



Commodity Derivatives Valuation with Autoregression and Moving Average in the Price Dynamics

Raphael Paschke
University of Mannheim
&
Marcel Prokopczuk
ICMA Centre, University of Reading

September 2009

ICMA Centre Discussion Papers in Finance DP2009-10

Copyright 2009 Paschke and Prokopczuk. All rights reserved.

ICMA Centre • University of Reading
Whiteknights • PO Box 242 • Reading RG6 6BA • UK
Tel: +44 (0)1183 787402 • Fax: +44 (0)1189 314741
Web: www.icmacentre.ac.uk
Director: Professor John Board, Chair in Finance
The ICMA Centre is supported by the International Capital Market Association



Commodity Derivatives Valuation with Autoregression and Moving Average in the Price Dynamics

Raphael Paschke* and Marcel Prokopczuk†

First version: May 2009

This version: September 2009

Abstract

In this paper we develop a continuous time factor model of commodity prices that allows for higher order autoregression and moving average components. The need for these components is documented by analyzing the convenience yield's time series dynamics. Making use of the affine model structure, closed-form pricing formulas for futures and options are derived. Empirically, a parsimonious version of the general model is estimated for the crude oil market using futures data. We demonstrate the model's superior performance in pricing nearby futures contracts in- and out-of-sample. Most notably, the model improves the pricing of long horizon contracts with information from the short end of the futures curve substantially.

JEL classification: G13, C50, Q40

Keywords: Commodity Pricing, *CARMA*, Futures, Crude Oil

*Department of Finance, University of Mannheim, D-68131 Mannheim, Germany. e-Mail: paschke@uni-mannheim.de.

†ICMA Centre, Henley Business School, University of Reading, Whiteknights, Reading, RG6 6BA, UK. e-Mail: m.prokopczuk@icmacentre.ac.uk.

I Introduction

Commodity prices and their stochastic behavior play a central role for many economic and financial decisions. Valuation and hedging of commodity-related securities and projects is an important problem, bringing forward the need for appropriate stochastic models. Brennan and Schwartz (1985) is one of the first works, proposing to employ financial modeling techniques for the commodity price, to evaluate natural resource-related projects.

More recent work extended this approach by recognizing that the inclusion of a second stochastic factor, a convenience yield, describing the benefits of holding the underlying commodity in stock¹ significantly improves the models' properties (see Gibson and Schwartz (1990), Schwartz (1997) (model 2), and Schwartz and Smith (2000)). These models have been extended for even more stochastic factors, e.g. Schwartz (1997) (model 3), Casassus and Collin-Dufresne (2005), and Geman and Nguyen (2005). It remains, however, controversial whether a third factor can improve the models' performance or merely yields overparameterization.

In this paper we take a different, more parsimonious approach than adding additional stochastic factors. All the studies mentioned above assume (explicitly or implicitly) that the convenience yield follows an Ornstein-Uhlenbeck type process, which is the continuous limit of a discrete $AR(1)$ process. Notwithstanding, when analyzing the convenience yield, we find that this assumption is not very satisfactory from an empirical point of view (see Section IV.B). Adding a moving average component yielding an $ARMA(1,1)$

¹See Kaldor (1939), Working (1949), or Brennan (1958) for a detailed discussion of the theory of storage and the arising convenience yield.

model, however, improves the statistical description of the convenience yield's dynamics significantly. Consequently, we propose to include this empirical feature in a continuous time commodity pricing model.

Our main contribution is thus twofold. First, theoretically, we develop a continuous time commodity pricing model which is able to incorporate higher order autoregression and moving average terms. This enables us to capture the stylized facts observed for the convenience yield without the need to add additional risk factors or leave the Gaussian world. The latter fact allows us to derive closed-form futures and options valuation formulas. Second, empirically, we implement a parsimonious specification of our model for the crude oil futures markets. A comparison with the benchmark model of Schwartz and Smith (2000), shows that the proposed model greatly improves the futures pricing at the short end of the futures curve in-sample and out-of-sample. Most notably, the model also improves the pricing of long horizon contracts with information from the short end of the futures curve substantially.

The model in this paper can be regarded as a generalization of the well-known model of Schwartz and Smith (2000). We follow their approach and do not consider an explicit convenience yield but formulate the model in a latent factor form which facilitates empirical implementation. Schwartz and Smith (2000) assume in their model that the second factor, describing short-term deviations from the long-term equilibrium price, follows an Ornstein-Uhlenbeck process. We generalize this approach by replacing it with a continuous autoregressive moving-average (*CARMA*) process. *CARMA* processes have been studied in the statistical literature for a long time (see Tsai and Chan (2000) or Brockwell (2001) and the references therein) but have received very little attention in

financial modeling. Benth et al. (2008) have recently proposed using *CARMA* processes for interest rate modeling and discuss the merits of this approach.

The properties of the *CARMA* process, as opposed to the simple Ornstein-Uhlenbeck process, are very desirable to model commodity futures prices.² First, adding higher order autoregression and, more importantly, moving average components to the model, allows much more flexibility with respect to the shape of the futures curve and, second, the term structure of volatilities. As a consequence, it is able to yield a much better pricing performance. This is especially true for the short end of the futures curve, usually the worst part of the curve with respect to pricing accuracy, due to the very high volatility of the nearby contracts.

The remainder of this paper is structured as follows. In Section II we first introduce the *CARMA* process in general and derive subsequently our commodity pricing model and discuss its properties. In Section III we describe the Kalman filter-based estimation approach of the model. Section IV presents our empirical study of crude oil futures. Concluding remarks are provided in Section V.

II *CARMA* Dynamics and Valuation

A. *CARMA*(p,q) Processes

When deciding to include higher order autoregressive as well as moving average terms in a model, many authors switch to discrete *ARMA*(p,q) models. However, a discrete time approach has the big disadvantage of losing analytical tractability, especially for

²Note, that the Ornstein-Uhlenbeck process is a special case of a *CARMA* process.

derivatives pricing. Therefore, we propose using continuous autoregressive moving-average ($CARMA(p, q)$) models for commodity price modeling (see Tsai and Chan (2000) or Brockwell (2001) for more detailed coverage of $CARMA$ processes). A $CARMA(p, q)$ process (with $0 \leq q < p$) is defined as the solution of the differential equation of order p :

$$\bar{Y}_t^{(p)} - \alpha_p \bar{Y}_t^{(p-1)} - \dots - \alpha_1 \bar{Y}_t - \alpha_0 = \sigma [\bar{W}_t^{(1)} + \beta_1 \bar{W}_t^{(2)} + \dots + \beta_q \bar{W}_t^{(q+1)}]. \quad (1)$$

The superscript denotes j -fold differentiation with respect to t . \bar{W}_t is a standard Brownian motion; $\alpha_i, i = 0, \dots, p; \beta_k, k = 1, \dots, p-1$ with $\beta_k = 0$ for $k > q$; and $\sigma > 0$ are constants. As a Brownian motion is nowhere differentiable, the derivatives $\bar{W}_t^{(j)}$ do not exist in the usual sense. As discussed by Tsai and Chan (2000), they can, however, be interpreted as observation and state equations:

$$\bar{Y}_t = \boldsymbol{\beta}' \bar{\mathbf{X}}_t, \quad (2)$$

$$d\bar{\mathbf{X}}_t = (\mathbf{A}\bar{\mathbf{X}}_t + \alpha_0 \boldsymbol{\omega})dt + \sigma \boldsymbol{\omega} dW_t, \quad (3)$$

with $d\bar{\mathbf{X}}_t = [X_t, X_t^{(1)}, \dots, X_t^{(p-1)}]'$; $\boldsymbol{\omega} = [0, 0, \dots, 0, 1]'$; $\boldsymbol{\beta} = [1, \beta_1, \dots, \beta_{p-1}]'$; and

$$\mathbf{A} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \\ \alpha_1 & \alpha_2 & \cdots & \cdots & \alpha_p \end{bmatrix}.$$

If $p = 1$, we set $\mathbf{A} = a_1$. The solution of (3) is given as (see Tsai and Chan (2000))

$$\bar{\mathbf{X}}_t = e^{\mathbf{A}t} \bar{\mathbf{X}}_0 + \alpha_0 \int_0^t e^{\mathbf{A}(t-u)} \boldsymbol{\omega} du + \sigma \int_0^t e^{\mathbf{A}(t-u)} \boldsymbol{\omega} d\bar{W}_u, \quad (4)$$

where $e^{\mathbf{M}}$ denotes the usual matrix exponential function.

B. Price Dynamics

The *CARMA* commodity pricing model presented in this article generalizes the well-known two-factor models of Schwartz and Smith (2000). A further extension to more factor models is straightforward. However, we consider the two-factor case, as we are convinced that parsimony is a very desirable model property.

The long-term/short-term model of Schwartz and Smith (2000) assumes that the log spot price of a commodity can be characterized by the sum of two stochastic factors, namely:

$$\ln S_t = Y_t + Z_t. \quad (5)$$

In this model, the factor Z_t denotes the long-term (non-stationary) equilibrium (log-)price level, following a standard arithmetic Brownian motion:

$$dZ_t = \mu dt + \sigma_2 dW_{2,t}, \quad (6)$$

where μ denotes the drift, $\sigma_2 > 0$ the volatility parameter, and $W_{2,t}$ a standard Wiener process.

The variable Y_t , represents short-term deviation from the equilibrium price level, and is

governed by an Ornstein-Uhlenbeck process as:

$$dY_t = -a_1 Y_t dt + \sigma_1 dW_{1,t}, \quad (7)$$

with parameters $a_1 > 0$, $\sigma_1 > 0$, and another Wiener process $W_{1,t}$ which can be correlated with $W_{2,t}$. As the process Y_t is mean-reverting towards zero, the log spot price will follow the process Z_t in the long term.

In this paper we focus on the mean-reverting factor Y_t and replace the simple Ornstein-Uhlenbeck process with a $CARMA(p, q)$ dynamics described in the previous section. The model we propose can be written similarly to (5) by replacing the dynamics of Y_t with a $CARMA(p, q)$ dynamics described by (3). Under the equivalent martingale measure, the resulting $ABM-CARMA(p, q)$ model can be compactly written as:

$$\ln S_t = \boldsymbol{\beta}' \mathbf{X}_t. \quad (8)$$

with $\boldsymbol{\beta} = [1, \beta_1, \dots, \beta_{p-1}, 1]'$, $\beta_k = 0$ for $k > q$, and

$$d\mathbf{X}_t = (\mathbf{A}\mathbf{X}_t + \boldsymbol{\mu})dt + \mathbf{V}d\mathbf{W}_t, \quad (9)$$

where

$$\mathbf{A} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ \alpha_1 & \alpha_2 & \cdots & \cdots & \alpha_p & 0 \\ 0 & 0 & 0 & \cdots & 0 & 0 \end{bmatrix},$$

$$\boldsymbol{\mu} = [0, 0, \dots, 0, \mu]', \quad d\mathbf{W}_t = [0, \dots, 0, dW_1, dW_2]',$$

$$\mathbf{V}\mathbf{V}' = \begin{bmatrix} 0 & \cdots & 0 & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & \cdots & 0 & 0 & 0 \\ 0 & \cdots & 0 & \sigma_1^2 & \sigma_1\sigma_2\rho \\ 0 & \cdots & 0 & \sigma_1\sigma_2\rho & \sigma_2^2 \end{bmatrix},$$

and $d\mathbf{X}_t = [X_t^{(0)}, X_t^{(1)}, \dots, X_t^{(p-1)}, Z_t]'$. It is worth mentioning that the model of Schwartz and Smith (2000) is a special case of the *ABM-CARMA*(p, q) model, given by $p = 1$ and $q = 0$. Furthermore, one should keep in mind that in contrast to the discrete *ARMA*(p, q) model, p must always be greater than q . Thus, for instance, a *ABM-CARMA*($1, 1$) model is not possible. The most parsimonious model variant including a moving average component is therefore the *ABM-CARMA*($2, 1$) specification.

The solution to the process is analogous to Equation (4):

$$\mathbf{X}_t = e^{\mathbf{A}t} \mathbf{X}_0 + \int_0^t e^{\mathbf{A}(t-u)} \boldsymbol{\mu} \, du + \sigma \int_0^t e^{\mathbf{A}(t-u)} \mathbf{V} \, dW_u. \quad (10)$$

Note that we formulate the model directly under the equivalent martingale measure, and, therefore, no risk premia are needed. This approach is taken, as previous studies such as Schwartz and Smith (2000) and Geman and Nguyen (2005) have shown that the drift under the physical measure and the risk premia can only be estimated with very low precision. Since our final goal is derivatives pricing, modeling and estimation directly under the equivalent martingale measure is favored, decreasing the error induced by unprecise estimates of these parameters.

C. Futures and Options Valuation

Standard theory within affine frameworks implies that futures prices are equal to the risk neutral expectation of the spot price at maturity³, i.e. conditional on information at time t , the log futures price $\ln F_t = \ln F(\mathbf{X}_t, t; T)$ is of the form:

$$\begin{aligned} \ln F_t &= \mathbb{E}_t[\ln S_T] + \frac{1}{2}\mathbb{V}_t[\ln S_T] \\ &= \boldsymbol{\beta}'\left(\mathcal{A}(t, T)\mathbf{X}_t + \mathcal{B}_\mu(t, T)\right) + \frac{1}{2}\boldsymbol{\beta}'\mathcal{B}_{\sigma^2}(t, T)\boldsymbol{\beta}, \end{aligned} \tag{11}$$

where $\mathbb{E}_t[\cdot]$ and $\mathbb{V}_t[\cdot]$ denote the conditional expected value and variance under the risk neutral measure respectively, and:

$$\begin{aligned} \mathcal{A}(t, T) &= e^{\mathbf{A}(T-t)}, \\ \mathcal{B}_\mu(t, T) &= \boldsymbol{\mu}(T-t), \\ \mathcal{B}_{\sigma^2}(t, T) &= \int_t^T e^{\mathbf{A}(T-u)}\mathbf{V}\mathbf{V}'e^{\mathbf{A}'(T-u)}du. \end{aligned}$$

³Strictly speaking, this is true for forward prices only. We are aware of the fact that futures and forwards may have different values in certain economic environments. For a clear-cut exposition of the differences in a similar framework, see, e.g., Miltersen and Schwartz (1998). In what follows we abstract from these differences and treat the two instruments as equal.

Note that the last term $\mathcal{B}_{\sigma^2}(t, T)$ can be evaluated in closed form.

Similarly, the price of a European option is the expected discounted pay-off at maturity of the option under the risk neutral measure. There are three points in time of interest, which are without loss of generality: today, which we normalize to zero, the maturity of the option t , and the maturity of the underlying futures contract T . The price of a call option $C_0 = C_0(\mathbf{X}_0, 0; t, T)$ with strike K is:

$$C_0 = e^{-rt} \mathbb{E}_0[\max\{F(\mathbf{X}_t, t; T) - K, 0\}]. \quad (12)$$

As the latent state variables \mathbf{X}_t are jointly normal and the log futures price in equation (11) is an affine-linear function of the states, the futures price at the option's maturity t is log-normally distributed with mean:

$$\begin{aligned} \mu_0(t, T) &= \mathbb{E}_0[\ln F(\mathbf{X}_t, t; T)] \\ &= \boldsymbol{\beta}' \mathcal{A}(t, T) \mathbb{E}_0[\mathbf{X}_t] + \boldsymbol{\beta} \mathcal{B}_\mu(t, T) \boldsymbol{\mu} + \frac{1}{2} \boldsymbol{\beta}' \mathcal{B}_{\sigma^2}(t, T) \boldsymbol{\beta} \\ &= \boldsymbol{\beta}' \mathcal{A}(0, T) \mathbf{X}_0 + \boldsymbol{\beta}' \left(\mathcal{A}(t, T) \mathcal{B}_\mu(0, t) + \mathcal{B}_\mu(t, T) \right) \boldsymbol{\mu} + \frac{1}{2} \boldsymbol{\beta}' \mathcal{B}_{\sigma^2}(t, T) \boldsymbol{\beta}, \end{aligned}$$

and deterministic variance:

$$\begin{aligned} \sigma^2(0, t, T) &= \boldsymbol{\beta}' \mathcal{A}(t, T) \mathbb{V}_0[\mathbf{X}_t] \mathcal{A}'(t, T) \boldsymbol{\beta} \\ &= \boldsymbol{\beta}' \mathcal{A}(t, T) \mathcal{B}_{\sigma^2}(0, t) \mathcal{A}'(t, T) \boldsymbol{\beta}. \end{aligned} \quad (13)$$

Therefore, the call option in Equation (12) as well as the corresponding put P_0 can be

evaluated by the use of the Black (1976) option pricing formula:

$$\begin{aligned} C_0 &= e^{-rt} \left(F_0 \Phi(d_1) - K \Phi(d_2) \right) \\ P_0 &= e^{-rt} \left(K \Phi(-d_2) - F_0 \Phi(-d_1) \right), \end{aligned} \tag{14}$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function and:

$$d_{1/2} = \frac{\ln(F_0/K) \pm \frac{1}{2}\sigma^2(0, t, T)}{\sigma(0, t, T)}.$$

D. Model Discussion

The futures price curve at some point in time t is:

$$\ln F(\mathbf{X}_t, t; T) = \boldsymbol{\beta}' \mathcal{A}(t, T) \mathbf{X}_t + \boldsymbol{\beta}' \mathcal{B}_\mu(t, T) \boldsymbol{\mu} + \frac{1}{2} \boldsymbol{\beta}' \mathcal{B}_\sigma(t, T) \boldsymbol{\beta},$$

directly showing that the (log) futures curve is affine-linear in the state variables.

Therefore, the futures curve of the $ABM-CARMA(p, q)$ model can be decomposed into $p+2$ parts, namely one for each state variable and the constant term. Moreover, it is also possible to disentangle the autoregressive and moving average components.

In the following we discuss the properties of the $ABM-CARMA(2, 1)$ model. This choice is motivated by the fact that it is the most simple specification incorporating both autoregressive and moving average components. The additive decompositions of the

futures curve reads as follows:⁴

$$\ln F(\mathbf{X}_t, t; T) = \underbrace{Z_t + A}_{ABM} + \underbrace{B\dot{X}_t + CX_t + D}_{CARMA}. \quad (15)$$

The coefficients depend on (t, T) only through the time to maturity $\tau = (T - t)$.

Furthermore, we can decompose these coefficients as:

$$\begin{aligned} B &= B^{AR} + B^{MA} \\ C &= C^{AR} + C^{MA} \\ D &= D^{AR} + D^{MA}, \end{aligned}$$

where the precise formulas are provided in the Appendix.

To be able to better interpret the different components, we illustrate the resulting curves, using the parameters estimated in the subsequent sections (see Table 2), in Figure 1.

The left part of Figure 1 shows the effects within the $ABM-CARMA(2,1)$ model, resulting from only the AR components; the right part displays the moving average terms. The upper Panels 1a and 1b depict the components of the futures curves. Panel 2a and 2b represent the coefficients (without the state variables) with respect to the underlying process short-term deviation process \dot{X}_t , and the lower Panels, 3a and 3b, give the coefficients (without the state variables) with respect to the integrated process X_t .

Panel 2a graphs the coefficient B^{AR} of the stochastic process \dot{X}_t . It starts at zero with slope one and approaches zero for long horizons. In between it attains a maximum at

⁴We commit a slight misuse of notation here, as A already denotes the coefficient matrix in the $CARMA$ process. No confusion should arise, as the meaning will be clear from the context.

$\tau = -\frac{\ln\left(\frac{\lambda_2}{\lambda_3}\right)}{\lambda_2 - \lambda_3}$, where λ_i denote the eigenvalues⁵ of the matrix \mathbf{A} . Consequently, this component is influential in shaping the medium-term behavior of the term structure. In the model of Schwartz and Smith (2000), this coefficient is always a monotonic decreasing function, limiting the model's flexibility.

The locally non-stochastic state variable X_t enters into the futures curve in a different manner. The coefficient C^{AR} originates at one and vanishes for long horizons. As X_t is 'slowly' moving relative to the underlying stochastic factor \dot{X}_t , its impact on the futures price regarding an infinitesimal longer maturity is perfectly predictable. Technically, this property results from the fact that the slope at the front end of the curve of C^{AR} is zero. The constant term D^{AR} is not plotted separately. It is zero for $\tau = 0$ and approaches a constant for $\tau \rightarrow \infty$.

Besides the autoregressive components, the futures curve consists of the non-stationary long-term (equilibrium) process Z_t . The futures curve resulting from this component alone ($Z_t + A$) is shown in Panel 1a as a dotted line. It starts at Z_t and has a constant slope of $\mu + \frac{1}{2}\sigma_2^2$. The dotted-dashed line adds the short-term deviation, $(Z_t + A + B^{AR}\dot{X}_t)$, stemming from the state variable \dot{X}_t , whereas the dashed line comprises the effect of all three state variables, $(Z_t + A + B^{AR}\dot{X}_t + C^{AR}X_t)$. Finally, the solid line in Panel 1a shows the entire future curve from the non-stationary and autoregressive parts of the model $(Z_t + A + B^{AR}\dot{X}_t + C^{AR}X_t + D^{AR})$. Note that the resulting futures curve represents only one point in time. As the state variables (Z_t, \dot{X}_t, X_t) evolve over time, the futures curve will change its shape.

⁵In general, the eigenvalues are possibly complex. With our parametrization, they are real. Most properties also hold in the general case.

Panel 2b depicts the moving average coefficient B^{MA} . It starts at β_1 with a negative slope, is of a humped shape with a minimum at $\tau = \frac{-2\ln(\frac{\lambda_2}{\lambda_3})}{\lambda_2 - \lambda_3}$, and converges to zero. It is evident from comparing the graphs in Panel 2a and 2b that only the moving average part of the model is able to explain fast movements of the short end of the futures curve. As the autoregressive coefficient B^{AR} of the 'fast' moving factor \dot{X}_t goes to zero for short maturities, the AR component is not able to follow fast changes of the short end.

In Panel 3b the dependence of the futures curve on the 'slowly' moving factor X_t is illustrated. The influence on short maturities diminishes and the slope is negative. As it converges to zero in the limit, it has a (negative) hump. The constant terms D^{MA} behave similarly to the constant autoregressive part D^{AR} .

Taking everything together results in the entire futures curve, which is shown in Panel 1b. As in Panel 1a, the dotted line represents the non-stationary model part ($Z_t + A$). The dotted dashed line represents the effects arising due to the 'fast' moving state variable \dot{X}_t ($Z_t + A + B\dot{X}_t$), whereas the dot-dashed line also includes the impact of the integrated state variable X_t ($Z_t + A + B\dot{X}_t + CX_t$). Finally, the solid line gives the complete log futures curve.

We wish to briefly discuss the role of the correlation coefficient ρ between the two stochastic factors (Z_t and \dot{X}_t). Correlation has no impact on the coefficients B and C . However, the correlation changes the slope in the medium/long term by adding an almost constant term in D . For very long futures horizons, all model specifications converge to the futures curve of a pure non-stationary one-factor model for log spot prices, namely into the ABM part $Z_t + A$, since A is linear in τ .

Finally, the model's instantaneous volatility curve of futures

$$\sigma^2(t, T) = \boldsymbol{\beta}' \mathcal{A}(t, T) \mathbf{V} \mathbf{V}' \mathcal{A}(t, T)' \boldsymbol{\beta}$$

is examined. Figure 2 shows the volatility curves of the *ABM-CARMA(2,1)* and also of the Schwartz and Smith (2000) model. As in all Gaussian models, the volatility curve is independent of the state variables and therefore constant over time. Moreover, in the *ABM-CARMA(2,1)* case it can be decomposed into

$$\sigma^2(t, T) = \sigma_2^2 + 2\sigma_2\rho\sigma_1 B + \sigma_1^2 B^2.$$

As $B \rightarrow 0$ for $\tau \rightarrow \infty$, the volatility of a very distant future is σ_2 , which is the volatility of the equilibrium process Z_t . If there is no moving average component, the volatility at the short end will also be σ_2 , as $B^{AR} \rightarrow 0$ for $\tau \rightarrow 0$. Thus, to match a behavior of the volatility curve which is in line with the Samuelson effect (i.e. a decreasing volatility curve), the moving average part is necessary, since $B^{MA}(t, t) = \beta_1$. Inspecting the functions B^{AR} and B^{MA} which are displayed in Panel 2a and 2b of Figure 1 one observes that the moving average part shapes the short-term volatility, whereas the *AR* part the medium-term structure.

III Estimating the State Variables and Process Parameters

As the latent state variables are unobservable we cannot estimate the processes' parameters directly using spot price data, as there are more factors than spot price observations per date. Furthermore, it was pointed out by Schwartz (1997) that spot price information of a commodity is often so uncertain that it is preferred to use the futures contract closest to maturity as a proxy for the spot price.

Therefore, we formulate our model in state space form and estimate it by employing the Kalman filter methodology. For a rigorous treatment of Kalman filtering see, e.g., Harvey (1989) and the references therein.⁶ The Kalman filter is a recursive procedure for computing the optimal estimator of some unobserved state variables based on observations of related quantities (in our case the futures prices). The observed quantities, i.e. the futures prices, are assumed to be measured with some noise, taking into account bid-ask spreads, price limits, nonsimultaneity of data, and errors in the data etc. (see Schwartz (1997)).

The *measurement equation* of the state space representation is obtained by adding serially and cross-sectionally uncorrelated zero mean noise to the futures valuation formula. The unobservable state variables follow the *transition equation*, which can be deduced from the assumed factor dynamics. When the factor dynamics are driven by Gaussian noise, and the observations are measured with Gaussian errors, the Kalman filter allows the estimation of the processes' parameters via maximum likelihood methods. The

⁶The Kalman filter approach has been applied to models of commodity derivatives by Schwartz (1997), Schwartz and Smith (2000), Geman and Nguyen (2005), and also Cortazar and Naranjo (2006).

log-likelihood function can be written as the sum of conditional log-likelihood terms, which remain Gaussian.

It is worth noting that using the Kalman filter approach we are able to explore time series as well as cross-sectional properties of the data at the same time. Furthermore, we can make use of all observations available and do not have to decide which two contracts to use when trying to estimate the parameters via an inversion of the measurement equation. It is also worth noting that in linear and Gaussian models the Kalman filter is the optimal filter.

From Equation (9) and the solution in (10), the exact *transition equation* is given by:

$$\mathbf{x}_{t+\Delta t} = \mathbf{G}\mathbf{x}_t + \mathbf{c} + \boldsymbol{\eta}_{\Delta t} , \quad (16)$$

for time step Δt and $\boldsymbol{\eta}_{\Delta t}$ serially uncorrelated, normally distributed disturbances with zero mean and constant variance and

$$\begin{aligned} \mathbf{c} &= \boldsymbol{\mu}\Delta t , \\ \mathbf{G} &= e^{\mathbf{A}\Delta t} , \\ \mathbf{E}[\boldsymbol{\eta}_{\Delta t}] &= \mathbf{0} , \\ \mathbf{V}[\boldsymbol{\eta}_{\Delta t}] &= \int_{\Delta t} e^{\mathbf{A}(\Delta t-u)} \mathbf{V}\mathbf{V}' e^{\mathbf{A}'(\Delta t-u)} du , \end{aligned} \quad (17)$$

where Δ_t demotes the length of the time steps as a fraction of one year. Note that the last integral can be evaluated analytically.

Writing the observed log futures prices $\ln F(t, T_i)$ at time t for maturities T_i , $i = 1, \dots, k$, as $\mathbf{y}_t = [\ln F(t, T_1), \dots, F(t, T_k)]'$, the *measurement equation* at time t is given by adding

measurement errors $\boldsymbol{\varepsilon}_t$ to the futures valuation formula (11). Hence

$$\mathbf{y}_t = \mathbf{d} + \mathbf{H}\mathbf{x}_t + \boldsymbol{\varepsilon}_t, \quad (18)$$

with

$$\begin{aligned} \mathbf{d} &= [\boldsymbol{\beta}'\mathcal{B}_\mu(t, T_1) + \boldsymbol{\beta}'\mathcal{B}_{\sigma^2}(t, T_1)\boldsymbol{\beta}, \dots, \boldsymbol{\beta}'\mathcal{B}_\mu(t, T_k) + \boldsymbol{\beta}'\mathcal{B}_{\sigma^2}(t, T_k)\boldsymbol{\beta}]', \\ \mathbf{H} &= [\boldsymbol{\beta}'\mathcal{A}(t, T_1), \dots, \boldsymbol{\beta}'\mathcal{A}(t, T_k)]', \\ \mathbf{E}[\boldsymbol{\varepsilon}_t] &= \mathbf{0}, \\ \mathbf{V}[\boldsymbol{\varepsilon}_t] &= \boldsymbol{\Xi}, \end{aligned} \quad (19)$$

and $\boldsymbol{\Xi}$ being a diagonal matrix of k serially and cross-sectionally uncorrelated error terms, i.e. $\boldsymbol{\Xi} = \text{diag}([\xi_1^2, \dots, \xi_k^2]')$. All parameters are collected in the set Ψ .

In general, one could allow the error terms to be cross-sectionally correlated. However, this would greatly complicate the estimation procedure. Therefore, we follow Schwartz (1997) and Schwartz and Smith (2000) and assume a diagonal covariance matrix for the measurement errors.

To start the Kalman filter-based estimation, one has to supply starting values Ψ_0 , as well as initial values for the (unobserved) state vector \mathbf{x}_0 . We follow Harvey (1989) and include the state vector at $t = 0$ in the set of parameters to be estimated.

IV Estimation Results and Model Comparison

In this section we describe our empirical study. Motivated by our preliminary analysis of the convenience yield in part *B* of this section, we implement the most parsimonious

ABM-CARMA model including a moving average component, i.e. an *ABM-CARMA(2,1)* model. We then analyze the model's futures pricing ability in and out-of-sample by comparing it to the benchmark model of Schwartz and Smith (2000).

We would like to mention that we also implemented and estimated a *ABM-CARMA(2,0)* model, although our theoretical analysis already showed some severe disadvantages of not including a moving average component. This was affirmed by the empirical results. The *ABM-CARMA(2,0)* performed relatively poorly, thus we do not report the results to keep the focus on the most interesting. Note that an *ABM-CARMA(1,1)* is not feasible, as the *CARMA* model requires $p > q$.

A. Data

Our data set consists of prices of crude oil futures contracts traded at the New York Mercantile Exchange (NYMEX), which is one of the most heavily traded commodity contracts worldwide. The short position in this contract commits the holder to deliver 1,000 barrels of domestic crude oil in Cushing, Oklahoma.⁷ We consider weekly observations, sampling Wednesday settlement prices between 01/01/1996 and 12/10/2008, yielding 676 observation dates. Crude oil futures are listed 9 years forward with monthly maturity for the first 6 years, and semiannual maturity thereafter. As the liquidity is rather low for longer-term contracts, we consider only the first 24 (i.e. the first two years) in our analysis. Thus, we yield a total number of 16,224 futures prices. We conduct our study using settlement values of futures prices, as they are classically considered to be

⁷The following domestic oil grades are deliverable: West Texas Intermediate, Low Sweet Mix, New Mexican Sweet, North Texas Sweet, Oklahoma Sweet, South Texas Sweet. Specific foreign crudes may also be deliverable, however, at a discount. For details on the specification of deliverable crudes and delivery locations see www.nymex.com.

representative for a trading day (see Geman and Nguyen (2005)). As maturity, we take the last day of trading.⁸ All data is obtained via Bloomberg.

Table 1 contains summary statistics for the futures price data, where F01 is the contract closest to maturity, F02 the second contract closest to maturity and so on. In line with prior research, the average futures curves is in backwardation, although on a much higher level, which is mainly due to the peak of the most recent observations (see Figure 3 for a time series plot of the closest to maturity future F01).

B. Preliminary Analysis of the Convenience Yield

In this subsection we conduct a preliminary analysis of the convenience yield in the crude oil market. The analysis is complicated by the fact that the convenience yield is not observable. Thus, we have to rely on some approximation. Using the well-known relationship between spot and futures price when storage costs s_t , interest rates r_t , and net convenience yields c_t are non-stochastic

$$F(t, T) = S_t e^{(s_t + r_t - c_t)(T-t)} = S_t e^{\delta_t(T-t)}, \quad (20)$$

enables us to estimate monthly forward total convenience yields δ_t . This total convenience yield already includes the costs of storage and capital. We can estimate this quantity as:

$$\delta_{t, T-1, T} = \ln \left(\frac{F(t, T)}{F(t, T-1)} \right). \quad (21)$$

⁸Trading ends at the close of business on the third business day prior to the 25th calendar day of the month preceding the delivery month. If the 25th calendar day of the month is a non-business day, trading shall cease on the third business day prior to the business day preceding the 25th calendar day.

As we do not have spot price data corresponding to the futures data, we use the two futures contracts closest to maturity. This procedure has been proposed by Gibson and Schwartz (1990). It is well known that the convenience yield exhibits a mean-reverting behavior which motivated previous research, for example Gibson and Schwartz (1990) and Schwartz and Smith (2000), to model the second stochastic factor in their models as Ornstein-Uhlenbeck processes, which is the continuous time equivalent of a discrete $AR(1)$ process. Having a time series of convenience yields at hand, it is easy to test whether this kind of model provides as satisfactory fit to the data.

We fit an $AR(1)$ model to the convenience yield time series δ_t and analyze the residuals. The Ljung-Box statistic testing the null hypothesis of independence yields 27.63 corresponding to a p-value smaller than 0.001. It seems, therefore, desirable to include higher order autoregressive and/or moving average components to describe the time series behavior of the convenience yield. In the spirit of parsimonious modeling, the next obvious step is to include a moving average term and to estimate an $ARMA(1,1)$ model.⁹ Repeating the Ljung-Box test with the resulting residuals yields a test statistic of 0.098 and a corresponding p-value of 0.75, providing clear evidence for no dependence in the remaining residuals. The inclusion of the moving average term clearly improves the discrete modeling of the convenience yield and serves as motivation to also include such a component in a continuous time pricing model.

⁹The alternative was to include higher order autogression terms in the model. We fitted $AR(p)$ models up to $p = 8$, all providing a fit inferior to the $ARMA(1,1)$ model. Besides providing a worse representation of the data, this approach also conflicts for higher values of p with the principle of parsimony.

C. *Parameter Estimates*

We estimate the most parsimonious *ABM-CARMA* model incorporating a moving average component, namely *ABM-CARMA*(2,1) and, as a benchmark, the model of Schwartz and Smith (2000). Thus, we have to estimate seven parameters in the model of Schwartz and Smith (2000), $\Psi = (\kappa, \mu, \sigma_1, \sigma_2, \rho, X_0, Z_0)$, and ten parameters for the *ABM-CARMA*(2,1) model $\Psi = (a_1, a_2, \mu, \sigma_1, \sigma_2, \rho, \beta, X_0, \dot{X}_0, Z_0)$, plus the k terms in the covariance matrix for the measurement errors Ξ in both cases. Table 2 reports the maximum likelihood parameter estimates of the two considered models.¹⁰ We do not report the estimated variance parameters of the measurement errors ξ_i , $i = 1, \dots, 24$. The average value of these is $0.15 \cdot 10^{-3}$ for the *ABM-CARMA*(2,1) model and $0.2 \cdot 10^{-3}$ for the Schwartz and Smith (2000) model.

The first thing to note is that all parameter estimates of the *ABM-CARMA*(2,1) model are highly significant. This suggests that the model is able to improve the ability to describe the underlying price dynamics. One should keep in mind that estimation was conducted directly under the equivalent martingale measure, and thus no risk premia were needed.

In both models, the short-term volatility is substantially higher than the long-term volatility. The correlation between the two factors is, for both models, negative. The coefficient β , which weights the differentiated process \dot{X}_t , is estimated as 0.96, indicating the gain of adding a moving average term to the model (β equals zero by definition in the *CARMA*(2,0) model).

¹⁰The model of Schwartz and Smith (2000) was also estimated by means of Kalman filtering and maximum likelihood.

D. In-Sample Model Comparison

Comparing the log-likelihoods (reported in Table 2) of the *ABM-CARMA(2,1)* and the Schwartz and Smith (2000) model, 78,676.52 against 76,498.50, indicates that the former provides a significantly better fit to the data. As the two models are not nested in the usual sense, however, a standard likelihood-ratio test cannot be applied. Therefore, we conduct a series of other tests to compare the models' pricing abilities.

When comparing the two models, one should take into account that both contain different numbers of parameters. Therefore, we compare the models with respect to two different goodness of fit criteria, namely the Akaike information criterion (AIC) and the Schwarz information criterion, which take the number of model parameters explicitly into account:

$$AIC = 2K - 2\ln(L(\Psi_i)), \quad (22)$$

$$SIC = K \ln T - 2\ln(L(\Psi_i)), \quad (23)$$

where $\ln(L(\Psi_i))$ denotes the log-likelihood values, T the number of observations, and K the number of parameters of the respective model, i.e. $K = 34$ for the *ABM-CARMA(2,1)* model, and $K = 31$ for the Schwartz and Smith (2000) model. The Schwarz information criterion penalizes a greater number of parameters more heavily than the Akaike information criterion.

Plugging the respective values into (22) and (23) yields

$$AIC_{ABM-CARMA} = -157,285, \quad AIC_{SS2000} = -152,935,$$

$$SIC_{ABM-CARMA} = -157,131, \quad SIC_{SS2000} = -152,795.$$

Both criteria attain the best (i.e. lowest) values for the $ABM-CARMA(2,1)$ model, providing more evidence for the better in-sample performance of this model.

In the real world, pricing accuracy is one of the most important features of a pricing model. We therefore compare the models with respect to their pricing precision. Table 3 reports the pricing errors of the (log-) futures prices.¹¹ We report root mean squared errors (RMSE) and mean absolute errors (MAE) for each maturity on an absolute (Panel A) and relative (Panel B) basis. Furthermore, we report the overall pricing errors in the last row of Table 3.

Considering the overall fit first, one can observe that the RMSE (MAE) decreases from 0.0141 (0.0065) to 0.0122 (0.0057) when switching from the Schwartz and Smith (2000) to the $ABM-CARMA(2,1)$ model. This corresponds to a reduction of the absolute RMSE (MAE) of 13.3% (12.4%).

For the individual contracts, the highest pricing error is observed for the closest future for both the $ABM-CARMA(2,1)$ model and the model of Schwartz and Smith (2000). The former, however, improves the in-sample fit clearly, reducing the absolute RMSE (MAE) from 0.0486 (0.0370) to 0.0409 (0.0318).

For some mid-term futures, some very small pricing errors are observed for the Schwartz and Smith (2000) case. The nine months and seventeen months futures are even priced almost perfectly. This is a well-known feature of the model, already observed by the authors. It stems from the fact that by having two state variables, it is possible to perfectly match two futures prices. We can observe a similar, but less pronounced pattern for the $ABM-CARMA(2,1)$ model.

¹¹We analyze log prices due to their use in the Kalman filter estimation.

The pricing at the long end of the futures curve also improves, although on a smaller absolute level, as pricing errors are small for both models anyway. Overall, it can be concluded that the inclusion of a moving average term does indeed substantially improve the model's ability to explain the commodity futures prices, especially the nearby futures contracts.

E. Out-of-Sample Model Comparison

To perform an out-of-sample comparison between the two models, two types of tests are performed, one based on the cross section of futures prices, the other one based on the time series of futures prices.

In the cross section-based test, we split the data set into two samples. The first one, containing the first twelve nearby futures (F01 - F12), is used for estimation. The second, containing the longer-term futures (F13 - F24), is used to evaluate the models' performance. In some sense, this is not a true out-of-sample test, as information of the entire observation period is used to value the futures at each date. A similar procedure, namely estimating parameters with every second contract (F01, F03, etc.) and then evaluating the pricing performance of the other ones (F02, F04, etc.), has been employed by Schwartz (1997). We consider our procedure as much more useful to evaluate the models' performance as, in reality, it is a frequently occurring case that short-term futures contracts are traded liquidly in the market, whereas longer-term contracts are only thinly traded and prices have to be determined by some kind of model. It is thus a critical question whether a model is useful in pricing longer-term contracts employing information from the short end of the futures curve.

Table 4 shows the RMSE and MAE for the out-of-sample contracts F13 to F24. The RMSE (MAE) for all twelve contracts amounts to 0.0179 (0.0126) for the model of Schwartz and Smith (2000), and to 0.0144 (0.0100) for the *ABM-CARMA(2,1)* model, improving the pricing accuracy by 19.6 % (20 %). Inspecting the pricing performance of the individual contracts, it can be seen that the improved model outperforms the standard model in every instance, i.e. for every maturity month.

As a second test, we perform a truly out-of-sample procedure. We split the entire data set of 676 weekly observations into two equally sized subsamples. The first one, containing the time series of all 24 futures prices of the first 338 weeks, is used for parameter estimation. The second subsample, containing the time series of all 24 futures prices of the weeks 339 to 676 is then used to evaluate the models' out-of-sample pricing performance.

Table 5 contains the results of this test. Considering the overall pricing performance, a rather small improvement of the *ABM-CARMA(2,1)* model is observed. The RMSE (MAE) considering all contracts decreases from 0.0381 (0.0290) to 0.0375 (0.0288) which equals a relative decrease of 1.74 % (0.66 %). However, inspecting the differences in pricing accuracy more thoroughly, it can be seen that a big difference exists for the closest to maturity futures. The RMSE (MAE) decreases from 0.0627 (0.0509) to 0.0564 (0.0451), a reduction of more than 10 %. The second and third closest futures show an improvement of about 6 % and 3 %, respectively.

Overall, the *ABM-CARMA(2,1)* model improves the pricing of long term contracts using information from the short end of the futures curve by 20 %. The out-of-sample pricing of the short term maturity futures improves by more than 10 %. We consider both of these improvements to be economically substantial.

V Conclusion

Assumptions regarding the underlying stochastic factors play a central role in financial modeling of commodity prices and commodity derivatives. In this paper we have argued that a simple $AR(1)$ representation for the short-term factor (i.e. the convenience yield) is not sufficient to model the futures curve. We therefore develop a new model, which we label $ABM-CARMA(p,q)$ model, relying on continuous time autoregressive moving average ($CARMA$) models. Closed-form futures and options valuation formulas were derived.

We then implement the $ABM-CARMA(2,1)$ model for the crude oil futures market and find that the inclusion of the moving average component considerably improves the quality of short-term futures pricing in- and out-of-sample. Moreover, the model improves the ability to price long maturity contracts using information from the short end of the futures curve.

Finally, we wish to conclude the paper by outlining some future research. Theoretically, the $ABM-CARMA(p,q)$ model allows for much greater flexibility of the term structure of volatilities. The volatility curve is especially important for options pricing. Having already derived European options pricing formulas, it is the obvious next step to evaluate the model's options pricing performance empirically. This aim is, however, complicated by the fact that most commodity options are of the American type.

Another direction of future work is to broaden the empirical basis by considering different commodity markets, e.g. the markets for agricultural commodities. As shown by Fama and French (1987), commodity markets with different seasonalities of supply and

demand may lead to different behavior of the convenience yield, and thus different model performance.

A Appendix

For $CARMA(2,1)$ we have

$$A = \begin{pmatrix} 0 & 1 & 0 \\ \alpha_1 & \alpha_2 & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

which can be decomposed into

$$\Lambda D(\lambda) \Lambda^{-1},$$

where

$$\begin{aligned} (\lambda_1, \lambda_2, \lambda_3)' &= \left(0, \frac{1}{2} \left(\alpha_2 - \sqrt{\alpha_2^2 + 4\alpha_1} \right), \frac{1}{2} \left(\alpha_2 + \sqrt{\alpha_2^2 + 4\alpha_1} \right) \right)' \\ &= (0, x + iy, x - iy)^\top, \end{aligned}$$

are the eigenvalues, Λ is the matrix of corresponding eigenvectors, and $i^2 = -1$.

In calculating the matrices $\mathcal{A}(t, T)$, $\mathcal{B}_\mu(t, T)$ and $\mathcal{B}_{\sigma^2}(t, T)$ the following expressions are substituted

$$\begin{aligned} e^{\lambda_i(T-t)} &= \psi(\lambda_i, t, T), \\ -\frac{1 - e^{\lambda_i(T-t)}}{\lambda_i} &= \phi(\lambda_i, t, T). \end{aligned}$$

With these substitutions, the decomposition in Section 2.D reads

$$\begin{aligned} A &= \mu(T-t) + \frac{1}{2}\sigma_2^2(T-t) \\ B^{AR} &= \frac{\psi(\lambda_3, t, T) - \psi(\lambda_2, t, T)}{\lambda_3 - \lambda_2} & B^{MA} &= \frac{\beta_1 \lambda_3 \psi(\lambda_3, t, T) - \beta_1 \lambda_2 \psi(\lambda_2, t, T)}{\lambda_3 - \lambda_2} \\ C^{AR} &= \frac{\lambda_2 \psi(\lambda_3, t, T) - \lambda_3 \psi(\lambda_2, t, T)}{\lambda_2 - \lambda_3} & C^{MA} &= \frac{\beta_1 \lambda_2 \lambda_3 (\psi(\lambda_3, t, T) - \psi(\lambda_2, t, T))}{\lambda_2 - \lambda_3} \\ D^{AR} &= \frac{1}{2} (h_2^{AR} \sigma_1^2 + 2h_1^{AR} \sigma_1 \rho \sigma_2) & D^{MA} &= \frac{1}{2} (h_2^{MA} \sigma_1^2 + 2h_1^{MA} \sigma_1 \rho \sigma_2) \end{aligned}$$

where

$$\begin{aligned} h_1^{AR} &= \frac{\phi(\lambda_2, t, T) - \phi(\lambda_3, t, T)}{\lambda_2 - \lambda_3} & h_1^{MA} &= \frac{\beta_1 (\lambda_2 \phi(\lambda_2, t, T) - \lambda_3 \phi(\lambda_3, t, T))}{\lambda_2 - \lambda_3} \\ h_2 &= \frac{\phi(2\lambda_2, t, T) (\beta_1 \lambda_2 + 1)^2}{(\lambda_3 - \lambda_2)^2} - \frac{2(\beta_1 \lambda_3 + 1) \phi(\lambda_2 + \lambda_3, t, T) (\beta_1 \lambda_2 + 1)}{(\lambda_3 - \lambda_2)^2} + \frac{(\beta_1 \lambda_3 + 1)^2 \phi(2\lambda_3, t, T)}{(\lambda_3 - \lambda_2)^2} \\ h_2^{AR} &= h_2|_{\beta_1=0} & h_2^{MA} &= h_2 - (h_2|_{\beta_1=0}) \end{aligned}$$

Using the boundaries

$$\begin{aligned}\psi(\lambda, t, T) &\xrightarrow{T \rightarrow \infty} 0 & \psi(\lambda, t, T) &\xrightarrow{T \rightarrow t} 1 \\ \phi(\lambda, t, T) &\xrightarrow{T \rightarrow \infty} -\frac{1}{\lambda} & \psi(\lambda, t, T) &\xrightarrow{T \rightarrow t} 0\end{aligned}$$

and the derivatives

$$\begin{aligned}\frac{d\phi(\lambda, t, T)}{dT} &= \psi(\lambda, t, T) \\ \frac{d\psi(\lambda, t, T)}{dT} &= \lambda\psi(\lambda, t, T)\end{aligned}$$

the results in the model discussion in Section II.D can be easily verified.

References

- F.E. Benth, S. Koekebakker, and V. Zakamouline. The CARMA interest rate model. Working Paper, 2008.
- F. Black. The pricing of commodity contracts. *Journal of Financial Economics*, 3: 167–179, 1976.
- M.J. Brennan. The supply of storage. *American Economic Review*, 48:50–72, 1958.
- M.J. Brennan and E.S. Schwartz. Evaluating natural resource investments. *Journal of Business*, 58(2):135–157, 1985.
- P.J. Brockwell. Continuous-time ARMA processes. In *Handbook of Statistics*, volume 19, pages 249–276. Elsevier Science Publishers, 2001.
- J. Casassus and P. Collin-Dufresne. Stochastic convenience yield implied from commodity futures and interest rates. *Journal of Finance*, 60(5):2283–2331, 2005.
- G. Cortazar and L. Naranjo. An n-factor gaussian model of oil futures prices. *Journal of Futures Markets*, 26:243–268, 2006.
- E.F. Fama and K.R. French. Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *Journal of Business*, 60(1):55–73, 1987.
- H. Geman and V.-N. Nguyen. Soybean inventory and forward curve dynamics. *Management Science*, 51(7):1076–1091, 2005.
- R. Gibson and E.S. Schwartz. Stochastic convenience yield and the pricing of oil contingent claims. *Journal of Finance*, 45(3):959–976, 1990.
- A.C. Harvey. *Forecasting, structural time series models and the Kalman filter*. Cambridge University Press, 1989.
- N. Kaldor. Speculation and economic stability. *Review of Economic Studies*, 7:1–27, 1939.

- K.R. Miltersen and E.S. Schwartz. 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, 1998.
- E.S. Schwartz. The stochastic behavior of commodity prices: Implications for valuation and hedging. *Journal of Finance*, 52(3):923–973, 1997.
- E.S. Schwartz and J.E. Smith. Short-term variations and long-term dynamics in commodity prices. *Management Science*, 46(7):893–911, 2000.
- H. Tsai and K.S. Chan. A note on the covariance structure of a continuous-time ARMA process. *Statistica Sinica*, 10:989–998, 2000.
- H. Working. The theory of price of storage. *American Economic Review*, 39:1254–1262, 1949.




Figure 1: **Futures Curve Decomposition**

This figure shows the different components of the ABM-CARMA(2,1) log futures curve using the parameters estimated in Section IV. The entire resulting futures curve is shown as a solid line in Panel 1b. The maturity $\tau = T - t$ is given in years.




Figure 2: **Volatility Term Structure**

This figure shows the instantaneous volatility of futures prices for different maturities, ranging from zero to five years. The solid line shows the term structure of volatilities for the ABM-CARMA(2,1) model, the dashed line for the model of Schwartz and Smith (2000).




Figure 3: **Futures Prices**

This figure shows weekly future prices of the closest to maturity future F01 from 01/01/1996 to 12/10/2008. Prices are in US dollars per barrel.

Table 1: **Statistics of Crude Oil Futures Contracts**

This table reports statistics for weekly observations of crude oil futures contracts from January 3, 1996 to December 10, 2008. Prices are in dollars per barrel. F01 denotes the one month futures contract, F02 the two months contract and so on.

	Mean Price	SE	Maturity	SE
F01	40.43	26.48	0.0450	0.0243
F02	40.42	26.65	0.1284	0.0243
F03	40.35	26.79	0.2119	0.0244
F04	40.24	26.91	0.2951	0.0244
F05	40.11	27.01	0.3785	0.0243
F06	39.98	27.09	0.4620	0.0244
F07	39.84	27.16	0.5453	0.0244
F08	39.71	27.22	0.6288	0.0244
F09	39.58	27.26	0.7121	0.0245
F10	39.45	27.29	0.7954	0.0244
F11	39.33	27.31	0.8789	0.0243
F12	39.22	27.33	0.9623	0.0244
F13	39.10	27.34	1.0456	0.0244
F14	38.99	27.34	1.1291	0.0244
F15	38.89	27.35	1.2124	0.0244
F16	38.79	27.35	1.2958	0.0243
F17	38.70	27.34	1.3793	0.0244
F18	38.61	27.33	1.4626	0.0245
F19	38.53	27.32	1.5460	0.0244
F20	38.45	27.31	1.6295	0.0244
F21	38.38	27.30	1.7128	0.0245
F22	38.31	27.28	1.7961	0.0244
F23	38.25	27.26	1.8796	0.0244
F24	38.19	27.25	1.9628	0.0243

Table 2: **Kalman Filter Parameter Estimates**

This table reports the estimated parameters and their standard errors estimated from weekly data using the Kalman filter maximum likelihood methodology.

Parameter	<i>ABM-CARMA(2,1)</i>		Schwartz/Smith (2000)	
	Estimate	Standard error	Estimate	Standard error
μ	0.0260	0.0045	-0.0179	0.0013
a_1	-0.5339	0.0207	0.8593	0.0058
a_2	-1.9636	0.0303	-	-
σ_1	0.7146	0.0343	0.3129	0.0100
σ_2	0.2359	0.0102	0.1918	0.0052
ρ	-0.5910	0.0448	-0.0894	0.0427
β	0.9552	0.0201	-	-
X_0	-0.5315	0.0452	0.1089	0.0439
\dot{X}_0	0.8902	0.1076	-	-
Z_0	2.6910	0.0442	2.7999	0.0269
Log-likelihood	78676.52		76498.50	

Table 3: In-Sample Pricing Errors

This table reports in-sample pricing for the $ABM-CARMA(2,1)$ and the model of Schwartz and Smith (2000). Panel A displays absolute Root Mean Squared Errors (RMSE), and Mean Absolute Errors (MAE). Panel B displays the corresponding statistics on a relative basis.

	RMSE			MAE			RMSE			MAE		
	<i>Panel A: Absolute</i>			<i>Panel A: Absolute</i>			<i>Panel B: Relative</i>			<i>Panel B: Relative</i>		
	$ABM-CARMA(2,1)$	SS (2000)	$ABM-CARMA(2,1)$	$ABM-CARMA(2,1)$	SS (2000)	$ABM-CARMA(2,1)$	$ABM-CARMA(2,1)$	SS (2000)	$ABM-CARMA(2,1)$	$ABM-CARMA(2,1)$	SS (2000)	SS (2000)
F01	0.0409	0.0486	0.0318	0.0370	1.26%	1.49%	0.95%	1.10%				
F02	0.0283	0.0330	0.0228	0.0260	0.87%	1.02%	0.68%	0.77%				
F03	0.0207	0.0230	0.0167	0.0186	0.64%	0.71%	0.50%	0.55%				
F04	0.0157	0.0161	0.0123	0.0132	0.48%	0.50%	0.37%	0.39%				
F05	0.0116	0.0111	0.0090	0.0091	0.36%	0.34%	0.27%	0.27%				
F06	0.0083	0.0073	0.0063	0.0059	0.25%	0.22%	0.19%	0.18%				
F07	0.0055	0.0043	0.0042	0.0034	0.17%	0.13%	0.13%	0.10%				
F08	0.0033	0.0020	0.0025	0.0015	0.10%	0.06%	0.08%	0.05%				
F09	0.0014	0.0000	0.0011	0.0000	0.04%	0.00%	0.03%	0.00%				
F10	0.0003	0.0014	0.0002	0.0011	0.01%	0.04%	0.01%	0.03%				
F11	0.0010	0.0021	0.0008	0.0016	0.03%	0.07%	0.02%	0.05%				
F12	0.0015	0.0025	0.0012	0.0019	0.05%	0.08%	0.04%	0.06%				
F13	0.0017	0.0025	0.0013	0.0019	0.05%	0.08%	0.04%	0.06%				
F14	0.0016	0.0023	0.0013	0.0018	0.05%	0.07%	0.04%	0.05%				
F15	0.0013	0.0018	0.0010	0.0014	0.04%	0.06%	0.03%	0.04%				
F16	0.0008	0.0011	0.0006	0.0008	0.03%	0.03%	0.02%	0.03%				
F17	0.0002	0.0000	0.0001	0.0000	0.01%	0.00%	0.00%	0.00%				
F18	0.0009	0.0012	0.0007	0.0009	0.03%	0.04%	0.02%	0.03%				
F19	0.0019	0.0026	0.0015	0.0020	0.06%	0.08%	0.05%	0.06%				
F20	0.0030	0.0040	0.0024	0.0031	0.09%	0.13%	0.07%	0.10%				
F21	0.0042	0.0055	0.0033	0.0044	0.13%	0.18%	0.10%	0.14%				
F22	0.0055	0.0070	0.0043	0.0056	0.17%	0.23%	0.13%	0.17%				
F23	0.0067	0.0085	0.0053	0.0069	0.21%	0.27%	0.16%	0.21%				
F24	0.0080	0.0100	0.0063	0.0081	0.25%	0.32%	0.19%	0.25%				
All	0.0122	0.0141	0.0057	0.0065	0.38%	0.43%	0.17%	0.20%				

Table 4: Out-of-Sample Pricing Errors: Cross Section

This table reports out-of-sample pricing errors for the $ABM-CARMA(2,1)$ and the model of Schwartz and Smith (2000) for the long end of the futures curve (contracts 13-24). Parameters were estimated using the short end (contracts 1-12). Panel A displays absolute Root Mean Squared Errors (RMSE), and Mean Absolute Errors (MAE). Panel B displays the corresponding statistics on a relative basis.

	RMSE		MAE		RMSE		MAE	
	$ABM-CARMA(2,1)$	SS (2000)	$ABM-CARMA(2,1)$	SS (2000)	$ABM-CARMA(2,1)$	SS (2000)	$ABM-CARMA(2,1)$	SS (2000)
	<i>Panel A: Absolute</i>				<i>Panel B: Relative</i>			
F13	0.0038	0.0054	0.0028	0.0041	0.12%	0.17%	0.08%	0.13%
F14	0.0052	0.0071	0.0038	0.0054	0.16%	0.23%	0.12%	0.17%
F15	0.0068	0.0090	0.0050	0.0069	0.21%	0.29%	0.15%	0.21%
F16	0.0083	0.0109	0.0061	0.0083	0.26%	0.35%	0.19%	0.26%
F17	0.0098	0.0128	0.0074	0.0098	0.31%	0.41%	0.23%	0.31%
F18	0.0115	0.0149	0.0088	0.0114	0.36%	0.48%	0.27%	0.36%
F19	0.0133	0.0171	0.0102	0.0131	0.42%	0.56%	0.31%	0.41%
F20	0.0152	0.0194	0.0117	0.0148	0.48%	0.63%	0.36%	0.47%
F21	0.0173	0.0216	0.0134	0.0166	0.55%	0.71%	0.41%	0.53%
F22	0.0194	0.0239	0.0150	0.0184	0.62%	0.78%	0.47%	0.58%
F23	0.0215	0.0262	0.0168	0.0202	0.68%	0.86%	0.52%	0.64%
F24	0.0237	0.0284	0.0185	0.0219	0.75%	0.93%	0.57%	0.69%
All	0.0144	0.0179	0.0100	0.0126	0.46%	0.59%	0.31%	0.40%

Table 5: **Out-of-Sample Pricing Errors: Time Series**

This table reports out-of-sample pricing errors for the $ABM-CARMA(2,1)$ and the model of Schwartz and Smith (2000) for the second half of the entire data set. Parameters were estimated using the first half. Panel A displays absolute Root Mean Squared Errors (RMSE), and Mean Absolute Errors (MAE). Panel B displays the corresponding statistics on a relative basis.

	RMSE			MAE			RMSE			MAE		
	$ABM-CARMA(2,1)$	SS (2000)	$ABM-CARMA(2,1)$	$ABM-CARMA(2,1)$	SS (2000)	$ABM-CARMA(2,1)$	$ABM-CARMA(2,1)$	SS (2000)	$ABM-CARMA(2,1)$	SS (2000)	$ABM-CARMA(2,1)$	SS (2000)
	Panel A: Absolute						Panel B: Relative					
F01	0.0409	0.0486	0.0318	0.0370	1.26%	1.49%	0.95%	1.10%				
F02	0.0283	0.0330	0.0228	0.0260	0.87%	1.02%	0.68%	0.77%				
F03	0.0207	0.0230	0.0167	0.0186	0.64%	0.71%	0.50%	0.55%				
F04	0.0442	0.0449	0.0358	0.0359	1.14%	1.14%	0.91%	0.91%				
F05	0.0417	0.0420	0.0335	0.0335	1.06%	1.06%	0.85%	0.85%				
F06	0.0398	0.0399	0.0316	0.0315	1.00%	1.00%	0.80%	0.79%				
F07	0.0383	0.0384	0.0303	0.0301	0.96%	0.96%	0.76%	0.76%				
F08	0.0372	0.0372	0.0291	0.0290	0.92%	0.92%	0.73%	0.73%				
F09	0.0364	0.0364	0.0284	0.0283	0.90%	0.90%	0.71%	0.71%				
F10	0.0357	0.0357	0.0278	0.0277	0.88%	0.88%	0.70%	0.69%				
F11	0.0351	0.0351	0.0272	0.0272	0.86%	0.86%	0.68%	0.68%				
F12	0.0345	0.0346	0.0266	0.0267	0.84%	0.85%	0.67%	0.67%				
F13	0.0340	0.0341	0.0261	0.0262	0.83%	0.83%	0.65%	0.66%				
F14	0.0336	0.0336	0.0257	0.0258	0.82%	0.82%	0.64%	0.64%				
F15	0.0331	0.0332	0.0253	0.0253	0.80%	0.81%	0.63%	0.63%				
F16	0.0327	0.0327	0.0249	0.0249	0.79%	0.79%	0.62%	0.62%				
F17	0.0323	0.0323	0.0245	0.0245	0.78%	0.78%	0.61%	0.61%				
F18	0.0320	0.0319	0.0242	0.0240	0.77%	0.77%	0.60%	0.60%				
F19	0.0317	0.0316	0.0239	0.0237	0.77%	0.76%	0.60%	0.59%				
F20	0.0315	0.0313	0.0237	0.0234	0.76%	0.75%	0.59%	0.58%				
F21	0.0315	0.0312	0.0238	0.0233	0.76%	0.75%	0.59%	0.58%				
F22	0.0315	0.0311	0.0239	0.0233	0.76%	0.75%	0.60%	0.58%				
F23	0.0318	0.0312	0.0242	0.0234	0.77%	0.75%	0.61%	0.58%				
F24	0.0322	0.0315	0.0248	0.0237	0.79%	0.76%	0.63%	0.59%				
All	0.0375	0.0381	0.0288	0.0290	0.94%	0.95%	0.72%	0.73%				