

QUANTITATIVE FINANCE
RESEARCH CENTRE



UNIVERSITY OF
TECHNOLOGY SYDNEY



QUANTITATIVE FINANCE RESEARCH CENTRE

Research Paper 260

July 2009

Modelling and Estimating the Forward Price Curve in the Energy Market

Carl Chiarella, Les Chewlow and Boda King

ISSN 1441-8010

www.qfrc.uts.edu.au

MODELLING AND ESTIMATING THE FORWARD PRICE CURVE IN THE ENERGY MARKET

CARL CHIARELLA*, LES CLEWLOW† AND BODA KANG‡

ABSTRACT. The stochastic or random nature of commodity prices plays a central role in models for valuing financial contingent claims on commodities. In this paper, by enhancing a multi factor framework which is consistent not only with the market observable forward price curve but also the volatilities and correlations of forward prices, we propose a two factor stochastic volatility model for the evolution of the gas forward curve. The volatility is stochastic due to a hidden Markov Chain that causes it to switch between “on peak” and “off peak” states. Based on the structure functional forms for the volatility, we propose and implement the Markov Chain Monte Carlo (MCMC) method to estimate the parameters of the forward curve model. Applications to simulated data indicate that the proposed algorithm is able to accommodate more general features, such as regime switching and seasonality. Applications to the market gas forward data shows that the MCMC approach provides stable estimates.

1. INTRODUCTION

The stochastic or random nature of commodity prices plays a central role in models for valuing financial contingent claims on commodities, and in procedures for evaluating investments to extract or produce the commodity. There are currently two approaches to modelling forward price dynamics in the literature. The first starts from a stochastic representation of the energy spot asset and other key variables, such as the convenience yield on the asset and interest rates (see for example Gibson & Schwartz (1990) and Schwartz (1997)), and derives the prices of energy contingent claims consistent with

Date: Current Version July 23, 2009.

* carl.chiarella@uts.edu.au; School of Finance and Economics, University of Technology, Sydney, Australia.

† les@lacimagroup.com; Lacima Group, Sydney, Australia.

‡ Corresponding author: boda.kang@uts.edu.au; School of Finance and Economics, University of Technology, Sydney, PO Box 123, Broadway, NSW 2007, Australia.

the spot process. However, one of the problems in implementing these models is that often the state variables are unobservable - even the spot price is hard to obtain, with the problems being exacerbated if the convenience yield has to be jointly estimated.

The second stream of literature models the evolution of the forward curve. Forward contracts are widely traded on many exchanges with prices easily observed - often the nearest maturity forward price is used as a proxy for the spot price with longer dated contracts used to imply the convenience yield. Clewlow & Strickland (1999a) work in this second stream, simultaneously modelling the evolution of the entire forward curve conditional on the initially observed forward curve and as such define a unified approach to the pricing and risk management of a portfolio of energy derivative positions.

These authors go on to develop a single-factor modelling framework which is consistent with market observable forward prices and volatilities and which leads to analytical pricing formulae for standard options, caps, floors, collars and convenient numerical schemes for swaptions. They also show how American style and exotic energy derivatives can be priced using trinomial trees which are constructed to be consistent with the forward curve and volatility structure.

However single-factor forward curve models leave very little room for complex movements of the forward price curve due to an underlying assumption of perfect correlation between forward prices of different maturities. This is inherently unrealistic, and becomes a critical flaw when the derivative in question depends on the relative values of points at different maturities along the forward price curve. Another shortcoming of the single-factor model of the forward price curve manifests itself in pricing long-term contracts on the spot price.

Clewlow & Strickland (1999b) develop a general framework with a multi-factor model for the risk management of energy derivatives. This framework is designed to be consistent not only with the market observable forward price curve but also the volatilities and correlations of forward prices. Breslin, Clewlow, Kwok & Strickland (2008) further generalize this framework to accommodate a more general multi-factor, multi-commodity (MFMC) model and also describe a process for estimating parameters from historical data.

There has been a great deal of fruitful research on the models of the interest rate term structure. Since the forward price curve shares similar patterns with forward interest rate dynamics, in particular the Heath, Jarrow & Morton (1992) (HJM) model, many of these ideas could potentially work in modelling forward energy prices. The instantaneous forward price volatility is not directly observable. However the observed implied volatility (obtained from the prices of various derivative contracts) is closely related to it and indicates that the market is also changing its belief about the instantaneous volatilities discontinuously. In deterministic volatility HJM models, the volatility curve is fixed and the volatility of a specific forward rate can change deterministically only with maturity. In order to properly describe the actual evolution of the volatility curve, one needs a process consisting of both deterministic factors and random factors. The drawback of diffusion models is that they cannot generate sudden and sufficiently large shifts of the volatility curve. If one augments that feature by adding traditional type jump processes, for example Poisson jumps, one finds that the frequency of the jumps is too large while the magnitude of the jumps is too small.

It seems that the class of piecewise-deterministic processes provide an appropriate framework for modelling the dynamics of the term structure of volatilities since they allow volatility to follow an almost deterministic process between two random jump times. Davis (1984) claims that this class covers almost all important non-diffusion applications. The simplest process in this class is the continuous-time homogeneous Markov chain with a finite number of jump times. Modelling with such a process approximates the actual jumps in volatility with jumps over a finite set of values but allows the well-developed stochastic calculus for continuous Markov chains (Elliott, Aggoun & Moore (1995) and Aggoun & Elliott (2005)) to be used.

In energy markets, forward contracts are widely traded on many exchanges with prices easily observed - often the nearest maturity forward price is used as a proxy for the spot price with longer dated contracts used to imply the convenience yield. Two volatility functions are proposed to model the gas forward curve in this paper. Both volatility functions contain a common seasonality adjustment term, but one volatility function is declining with increasing maturity to a constant; another volatility function captures the

overall tilting of the forward curve where the short maturity contracts move in the opposite direction to the longer maturity contracts. We allow the parameters (including the spot volatilities and the attenuation parameters) of the volatility functions to take different values in different states of the world. The dynamics of the “states of the world”, for example an “on-peak” or “off-peak” time for gas, or a “good” or “bad” economic environment are represented by a Markov Chain. The evolution of the two volatility functions depend on the transition of the Markov Chain.

The rest of the paper is organized as follows. A two factor regime switching model for the gas forward price curve is discussed in Section 2. In Section 3 a Markov Chain Monte Carlo approach including several detailed algorithms is analysed and implemented to estimate the parameters in the model. A number of numerical examples and the calibration results from the real market data are discussed in Section 4. Finally a conclusion is drawn in Section 5.

2. TWO-FACTOR REGIME SWITCHING FORWARD PRICE CURVE MODEL

The single-factor models proposed in Clewlow & Strickland (1999a) have a wide range of applicability in energy valuation and risk management, and are relatively simple to understand and parameterise. While these models can capture much of the dynamics of actual life processes in many circumstances, by definition they only use a small amount of the potential information available from the market. In particular, one major drawback of single-factor models is that they imply that instantaneous changes in forward prices at all maturities are perfectly correlated. Increasingly, energy risk practitioners are attracted to modelling frameworks that avoid such simplifications. Where enough data is available, a more general multi-factor model can be used to capture extra information about the price dynamics, and this is the modelling framework that we concentrate on here. It is also relatively straightforward to extend such a multi-factor model to incorporate multiple commodities.

Following the spirit of both Clewlow & Strickland (1999b) and Breslin et al. (2008) and the idea of the piecewise deterministic processes discussed in the previous section, rather than a general form of the volatility functions in the multi-factor model, we propose a two factor model in which we specify the explicit form of the volatility function. Benth & Koekebakker (2008) compared

several different volatility dynamics of the forward curve models under the HJM framework. In this paper, incorporating both the volatility functional form of these latter authors and the Markov switching of the coefficients, we propose the following two-factor regime switching model for the gas forward curve:

$$\frac{dF(t, T)}{F(t, T)} = \sigma_1(t, T)dW_1(t) + \sigma_2(t, T)dW_2(t), \quad (1)$$

with

$$\begin{aligned} \sigma_1(t, T) &= \langle \sigma_1, X_t \rangle c(t) \left(e^{-\langle \alpha_1, X_t \rangle (T-t)} \cdot (1 - \sigma_{l_1}) + \sigma_{l_1} \right), \\ \sigma_2(t, T) &= \langle \sigma_2, X_t \rangle c(t) \left(\sigma_{l_2} - e^{-\langle \alpha_2, X_t \rangle (T-t)} \right), \\ c(t) &= c + \sum_{j=1}^J (d_j (1 + \sin(f_j + 2\pi j t))), \end{aligned}$$

where:

- $F(t, T)$ is the price of the gas forward at time t with a maturity at time T .
- X_t is a finite state Markov chain with state space $S = \{e_1, e_2, \dots, e_N\}$ where e_i is a vector with length N and 1 at the i -th position and 0 elsewhere, that is

$$e_i = (0, \dots, 0, 1, 0, \dots, 0)' \in \mathbb{R}^N.$$

- $P = (p_{ij})_{N \times N}$ is the transition probability matrix of the Markov Chain X_t . For all $i = 1, \dots, N, j = 1, \dots, N$, p_{ij} is the conditional probability that the Markov Chain X_t transit from state e_i at time t to state e_j at time $t + 1$, that is,

$$p_{ij} = \Pr(X_{t+1} = e_j | X_t = e_i).$$

- For $i = 1, 2$,

$$\sigma_i = (\sigma_{i1}, \sigma_{i2}, \dots, \sigma_{iN}), \quad \alpha_i = (\alpha_{i1}, \alpha_{i2}, \dots, \alpha_{iN}),$$

are the different values of the volatilities and attenuation parameters which evolve following the rule of the Markov Chain X_t .

- $\langle \cdot, \cdot \rangle$ denotes the scalar product in \mathbb{R}^N , if $u = (u_1, \dots, u_N)$ then

$$\langle u, X_t \rangle = \sum_{i=1}^N u_i \mathbb{I}_{(X_t=e_i)};$$

where $\mathbb{I}_{(X_t=e_i)}$ is the indicator function:

$$\mathbb{I}_{(X_t=e_i)} = \begin{cases} 1 & X_t = e_i \\ 0 & \text{otherwise.} \end{cases}$$

- σ_{l_1} and σ_{l_2} describe some features of the long run volatility;
- $c(t)$ - the seasonal part which is modelled as a truncated Fourier series;
- $W_1(t)$ and $W_2(t)$ are independent Brownian motions.

In the above model, both volatility functions contain a common seasonality adjustment term $c(t)$, but $\sigma_1(t, T)$ is declining with increasing maturity to a constant and $\sigma_2(t, T)$ captures the overall tilting of the forward curve where the short maturity contracts move in the opposite direction to the longer maturity contracts.

In the above model, some parameters of the proposed volatility function will change from one value to another depending on the state of the world or the change of the market structure. The transition of the state is characterized by a finite state Markov Chain. The element of the transition probability matrix P will form part of the parameter set to be estimated.

3. ESTIMATION ALGORITHMS AND IMPLEMENTATIONS

Hahn, Frühwirth-Schnatter & Sass (2007) use a Bayesian approach to estimate a Markov Switching Model (MSM). They derive a Markov Chain Monte Carlo Method (MCMC) for the MSM which yields better results than the corresponding EM algorithm.

Based on the structure of our volatility functional forms, we propose and implement the MCMC to estimate the parameters of the above forward curve model. Applications to simulated data indicate that the proposed algorithm is able to accommodate more general features, such as regime switching, seasonality, certain functional form of forward volatility functions. Applications to the market gas forward data shows that the MCMC approach provides stable estimates.

In this section, we describe an algorithm to estimate the set of parameters

$$\theta = \{\sigma_i, \alpha_i, \sigma_{l_i} (i = 1, 2), c, d_j, f_j, (j = 1, 2, \dots, J)\}$$

and the elements of the transition probability matrix P given the forward prices at fixed observation times $\Delta t, 2\Delta t, \dots, M\Delta t = t$ assuming that the state process can jump only at the discrete observation times.

3.1. Data transformation. For any maturity time T , we observe forward prices $F(k\Delta t, T), k = 1, \dots, M$ from the market. Setting $\mathbf{y}_T = \{y_{t,T}, t \geq 0\}$, we rewrite the forward dynamics (1) in discrete time as

$$\begin{aligned} y_{t,T} &:= \log(F(t + \Delta t, T)) - \log(F(t, T)) \\ &= -\frac{1}{2} \sum_{i=1}^2 \sigma_i(t, T)^2 \Delta t + \sum_{i=1}^2 \sigma_i(t, T) \Delta W_i(t). \end{aligned}$$

Apparently, given all parameters and the state of the Markov Chain at each time, the difference of the log of forward price, namely $y_{t,T}$ is normally distributed and the random variables are independent for different t . Hence, $y_{t,T}$ is normal distributed with a mean of $\mu_{t,T} = -\frac{1}{2} \sum_{i=1}^2 \sigma_i(t, T)^2 \Delta t$ and standard deviation $\sigma_{t,T} = \sqrt{\Delta t \sum_{i=1}^2 \sigma_i(t, T)^2}$.

3.2. Prior Distributions. Prior distributions have to be chosen for θ and X_0 .

Assumption 3.1. Assume that $\theta = \{\sigma_i, \alpha_i, \sigma_{l_i} (i = 1, 2), c, d_j, f_j, (j = 1, 2, \dots, J)\}, P$ and X_0 are independent, that is

$$Pr(\theta, P, X_0) = Pr(P) \prod_{i=1}^2 (Pr(\sigma_i) \cdot Pr(\alpha_i) \cdot Pr(\sigma_{l_i})) Pr(c) \prod_{j=1}^J (Pr(d_j) Pr(f_j)) Pr(X_0).$$

In addition to Assumption 3.1 we further assume that, for $l = 1, \dots, N$ and $i = 1, 2, j = 1, \dots, J$:

- The rows P_l of P are assumed to be independent and to follow a Dirichlet distribution

$$P_l \sim D(g_{l1}, \dots, g_{lN}).$$

- The priors of σ s and α s are

$$\sigma_{il} \sim U(a_{il}^\sigma, b_{il}^\sigma), \alpha_{il} \sim U(a_{il}^\alpha, b_{il}^\alpha).$$

- The priors of σ_{l_i} are

$$\sigma_{l_i} \sim U(0, 1).$$

- The priors of c, d, f are

$$c, d, f \sim U(a^c, b^c).$$

- The prior of the initial state of the Markov chain

$$X_0 \sim U(\{1, \dots, N\}).$$

All parameters $g_{l1}, \dots, g_{lN}, a_{il}^\sigma, b_{il}^\sigma, a_{il}^\alpha, b_{il}^\alpha, a^c, b^c$ in the prior distributions need to be chosen carefully before the calibrations process.

3.3. Full conditional posterior distributions. To sample from the joint posterior distribution of $\{\sigma_i, \alpha_i, \sigma_{l_i} (i = 1, 2), c, d_j, f_j, (j = 1, 2, \dots, J)\}, P$ and X given the observed data \mathbf{y}_T , we partition the unknowns into three blocks θ, P and X , and draw each of them from the appropriate conditional density.

3.3.1. Complete-data likelihood function. Given θ, P and X , the log return process \mathbf{y}_T is independent, and the likelihood function is given by

$$P(\mathbf{y}_T | \theta, P, X) = \prod_{k=1}^M \phi(y_{k\Delta t, T}, \mu_{k\Delta t, T}, \sigma_{k\Delta t, T}), \quad (2)$$

where $\phi(y, \mu, \sigma)$ denotes the density of a normal distribution with mean $\mu_{k\Delta t, T}$ and standard deviation $\sigma_{k\Delta t, T}$, given here by

$$\mu_{t, T} = -\frac{1}{2} \sum_{i=1}^2 \sigma_i(t, T)^2 \Delta t, \quad \sigma_{t, T} = \sqrt{\Delta t \sum_{i=1}^2 \sigma_i(t, T)^2}.$$

3.3.2. Drift and volatility. The conditional joint distribution of the parameters associated with drift and volatility is given by

$$\Pr(\theta | \mathbf{y}_T, P, X) \propto \Pr(\mathbf{y}_T | \theta, X) \Pr(\theta).$$

3.3.3. State process. The prior distribution of the state process $X_{l\Delta t}, l = 1, 2, \dots, M$ is determined by the distribution of X_0 , and the transition probabilities P , and is independent of θ . Therefore the full conditional posterior is

$$\Pr(X | \mathbf{y}_T, \theta, P) \propto \Pr(\mathbf{y}_T | \theta, X) P(X | P).$$

Defining $N_{lq} = \sum_{k=1}^M \mathbb{I}_{\{X_{k-1}=l, X_k=q\}}$, the number of transitions from state l to q , the probability of X given P is:

$$\Pr(X|P) = P(X_0|P) \prod_{k=1}^M \Pr(X_k|X_{k-1}, P) = P(X_0|P) \prod_{l,q=1}^N P_{lq}^{N_{lq}}.$$

By forward-filtering-backward-sampling, we update X by drawing from the full conditional posterior $\Pr(X|\mathbf{y}_T, \theta, P)$.

3.4. Proposal distributions.

3.4.1. *Drift and volatility.* For the update of θ , we use the normal random walk

$$\theta' = \theta + r^\theta \psi, \tag{3}$$

where ψ is a matrix of independent standard normal random variables and r^θ are parameters scaling the step widths. Hence, we have a Metropolis step (see Hahn et al. (2007)) with acceptance probability $\alpha_\theta = \min\{1, \bar{\alpha}_\theta\}$, where

$$\bar{\alpha}_\theta = \frac{\Pr(\mathbf{y}_T|\theta', X)\Pr(\theta')}{\Pr(\mathbf{y}_T|\theta)P(\theta)}.$$

3.4.2. *Transition matrix.* For the update of the transition matrix, we use the method of sampling from a Dirichlet distribution as described in Frühwirth-Schnatter (2006). For each row $l = 1, \dots, N$ of the transition matrix, the proposal

$$P'_l \sim D(g_{l1} + N_{l1}, \dots, g_{lN} + N_{lN}),$$

is used. If the initial distribution of the state process $\Pr(X_0)$ is independent of P , then P'_l is a sample from the appropriate full conditional distribution, and we obtain a Gibbs step (see Hahn et al. (2007)) with acceptance rate 1.

3.5. **Algorithm.** Start with some state process \mathbf{X}^0 and repeat the following steps for $k = 1, \dots, K_0, \dots, K + K_0$.

- (1) **Parameter simulation** conditional on the states $\mathbf{X}^{(k-1)}$:
 - **Sample the transition matrix** $P^{(k)}$ from the complete-data posterior distribution $p(\cdot|\mathbf{X}^{(k-1)})$.
 - **Sample the model parameters** $\theta^{(k)}$ from the complete-data posterior $p(\theta|\mathbf{y}_T, \mathbf{X}^{(k-1)})$ (Metropolis-Hastings algorithm). Store the actual values of all parameters as $(\theta^{(k)}, P^{(k)})$.

- (2) **Markov Chain State simulation** conditional on $(\theta^{(k)}, P^{(k)})$ by sampling a path X of the hidden Markov chain from the conditional posterior $p(\mathbf{X} | (\theta^{(k)}, P^{(k)}), \mathbf{y}_T)$ (forward-filtering-backward-sampling). Store the actual values of all states as $\mathbf{X}^{(k)}$, increase k by one, and return to step (1). Finally, the first K_0 draws are discarded.

4. APPLICATIONS

In this section, we describe some details on the implementation, then we present numerical results of the proposed algorithms both for simulated data and historical gas forward prices.

4.1. Implementations on gas forward data. We have a number of time series of daily NBP Natural Gas forward price data denoted by $F(t, T)$. We have total 24 different maturities and for every maturity T which ranges from March 2008 to February 2010, we have a little more than one year of data with t from 29/09/2006 to 19/02/2008. Following the previous section, we apply the log difference transform to the original data $F(t, T)$ to obtain the data series \mathbf{y}_T for every maturity T . Figures 1 and 2 below show the behavior of $F(t, T)$ and \mathbf{y}_T respectively.




FIGURE 1. Forward price curves of 24 different maturities from March 2008 to February 2010 with data from 29/09/2006 to 19/02/2008.




FIGURE 2. The differences of the log of forward prices y_T with 24 different maturities from March 2008 to February 2010 with data from 30/09/2006 to 19/02/2008.

We use a two-state Markov Chain in this implementation, hence we set $N = 2$.

4.1.1. *Choosing the prior.* The prior distribution should be independent of the data and the posterior distribution. However in order to obtain a relatively accurate and relevant prior distribution, we obtain some information from both the data and the historical estimation and information about the range of the parameters. The following priors are used in the implementation:

- (1) For the transition matrix P , the vector g_l equals the prior expectation of P_l times a constant that determines the variance. If P^0 denotes our prior expectation of P , we may set $g_l = P_l^0 c_p$. Then c_p can be interpreted as the number of observations of jumps out of state l in the prior distribution. Hence, the matrix of parameters for the Dirichlet distribution is set to

$$g = \begin{pmatrix} 0.64 & 0.36 \\ 0.16 & 0.84 \end{pmatrix} \cdot c_p. \quad (4)$$

where $c_p = 15$.

- (2) The priors of σ_s and α_s ($i, l = 1, 2$) are

$$\sigma_{il} \sim U(0, 1.5), \quad \alpha_{il} \sim U(0, 5).$$

(3) The priors of σ_{l_i} are

$$\sigma_{l_i} \sim U(0, 1).$$

(4) The priors of c, d, f ($J = 1$) are

$$c, d_1, f_1 \sim U(0, 1).$$

(5) The prior of the initial state of the Markov chain are

$$X_0 \sim U(\{1, 2\}).$$

4.1.2. *Running MCMC.* To implement the Metropolis-Hastings algorithm, the scaling factors r^θ have to be selected in (3). We found that selecting r^θ around 1% to 5% of the minimal difference between two adjacent initial values among θ worked well.

In order to use all available information about forward prices on the market, we implemented a “randomised” algorithm to calculate the complete-data likelihood function.

In fact, before the calibration, we choose a sequence $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_M)$ with $\lambda_k \sim U(\{1, 2, \dots, 24\})$, $k = 1, \dots, M$ where λ_k ranges from 1 to 24 to label 24 different maturities. At each time $k \cdot \Delta t$ in the data series, we choose the data y_{t, T_j} with the maturity T_j if $\lambda_k = j$. Hence the complete-data likelihood function (2) under this setting is

$$P(\mathbf{y}_{T_\lambda} | \theta, P, X) = \prod_{k=1}^M \phi(y_{k\Delta t, T_{\lambda_k}}, \mu_{k\Delta t, T_{\lambda_k}}, \sigma_{k\Delta t, T_{\lambda_k}}), \quad (5)$$

where ϕ, μ and σ have the same definition as before.

4.2. **Numerical results.** For a two state regime, two factors model as above, about 150,000 steps were run and the following is the estimates of the parameters:

- Number of States $N = 2$

j	1	2	j	1	2
σ_{1j}	0.3057	0.8429	σ_{2j}	0.4762	1.0292
α_{1j}	2.0464	1.8932	α_{2j}	1.1533	3.2536

- $c = 0.8121$, $d = 0.0781$, $f = 0.3070$
- $\sigma_{l_1} = 0.4869$, $\sigma_{l_2} = 0.6203$;

- Transition probabilities

j	1	2
1	0.8516	0.1484
2	0.7080	0.2920

which shows that with high probability the Markov chain will stay in regime 1 most of the time but with some (low) probability will jump from regime 1 to regime 2 occasionally.

Figure 3 shows simulations of the two volatility functions $\sigma_1(t, T)$ and $\sigma_2(t, T)$ and one trajectory of the Markov Chain X_t .




FIGURE 3. A simulation path of two volatility functions $\sigma_1(T-t)$ and $\sigma_2(T-t)$ as a function of $(T-t)$; but X_t evolve forward in time with t .

Figure 4 demonstrates the shape of two volatility functions in different regimes after the seasonal components $c(t)$ has been removed. Those results emphasize that one volatility function $\sigma_1(t, T)$ is declining with increasing maturity to a constant and another volatility function $\sigma_2(t, T)$ captures the overall tilting of the forward curve where the short maturity contracts move in the opposite direction to the longer maturity contracts.

A number of MCMC estimation procedures have been run starting from different sets of parameters with different initial samples implemented by




FIGURE 4. After removing the seasonal effect, volatility functions when the state is either in regime 1 or in regime 2 all the time.

different random seeds. We have obtained fairly similar results from those estimations, which is a testimony that the algorithm converges.

5. CONCLUSION

We have proposed a two factor regime switching volatility model for the forward price curve in the energy market. The regime change will correspond to the change of the demand or the change of the market structure which would apply to all of the forward curves in the market.

Using the approaches suggested by Hahn et al. (2007) and Frühwirth-Schnatter (2006), we implemented a MCMC approach to estimate the parameters of the model including the transition probabilities of the Markov Chain from all available forward curves on the market. The approach we proposed and implemented is able to accommodate more complicated functional forms for the volatility functions.

This is an “off line” approach which is not efficient enough to update the estimates when new data arrives, so in future research we will look into some “online” approach, such as the particle filter, to estimate the parameters of




FIGURE 5. A number of simulation paths of forward prices with different maturities. In each of the graph above, the dark blue line represent the original forward price data and the others are the simulated paths based on the parameters estimated.

those complicated volatility functions when the observations are constantly arriving.

REFERENCES

- Aggoun, L. & Elliott, R. (2005), *Measure Theory and Filtering Introduction with Applications*, Cambridge University Press.
- Benth, F. & Koekebakker, S. (2008), ‘Stochastic modeling of financial electricity contracts’, *Energy Economics* **30**, 1116–1157.
- Breslin, J., Clewlow, L., Kwok, C. & Strickland, C. (2008), ‘Gaining from complexity: MFMC models’, *Energy Risk* **April**, 60–64.

- Clewlow, L. & Strickland, C. (1999a), ‘Valuing Energy Options in a One Factor Model Fitted to Forward Prices’, *QFRC Research Paper Series 10, University of Technology, Sydney*.
- Clewlow, L. & Strickland, C. (1999b), ‘A Multi-Factor Model for Energy Derivatives’, *QFRC Research Paper Series 28, University of Technology, Sydney*.
- Davis, M. (1984), ‘Piecewise-deterministic Markov processes: a general class of non-diffusion stochastic models’, *Journal of Royal Statistical Society B* **46**, 353–388.
- Elliott, R., Aggoun, L. & Moore, J. (1995), *Hidden Markov Models: Estimation and Control, Applications of Mathematics*, Springer-Verlag, New York.
- Frühwirth-Schnatter, S. (2006), *Finite Mixture and Markov Switching Models*, Springer, New York.
- Gibson, R. & Schwartz, E. S. (1990), ‘Stochastic Convenience Yield and the Pricing of Oil Contingent Claims’, *The Journal of Finance* **45**, 959–976.
- Hahn, M., Frühwirth-Schnatter, S. & Sass, J. (2007), Markov chain monte carlo methods for parameter estimation in multidimensional continuous time markov switching models, Ricam report, RICAM, Linz.
- Heath, D., Jarrow, R. & Morton, A. (1992), ‘Bond pricing and the term structure of interest rates: a new methodology’, *Econometrica* **60**, 77–105.
- Schwartz, E. S. (1997), ‘The Stochastic Behaviour of Commodity Prices: Implications for Pricing and Hedging’, *The Journal of Finance* **3**, 923–973.