Multi-factor Models and the Risk Premiums: A Simulation Study

Arne Andresen and Johan M. Sollie

February 23, 2013

Abstract

The estimation of commodity spot price models often involves the estimation of risk premiums. We show in a simulation study that the market prices of risk cannot be accurately estimated using two popular estimation techniques; the Kalman filter and an iterative routine. Risk premium parameters may be dependent on the starting value for the iterative routine, and cannot be accurately estimated using the Kalman filter technique. We conclude with a short analysis of results from the spot price model literature by examining the implied volatility term structure from other published research papers.

1 Introduction

The rate of development of commodity spot price models has increased during the last decades. Incomplete commodity markets give rise to a wide variety of specifications ranging from the mean reverting model of Schwartz (1997), the two factor model of Schwartz and Smith (2000), to the more general model by Cortazar and Naranjo (2006).

Spot price models are often fitted to data using the Kalman filter, since they include some possibly unobservable state variables. Another approach is an iterative routine as suggested by Lucia and Schwartz (2002), and further examined in Cortazar and Schwartz (2003). The algorithm offers a way of calibrating factor models to observed spot and forward prices.

When calibrating spot price models, estimates of the risk premium parameters are routinely reported. See for example Schwartz and Smith (2000), Wilkens and Wimschulte (2007), Cartea and Williams (2008), Bhar and Lee (2011) and Nomikos and Soldatos (2010). If the Kalman filter is applied, standard deviations of parameter estimates are readily available. The iterative algorithm, however, does not provide any measure of the precision of the parameter estimates.

As discussed by Schwartz and Smith (2000), the risk premium parameters cannot be estimated with much accuracy. They argue that a possible solution to the problem is using a much longer time series of forward prices. The aim of this paper is to shed light on the accuracy of the risk premium parameter estimates obtained using the Kalman filter or the iterative algorithm. That is, how the accuracy differs between the two methods, how the accuracy depends on the sample size and the volatility parameters. We focus on the estimation of the risk premium parameters in the two-factor model by Schwartz and Smith (2000). This model is chosen based on its popularity as a basis for commodity pricing models. In addition, if the risk premium parameters cannot be accurately estimated in this simple and intuitively appealing model, it is unlikely that risk premium parameters are identifiable in more complicated models.

The finite-sample properties of some of the standard techniques used to estimate term structure models are studied in Duffee and Stanton (2012). One finding is that maximum likelihood produces strongly biased parameter estimates when the model includes a flexible specification of the dynamics of interest rate risk. The authors argue that this result underscore the importance of performing detailed Monte Carlo analysis to study the small-sample properties of new estimators.

We extend the literature by examining the properties of the Kalman filter and the iterative estimator applied to a commodity spot price model with constant market price of risk. The experiment is a simulation study where samples are generated from the model and the parameters of the model are re-estimated. By using simulated data, and not real observed data, the true values of the process are known a priori, and effects of market imperfections are minimized. Examples of market imperfections can be poor liquidity, resulting in noisy data, transaction costs or investor sentiment. In a real world setting, it is not plausible that the parameters of a two factor model are stable over time. Risk aversion tends to fluctuate, see Lee et al. (1990) for a discussion on time varying risk aversion, and De Long et al. (1990) for a broader discussion of investor behavior. In addition, the volatility of commodity prices is known to be non-constant, see for example Pindyck (2004). Using a simulation study, the problem of estimating the risk premiums can be isolated, and the results will not be subject to other sources of randomness.

The paper is organized as follows: Section 2 introduces the two-factor model used in the study. Section 3 demonstrates how the model can be set in state space form and estimated with the Kalman filter, while Section 4 explains the alternative iterative estimation routine. Section 5 outlines how the model can be used for simulation. The results of the simulation study are discussed in Section 6 together with a short analysis of the implied volatility term structure of other published research papers. Section 7 concludes the paper.

2 Model formulation

This study examines the estimation of the short- and long-term risk premiums implied by the Schwartz and Smith (2000) model. In this model the spot price $S(t)$ is decomposed into two components,

$$s(t) = \ln S(t) = \chi_t + \xi_t. \tag{1}$$

χ_t is assumed to follow an Ornstein-Uhlenbeck process and ξ_t is assmued to follow a Wiener process with drift.

$$d\chi_t = -\kappa\chi_t dt + \sigma_\chi dW_\chi, \tag{2}$$
$$d\xi_t = \mu_\xi dt + \sigma_\xi dW_\xi, \tag{3}$$

where dW_ξ and dW_χ are increments of standard Wiener processes, and $dW_\chi dW_\xi = \rho dt$. Given χ_0 and ξ_0, the vector of expected values and covariance matrix is given by

$$\mathrm{E}\left[\chi_t, \quad \xi_t\right] = \left[e^{-\kappa t}\chi_0, \quad \xi_0 + \mu_\xi t\right], \tag{4}$$
$$\mathrm{Cov}\left[\chi_t, \quad \xi_t\right] = \begin{bmatrix} (1-e^{-2\kappa t})\frac{\sigma_\chi^2}{2\kappa} & (1-e^{-\kappa t})\frac{\rho\sigma_\chi\sigma_\xi}{\kappa} \\ (1-e^{-\kappa t})\frac{\rho\sigma_\chi\sigma_\xi}{\kappa} & \sigma_\xi^2 t \end{bmatrix}. \tag{5}$$

The log of the future spot price is then normally distributed with expectation given by

$$\ln S(t) = e^{-\kappa t}\chi_0 + \xi_0 + \mu_\xi t.$$

The latent stochastic factors of the model can be interpreted as short- and long-term components as follows. The influence of χ_t on the expected future spot price decays monotonically to 0. Hence it may be viewed as a short-term factor. The short-term factor typically captures temporary price changes that are not expected to persist. The mean reversion parameter κ describes the rate at which short-term deviations are expected to persist. ξ_t may be viewed as a long-term factor since changes to this factor will have a permanent effect on the expected future spot price.

To obtain the risk neutral dynamics two additional parameters, λ_χ and λ_ξ, are introduced.

$$d\chi_t = (-\kappa\chi_t - \lambda_\chi)dt + \sigma_\chi dW_\chi^*, \tag{6}$$
$$d\xi_t = (\mu_\xi - \lambda_\xi)dt + \sigma_\xi dW_\xi^*, \tag{7}$$

where dW_ξ^* and dW_χ^* are increments of standard Wiener processes, and $dW_\chi^* dW_\xi^* = \rho dt$. That is, the risk-neutral process for the short-term deviation is now an Ornstein-Uhlenbeck process reverting to λ_χ/κ, and the long-term factor has drift $\mu_\xi^* = \mu_\xi - \lambda_\xi$.

The tractability of the model allows us to solve for the log forward prices implied by the model, as given by

$$\begin{aligned} \ln(F_{T,0}) &= \ln(\mathrm{E}^*[S_T]) \\ &= e^{-\kappa T}\chi_0 + \xi_0 + A(T), \end{aligned} \tag{8}$$
$$A(T) = \mu_\xi^* T - (1-e^{-\kappa T})\frac{\lambda_\chi}{\kappa} + \frac{1}{2}\left((1-e^{-2\kappa T})\frac{\sigma_\chi^2}{2\kappa} + \sigma_\xi^2 T + 2(1-e^{-\kappa T})\frac{\rho\sigma_\chi\sigma_\xi}{\kappa}\right).$$

The rest of the paper will focus on the estimation of the short- and long-term risk premiums.

3 State space form

The state variables, χ_t and ξ_t, are not directly observable, but can be estimated by casting the model into state space form and employing a Kalman filter procedure, see Schwartz and Smith (2000). The Kalman filter facilitates the calculation of the likelihood, which allows us to estimate parameters using maximum likelihood techniques. A well-known property of maximum likelihood estimators is that they are efficient when the sample size tends to infinity. That is, for large sample sizes no other unbiased estimator has a lower mean squared error.

Let the state equation be given by

$$\mathbf{x}_t = \mathbf{c} + \mathbf{G}\mathbf{x}_{t-1} + \boldsymbol{\omega}_t \quad t = 1, \ldots, N \tag{9}$$

where

$$\mathbf{x}_t = \begin{bmatrix}\chi_t, & \xi_t\end{bmatrix}'$$
$$\mathbf{c} = \begin{bmatrix}0, & \mu_\xi \Delta t\end{bmatrix}'$$
$$\mathbf{G} = \begin{bmatrix} e^{-\kappa \Delta t} & 0 \\ 0 & 1 \end{bmatrix},$$

and $\boldsymbol{\omega}_t \sim N(\mathbf{0}, \boldsymbol{\Sigma})$, where $\boldsymbol{\Sigma}$ is given by Eq. (5). The measurement equation is given by

$$\mathbf{y}_t = \mathbf{d} + \mathbf{F}\mathbf{x}_t + \mathbf{v}_t, \quad t = 1, \ldots, N, \tag{10}$$

where

$$\mathbf{y}_t = \begin{bmatrix}\ln F_{T_1}, \ldots \ln F_{T_n}\end{bmatrix}'$$
$$\mathbf{d} = \begin{bmatrix}A(T_1) \ldots A(T_n)\end{bmatrix}'$$
$$\mathbf{F} = \begin{bmatrix} e^{-\kappa T_1} & 1 \\ \vdots & \vdots \\ e^{-\kappa T_n} & 1 \end{bmatrix}.$$

$\mathbf{v}_t \sim N(\mathbf{0}, \mathbf{V})$, $\mathbf{V} = \sigma_f^2 \mathbf{I}_n$, where n equals the number of different forward contracts.

The Kalman filter is a recursive procedure for computing the optimal estimate of the unknown state vector $\mathbf{x}_t$, $t = 1, 2, ..., N$, assuming that the model parameters are known. Let $\hat{\mathbf{x}}_{t|t-1}$ and $\hat{\mathbf{x}}_{t|t}$ denote the estimate of the state vector based on available information up to time $t-1$ and t, respectively. Let $\mathbf{P}_{t|t-1}$ and $\mathbf{P}_{t|t}$ denote the covariance matrix of the estimate of the state vector based on information up to time $t-1$ and t, respectively. The Kalman filter is described be the following equations:

1. $\hat{\mathbf{x}}_{t|t-1} = \mathbf{c} + \mathbf{G}\hat{\mathbf{x}}_{t-1|t-1}$

2. $\mathbf{P}_{t|t-1} = \mathbf{G}\mathbf{P}_{t-1|t-1}\mathbf{G}' + \mathbf{\Sigma}$

3. $\mathbf{v}_{t|t-1} = \mathbf{y}_t - \left(\mathbf{d} + \mathbf{F}\hat{\mathbf{x}}_{t|t-1}\right)$

4. $\mathbf{H}_{t|t-1} = \mathbf{F}\mathbf{P}_{t|t-1}\mathbf{F}' + \mathbf{V}$

5. $\hat{\mathbf{x}}_{t|t} = \hat{\mathbf{x}}_{t|t-1} + \mathbf{K}_t\mathbf{v}_{t|t-1}$

6. $\mathbf{P}_{t|t} = \mathbf{P}_{t|t-1} - \mathbf{K}_t\mathbf{F}\mathbf{P}_{t|t-1}$

where $\mathbf{K}_t = \mathbf{P}_{t|t-1}\mathbf{F}'\mathbf{H}_{t|t-1}^{-1}$ is the Kalman gain.

Assume $\hat{\mathbf{x}}_1$ is $N(\mathbf{x}_{1|0}, \mathbf{P}_{1|0})$ where $\mathbf{x}_{1|0}$ and $\mathbf{P}_{1|0}$ are known. The log likelihood function of the data is

$$\ln L = -\frac{1}{2}\sum_{t=1}^{n} \ln(2\pi\mathbf{H}_{t|t-1}) - \frac{1}{2}\sum_{t=1}^{n} \mathbf{v}'_{t|t-1}\mathbf{H}_{t|t-1}^{-1}\mathbf{v}_{t|t-1}. \tag{11}$$

The parameters of the model can now be estimated by maximizing the likelihood of the model with respect to the model parameters. The likelihood is evaluated using R and the package KFAS by Helske (2010). See Durbin and Koopman (2001) for details on the Kalman filter and the likelihood function.

Both the model parameters and the distribution of the state vector $\hat{\mathbf{x}}_1$ have to be initialized. It is common to provide a diffuse initialization of the Kalman filter, that is setting the variance to a large value, for example $\mathbf{P}_{1|0} = 10^6 \cdot \mathbf{I}$. When estimating the model to real data, we can obtain starting values of the model parameters by utilizing that the long-term factor will approximately follow the dynamics of a long forward contract, while the short-term factor will closely follow the spread between a short and long forward contract. Since these quantities are observable, the model parameters can be estimated using standard techniques. To ensure that the maximum likelihood estimation routine has reached a global maximum, it is also advised to run the optimizing routine using several different initial parameter values.

4 Iterative procedure

The unobserved state vector can also be recovered using an iterative procedure, as first described in Lucia and Schwartz (2002). For a given parameter configuration, for each day, the squared distance between model and observed log prices is minimized with respect to the state variables. From the estimated state variables, the covariance matrix and the long-term drift parameter are obtained. The last step in each iteration of the algorithm is once again minimizing the squared distance between model implied prices and observed prices, but with respect to the risk premiums and mean reversion parameters.

The state vector is initialized with vectors of zeros. After the algorithm has completed one iteration, the previous values of the state variables are used as starting values in the optimization to provide faster convergence. i.e., $\chi_{i,j-1}$ and $\xi_{i,j-1}$ is used as starting values for $\chi_{i,j}$ and $\xi_{i,j}$ respectively, with $\chi_{i,0} = \xi_{i,0} = 0$.

The minimization is done using the Nelder-Mead algorithm, as implemented in *optim* in R. A convergence criteria must be set with regards to the sum of pricing errors, e. The convergence criteria is set to

$$\ln\left(\frac{e_{j-1}}{e_j}\right) < \delta,$$

and $\delta = 0.0001$.

The details are as follows;

$j \leftarrow 1$
while $(\ln e_{j-1} - \ln e_j) > \delta$ **do**
  for $i = 1$ to N **do**
    $\{\hat{\chi}_{i,j}, \hat{\xi}_{i,j}\} \leftarrow \underset{\chi,\xi}{\text{argmin}}[\mathbf{Y}_i - \mathbf{f}_i(\chi, \xi)]^2$
  end for
  $\boldsymbol{\psi} \leftarrow (\hat{\boldsymbol{\xi}}_j - \hat{\boldsymbol{\xi}}_{-1,j} - \hat{\boldsymbol{\mu}}_{\xi,j-1}\Delta t)/\sqrt{\Delta t}$
  $\boldsymbol{\pi} \leftarrow \frac{\hat{\boldsymbol{\chi}}_j - \hat{\boldsymbol{\chi}}_{-1,j} exp(-\hat{\kappa}_{j-1}\Delta t)}{\sqrt{(1-exp(-2\hat{\kappa}_{j-1}\Delta t))/2\hat{\kappa}_{j-1}}}$
  $\hat{\mu}_{\xi,j} \leftarrow N^{-1}\sum_N(\hat{\boldsymbol{\xi}}_j - \hat{\boldsymbol{\xi}}_{-1,j})/\Delta t$
  $\hat{\sigma}_{\xi,j} \leftarrow \sqrt{\text{Var}[\boldsymbol{\psi}]}$
  $\hat{\sigma}_{\chi,j} \leftarrow \sqrt{\text{Var}[\boldsymbol{\pi}]}$
  $\hat{\rho}_j \leftarrow \text{Cor}[\boldsymbol{\psi}, \boldsymbol{\pi}]$
  $\{\hat{\kappa}_j, \hat{\mu}^*_{\xi,j}, \hat{\lambda}_{\chi,j}\} \leftarrow \underset{\kappa,\mu^*_\xi,\lambda_\chi}{\text{argmin}}[\mathbf{Y} - \mathbf{f}(\kappa, \mu^*_\xi, \lambda_\chi)]^2$
  $e_j \leftarrow \sum_i [\mathbf{Y}_i - \mathbf{f}_i(\hat{\chi}_j, \hat{\xi}_j, \hat{\kappa}_j, \hat{\mu}_{\xi,j}, \hat{\mu}^*_{\xi,j}, \hat{\lambda}_{\chi,j}, \hat{\sigma}_{\chi,j}, \hat{\sigma}_{\xi,j}, \hat{\rho}_j)]^2$
  $j \leftarrow j + 1$
end while

where $\mathbf{Y}$ is panel of simulated log forward prices, $\mathbf{f}(\cdot)$ is a panel of model implied log forward prices, $\boldsymbol{\chi}_{-1,j}$ and $\boldsymbol{\xi}_{-1,j}$ denote the lagged short- and long-term state vectors, respectively.

5 Simulation procedure

The model allows for direct simulation of the log forward price curve via Eq. (8). We simulate the log forward price curve consisting of 20 log forward prices equally spaced in time, reaching up to five years ahead. The prices are computed daily for 2500 days, totaling 50000 observations.

Let $\boldsymbol{\epsilon}$ be a matrix of correlated normal random variables with expectation $\mathbf{0}$ and covariance matrix given by

$$\begin{bmatrix} \sigma_\chi^2 & \rho\sigma_\chi\sigma_\xi \\ \rho\sigma_\chi\sigma_\xi & \sigma_\xi^2 \end{bmatrix}.$$

The details of the simulation are as follows[1];

[1] ξ is initialized with an arbitrarily chosen value of 3. The initial value of the state variables do not influence the estimation of the model parameters.

Generate $\boldsymbol{\epsilon}$
$\chi_0 \leftarrow 0$
$\xi_0 \leftarrow 3$
for $i = 1$ to 2500 **do**
$\chi_i \leftarrow \chi_{i-1} exp(-\kappa\Delta t) + \sqrt{\frac{1-exp(-2\kappa\Delta t)}{2\kappa}}\epsilon_{i,1}$
$\xi_i \leftarrow \xi_{i-1} + \mu_\xi \Delta t + \sqrt{\Delta t}\epsilon_{i,2}$
end for

Given the $\boldsymbol{\chi}$ and $\boldsymbol{\xi}$ vectors, Eq. (8) is evaluated with the suitable time to maturity to get a panel $\mathbf{Y}$, $dim(\mathbf{Y}) = 20 \times 2500$ of forward prices, 50000 observations. A small measurement error, $\epsilon \sim N(0, \sigma_f^2)$ is added to each of the 50000 observations to avoid numerical issues in the Kalman filter.[2] The procedure is repeated each time a sample is generated.

The simulation study is conducted by choosing a set of reasonable parameter values as configuration for the simulation. These values are not changed during the study. Step $1-3$ in the procedure below is repeated five times, each time with a different set of starting values.

1. For the given configuration, a set of forward prices is simulated according to the previously described algorithm.

2. For a given set of starting values, the model is fitted to the simulated data using both the Kalman filter and the iterative routine.

3. Step $1-2$ repeated 100 times, and the estimated parameters stored.

The parameter configuration of the simulated process is chosen to be the same as estimated for the Enron data in Table 2 of Schwartz and Smith (2000); $\mu_\xi = -0.0386, \mu_\xi^* = 0.0161, \kappa = 1.19, \rho = 0.189, \sigma_\chi = 0.158, \sigma_\xi = 0.115, \sigma_f = 0.001, \lambda_\chi = 0.014$. The starting values $\mu^s, \kappa^s, \rho^s, \sigma_\chi^s, \sigma_\xi^s, \sigma_f^s$ are all given the true value[3]. First, the starting value of the short-term risk premium is changed, $\lambda_\chi^s = \lambda_{\chi,k}$, where $\lambda_{\chi,k} = -0.114 + (k \times 0.05)$ for $k = \{0, 1, 2, 3, 4\}$. I.e., only the value of the short-term risk premium λ_χ is changed, and the other model parameters are initialized at the same configuration as used for the simulation. Then the same procedure is followed for the long-term risk premium, λ_ξ. Since λ_ξ is given implicitly by the two μ parameters, we let $\mu_\xi^s = \mu_{\xi,m}$, where $\mu_{\xi,m} = -0.1547 + (m \times 0.05)$ for $m = \{0, 1, 2, 3, 4\}$. When examining the short-term risk premium, the long-term risk premium is initiated at its true value in the numeric search and vice versa.

[2]While the price of financial assets usually are observed without measurement error, the model can be slightly misspecified, giving rise to an error term. Regardless of the economical interpretation of the error term, the variance of the innovation in this experiment is given such a low value that it does not affect the estimated model parameters. We set $\sigma_f = 0.001$.

[3]To better isolate the problems that occur in the estimation of the risk premium parameters the initial configuration of the model parameters is fixed at their true value.

6 Results

In this section the results of the study are presented together with a short analysis of the volatility term structure.

The results from the study show that while both the Kalman filter and the iterative routine recover the correct values for $\mu_\xi^*, \kappa, \rho, \sigma_\chi, \sigma_\xi$, both routines fail in estimating the risk premiums with sufficient accuracy. Summary statistics for the Kalman filter routine are shown in Table 1 and the corresponding results for the iterative routine in Table 2.

The estimated value of the short-term risk premium is shown in Figure 1. The starting values clearly affect the parameter estimate using the iterative routine. This means the model will always give reasonable estimates of this parameter if it is given a reasonable starting value. This is not the case for the Kalman filter routine, but this routine is also not able to recover the true parameter value with any reasonable precision. The long-term risk premium results are shown in Figure 2 where the results from the estimation routines are plotted against each other. Two insights can be drawn from the plot; none of the routines result in accurate parameter estimates, and estimated parameters are highly correlated.

It is interesting to see how the uncertainty of the parameter estimates is affected by the sample size and volatility parameters. We choose two different parameter configurations. The first is the Enron data in Table 2 of Schwartz and Smith (2000) referred to above, while the second is the Nord Pool data estimates in Table 1 of Nomikos and Soldatos (2010); $\mu_\xi = 0.097, \mu_\xi^* = 0.0623, \kappa = 1.32, \rho = -0.42, \sigma_\chi = 1.00, \sigma_\xi = 0.3248, \sigma_f = 0.001, \lambda_\chi = 0.5416$. Electricity prices are know to be highly volatile, which is reflected in the σ parameters of the Nord Pool data. For each parameter configuration we generate 9×250 samples and estimate the parameters using the Kalman filter. In each sample we simulate the log forward price curve consisting of 20 log forward prices equally spaced in time, reaching up to five years ahead. Years of simulated data is chosen to be $1, 2.5, 5$, 10, 15, 20, 25, 35 and 50 years.

Figures 3 and 4 show how the uncertainty of the estimate is reduced with the sample size. It is seen that the risk premium parameters cannot be accurately estimated even with 50 years of data (250000 data points), but the uncertainty is greatly reduced from the sample of one year. The uncertainty of the estimates is highest for Nord Pool data, especially for the short-term risk premium. This is as expected, because the short-term volatility of the Nord Pool data is over six times higher compared to the Enron data. Both parameter estimates are centered on the true value, confirming that the Kalman filter yields unbiased parameter estimates.

As discussed by Schwartz and Smith (2000), the risk premiums determine the difference between the expected future spot price and forward prices. Because the expected future spot price is unobserved, the risk premiums are not accurately estimated. In more details, if we are given the risk-neutral drift μ_ξ^*, the equilibrium drift μ_ξ plays no role in the risk neutral price process. In addition, any shift in the short-term risk premium can be offset by a scaled (scaled

Maximum likelihood estimates

	Min	1st Qu.	Median	Mean	3rd Qu.	Max
μ_ξ	-0.1480	-0.0635	-0.0382	-0.0389	-0.0156	0.0977
μ_ξ^*	0.0158	0.0161	0.0161	0.0161	0.0161	0.0164
κ	1.1881	1.1902	1.1901	1.1902	1.1902	1.1921
ρ	0.1228	0.1759	0.1885	0.1886	0.2016	0.2378
σ_χ	0.1524	0.1569	0.1580	0.1580	0.1591	0.1623
σ_ξ	0.1122	0.1146	0.1148	0.1148	0.1149	0.1174
σ_f^2	9.816e-07	9.955e-07	1.000e-06	1.000e-06	1.005e-06	1.023e-06
λ_χ	-0.1376	-0.0220	0.0162	0.0158	0.0510	0.1461
λ_ξ	-0.1642	-0.0797	-0.0543	-0.0551	-0.0317	0.0815

Table 1: Descriptive statistics for the Maximum likelihood estimation routine. Notice the filtration of noise as given by σ_f^2, and the high dispersion of the risk premium parameters, λ_χ and λ_ξ. The parameter estimates are initiated at their true values, except for λ_χ and λ_ξ.

Iterative routine estimates

	Min	1st Qu.	Median	Mean	3rd Qu.	Max
μ_ξ	-0.1466	-0.0632	-0.0402	-0.0390	-0.0148	0.0990
μ_ξ^*	0.0154	0.0159	0.0161	0.0161	0.0162	0.0166
κ	1.1890	1.1900	1.1910	1.1911	1.1911	1.195
ρ	0.1170	0.1692	0.1814	0.1818	0.1956	0.2302
σ_χ	0.1525	0.1587	0.1602	0.1601	0.1614	0.1666
σ_ξ	0.1105	0.1141	0.1151	0.1152	0.1162	0.1214
λ_χ	-0.0876	-0.0364	0.0139	0.0139	0.0644	0.1159
λ_ξ	-0.1630	-0.0793	-0.0561	-0.0551	-0.0307	0.0828
Iterations	4.00	6.00	7.00	7.28	8.00	16.00

Table 2: Descriptive statistics for the iterative estimation routine. The parameter estimates are initiated at their true values, except for λ_χ and λ_ξ.

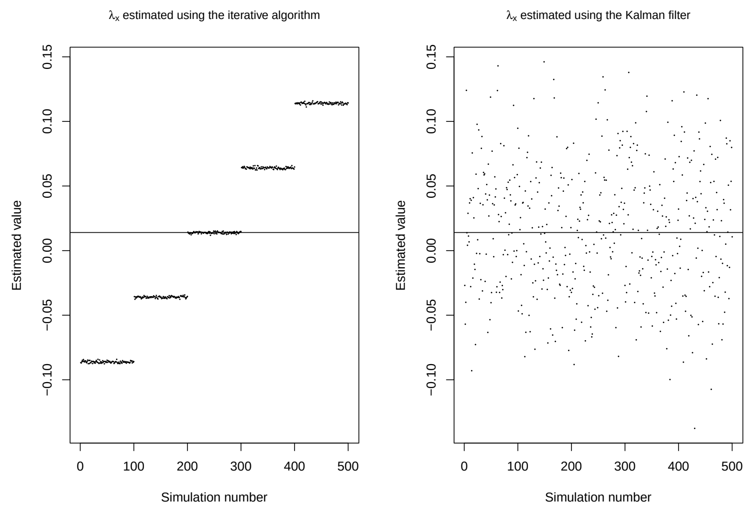

Figure 1: Estimated short-term risk premium, λ_χ. The solid line shows the correct value, $\lambda_\chi = 0.014$. None of the algorithms can recover the correct parameter value.

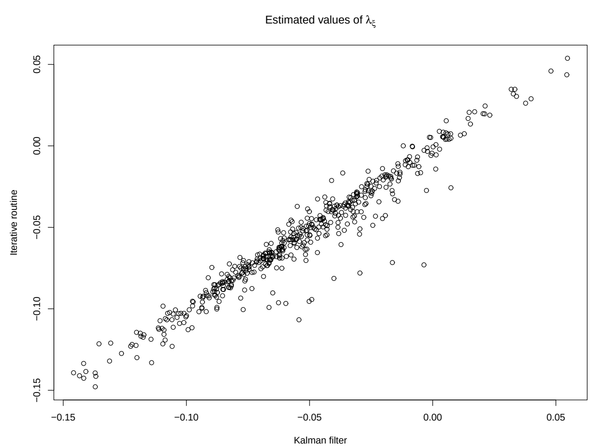

Figure 2: Estimated long-term risk premium, λ_ξ obtained from the Kalman filter routine plotted against the estimated parameter values given by the iterative routine. The samples are generated using $\lambda_\xi = -0.0547$.

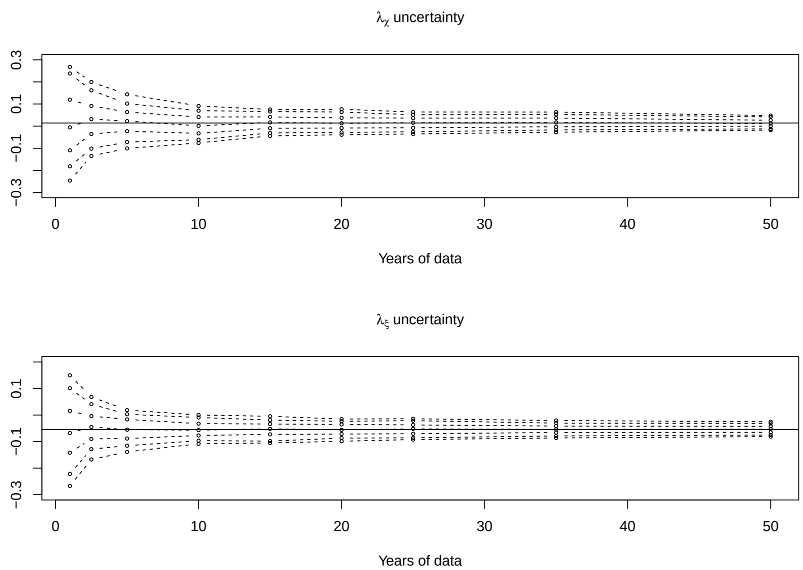

Figure 3: Uncertainty of the risk premium estimates of the Enron data. Dotted lines indicate the $95, 90, 75, 50, 25, 10$ and 5th percentile of the risk premium parameters estimates. Solid lines represent the true parameter values. The short-term risk premium λ_χ is shown in the upper plot, and the long-term risk premium, λ_ξ, is shown in the lower plot.

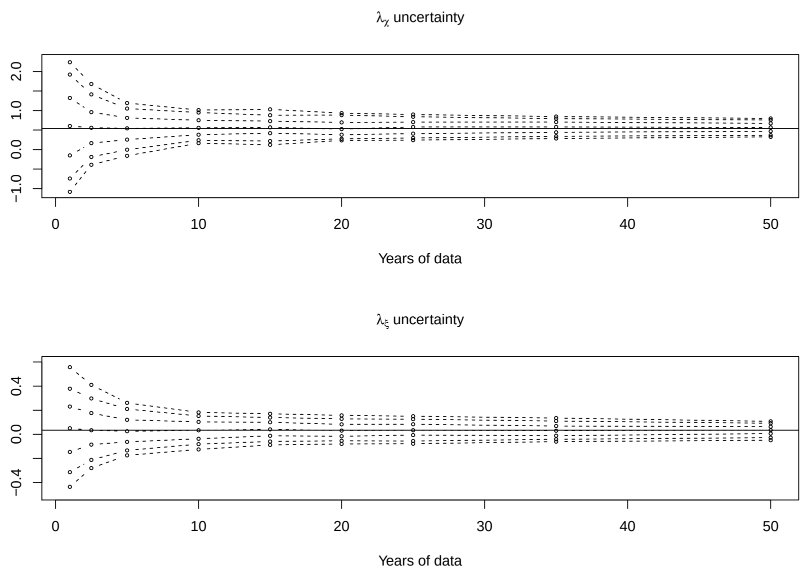

Figure 4: Uncertainty of the risk premium estimates of the Nord Pool data. Dotted lines indicate the $95, 90, 75, 50, 25, 10$ and 5th percentile of the risk premium parameters estimates. Solid lines represent the true parameter values. The short-term risk premium λ_χ is shown in the upper plot, and the long-term risk premium, λ_ξ, is shown in the lower plot.

with the mean reversion coefficient, κ) change in the initial value of the state vectors. Hence, we cannot precisely estimate the true process.

The estimate of the short-term risk premium provided by the iterative routine is dependent on the starting values for the same reason. In the iterative routine the state vectors are estimated conditional upon all model parameters, and the parameters are estimated conditional upon the state vectors. Assume we initialize the short-term risk premium with a perturbation δ. The unobserved short- and long-term vectors will then be estimated as $\chi - \delta/\kappa$ and $\xi + \delta/\kappa$, where χ and ξ are the true state vectors. Conditional upon these shifted state vectors the estimate of the short-term risk premium is $\lambda_\chi + \delta$, where is λ_χ the true value of the short-term risk premium which was used to generate the true state vectors.

The simulation results show that the uncertainty of the risk premium parameters increases if the volatility parameters increases. This lead to an economical reason for the difficulty of estimating statistically significant risk premiums. Risk premiums are premiums offered to hold financial risk. In the model, the financial risk is measured using volatilities, i.e., the short- and long-term diffusion coefficients. If the diffusion coefficients are large, market participants will require a large risk premium to compensate for the financial risk of holding the commodity. If the market is sufficiently integrated with other financial markets, no arbitrage conditions will ensure the higher risk premium as capital is allocated towards the highest risk adjusted rate of return. This implies that it will be challenging to obtain estimates of the risk premium parameters that are significant different from zero for any realistic parameter configuration of the model.

It should also be noted that the iterative algorithm cannot contrast the signal from noise, which can increase the uncertainty in parameter estimates in the presence of noisy observations. The Kalman filter approach does not suffer the same shortcoming, since measurement errors are filtered. The two-factor model is sometimes used in real options analysis, where the investment horizon tends to be quite long. In this setting, it is crucial to obtain accurate estimates of the long-run volatility.

Noisy observations may result in an inflated volatility estimate using the iterative routine, as there is no explicit filtering of measurement errors. The Kalman filter can separate the signal from the noise, and is recommended in the case of noisy observations given by low liquidity or other market imperfections.

6.1 Volatility term structure

Both algorithms described above are designed to minimize the squared pricing error. In commodity markets, the volatility of observed forward contracts is often found to be monotonically declining, also known as the Samuelson (1965) effect. If the model describes the data well, the volatility term structure implied from the model should also be in accordance to the observed volatility term structure.

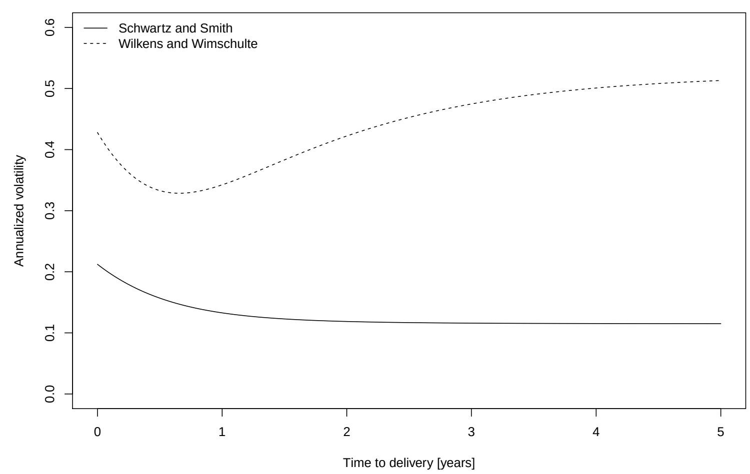

Figure 5: Volatility term structure implied by previous studies on EEX and oil data. The EEX results are not in line with the expected declining volatility term structure.

As shown, the iterative algorithm does not always converge to the true values, but is highly dependent on the starting values. During additional simulation studies, it was revealed that the iterative routine does converge to sensible values for all parameters except the risk premiums (for details, contact the authors). The algorithm however may need a large number of iterations to converge if it is initialized far from the correct values. In order to demonstrate the results the iterative routine can give, we graph the volatility term structure implied by previously published research papers. We choose to illustrate the volatility term structure with two examples from energy markets. Wilkens and Wimschulte (2007) use electricity swap prices from the European Energy Exchange (EEX), while Schwartz and Smith (2000) use data on oil forwards. The first example uses the iterative algorithm, while the last employs the Kalman filter. Figure 5 shows the volatility term structure implied by their results. The data from the EEX results in a u-shaped term structure, not in line with the expected shape. On the other hand, the data from the oil market seem to result in a reasonable volatility term structure.

The unrealistic volatility term structure is taken as an indication of the problems that can occur using the iterative routine. Also the Kalman filter routine does not always provides a good fit to the volatility term structure, see for example Manoliu and Tompaidis (2002) and Cortazar and Naranjo (2006). Note however, that the Kalman filter routine can easily be extended to provide a good fit to the volatility term structure by including the empirical volatility in the observation vector.

7 Conclusion

The paper analyzes the estimation of the risk premium parameters in the two-factor model by Schwartz and Smith (2000) by the iterative algorithm first described by Lucia and Schwartz (2002) and by employing the Kalman filter. The study shows that the risk premiums cannot be estimated with any reasonable precision, even under a controlled experiment such as a simulation study. The iterative procedure is highly dependent on its starting values for the short-term risk premium, while the Kalman filter approach does not return estimates within a reasonable range. This parameter indeterminacy does not affect the robustness of the model for use of valuation purposes, but it does affect its robustness regarding forecasting purposes (Schwartz and Smith (2000)).

The Kalman filter facilitates the calculation of the likelihood, which allows us to estimate parameters using maximum likelihood techniques. We examine the relationship between volatility parameters, sample size and accuracy of parameter estimates using the maximum likelihood estimator. It is shown that the uncertainty of the risk premium parameters increases if the volatility parameters increases. The sample size is varied between 5000 and 250000 observations. The accuracy increases as the sample sizes increases, but even 250000 data points are not enough to achieve precise estimates of risk premium parameters. However, if the risk premium parameters are relatively far from zero, we should be able to tell the sign of the risk premium using a large sample size.

Maximum likelihood estimators are efficient when the sample size tends to infinity, hence our results indicates a lower bound for the uncertainty of the risk premium parameter estimates. The iterative method does not provide any measure of the precision of the parameter estimates, therefore our results can serve as an useful indication of the minimum uncertainty of parameters estimates obtained using this method.

References

Bhar, R., Lee, D., 2011. Time-varying market price of risk in the crude oil futures market. Journal of Futures Markets 31 (8), 779–807.

Cartea, A., Williams, T., 2008. UK gas markets: the market price of risk and applications to multiple interruptible supply contracts. Energy Economics 30 (3), 829–846.

Cortazar, G., Naranjo, L., 2006. An N-factor Gaussian model of oil futures prices. Journal of Futures Markets 26 (3), 243–268.

Cortazar, G., Schwartz, E. S., 2003. Implementing a stochastic model for oil futures prices. Energy Economics 25 (3), 215–238.

De Long, J. B., Shleifer, A., Summers, L. H., Waldmann., R. J., 1990. Noise trader risk in financial markets. The Journal of Political Economy 98 (4), 703–738.

Duffee, G., Stanton, R., 2012. Estimation of dynamic term structure models. The Quarterly Journal of Finance 2 (02).

Durbin, J., Koopman, S. J., 2001. Time Series Analysis by State Space Methods. Oxford University Press.

Helske, J., 2010. KFAS: Kalman filter and smoothers for exponential family state space models. R package version 0.6.0.
URL `http://CRAN.R-project.org/package=KFAS`

Lee, C. M., Shleifer, A., Thaler, R. H., 1990. Closed-end mutual funds. Journal of Economic Perspectives 4 (4), 153–164.

Lucia, J. J., Schwartz, E. S., 2002. Electricity prices and power derivatives: Evidence from the nordic power exchange. Review of Derivatives Research 5 (1), 5–50.

Manoliu, M., Tompaidis, S., 2002. Energy futures prices: term structure models with Kalman filter estimation. Applied Mathematical Finance 9 (1), 21–43.

Nomikos, N. K., Soldatos, O. A., 2010. Modelling short and long-term risks in power markets: Empirical evidence from nord pool. Energy Policy 38 (10), 5671–5683.

Pindyck, R., 2004. Volatility and commodity price dynamics. Journal of Futures Markets 24 (11), 1029–1047.

Samuelson, P. A., 1965. Proof that properly anticipated prices fluctuate randomly. Industrial Management Review 6 (2), 41–49.

Schwartz, E., Smith, J. E., 2000. Short-term variations and long-term dynamics in commodity prices. Management Science 46 (7), 893–911.

Schwartz, E. S., 1997. The stochastic behavior of commodity prices: Implications for valuation and hedging. Journal of Finance 52 (3), 923–974.

Wilkens, S., Wimschulte, J., 2007. The pricing of electricity futures: Evidence from the european energy exchange. Journal of Futures Markets 27 (4), 387–410.