Identifying Uncertainties from Multiple Factors: A Study on Electricity Price*

Wei Wei[a] and **Asger Lunde**[a]

[a] *Aarhus University & CREATES*

November 30, 2015

Abstract

Using a multi-factor model, we separate the uncertainties in electricity spot prices into three risk factors: spikes, base-signals, and stochastic volatility. The model is estimated using the particle Markov chain Monte Carlo method and is applied to the Germany/Austria electricity market. Our results indicate that spike shocks are large and infrequent, and they usually die out within a day, while base-signals are more persistent than spikes. Moreover, the observed clustering of large price movements is explained by stochastic volatility. We apply our estimates in the spot market to the futures market and find that different risk factors have distinct impacts on futures prices. In particular, the base-signal dynamics govern the futures price dynamics. Lastly, we find evidence that the market price of both spike risk and base-signal risk are time-varying.

JEL Codes: C51, G13, Q4

Keywords: Risk factors, Stochastic volatility, Jumps, Particle filters

*The authors acknowledge support from the project entitled "Stochastic and Econometric Analysis of Commodity Markets" funded by Aarhus University Research Foundation. Both authors are affiliated with CREATES - Center for Research in Econometric Analysis of Time Series (DNRF78), funded by the Danish National Research Foundation. Furthermore, Wei Wei is grateful for financial support from the The Danish Council for Independent Research, Social Sciences (4003-00106B/FSE).

1 Introduction

Electricity is an essential part of economic activity. In recent decades, electricity markets have switched from being highly regulated to being more competitive and open. The deregulation has increased efficiency and transparency, but also introduced a great amount of price uncertainty. The spot price of electricity is determined by the orders of market players, placed on the day before delivery, in the day-ahead spot market. Besides the traditional market participants (consumers and producers) , financial institutions such as energy trading companies and banks have begun to play an increasing role. The spot market is therefore a mixture of different types of market players, and likely faces different sources of uncertainties, or risk factors, that have distinct dynamics.

While there is a large literature concerning spot price dynamics, most models are silent about the individual risk factors and how the market prices these different types of risk.[1] Those models that do allow multiple risk factors and admit close-form derivative pricing are typically challenged by their estimation complexity, and hence often rely on moment matching. Therefore, there is a trade-off between the model's ability to disentangle the sources of risk and the difficulty of estimation. The lack of efficiently estimated risk factors, or the omission of multiple risk factors, hinders the usefulness of the current literature for risk management; electricity producers need to know what risk they are facing before they can hedge their risk exposure.

In this paper, we efficiently estimate and disentangle the risk factors in the spot prices using an likelihood-based procedure, particle Markov chain Monte Carlo (PMCMC). The multi-factor model we adopt allows analytical futures prices. Applying our estimates of the spot price dynamics to the futures market, we are able to analyze the risk premium associated with each risk factor.

Our model have three main sources of price uncertainty, or risk factors: spike shocks, base-signal shocks, and volatility shocks. Spike shocks are large and infrequent price movements that die out fast, base-signals describe the "normal" fluctuations, and volatility shocks modulate the magnitude of base-signal shocks. The decomposition of the shocks is not only necessary for capturing important features of the data, it is also intuitive given the production structure of electricity. The supply curve of electricity is characterized by a "hockey-stick" shaped function, gradually increasing below a threshold and steep sloped above the threshold, see e.g. Kanamura and Ōhashi (2007). This shape emerges because electricity producers can not adjust the output above or below certain thresholds without incurring

[1] Weron (2014) provides a comprehensive summary of this literature and comparison of competing models.

adjustment costs for restarting or shutting down a power generation unit, and storing electricity is costly on a large scale. The shape of the supply curve provides the intuition to the different formation of spikes and base-signals. When there is a large increase in demand, producers switch on additional generation units, and electricity prices tend to jump up then revert back in the next day or two. These short-lived price hikes are modeled by spikes. This is in contrast with the base-signals, which account for price fluctuations that occur in the gradually increasing part of the supply curve. Although base-signals are also considered mean-reverting, as the effect of a temporary demand or supply shock eventually dies out, it mean reverts in a slower speed compared to the spikes.

We model the base-signals and spikes using an additive structure of stochastic processes. In particular, we specify a non-Gaussian Ornstein-Uhlenbeck (OU) process for the spikes, and a Gaussian OU process for base-signals. The use of OU process accounts for the mean-reverting behavior of both. Base-signals are also modulated by stochastic volatility, modeled by an independent non-Gaussian Ornstein-Uhlenbeck (OU) process. Our arithmetic factor model is able to capture both the spike behavior and multi-scale autocorrelation structure of spot prices in a straightforward manner. Benth, Kallsen, and Meyer-Brandis (2007) propose an arithmetic factor model where each factor is modeled as a non-Gaussian Ornstein-Uhlenbeck (OU) process. Our use of Gaussian OU process for the base-signals allows negative prices, which is an advantage in electricity market since negative prices can actually occur. Moreover, Benth, Kiesel, and Nazarova (2012) find that finite activity jumps can not fully explain the noise in base-signals for the electricity data and suggest Gaussian shocks instead. The stochastic volatility model we adopt was proposed in Barndorff-Nielsen and Shephard (2001) for the equity market. Benth (2011) and Brix, Lunde, and Wei (2015) consider these stochastic volatile processes for energy market in a geometric factor model. Our arithmetic model is more tractable than geometric models when it comes to derivative pricing. The underlying of an electricity derivative contract is usually the electricity with delivery period of a week or a month. In this case, arithmetic models still admit explicit pricing formula while this property is only present in geometric models when the delivery period of the contract is a day.

The three risk factors (spikes, base-signals and stochastic volatility) in our model have different implications in risk management. If a power plant seeks to reduce its exposure from the volatile market, it can utilize financial derivatives such as futures or forward contracts. The relationship between spot and forward/futures price is particularly interesting in the case of electricity because it is a non-storable

commodity (see for instance Bessembinder and Lemmon, 2002 and Longstaff and Wang, 2004). We follow the strand of literature which obtains futures price based on the arbitrage pricing theory (Benth, Benth, and Koekebakker, 2008; Veraart and Veraart, 2014). The no-arbitrage futures price is equal to the expected spot prices under a risk neutral measure, and it can be decomposed into a predicted spot price (expected spot under the physical measure) and a market risk premium. If spot prices follow a random walk, the predicted spot prices is simply the current price. For a mean-reverting process, the predicted spot is a weighted average between current price and the long-run mean, where the weights depend on the speed of mean-reversion and time to maturity. Since spikes mean-revert faster than base-signals, the current level of spike process contributes less to futures price than the current level of base-signals. In the arithmetic model, volatility does not directly affect the predicted spot, but it impacts spikes identification, which in turn affects futures price.

The market risk premium, defined as the difference between the futures price and the predicted spot price, measures what the market demands to be compensated for bearing risk. Since the market has different aversions to different sources of uncertainty, spike shock, base-signal shock and volatility shock could all have their own price of risk. Separating the risk premium associated with different risk factors has been an active area of research in the equity market, see for instance Pan (2002) and Bollerslev and Todorov (2011). The risk premium in the electricity markets can be either positive or negative. Benth, Álvaro Cartea, and Kiesel (2008) connects the sign of the risk premium to the risk preference of market players. If the producers are more risk-averse about base-signal shocks than the consumers, there will be excess supply on the futures market. The excess supply lowers the futures prices compared to predicted spot prices, which implies negative base-signal risk premium. On the other hand, consumers would want to be protected against positive spikes, hence spikes are likely to increase futures price, in other words, spike risks are likely to carry a positive price of risk.

We estimate our spot price models using the particle Markov chain Monte Carlo (PMCMC) methods developed in Andrieu, Doucet, and Holenstein (2010). Flury and Shephard (2011) illustrates the usefulness of this method in economic models. Brix, Lunde, and Wei (2015) extend the method to the estimation of factor models with stochastic volatility. In Brix, Lunde, and Wei (2015), the logarithm of spike process are driven by a compound Poisson process with normally distributed jump sizes. This specification allows them to use Rao-Blackwellization and obtain an efficient particle filter. We adapt their method to accommodate non-Gaussian jump size distribution in the spike process. PMCMC is a

relatively easy estimation method for non-Gaussian OU stochastic volatility models, and parameters are estimated from a likelihood-base approach. In addition, the latent risk factors are obtained directly from the particle filters while accounting for parameter uncertainty.

The contribution of this paper is threefold. First, we separate spikes, base-signals and stochastic volatility in the spot prices using a likelihood-based approach. Second, we analyze how risk factors affect the futures prices and their risk premium. Third, we contribute to the jump filtering literature when the jump size distribution is not Gaussian.

The rest of the paper is organized as follows. Section 2 describes our model for spot prices as well as some nested alternatives we use to help identify the significance of each risk factor. In section 3, we detail how to compute futures prices based on the arithmetic factor model. Section 4 outlines the estimation method. An application with Germany/Austria electricity prices in presented in Section 5. Section 6 concludes.

2 Models for the Spot Market

Electricity prices have a trend and seasonality component which is commonly considered deterministic, see e.g. Koopman, Ooms, and Carnero (2007) and Janczura, Trück, Weron, and Wolff (2013). We adopt a multiplicative specification for the trend and seasonality component $\Lambda(t)$ and the stochastic component $Z(t)$, i.e., the spot price $S(t)$ equals $\Lambda(t)Z(t)$.

We focus on the stochastic component $Z(t)$ which accounts for the risk factors: spikes, base-signals and stochastic volatility. Several classes of models have been proposed to separate spikes and base-signals. Following Schwartz (1997) who models the mean-reverting behavior in commodity prices using a Ornstein-Uhlenbeck (OU) process, Cartea and Figueroa (2005) suggest a mean-reverting jump-diffusion driven by both a Brownian motion and a compound Poisson process. Knittel and Roberts (2005) refine the model to allow the jump intensity to depend on the season and time of day. However, jump-diffusion models impose the same mean-reversion speed for base-signals and spikes. Regime-switching models offer a more flexible way to deal with spikes. Spike regime and base-signal regime can be driven by the same Gaussian shocks but differ in parameters (Huisman and Mahieu, 2003; Mount, Ning, and Cai, 2006), or they can have different shock distributions (Bierbrauer, Menn, Rachev, and Trück, 2007; Weron, Bierbrauer, and Trück, 2004). The latter allows spike regime to have a positive distribution such as lognormal or exponential, but it complicates the estimation procedure. Also, to

capture the different persistence of a spike shock and a base-signal shock, these models need three regimes with restricted transition probabilities. Instead, we adopt the arithmetic factor model so that spikes and base-signals are explained by a simple additive structure.

Stochastic volatility is an important feature in financial markets in general. Heston (1993) suggest the square-root volatility process whose affine structure allows for close-form derivative pricing. The Heston model has only diffusion terms, and it sometime requires undesirably high "volatility of volatility" parameter when calibrated to market data (Bates, 2000). Barndorff-Nielsen and Shephard (2001) propose the class of non-Gaussian OU processes. These processes are driven up entirely by jumps then tails off exponentially. Benth (2011) compare the non-Gaussian OU model to Heston model and find that non-Gaussian OU model offers more flexible forward dynamics. We follow Benth (2011) and adopt the non-Gaussian OU processes to model the base-signal variance.

We describe our continuous-time model for spikes, base-signals and stochastic volatility in subsection 2.1. Subsection 2.2 and 2.3 outline the discretized model and a few nested models that we also consider.

2.1 The Arithmetic Model

The stochastic component $Z(t)$ is composed of the base-signal $X(t)$ and the spike process $Y(t)$ in an additive structure, $Z(t) = X(t) + Y(t)$. Both $X(t)$ and $Y(t)$ are modeled by the class of OU processes:

$$\begin{aligned} \mathrm{d}X(t) &= -\alpha_x(X(t) - \mu_x)\mathrm{d}t + \sigma(t)\mathrm{d}B(t) \\ \mathrm{d}Y(t) &= -\alpha_y Y(t)\mathrm{d}t + \mathrm{d}I(t). \end{aligned}$$

We impose $\mu_x = 1$ as μ_x and the constant in $\Lambda(t)$ are not separately identifiable. The spike process $Y(t)$ is driven by the compound Poisson process $I(t)$ with jump intensity λ_J. We specify jump size to follow an exponential distribution, $\mathrm{Exp}(\mu_J)$. This specification leads to a Gamma OU process. In other words, the marginal distribution of $Y(t)$ follows an Gamma distribution, $Y(t) \sim \mathrm{G}(\lambda_J/\alpha_y, \mu_J)$.

Variance in the base-signal is modeled by an superposition of inverse Gaussian OU processes. That is, $\sigma^2(t) = \sigma_1^2(t) + \sigma_2^2(t)$, and $\sigma_i^2(t)$ solves

$$\mathrm{d}\sigma_i^2(t) = -\lambda_i \sigma_i^2(t)\mathrm{d}t + \mathrm{d}L_i(\lambda_i t),$$

with marginal distribution of $\sigma_i^2(t)$ follows an inverse Gaussian (IG) distribution, $\sigma_i^2(t) \sim \text{IG}(\delta_i, \gamma)$. We restrict γ to be the same for the two variance components such that the superposition of $\sigma_1^2(t)$ and $\sigma_2^2(t)$ still follows IG marginal, $\sigma^2(t) \sim \text{IG}(\delta_1 + \delta_2, \gamma)$. The different mean-reversion speed, λ_1 and λ_2, give rises to flexible autocorrelation structure in the variance process,

$$\text{Corr}(\sigma^2(t), \sigma^2(t-u)) = w_1 \exp(-\lambda_1 u) + w_2 \exp(-\lambda_2 u),$$

where $w_i = \delta_i/(\delta_1 + \delta_2)$.

Both $\sigma_i^2(t)$ and $Y(t)$ belong in the class of Non-Gaussian OU processes. However, the IG-OU process $\sigma_i^2(t)$ allows for infinitely many jumps in finite time intervals, whereas the Gamma-OU process $Y(t)$ is characterized by finite activity jumps. Spike are rare events, hence the adoption of Poisson jumps. On the other hand, IG-OU allows more variation in the volatility path, and this helps explain the kurtosis of base-signal shocks.

2.2 Discretization

The discretization of the base-signal $X(t)$ leads to an AR(1) model:

$$X_{t+1} = e^{-\alpha_x} X_t + \epsilon_{t+1}, \tag{1}$$

where $\epsilon_{t+1} \sim N\left(0, \sigma^2(t) a_2\right)$ and $a_2 = \left(1 - e^{-2\alpha_x}\right)/(2\alpha_x)$. This discretization utilizes the integral representation of an OU process and the approximation $\int_t^{t+1} e^{-2\alpha(t+1-s)} \sigma^2(s) \mathrm{d}s \approx \sigma^2(t) a_2$.

For the Gamma-OU process $Y(t)$, we have the following shot noise representation:

$$Y(t+\Delta) = e^{-\alpha_y \Delta} Y(t) + \sum_{j=1}^{n} \xi_j e^{-\alpha_y u_j \Delta},$$

where the number of jumps n follows a Poisson distribution, $n \sim \text{Poiss}(\lambda_J \Delta)$. The jump size ξ_j follows an exponential distribution, $\xi_j \sim \text{Exp}(\mu_J)$, and $u_j \sim \text{Unif}(0,1)$, for $j = 1, ..., n$. Let $\phi = \xi_j e^{-\alpha_y u_j \Delta}$, we can derive an analytical expression for the density of ϕ:

$$p(\phi) = \frac{1}{\alpha_y \Delta \phi} \left(\exp\left(-\frac{\phi}{\mu_J}\right) - \exp\left(-\frac{\phi}{\mu_J c}\right) \right), \tag{2}$$

where $c = e^{-\alpha_y \Delta}$.Consider $\Delta = 1$. If λ_J is small, we can approximate n by a Bernoulli random

variable, $J_{t+1} \sim \text{Bern}(\lambda_J)$. In other words, we assume at most one jump occurs per day. In this case, we obtain the following discretization,

$$Y_{t+1} = e^{-\alpha_y} Y_t + J_{t+1}\phi_{t+1}, \tag{3}$$

where ϕ is distributed according to equation (2). The observed price is the sum of the two factors,

$$Z_{t+1} = X_{t+1} + Y_{t+1}. \tag{4}$$

For the volatility process, it can be simulated recursively from

$$\sigma_i^2(t+1) = e^{-\lambda_i}\sigma_i^2(t) + \epsilon_i^\sigma(t+1), \tag{5}$$

where $\epsilon_i^\sigma(t+1) \stackrel{d}{=} e^{-\lambda_i}\int_0^1 e^{\lambda_i u} dL_i(\lambda_i u)$. The density of $\epsilon_i(t+1)$ does not admit explicit expression, but our estimation procedure only requires us to simulate $\epsilon_i(t+1)$. Equation (1), (3), (4) and (5) form our proposed model. It features a two-factor price dynamics with jumps and a two-factor stochastic volatility. We denote this model TFJ-S2V.

2.3 Nested Models

We also consider a few nest models to help assess the importance of each factor. There models are derived from the proposed model TFJ-S2V with the restrictions presented in Table 1:

Table 1: Model Overview

Model	Restrictions
TFJ-S2V	None
TFJ-SV	$\lambda_1 = \lambda_2$
SFJ-SV	$\alpha_x = \alpha_y$ and $\lambda_1 = \lambda_2$
SF-SV	$\lambda_J = 0$ and $\lambda_1 = \lambda_2$
TFJ	$\sigma^2(t) = \sigma^2$
SFJ	$\alpha_x = \alpha_y$ and $\sigma^2(t) = \sigma^2$

Specifically, we implement a two-factor model with one-factor stochastic volatility, TFJ-SV; a single factor jump-diffusion model with stochastic volatility, SFJ-SV; a single factor diffusion model with stochastic volatility, SF-SV, which corresponds to the model in Benth (2011); a two-factor model with constant volatility, TFJ, which is similar to the model in Benth, Kallsen, and Meyer-Brandis (2007)

but with diffusive base-signals; and a single-factor jump-diffusion model with constant volatility, SFJ, which is the model proposed in Cartea and Figueroa (2005).

3 Pricing Futures Contracts

Futures contracts are widely traded financial instruments for hedging risk. Electricity futures have the distinctive feature that the delivery is usually made over a period of time such as a week or a month. Let $f(t,T)$ denote the futures price[2] at time t with maturity T, and $F(t,T_1,T_2)$ denote the futures price at time t with delivery period $[T_1,T_2]$. Suppose that the financial settlement of the contract takes place at the end of the delivery period, we obtain the following relationship,

$$F(t,T_1,T_2) = \frac{1}{T_2-T_1}\int_{T_1}^{T_2} f(t,T)\mathrm{d}T.$$

We first derive $f(t,T)$ using the arbitrage pricing theory, then compute $F(t,T_1,T_2)$ and the risk premium. Detailed derivations for a purely additive model can be found in Benth, Álvaro Cartea, and Kiesel (2008) and Benth, Benth, and Koekebakker (2008).

Since futures contracts are zero-cost securities, the arbitrage-free futures price is given by

$$f(t,T) = \mathrm{E}_Q\left[S(T)|\mathcal{F}_t\right] = \Lambda(T)\mathrm{E}_Q\left[X(T)+Y(T)|\mathcal{F}_t\right],$$

where Q is risk-neutral probability measure. The risk-neutral measure Q and physical measure P are equivalent measures, and we let $D(t)$ denote the density process of the Radon-Nikodym derivative $\mathrm{d}Q/\mathrm{d}P$. In an incomplete market such as the electricity market, the choice of Q is not unique. We follow Benth, Álvaro Cartea, and Kiesel (2008) and choose a Girsanov transform for $B(t)$ and an Esscher transform for $I(t)$. Specifically,

$$D(t) = \frac{\mathrm{d}Q}{\mathrm{d}P}|\mathcal{F}_t = D_B(t)\times D_I(t),$$

[2]We do not distinguish between futures contract which is mark-to-market with the forward prices.

with

$$D_B(t) = \exp\left(\int_0^t \frac{\theta_B}{\sigma(u)} \mathrm{d}B(u) - \frac{1}{2}\int_0^t \frac{\theta_B^2}{\sigma^2(u)} \mathrm{d}u\right)$$
$$D_I(t) = \frac{\exp\left(\int_0^t \theta_I \mathrm{d}I(u)\right)}{\mathrm{E}\left(\exp\left(\int_0^t \theta_I \mathrm{d}I(u)\right)\right)}.$$

The diffusion term $B(t)$ and the jump process $I(t)$ under Q can be obtained from a shift in the drift term:

$$\mathrm{d}B(t) = \mathrm{d}B^Q(t) + \frac{\theta_B}{\sigma(t)}\mathrm{d}t$$
$$\mathrm{d}I(t) = \mathrm{d}I^Q(t) + \eta_I \mathrm{d}t.$$

Here θ_B is the market price of diffusion risk per unit of volatility, and B^Q is a Brownian motion under Q. The Esscher transform preserves the Levy structure of $I(t)$, and $I^Q(t)$ is a Compound Poisson process under Q, with $\mathrm{E}_Q(I^Q(t)) = \mu_J \lambda_J t$. We denote η_I the price of jump risk , and it is related with θ_I through $\eta_I = \lambda_J \mu_J \theta_J / (1 - \mu_J \theta_I)$.

Using the the integral representation of an OU process, we obtain:

$$\begin{aligned} f(t,T) =& \Lambda(T)\left(X(t)e^{-\alpha_x(T-t)} + Y(t)e^{-\alpha_y(T-t)} + \mu_x\left(1 - e^{-\alpha_x(T-t)}\right)\right. \\ &\left. + \frac{\theta_B}{\alpha_x}\left(1 - e^{-\alpha_x(T-t)}\right) + (\lambda_J\mu_J + \eta_I)\frac{1}{\alpha_y}\left(1 - e^{-\alpha_y(T-t)}\right)\right) \\ =& \Lambda(T)\left(\mathrm{E}_P\left(X(T)|\mathcal{F}_t\right) + \mathrm{E}_P\left(Y(T)|\mathcal{F}_t\right) + \frac{\theta_B}{\alpha_x}\left(1 - e^{-\alpha_x(T-t)}\right) + \frac{\eta_I}{\alpha_y}\left(1 - e^{-\alpha_y(T-t)}\right)\right). \quad (6) \end{aligned}$$

The futures price formula has four components: the expected base-signal, the expected spike level, the diffusion risk premium, and the jump risk premium. The expected base-signal is a weighted average between current base-signal $X(t)$ and its long-run mean μ_x. As $(T-t)$ decrease, or as the futures contract gets closer to maturity, the weight of $X(t)$ increases, until the expectation converges to $X(t)$ at maturity. The same dynamics also applies to the expect spike process. However, since spikes mean revert faster, the weight of $Y(t)$ tends to be smaller than $X(t)$. In other words, the current level of base-signal matters more for futures prices than the current spike level. The first two components and the seasonality function $\Lambda(T)$ forms the futures price under P, $f^P(t,T)$, also referred to the expected

spot price. The difference between $f(t,T)$ and $f^P(t,T)$ is the market risk premium. If both θ_B and η_I are constant, the risk premium is a deterministic function of $(T-t)$. Also, the risk premium converges to zeros as $(T-t) \to 0$, and the futures price converges to the spot price.

For futures contract with delivery period $[T_1, T_2]$, we define

$$\bar{\Lambda}(T_1,T_2) = \frac{1}{T_2-T_1}\int_{T_1}^{T_2}\Lambda(T)\mathrm{d}T$$
$$\bar{\alpha}(\alpha,t,T_1,T_2) = \frac{\int_{T_1}^{T_2}\Lambda(T)e^{-\alpha(T-t)}\mathrm{d}T}{\int_{T_1}^{T_2}\Lambda(T)\mathrm{d}T},$$

then the futures price under P is the predicted spot price averaged over $[T_1, T_2]$:

$$\begin{aligned} F^P(t,T_1,T_2) = \bar{\Lambda}(T_1,T_2)\Big(&X(t)\bar{\alpha}(\alpha_x,t,T_1,T_2) + \mu_x\left(1-\bar{\alpha}(\alpha_x,t,T_1,T_2)\right) \\ &+Y(t)\bar{\alpha}(\alpha_y,t,T_1,T_2) + \frac{\lambda_J\mu_J}{\alpha_y}\left(1-\bar{\alpha}(\alpha_y,t,T_1,T_2)\right)\Big). \end{aligned} \tag{7}$$

The market risk premium is defined as the difference between the futures price under risk-neutral measure Q and the futures price under P measure:

$$\begin{aligned} R(t,T_1,T_2) &= F(t,T_1,T_2) - F^P(t,T_1,T_2) \\ &= \bar{\Lambda}(T_1,T_2)\left(\frac{\theta_B}{\alpha_x}\left(1-\bar{\alpha}(\alpha_x,t,T_1,T_2)\right) + \frac{\eta_I}{\alpha_y}\left(1-\bar{\alpha}(\alpha_y,t,T_1,T_2)\right)\right). \end{aligned} \tag{8}$$

Equation (7) and (8) have a similar structure as the futures price with point delivery in equation (6). The predicted spot is a weighted average of current $X(t)$ and $Y(t)$ and the long-run mean of $X(t)$ and $Y(t)$, and the risk premium is a deterministic function. However, $\bar{\alpha}$ does not approach zero as $T_1 - t \to 0$. In other words, the futures price does not converge to the spot price because of the uncertainty in the period $[T_1, T_2]$.

4 Estimation Method

We adopt the PMCMC algorithm, in particular the particle marginal Metropolis-Hastings (PMMH) sampler, to estimate the proposed model. The PMMH sampler is consist of a particle filter or sequential Monte Carlo (SMC) part and a MH part. The SMC part obtains an unbiased estimate of the likelihood conditional on model parameters, while the MH part obtains the posterior distribution of parameters

and latent factors. The MH algorithm is easy to implement and has been widely used in Bayesian likelihood-based inference. However, the likelihood obtained from SMC is a random variable, and the performance of MH is highly dependent on the variance of likelihood estimate (see Pitt, dos Santos Silva, Giordani, and Kohn, 2012). Hence, it is important to have an efficient SMC algorithm, for instance the auxiliary particle filter by Pitt and Shephard (1999), that minimizes the variance of the likelihood estimates.

SMC is a simulated-based algorithm that can compute likelihood that involves high-dimensional integral, which is common when a model has latent variables. The basic element of SMC involves using importance sampling to obtain a filtering distribution of the latent variables sequentially. The variance of the likelihood depends on the variance of the importance weights, which in turn depends on the proximity of importance density to the target filtering density (see Creal, 2012 for an extensive review on SMC). If the filtering distribution can be sampled from directly, we have the fully adapted SMC. If part of the latent variables can be sampled directly conditional on the other latent variables and the data, we can use a technique called Rao-Blackwellization. Brix, Lunde, and Wei, 2015 apply Rao-Blackwellization to a geometric factor model for energy prices. In their model, the spike shock has normally distributed jump size with Bernoulli distributed jump probability, and its filtering distribution can be sampled directly conditional on the stochastic volatility.

However, the spike shocks in our model does not follow a standard distribution, so we can not use Rao-Blackwellization to deal with jumps directly. Instead, we use an importance prior density where it is possible to adapt the jumps, then reweigh the particles to correct the difference. The idea shares some similarity to Durbin and Koopman (1997), where an approximating model is used to compute likelihood from kalman filter, then the residuals from a simulation smoother is used in importance sampling to correct the likelihood. Our method choose the approximating model in a sequential way, and it can be used for different variations of jump-diffusion models and stochastic volatility. We illustrate the SMC procedure using the jump-diffusion model SFJ in section 4.1, and briefly describe how to adapt it for the TFJ-S2V model in section 4.2. Section 4.3 outlines the PMMH algorithm.

4.1 SMC for SFJ

In the SFJ model, the state variables at time $t+1$, which we denote $K_{t+1} = \{\phi_{t+1}, J_{t+1}\}$, does not depend on past states, $K_{1:t}$. The measurement equation in this model is

$$Z_{t+1} = e^{-\alpha} Z_t + \epsilon_{t+1} + \phi_{t+1} J_{t+1},$$

where the prior distribution of the jump size parameter $p(\phi_{t+1})$ is given in equation (2), and $\epsilon_{t+1} \sim N\left(0, \sigma^2 a_2\right)$. We first specify an importance prior density for ϕ_{t+1}, $g(\phi_{t+1}) = \text{Exp}(\mu_\phi)$. Based on this importance prior, we evaluate an "importance likelihood" of Z_{t+1}, which we denote $g(Z_{t+1}|Z_t)$. Specifically, if $J_{t+1} = 1$, the sum of ϵ_{t+1} and ϕ_{t+1} follows a normal modified exponential distribution, so we can evaluate $g(Z_{t+1}|Z_t, J_{t+1} = 1)$. If $J_{t+1} = 0$, $g(Z_{t+1}|Z_t, J_{t+1} = 0)$ is simply normal. Hence, the density $g(Z_{t+1}|Z_t)$ based the importance prior $g(\phi_{t+1})$ can be computed from

$$g(Z_{t+1}|Z_t) = g(Z_{t+1}|Z_t, J_{t+1} = 1)\lambda_J + g(Z_{t+1}|Z_t, J_{t+1} = 0)(1 - \lambda_J). \tag{9}$$

When the state variables include jumps, the measurement density can be highly peaked. Hence, it is important to incorporate information from the observation when propagating particles. We generate N particles, $J_{t+1}^{(i)}$, $i = 1, ..., N$, using an adapted importance filtering density[3], $J_{t+1}^{(i)} \sim g(J_{t+1}|Z_{1:t+1}) = \text{Bern}(\lambda_J^*)$, where

$$\lambda_J^* = \frac{g(Z_{t+1}|Z_t, J_{t+1} = 1)\lambda_J}{g(Z_{t+1}|Z_t)}.$$

Next, the jump sizes $\phi_{t+1}^{(i)}$ are generated conditional on $J_{t+1}^{(i)} = 1$ and Z using

$$\begin{aligned} g(\phi_{t+1}|Z_{1:t+1}, J_{t+1}^{(i)} = 1) &= \frac{g(Z_{t+1}|Z_t, \phi_{t+1}, J_{t+1}^{(i)} = 1) g(\phi_{t+1}|J_{t+1}^{(i)} = 1)}{g(Z_{t+1}|Z_t, J_{t+1}^{(i)} = 1)} \\ &\sim \text{TN}(\mu_\phi^*, a_2\sigma^2), \end{aligned}$$

where TN denotes truncated normal, and $\mu_\phi^* = Z_{t+1} - e^{-\alpha} Z_t - a_2\sigma^2/\mu_\phi$. If $J_{t+1}^{(i)} = 0$, we do not gain any information about ϕ_{t+1} from Z, so $g(\phi_{t+1}|Z_{1:t+1}, J_{t+1}^{(i)} = 0) = g(\phi_{t+1}|J_{t+1}^{(i)} = 0) = p(\phi_{t+1})$.

If $p(\phi_{t+1}) = g(\phi_{t+1})$, we have a fully adapted SMC algorithm as $g(Z_{t+1}|Z_t)$ is analytically tractable.

[3] We use g to denote the the importance prior/likelihood/filtering, and p to denote the true prior/likelihood/filtering density.

However, this is not the case in our model, so we need adjust the difference between $p(\phi_{t+1})$ and the importance prior $g(\phi_{t+1})$. The importance weight is given by

$$\begin{aligned}\tilde{\omega}^{(i)} &= \frac{p(\phi^{(i)}, J^{(i)}|Z_{1:t+1})}{g(\phi^{(i)}, J^{(i)}|Z_{1:t+1})} \\ &= \frac{p(Z_{t+1}|Z_{1:t}, \phi_{t+1}^{(i)}, J_{t+1}^{(i)})p(\phi_{t+1}^{(i)})p(J_{t+1}^{(i)})}{g(Z_{t+1}|Z_t, \phi_{t+1}^{(i)}, J_{t+1}^{(i)})g(\phi_{t+1}^{(i)}|J_{t+1}^{(i)})g(J_{t+1}^{(i)})} \frac{g(Z_{t+1}|Z_{1:t})}{p(Z_{t+1}|Z_{1:t})}.\end{aligned}$$

We have specified $g(J_{t+1}) = p(J_{t+1})$ when generating $J_{t+1}^{(i)}$ and $\phi_{t+1}^{(i)}$. Also, $g(Z_{t+1}|Z_{1:t}, \phi_{t+1}^{(i)}, J_{t+1}^{(i)}) = p(Z_{t+1}|Z_{1:t}, \phi_{t+1}^{(i)}, J_{t+1}^{(i)})$ since Z_{t+1} has the same normal distribution conditional on the jumps. Hence, we are left with the incremental weights

$$\omega_{t+1}^{(i)} = \frac{p(\phi_{t+1}^{(i)})}{g(\phi_{t+1}^{(i)}|J_{t+1}^{(i)})} g(Z_{t+1}|Z_{1:t}),$$

and

$$\hat{p}(Z_{t+1}|Z_{1:t}) = \frac{1}{N}\sum_i \omega_{t+1}^{(i)}.$$

In summary, our algorithm specifies an importance prior density for the state variables that can be fully adapted. Using this importance prior density, we obtain an importance likelihood and an importance filtering density. Then the importance weight, defined as the ratio of target filtering density over importance filtering density, simplifies to the ratio of true prior density over the importance prior, multiplied by the importance likelihood. Note that the importance likelihood is constant for all particles, hence, the variance of the weights is significantly reduced.

4.2 SMC for TF-SVJ

We modify the SMC algorithm in Brix, Lunde, and Wei, 2015 using the idea illustrated in the last subsection to estimate TFJ-S2V model. The two-factor model admits a non-markovian measurement equation:

$$Z_{t+1} = e^{-\alpha_x} Z_t + Y_{t+1} - e^{-\alpha_x} Y_t + \epsilon_{t+1}.$$

Given particles $\{Y_t^{(i)}, \sigma^2(t-1,i)\}$ that approximates the filtering distribution, $p(Y_t, \sigma^2(t-1)|Z_{1:t})$, we first propagate the variance using its transition density, $\sigma^2(t,i) \sim p(\sigma^2(t)|\sigma^2(t-1,i))$. Then, conditional on $\sigma^2(t,i)$ and $Y_t^{(i)}$, we can evaluate $g(Z_{t+1}|Y_t^{(i)}, Z_{1:t}, \sigma^2(t,i))$ and propagate $J_{t+1}^{(i)} \sim g(J_{t+1}|Y_t^{(i)}, \sigma^2(t,i), Z_{1:t+1})$ and $\phi_{t+1}^{(i)} \sim g(\phi_{t+1}|Y_t^{(i)}, J_{t+1}{}^{(i)} = 1, \sigma^2(t,i),, Z_{1:t+1})$ based on the importance prior as in section 4.1. The incremental weights are given by

$$\begin{aligned}\omega_{t+1}^{(i)} &= \frac{p(Z_{t+1}|Z_{1:t}, Y_t^{(i)}, \phi_{t+1}^{(i)}, J_{t+1}^{(i)}, \sigma^2(t,i))p(\phi_{t+1}^{(i)})p(J_{t+1}^{(i)})}{g(\phi_{t+1}^{(i)}|Y_t^{(i)}, J_{t+1}^{(i)}, \sigma^2(t,i), Z_{1:t+1})g(J_{t+1}^{(i)}|Y_t^{(i)}, \sigma^2(t,i), Z_{t+1})} \\ &= \frac{p(\phi_{t+1}^{(i)})}{g(\phi_{t+1}^{(i)}|J_{t+1}^{(i)})} g(Z_{t+1}|Y_t^{(i)}, Z_{1:t}, \sigma^2(t,i)),\end{aligned}$$

with

$$\hat{p}(Z_{t+1}|Z_{1:t}) = \frac{1}{N}\sum_i \omega_{t+1}^{(i)}.$$

Let $\hat{\omega}_{t+1}^{(i)}$ denote the normalized weights, $\hat{\omega}_{t+1}^{(i)} = \omega_{t+1}^{(i)} / \sum_i \omega_{t+1}^{(i)}$. We resample the particles using $\hat{\omega}_{t+1}^{(i)}$ and let $Y_{t+1}^{(i)} = e^{-\alpha_y} Y_t^{(i)} + \phi_{t+1}^{(i)} J_{t+1}^{(i)}$, then $\{Y_{t+1}^{(i)}, \sigma^2(t,i)\}$ approximates the filtered distribution, $p(Y_{t+1}, \sigma^2(t)|Z_{1:t+1})$.

4.3 The PMMH algorithm

We outline the PMMH algorithm in Andrieu, Doucet, and Holenstein, 2010 with a random-walk MH sampler. For $g = 1, ..., G$, where G is the number of MCMC iterations:

1. Sample $\theta^{g+1} \sim q(\theta^{g+1}; \theta^g, \Sigma^g)$. We set the proposal density q to be truncated multivariate normal with mean θ^g and covariance matrix Σ^g.

2. Run the SMC algorithm to obtain the likelihood estimate $\hat{p}(Z_{1:T}|\theta^{g+1})$, and draw one realization of the state variables $K_{1:T}^{g+1}$.

3. Accept θ^{g+1} and $K_{1:T}^{g+1}$ with probability:

$$\alpha^{g+1} = \min\left\{\frac{\hat{p}(Z_{1:T}|\theta^{g+1)})p(\theta^{g+1})q(\theta^g; \theta^{g+1}, \Sigma^g)}{\hat{p}(Z_{1:T}|\theta^g)p(\theta^g)q(\theta^{g+1}; \theta^g, \Sigma^g)}, 1\right\},$$

   where $p(\theta)$ is the prior density of the model parameters.

4. The covariance matrix Σ^g is adjusted using the global adaptive scaling in Andrieu and Thoms (2008).

5 Empirical Application

5.1 The Electricity Market

We analyze the spot electricity price for the Germany/Austria market, also called the Phelix index, from the European Energy Exchange (EEX). In the spot market, market players place their bids for electricity with delivery in the following day in a 24 hour interval. For each hour of the next day, an equilibrium price is determined from an aggregated supply and demand curve based on the submitted bids. The base-load index is then computed as the average of the 24 hourly prices.

In the Germany/Austria market, major producers of electricity include nuclear, lignite, hard coal and natural gas plants. The "hockey-stick" shaped supply curve reflect not only the different marginal cost of production technology, but also the power plants' inflexibility to increase operation capacity. On the other hand, it is costly for some power plants, especially the ones with lower marginal cost of production such as nuclear and lignite plants, to decrease production. Since 2008, the EEX allows negative spot price. If the demand is too low, producers can pay consumer to avoid shutting down a power generation unit.

EEX also has an active futures market for the Phelix index. Market participates can enter a long term contract with weekly, monthly, quarterly or yearly delivery. The maturities of the contracts ranges from current week to next 6 years.

The data we use is the Phelix spot prices from 01/01/2002 to 10/21/2008. The cutoff point of 2008 is mostly due to futures price availability. Also, EEX has gone though major changes in the past few years. For example, Germany, France and Benelux started market coupling in 2010 to increase market integration and decrease price differences cross-border. In 2014, a price coupling system is launched covering much of North Western Europe including the Nordic countries and Great Britain. Time-series modeling should account for these structural changes in the market, but we leave this issue to future research.

5.2 Trend and Seasonality

The spot price $S(t)$ has a deterministic trend and seasonality component $\Lambda(t)$. We filter out this component before estimation following the procedure in Benth, Kiesel, and Nazarova (2012). We adopt the multiplicative structure, $S(t) = \Lambda(t)Z(t)$, and the following function for $\ln \Lambda(t)$:

$$\ln \Lambda(t) = a_0 + a_1 t + a_2 \cos(\frac{2\pi t}{261} + a_3).$$

Weekend prices are excluded from the series, leaving 1776 observations. To remove the effect of outliers on the estimation of $\ln \Lambda(t)$, we first obtain a smoothed log prices using a moving average of 21 days. Next, if the absolute difference between the smoothed and original log price is larger than 0.5, it is considered as an outlier and the smoothed value is substituted in. The $\ln \Lambda(t)$ we adopt has a trend and a 12 month period. We also fit the data with a 6 month period but the coefficients turn out to be insignificant at the 5% level. The detrended and deseasonalized spot price, $Z(t)$, is computed by inserting back the detected outliers and divided by the fitted $\Lambda(t)$ function. Figure 1 display the detrended and deseasonalized spot prices, $Z(t)$.

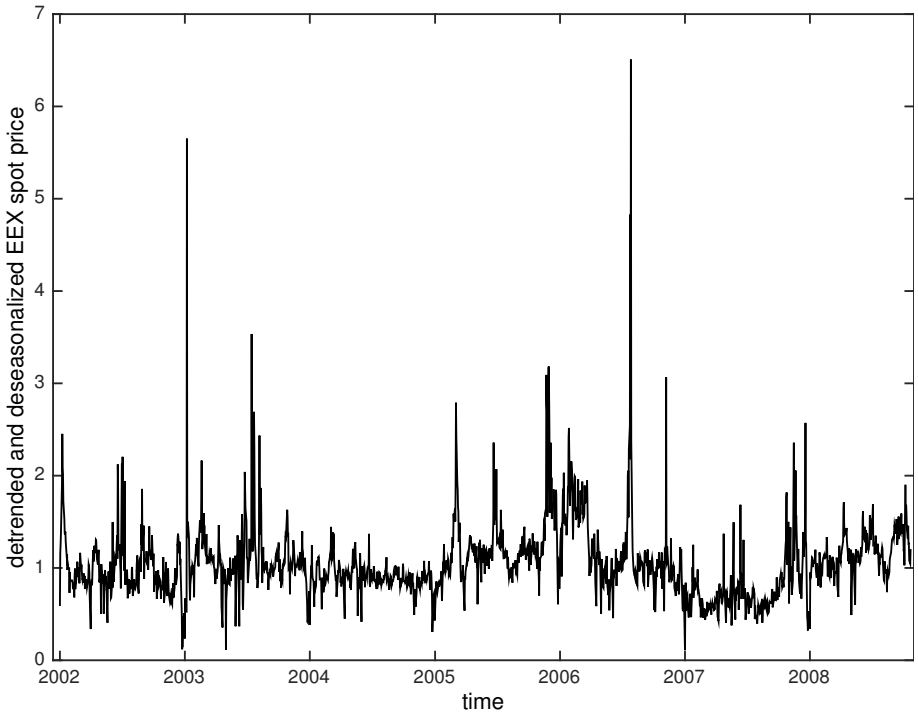

Figure 1: Detrended and Deseasonalized Spot Prices

5.3 Estimation Results

We apply the PMMH estimation algorithm from section 4 to the six models in Table 1. The same prior is specified for models with overlapping parameters. For the jump size parameter, we choose $\mu_J \sim G(2,2)$ which places low probability on small jumps. Other parameters have a diffusive prior. For the mean reversion parameters, we specify $\alpha_x \sim G(1, 0.5)$ and $\alpha_y \sim G(1,1)$. For stochastic volatility, the priors are $\lambda_1 \sim G(1, 0.2)$, $\lambda_2 \sim G(1,1)$, and $\delta_1, \delta_2, \gamma \sim G(1, \sqrt{50})$. For models with constant volatility, we choose $\sigma^2 \sim G(1,2)$. Finally, the priors for jump intensity is $\lambda_J \sim \text{Beta}(1,2)$. For the SMC part of the algorithm, the number of particles is set to three time the number of observations. For the MH part of the algorithm, we start with a preliminary run which uses an adaptive MCMC algorithm as in Brix, Lunde, and Wei (2015), then "freeze" the covariance matrix and scaling factor and run a further 20000 iterations to obtain the posterior distributions of θ and K.

To assess model fit, we report the likelihood evaluated at the posterior mean and median of θ from the SMC procedure with 300,000 particles. The large number of particles is used to minimize the randomness of likelihood estimated from SMC. We also compute marginal likelihood $p(Z|M)$ using the formula in Newton and Raftery (1994):

$$\hat{p}(Z|M) = \frac{\delta_p G/(1-\delta_p) + \sum_{g=1}^{G} p(Z|\theta^g, M)/\left(\delta \hat{p}(Z|M) + (1-\delta_p)\, p(Z|\theta^g, M)\right)}{\delta_p G/\left((1-\delta_p)\hat{p}(Z|M)\right) + \sum_{g=1}^{G} \left(\delta \hat{p}(Z|M) + (1-\delta_p)\, p(Z|\theta^g, M)\right)^{-1}}, \tag{10}$$

where δ_p is the fraction of parameter draws we assume to be from the priors. While $\delta_p = 0$ corresponds to the harmonic mean of the log-likelihood, choosing $\delta_p \neq 0$ stabilizes the marginal likelihood estimate, and we set $\delta_p = 0.01$. Marginal likelihood is directly related to Bayes factors: if all models are equally probable a priori, the logarithmic Bayes Factor between model M_1 and M_0 can be computed from $\text{LBF}(\text{M}_1, \text{M}_0) = -2\left(\ln p(Z|M_1) - \ln p(Z|M_0)\right)$. Bayes Factors compare model fit while taking parameter uncertainty into consideration. Kass and Raftery (1995) suggest the following interpretation for the scale: if $\text{LBF}(\text{M}_1, \text{M}_0)$ is between 2 to 6, it is viewed as positive evidence favoring M_0; between 6 to 10, it indicates strong evidence; and a value greater than 10 is interpreted as very strong evidence. Negative values are interpreted on the same scale favoring M_1.

Table 2 reports the posterior mean and standard deviation of model parameters, and Table 3 reports the logarithmic bayes factors. From Table 2, the simple jump-diffusion model SFJ estimates the mean-reversion rate to be 0.546. Hence, both spike shocks and base-signal shocks have a half-life

of 1.27 days. Spike occurs 10 times a year on average with a mean shock size of 1.01. The constant variance estimate for the base-signal is 0.066. Model TFJ also assumes constant base-signal variance, but the spikes can mean-revert faster than base-signals. The estimated half-life of a spike shock is about one-third of a day, compared to a half-life of 4.4 days for a base-signal shock. The variance of base-signal decreases to 0.028, while the jump intensity increases to 0.095 with a slightly larger mean shock size, $\mu_J = 1.21$. The TFJ model attributes more of the price uncertainty to spikes than the SFJ model, and Table 3 suggests that TFJ is strongly favored over SFJ.

Stochastic volatility is also an important source of uncertainty; models with stochastic volatility have higher likelihood than models with constant volatility. When volatility is allowed to jump, the jump-diffusion specification provides worse model fit than the simple no jump model, in the sense that LBF(SF-SV, SFJ-SV) is negative. Also, the posterior of jump size parameter μ_J is dominated by the prior, indicating that the data is uninformative about jump sizes. On the other hand, specify different mean-reversion rate for base-signals and spikes does improve model fit, with LBF(SF-SV, TFJ-SV) = 93.2. Compare TFJ-SV or TFJ-S2V with TFJ, we find the estimated half-life of base-signals is similar, while the half-life of spikes decreases to about a quarter of a day. The jump intensity in TFJ-SV/TFJ-S2V is smaller than TFJ, with larger jump size.

Spike identification also affect the estimation of volatility dynamics. In SF-SV and SFJ-SV model, almost all of the price variations are explained by volatility. The estimated volatility process is less persistent, and the half-life of a volatility shock is about 1.5 days. In the TFJ-SV model, large transient changes are explained by spikes, and volatility shock has a half-life of 2.5 days. The TFJ-S2V allows two factors in the volatility process. We identify a slower mean-reverting process with a half-life of 12.5 days and weight of 0.4 and a faster mean-reverting process with a half-life of 1.3 days and weight of 0.6.

5.4 Estimated Risk Factors

In this section, we compare the risk factors identified by the different models. Figure 2 presents the estimated spike processes from the models that allow jumps. In SFJ-SV model, most of the price variations are explained by volatility, while spike process is almost constant. Spikes identified by the SFJ model die out slower than TFJ model as SFJ imposes the same mean-reversion rate for spikes and base-signals. Spikes in the TFJ model better resembles the short-lived jumps we observe in the data,

Table 2: Parameter Estimates

	SFJ	TFJ	SFJ-SV	SF-SV	TFJ-SV	TFJ-S2V
α_x	0.546	0.157	0.158	0.158	0.149	0.157
	(0.024)	(0.015)	(0.016)	(0.016)	(0.015)	(0.016)
α_y		2.332			2.991	3.000
		(0.326)			(0.829)	(1.226)
σ^2	0.066	0.028				
	(0.003)	(0.001)				
λ_1			0.438	0.435	0.279	0.056
			(0.063)	(0.067)	(0.049)	(0.015)
λ_2						0.510
						(0.104)
δ_1			0.133	0.133	0.148	0.059
			(0.008)	(0.008)	(0.011)	(0.019)
δ_2						0.089
						(0.016)
γ			1.291	1.348	3.339	3.264
			(0.240)	(0.274)	(0.469)	(0.511)
μ_J	1.010	1.208		3.620	2.412	2.259
	(0.172)	(0.228)		(2.583)	(1.094)	(0.859)
λ_J	0.043	0.095		0.002	0.028	0.029
	(0.008)	(0.021)		(0.002)	(0.012)	(0.016)
LLMean	0.0	299.9	436.5	435.0	485.1	495.1
LLMedian	0.0	300.1	436.5	435.5	485.5	495.4
MarginalLL	0.0	299.6	437.0	435.2	483.6	491.1

The table reports the posterior mean and posterior standard deviation of model parameters. LLmean/LLmedian is the log-likelihood evaluated at the posterior mean/median using 300,000 particles. The marginalLL is computed using equation (10). All three log-likelihood are reported in excess of model SFJ.

Table 3: Bayes Factor

	SFJ	TFJ	SFJ-SV	SF-SV	TFJ-SV
TFJ	599.2				
SFJ-SV	874.0	274.8			
SF-SV	870.4	271.2	-3.6		
TFJ-SV	967.2	368.0	93.2	96.8	
TFJ-S2V	982.3	383.1	108.3	111.9	15.1

The table reports twice the natural logarithm of the Bayes factors. The entry (i,j) in the matrix compares the model in the ith row and the model in the jth column, with a positive value favoring the first model and a negative value favoring the latter.

see for example the large spikes in January 2003 and July 2006. However, spike innovations tend to be clustered, like in the beginning of 2006. TFJ-SV and TFJ-S2V both captures the transient nature of the spike process, and there is less clustering in the spike innovation.

Figure 3 presents the estimate volatility from model SF-SV and TFJ-SV. The estimated volatility from SFJ-SV is almost identical to SF-SV and we omit it here. Volatility identified by SF-SV is less persistent with a heavier tail, which is consistent with the parameter estimates in Table 2. Figure 4 plots the estimated volatility process from TFJ-S2V model as well as the individual volatility factors. The first volatility factor $\sigma_1(t)$ is slowly mean-reverting and it captures high volatility period such as months near the end of 2005. The second volatility factor $\sigma_1(t)$ accounts for the daily variations and contributes to the kurtosis of base-signal shocks.

5.5 Futures Prices and Risk Premium

In this section, we compare the market futures prices, $F(t, T_1, T_2)$, with the futures price under P measure, $F^P(t, T_1, T_2)$. Since $F^P(t, T_1, T_2)$ is computed from the predicted spot prices averaged over the delivery period $[T_1, T_2]$, it depends on the model parameters and estimated risk factors.

We select two monthly futures contract to illustrate the effect of different risk factors and model specifications on futures prices. Figure 5 presents the futures prices in July 2006 for the contract with delivery in August 2006. From Figure 5, we see that $F^P(t, T_1, T_2)$ becomes more volatile as it approaches delivery, which is known as the Samuelson effect. Recall from equation (7), $F^P(t, T_1, T_2)$ is

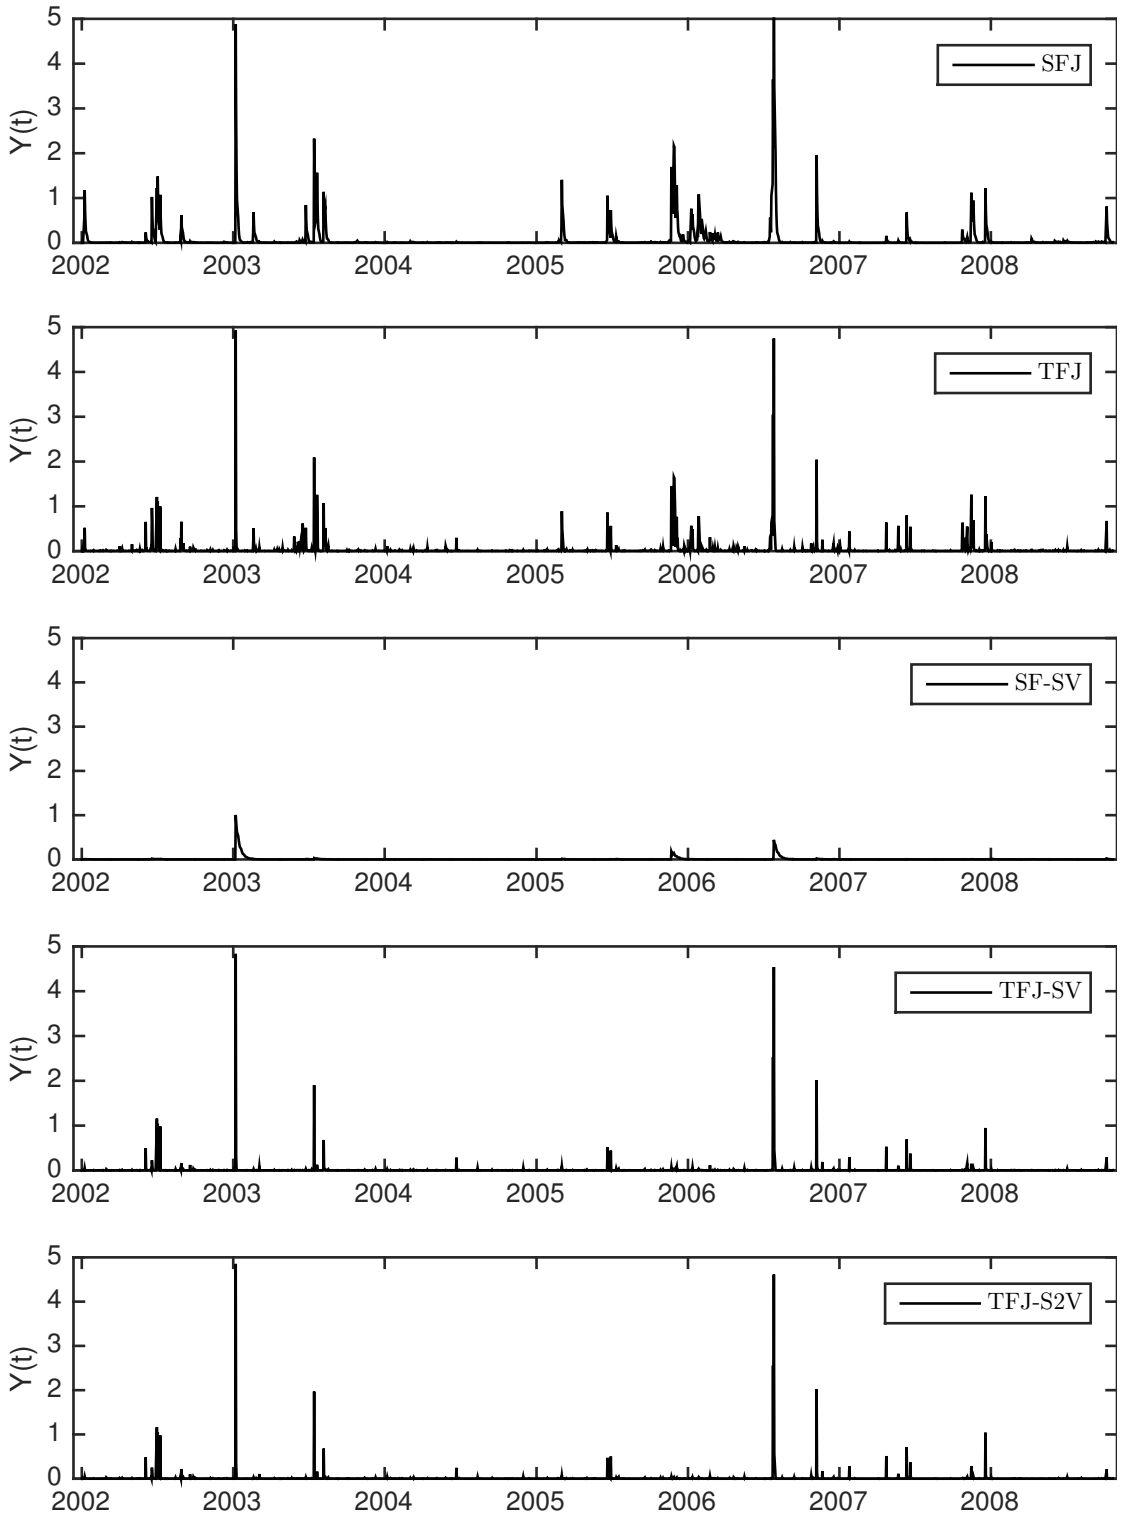

Figure 2: Estimated Spike Processes. The figure depicts the spike processes computed using the mean of the posterior distribution $p(Y_{1:T}|Z_{1:T})$.

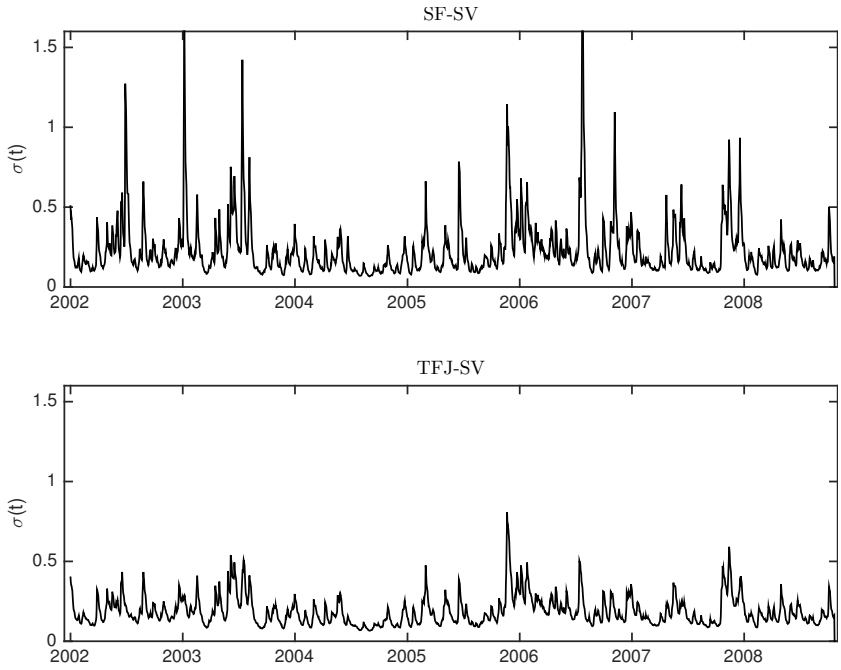

Figure 3: Estimated Volatility Processes. The figure depicts the volatility processes from model SF-SV and TFJ-SV, computed using the mean of the posterior distribution $p(\sigma(1:T-1)|Z_{1:T})$.

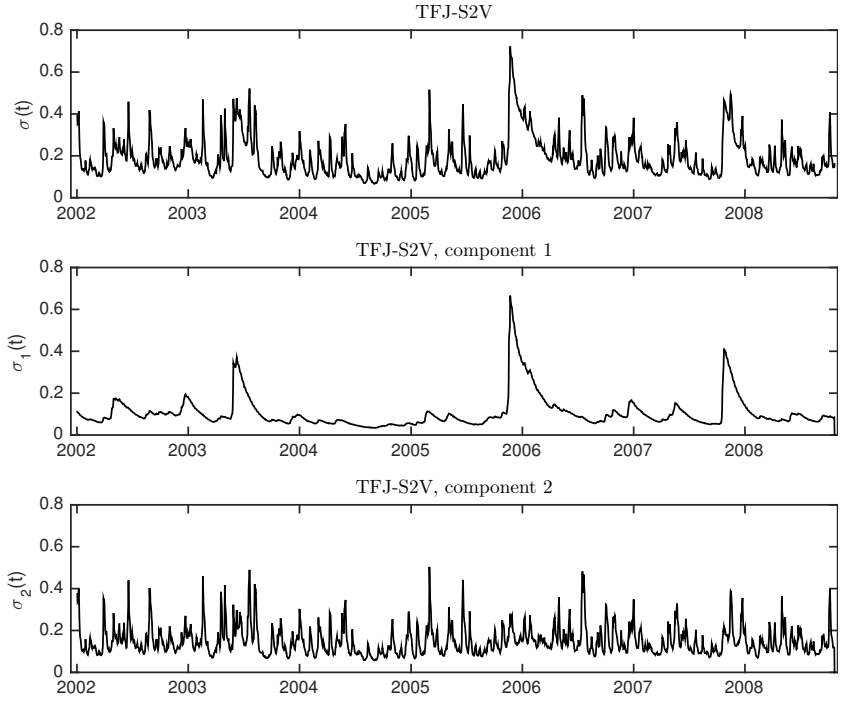

Figure 4: Estimated Volatility Processes. The figure depicts the volatility process and its two component from model TFJ-S2V, computed using the mean of the posterior distribution $p(\sigma(1:T-1)|Z_{1:T})$.

a weighted average of long-run mean and current level of $X(t)$ and $Y(t)$. As $T_1 - t$ decreases, the weight on the constant long-run mean decreases, and the weight on the stochastic $X(t)$ and $Y(t)$ increases. However, since the delivery period is $[T_1, T_2]$, $F^P(t, T_1, T_2)$ does not converge to $S(T_1)$ when $t \to T_1$.

For this contract, the predicted spot $F^P(t, T_1, T_2)$ is very different for different model specifications. Model SFJ predicts nearly constant futures price as the half-life of both spike shocks and base-signal are only about one day. Hence, shocks are transient and they have little influence on futures price. This is in contrast with model SF-SV, where both spike shocks and base-signal shocks have half-life of more than 4 days. Hence, the large spike on July 27th (3 business days before delivery) has a large impact on the futures price. For the two-factor models, spikes are transient and base-signals are more persistent. Hence, the predicted futures price has more variation than SFJ, whereas the large spike does not affect futures prices.

Next, we look at the futures price in August 2008 with maturity in September 2008 in Figure 6. During this month, $Y(t)$ is almost zero for all the models. Hence, the shape of the futures price is similar for all models except SFJ since their estimated α_x is similar. However, the level of the curves varies across model. This variation comes from the long-run mean of $Y(t)$, $\lambda_J \mu_J / \alpha_y$. TFJ model has the highest estimated jump intensity, while in model SF-SV λ_J is almost zero. The possibility of positive spikes drives up the futures price, hence $F^P(t, T_1, T_2)$ from TFJ lies above that of SF-SV. Compared to TFJ, the TFJ-SV/TFJ-S2V model has smaller jump probability with larger jump size, and slighter lower futures price. Although volatility does not directly impact the predicted spot, it affected the estimated spike parameters which in turn shifts the level of the futures prices.

The difference between $F(t, T_1, T_2)$ and $F^P(t, T_1, T_2)$ is the market risk premium. From equation (8), if the market price of jump risk η_I is constant and α_y is large, the spike risk premium is close to being constant as $\bar{\alpha}(\alpha_y, t, T_1, T_2)$ is close to zero. If the market price of diffusion risk per unit of volatility θ_B is also constant, the diffusion risk premium $\theta_B \left(1 - \bar{\alpha}(\alpha_x, t, T_1, T_2)\right) / \alpha_x$ is a monotonically increasing or decreasing function depending on the sign of θ. Figure 7 plots the risk premium from six different contracts, where $F^P(t, T_1, T_2)$ is computed from model TFJ-S2V. The risk premium varies in both shape and level, suggesting that neither θ_B nor η_I is constant over time.

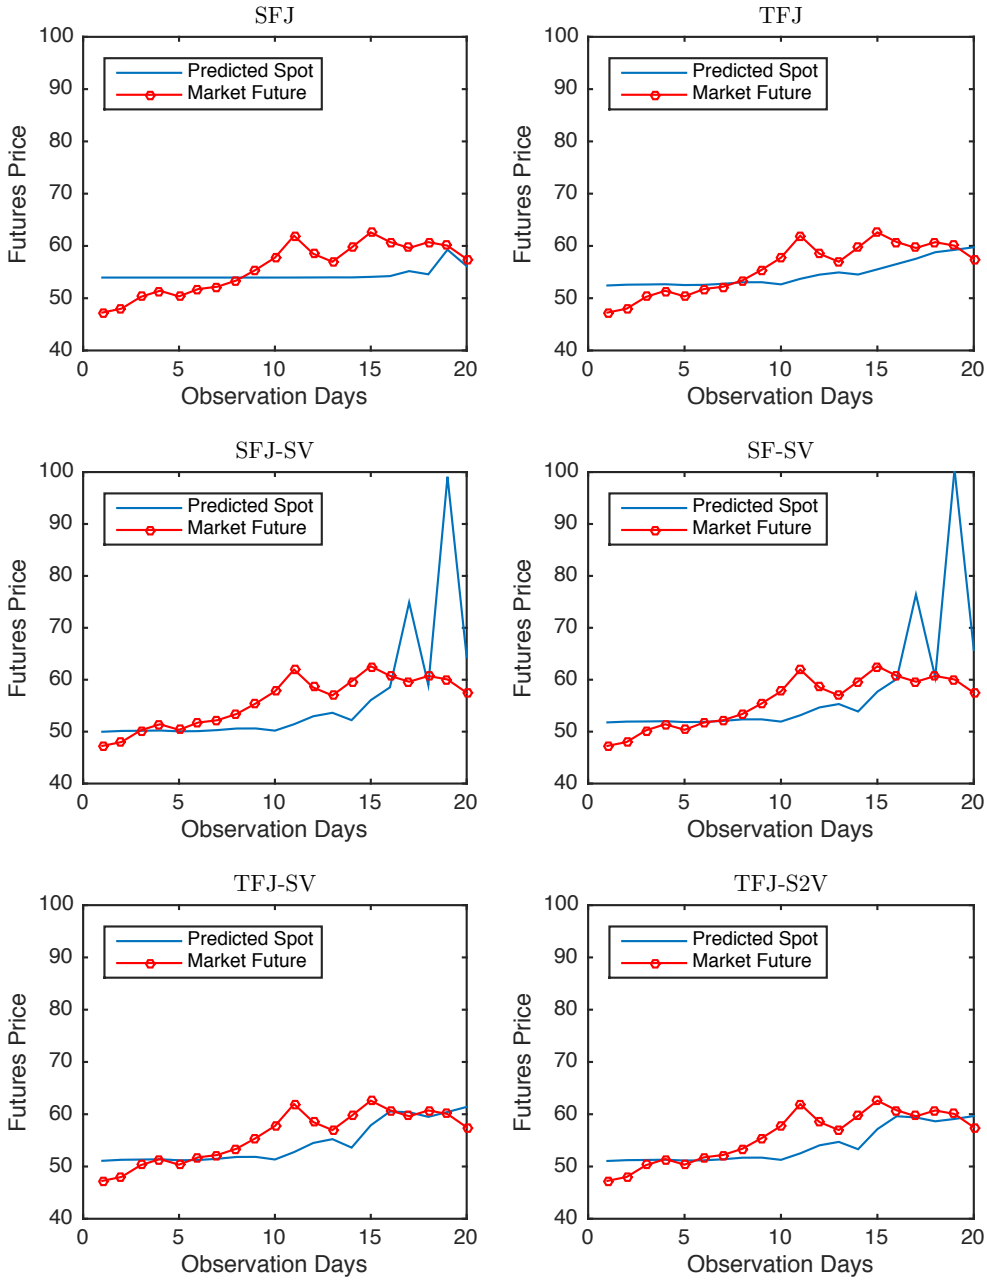

Figure 5: Predicted Spot vs Market Future Prices, Maturity in Aug 2006.

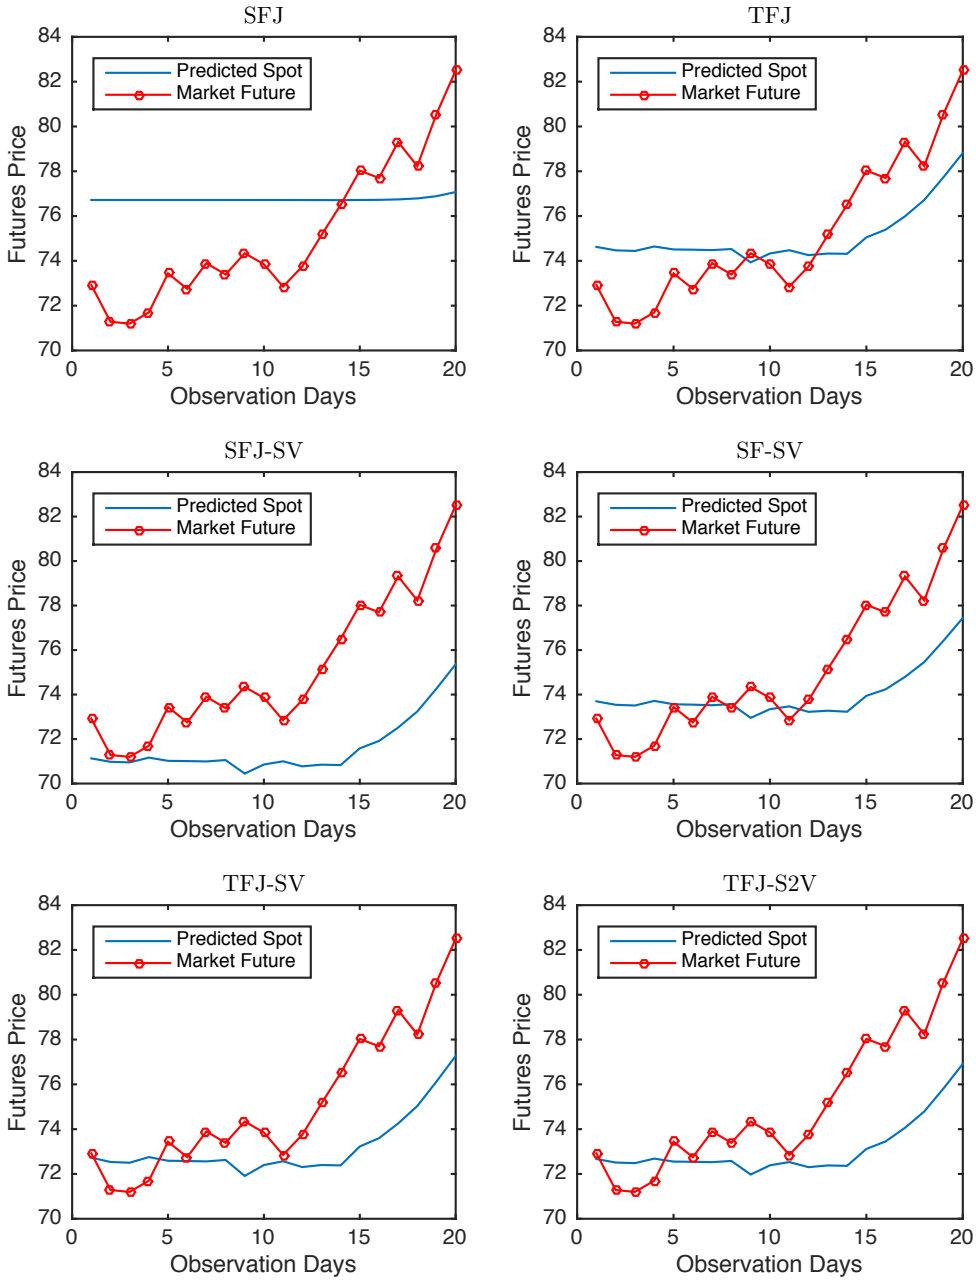

Figure 6: Predicted Spot vs Market Future Prices, Maturity in Sep 2008.

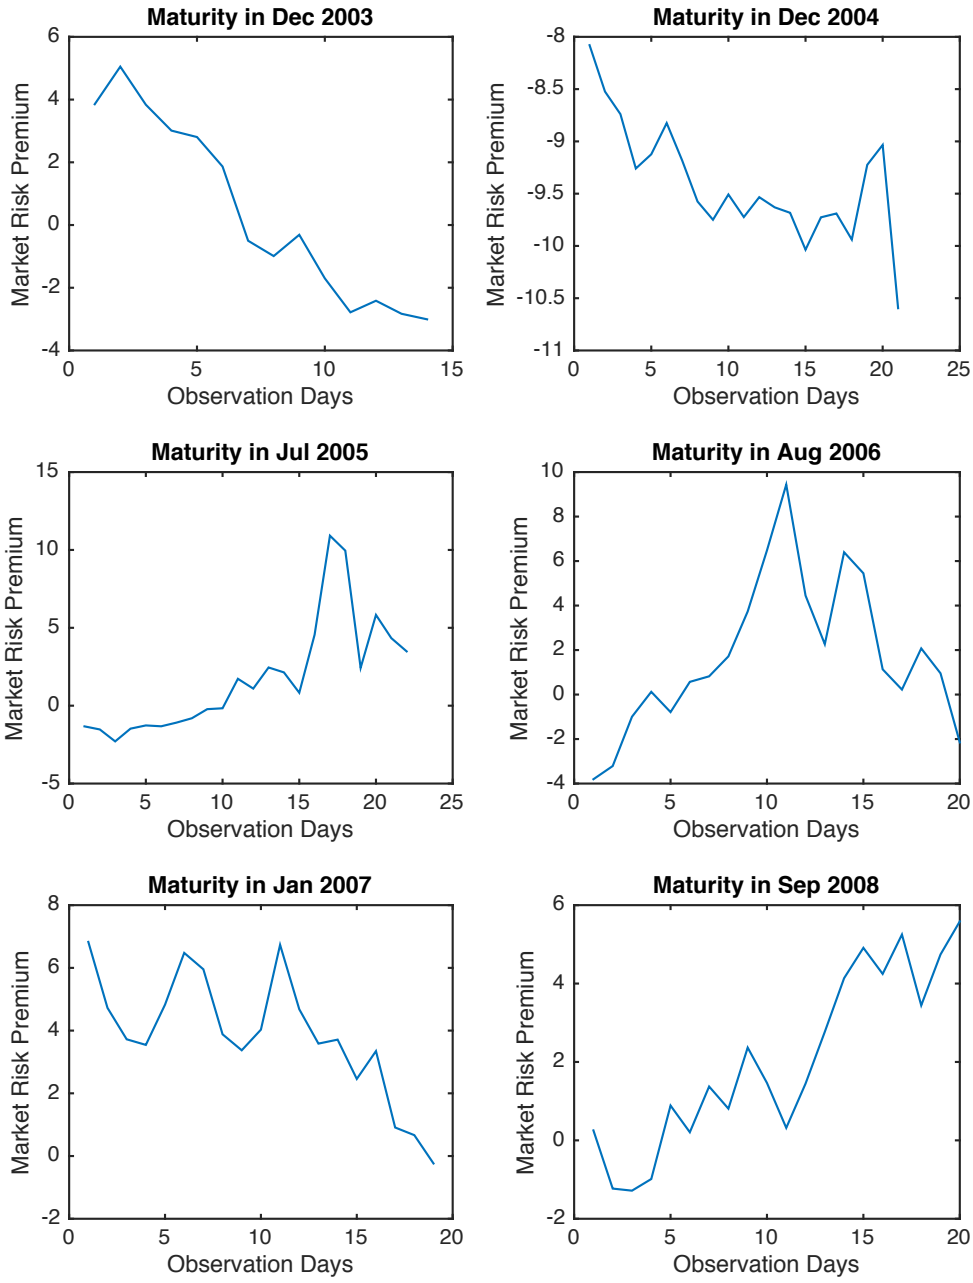

Figure 7: Market Risk Premium from Model TF-S2V.

6 Conclusions

We identify three main sources of uncertainty in electricity prices using an arithmetic multi-factor model. Specifically, electricity prices are subject to large spike shocks which die out fast, base-signal shocks which account for the more persistent variation, and volatility shocks which explains clustered large variations. The model is estimated using the PMCMC method. To accommodate non-Gaussian distributed spike shocks efficiently, we develop an adapted SMC algorithm which uses an importance distribution for adaption.

We analyze the EEX phelix index for the empirical study. To assess the importance of each factor, we also consider nested models which are restricted versions of our proposed model. The empirical application suggest both spikes and stochastic volatility are important factors in addition to base-signals.

The multi-factor model in this paper allows for analytically tractable pricing. Market futures price based on arbitrage pricing theory is consist of a predicted spot price and a market risk premium. Using the estimation results from the EEX spot market, we compute the predicted spot price and compare it the the market futures price. We find that spikes are transient, and the current level of spikes has little impact on futures price. Instead, the long-run mean of spikes shifts the level of futures price. The dynamics of the predicted spot is most determined by the base-signals. As the time to delivery approaches, the predicted spot price moves closer to the stochastic base-signals and hence become more volatile. Volatility does not directly impact futures price, but it impacts estimated spikes and hence affect future prices indirectly. We compute risk premium as the difference between the predicted spot price and the market futures price, and the results suggest that market price of both jump risk and diffusion risk are time-varying.

References

Andrieu, C., Doucet, A., and Holenstein, R. (2010). 'Particle Markov chain Monte Carlo methods', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(3): 269–342.

Andrieu, C., and Thoms, J. (2008). 'A tutorial on adaptive MCMC', *Statistics and Computing*, 18(4): 343–373.

Barndorff-Nielsen, O. E., and Shephard, N. (2001). 'Non-Gaussian Ornstein-Uhlenbeck-based models

and some of their uses in financial economics', *Journal Of The Royal Statistical Society Series B*, 63(2): 167–241.

Bates, D. S. (2000). 'Post-87 crash fears in the S&P 500 futures option market', *Journal of Econometrics*, 94(1-2): 181–238.

Benth, F. E. (2011). 'The Stochastic Volatility Model of Barndorff-Nielson and Shephard in Commodity Markets', *Mathematical Finance*, 21(4): 595–625.

Benth, F. E., Álvaro Cartea, and Kiesel, R. (2008). 'Pricing forward contracts in power markets by the certainty equivalence principle: Explaining the sign of the market risk premium', *Journal of Banking & Finance*, 32(10): 2006 – 2021.

Benth, F. E., Benth, J. S., and Koekebakker, S. (2008). *Statistical Modeling of Electricity and Related Markets*, Advanced Series on Statistical Science and Applied Probability. World Scientific.

Benth, F. E., Kallsen, J., and Meyer-Brandis, T. (2007). 'A Non-Gaussian Ornstein-Uhlenbeck Process for Electricity Spot Price Modeling and Derivatives Pricing', *Applied Mathematical Finance*, 14(2): 153–169.

Benth, F. E., Kiesel, R., and Nazarova, A. (2012). 'A critical empirical study of three electricity spot price models', *Energy Economics*, 34(5): 1589 – 1616.

Bessembinder, H., and Lemmon, M. L. (2002). 'Equilibrium Pricing and Optimal Hedging in Electricity Forward Markets', *The Journal of Finance*, 57(3): 1347–1382.

Bierbrauer, M., Menn, C., Rachev, S. T., and Trück, S. (2007). 'Spot and derivative pricing in the {EEX} power market', *Journal of Banking & Finance*, 31(11): 3462 – 3485, Risk Management and Quantitative Approaches in Finance.

Bollerslev, T., and Todorov, V. (2011). 'Tails, Fears, and Risk Premia', *The Journal of Finance*, 66(6): 2165–2211.

Brix, A. F., Lunde, A., and Wei, W. (2015). 'A Generalized Schwartz Model for Energy Spot Prices - Estimation using a Particle MCMC Method', .

Cartea, A., and Figueroa, M. (2005). 'Pricing in electricity markets: a mean reverting jump diffusion model with seasonality', *Applied Mathematical Finance*, 12(4): 313–335.

Creal, D. (2012). 'A Survey of Sequential Monte Carlo Methods for Economics and Finance', *Econometric Reviews*, 31(3): 245–296.

Durbin, J., and Koopman, S. J. (1997). 'Monte Carlo maximum likelihood estimation for non-Gaussian state space models', *Biometrika*, 84(3): 669–684.

Flury, T., and Shephard, N. (2011). 'Bayesian Inference Based Only On Simulated Likelihood: Particle Filter Analysis Of Dynamic Economic Models', *Econometric Theory*, 27(05): 933–956.

Heston, S. L. (1993). 'A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options', *Review of Financial Studies*, 6(2): 327–43.

Huisman, R., and Mahieu, R. (2003). 'Regime jumps in electricity prices', *Energy Economics*, 25(5): 425 – 434.

Janczura, J., Trück, S., Weron, R., and Wolff, R. C. (2013). 'Identifying spikes and seasonal components in electricity spot price data: A guide to robust modeling', *Energy Economics*, 38: 96 – 110.

Kanamura, T., and Ōhashi, K. (2007). 'A structural model for electricity prices with spikes: Measurement of spike risk and optimal policies for hydropower plant operation', *Energy Economics*, 29(5): 1010 – 1032.

Kass, R. E., and Raftery, A. E. (1995). 'Bayes Factors', *Journal of the American Statistical Association*, 90(430): pp. 773–795.

Knittel, C. R., and Roberts, M. R. (2005). 'An empirical examination of restructured electricity prices', *Energy Economics*, 27(5): 791 – 817.

Koopman, S. J., Ooms, M., and Carnero, M. A. (2007). 'Periodic Seasonal Reg-ARFIMA-GARCH Models for Daily Electricity Spot Prices', *Journal of the American Statistical Association*, 102(477): 16–27.

Longstaff, F. A., and Wang, A. W. (2004). 'Electricity Forward Prices: A High-Frequency Empirical Analysis', *The Journal of Finance*, 59(4): 1877–1900.

Mount, T. D., Ning, Y., and Cai, X. (2006). 'Predicting price spikes in electricity markets using a regime-switching model with time-varying parameters', *Energy Economics*, 28(1): 62 – 80.

Newton, M. A., and Raftery, A. E. (1994). 'Approximate Bayesian Inference with the Weighted Likelihood Bootstrap', *Journal of the Royal Statistical Society. Series B (Methodological)*, 56(1): pp. 3–48.

Pan, J. (2002). 'The jump-risk premia implicit in options: evidence from an integrated time-series study', *Journal of Financial Economics*, 63(1): 3–50.

Pitt, M. K., dos Santos Silva, R., Giordani, P., and Kohn, R. (2012). 'On some properties of Markov chain Monte Carlo simulation methods based on the particle filter', *Journal of Econometrics*, 171(2): 134 – 151, Bayesian Models, Methods and Applications.

Pitt, M. K., and Shephard, N. (1999). 'Filtering via Simulation: Auxiliary Particle Filters', *Journal of the American Statistical Association*, 94(446): 590–599.

Schwartz, E. S. (1997). 'The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging', *Journal of Finance*, 52(3): 923–73.

Veraart, A. E. D., and Veraart, L. A. M. (2014). 'Risk premiums in energy markets', *The Journal of Energy Markets*, 6(4): 91–132.

Weron, R. (2014). 'Electricity price forecasting: A review of the state-of-the-art with a look into the future', *International Journal of Forecasting*, 30(4): 1030 – 1081.

Weron, R., Bierbrauer, M., and Trück, S. (2004). 'Modeling electricity prices: jump diffusion and regime switching', *Physica A: Statistical Mechanics and its Applications*, 336(1-2): 39 – 48.