

Article

Bayesian Option Pricing Framework with Stochastic Volatility for FX Data

Ying Wang ¹, Sai Tsang Boris Choy ^{2,*} and Hoi Ying Wong ¹

¹ Department of Statistics, The Chinese University of Hong Kong, Hong Kong, China; rachelwang415@gmail.com (Y.W.); hywong@sta.cuhk.edu.hk (H.Y.W.)

² Discipline of Business Analytics, The University of Sydney, NSW 2006, Australia

* Correspondence: boris.choy@sydney.edu.au; Tel.: +61-2-9351-2787

Academic Editor: Qihe Tang

Received: 31 August 2016; Accepted: 9 December 2016; Published: 16 December 2016

Abstract: The application of stochastic volatility (SV) models in the option pricing literature usually assumes that the market has sufficient option data to calibrate the model's risk-neutral parameters. When option data are insufficient or unavailable, market practitioners must estimate the model from the historical returns of the underlying asset and then transform the resulting model into its risk-neutral equivalent. However, the likelihood function of an SV model can only be expressed in a high-dimensional integration, which makes the estimation a highly challenging task. The Bayesian approach has been the classical way to estimate SV models under the data-generating (physical) probability measure, but the transformation from the estimated physical dynamic into its risk-neutral counterpart has not been addressed. Inspired by the generalized autoregressive conditional heteroskedasticity (GARCH) option pricing approach by Duan in 1995, we propose an SV model that enables us to simultaneously and conveniently perform Bayesian inference and transformation into risk-neutral dynamics. Our model relaxes the normality assumption on innovations of both return and volatility processes, and our empirical study shows that the estimated option prices generate realistic implied volatility smile shapes. In addition, the volatility premium is almost flat across strike prices, so adding a few option data to the historical time series of the underlying asset can greatly improve the estimation of option prices.

Keywords: option pricing; volatility smile; Student-*t*; variance gamma; Markov chain Monte Carlo (MCMC)

1. Introduction

The constant volatility assumption in the original Black–Scholes model has been criticized over the years for its failure to produce the implied volatility smile. Time varying volatility offers a promising remedy to capture the smile. For instance, the autoregressive conditional heteroskedasticity (ARCH) models of [1] and the generalised ARCH (GARCH) models of [2] offer the possibility of capturing many stylized time-varying volatility facts in a time series perspective. However, they were mainly used in the investigation of economic series, until Duan ([3]) proposed a GARCH option pricing framework that enabled the GARCH model estimated from asset return series to be conveniently transformed into a risk-neutral process for option pricing purposes. More precisely, the concept of the locally risk-neutral valuation relationship (LRNVR) was established by [3]. Although this option pricing model successfully captures the volatility skewness of the equity option market, it is inadequate to generate the U-shaped volatility smiles in other option markets, such as commodity and foreign exchange (FX) markets.

Two major drawbacks of the GARCH-type models were pointed out by [4]. First, the constraints imposed on the parameters of these models to ensure a positive conditional variance are often violated

during estimation. Second, a random oscillatory behaviour of the conditional variance process is ruled out. For the purpose of option pricing, Hull and White [5] proposed the stochastic volatility (SV) model, which contains an additional unobserved random process in the volatility. Heston [6] extended the SV model to allow for a non-zero correlation between the asset return and its volatility, which is called the leverage effect in financial econometrics. Although SV models are often calibrated to observed option data under the risk-neutral measure, option data are not always available for certain underlying assets. When option data are unavailable, market practitioners must resort to estimating the model under the physical measure using historical asset returns, and then transforming the model into its risk-neutral counterpart for derivative pricing.

Unlike GARCH models, SV models do not admit a computable likelihood function, which makes the estimation a highly challenging task. While the maximum likelihood estimation (MLE) method is infeasible, Cordis and Korby [7] introduced discrete stochastic autoregressive volatility (DSARV) models, in which volatilities only take discrete random values. These models greatly reduce computational costs, because the MLE method can be applied. However, most SV models assume a continuous stochastic volatility, and the Bayesian framework has become a useful alternative for statistical inference. The implementation of Bayesian methods usually requires the construction of a Markov chain Monte Carlo (MCMC) simulation from the intractable joint posterior distribution. Jacquier et al. [8] analysed SV models with a leverage effect by adopting the Gibbs sampling scheme. Shephard and Pitt [9] employed the Metropolis–Hastings scheme for the same problem. Other generalizations include the work of [10–12], among others. Choy et al. [13] analysed various FX data using different heavy-tailed SV models and found that there were no leverage effects. Wang et al. [14] considered SV models with leverage (SVL) and a bivariate Student- t error distribution. The t -distribution is expressed in a scale mixture of normal (SMN) representation, which significantly simplifies the Gibbs sampler and dramatically reduces computational time. All of the methods presented above are useful in the estimation of the SV model under the physical measure, but the transformation into the risk-neutral process for option pricing can be a highly non-trivial task.

The risk-neutral valuation for derivative pricing is concerned with the appropriate drift term of the asset return process under the pricing (risk-neutral) measure. However, most SV estimation ignores the drift term, speculating that it can be well estimated by sample mean. An exceptional case is the work by [15], but their drift specification cannot be effectively transformed into a risk-neutral process for option pricing purposes. Bates [16] considered the transformation between the two measures and allowed the parameters of the volatility process to change, but he did not set up the locally risk-neutral valuation formula for option pricing with stochastic volatility.

Inspired by the concept of LRNVR in [3], we propose an SV model that enables a convenient transformation into the risk-neutral process. Our model permits the innovations of the return and volatility processes to have different heavy-tailed SMN distributions. Specifically, we choose a variance gamma (VG) error distribution in the return process and a Student- t error distribution in the volatility process. We shall refer to this SV model as the $VG-t$ SV model. For modelling FX data, no leverage effect is assumed. The use of SMN distributions permits an efficient Bayesian inference using MCMC algorithms, and provides a better fit to FX data. The VG distribution ([17]) has thicker tails than the Gaussian distribution, retains finite moments for all orders, and offers a good empirical fit to the data. Assuming a constant volatility, Madan et al. [18] applied the VG process to option pricing and showed that the VG model fitted the volatility smile well. In this paper, we allow the volatility to have a stochastic process. We derive the physical process that can be efficiently transformed into a risk-neutral process for the derivative pricing of FX rates.

This paper also contributes to the literature by the use of the Bayesian framework in option pricing with empirical data. We perform the estimation on an FX rate and then simulate option prices based on the proposed model. We find evidence that heavy-tailed distributions for the FX return and its volatility have important implications. When both the FX return and its volatility marginally follow SMN distributions, the generated implied volatility smile matches the shape of the market observed

implied volatility very well. The small gap between the model and market implied volatility smile is almost flat across strike prices. The difference between the model and market implied volatility is often known as the volatility premium, and can be managed if a few option data are available. This is because the volatility premium is quite flat across strike prices.

The remainder of this paper is organized as follows. Section 2 presents the model and problem formulation. Section 3 discusses the Bayesian estimation framework using SMN representations and MCMC methods. The simulations of option prices are also detailed. An empirical study on FX options is shown to justify the potential use of the framework. Specifically, the out-of-sample fit from our framework outperforms the standard Black–Scholes model when compared with market option data. We also discuss how to further improve our framework by considering the volatility premium once some (but not necessarily many) option data are available. Finally, concluding remarks are made in Section 4.

2. Stochastic Volatility Model and Option Pricing

2.1. Problem Formulation

Consider a fixed filtered complete probability space $(\Omega, \mathcal{F}, \mathbb{P}, \mathcal{F}_{t \geq 0})$, where \mathcal{F}_t is the filtration generated by independent processes $\{\varepsilon_t\}$ and $\{\eta_t\}$, augmented by the null sets of the data-generating (physical) probability measure \mathbb{P} . In addition, $\mathbb{E}^{\mathbb{P}}[\varepsilon_t] = \mathbb{E}^{\mathbb{P}}[\eta_t] = 0$ for all $t \in \mathbb{N}$. Let S_t be the underlying asset price at time t . We construct the following discrete-time SV model under \mathbb{P} :

$$y_t = f(\sigma_t^2) + \sigma_t \varepsilon_t, \quad (1)$$

$$h_t = \mu + \phi(h_{t-1} - \mu) + \tau \eta_t, \quad (2)$$

where $y_t = \ln S_t - \ln S_{t-1}$ is the geometric return of an asset at time t , $\sigma_t^2 = e^{h_t}$ and h_t are the volatility and log-volatility at time t , and $f(\sigma_t^2)$ is the drift of the return process, which is a function of the volatility in general. For the log-volatilities, the conditional mean and conditional variance of h_t are given by

$$\mathbb{E}[h_t | h_{t-1}] = \mu + \phi(h_{t-1} - \mu), \quad \text{and} \quad \mathbb{V}[h_t | h_{t-1}] = \tau^2,$$

while the unconditional mean and unconditional variance can be derived as

$$\mathbb{E}[h_t] = \mu, \quad \text{and} \quad \mathbb{V}[h_t] = \frac{\tau^2}{1 - \phi^2}.$$

Obviously, μ is the unconditional mean of h_t , and τ is the conditional standard deviation of h_t . We assume that the persistence ϕ in the volatility equation satisfies $|\phi| < 1$ to ensure that h_t is stationary.

It remains to specify the drift function $f(\sigma_t^2)$ to complete the entire SV model for the asset return process. We choose the drift function that meets the following three requirements.

1. It should be a simple model that enables us to apply the standard Bayesian inference via MCMC;
2. It should be convenient to transform the model under \mathbb{P} to the model under \mathbb{Q} , the risk-neutral probability;
3. The resulting \mathbb{Q} -process should generate option prices close to market prices.

While the third requirement is justified empirically in a later section, we focus on the former two requirements.

We use SMN distributions to model ε_t and η_t so that the Bayesian inference for SV models reported in the literature can be easily applied. For the second requirement, our model is developed using the concept of LRNVN [3], which requires the following properties.

- (P₁) Equivalent probability measures: $\mathbb{P}(A) = 0$ iff $\mathbb{Q}(A) = 0$ for any event A ;
- (P₂) Martingale property: $E^{\mathbb{Q}}[S_{t+1} | \mathcal{F}_t] = S_t e^r$, where r is the one-step continuously compounded interest rate;
- (P₃) Equivalent local variance: $V^{\mathbb{Q}}(y_{t+1} | \mathcal{F}_t) = V^{\mathbb{P}}(y_{t+1} | \mathcal{F}_t)$,

where \mathcal{F}_t is the information accumulated up to time t , and $V(\cdot)$ denotes the variance. Duan [3] offered a sound economic interpretation of these requirements, and showed that they ruled out the local arbitrage opportunities within the GARCH models. In fact, (P₁) is a compatibility condition that preserves unlikely events to remain unlikely in both probability measures. (P₂) ensures that forward contracts—the simplest derivative contracts—are correctly priced so that no arbitrage opportunity exists in the risk-neutral measure \mathbb{Q} , consistent with standard financial econometric theory. (P₃) maintains the conditional risk level so that no additional risk is generated or reduced from issuing a derivative.

Inspired by the Girsanov theorem on Itô’s processes, we consider the change of measure by linearly shifting the drift term in the asset return process (1) so that it guarantees (P₁) and does not affect the Bayesian inference for SV models. The (P₂) of LRNVR implies that the drift term contains the interest rate r under \mathbb{Q} . Combining these two considerations, we have

$$\begin{aligned} y_t &= r + \beta\sigma_t + \Phi(\sigma_t^2) + \sigma_t\varepsilon_t, \\ h_t &= \mu + \phi(h_{t-1} - \mu) + \tau\eta_t, \end{aligned}$$

where r is the one-period continuously compounded risk-free interest rate, β is a constant, and $\Phi(\sigma^2)$ is the compensator function enforcing (P₂) to hold. In addition, (P₃) is satisfied, as the conditional variance is evaluated as σ^2 after the mean shift.

To illustrate our idea, we consider the Gaussian innovations for both return and conditional volatility processes. Set $\Phi = -\sigma_t^2/2$. The unconditional distribution of h_t remains Gaussian. Under \mathbb{P} ,

$$\begin{aligned} y_t &= r + \beta\sigma_t - \frac{\sigma_t^2}{2} + \sigma_t\varepsilon_t, \\ h_t &= \mu + \phi(h_{t-1} - \mu) + \tau\eta_t, \end{aligned} \tag{3}$$

and a pair of \mathbb{Q} -processes $(\varepsilon_t^*, \eta_t^*)$ satisfying the LRNVR can then be identified. Specifically, $\varepsilon_t^* = \varepsilon_t + \beta$ and $\eta_t^* = \eta_t$. Clearly, \mathbb{P} and \mathbb{Q} are equivalent. The SV model under \mathbb{Q} is deduced as

$$y_t^* = r - \frac{\sigma_t^2}{2} + \sigma_t\varepsilon_t^*, \tag{4}$$

$$h_t^* = \mu^* + \phi(h_{t-1}^* - \mu) + \tau\eta_t^*. \tag{5}$$

It is easy to verify that

$$\begin{aligned} E^{\mathbb{Q}}[S_{t+1} | \mathcal{F}_t] &= E^{\mathbb{Q}}[S_t e^{y_{t+1}} | \mathcal{F}_t] = S_t e^r E^{\mathbb{Q}}\left[e^{-\frac{1}{2}\sigma_t^2 + \sigma_t\varepsilon_t^*} \middle| \mathcal{F}_t\right] = S_t e^r; \\ V^{\mathbb{Q}}[y_{t+1} | \mathcal{F}_t] &= \sigma_t^2 = V^{\mathbb{P}}[y_{t+1} | \mathcal{F}_t]. \end{aligned}$$

The following theorem presents a more general result.

Theorem 1. Consider the SV model in (1)–(2). Let $\varepsilon_t^* = \varepsilon_t + \beta$ and $\eta_t^* = \eta_t$ for any fixed constant β . If there exists an equivalent measure \mathbb{Q} under which $E^{\mathbb{Q}}[\varepsilon_t^*] = E^{\mathbb{Q}}[\eta_t^*] = 0$ and $M(s) = E^{\mathbb{Q}}[e^{s\varepsilon_t^*}] < \infty$ for all $s > 0$, $t \in \mathbb{N}$ is the moment generating function (MGF) of ε_t^* , then the SV model under \mathbb{P} :

$$\begin{aligned} y_t &= r + \beta\sigma_t - \ln M(\sigma_t) + \sigma_t\varepsilon_t, \\ h_t &= \mu + \phi(h_{t-1} - \mu) + \tau\eta_t, \end{aligned} \tag{6}$$

admits an SV model under \mathbb{Q} satisfying the LRNVR:

$$\begin{aligned} y_t^* &= r - \ln M(\sigma_t) + \sigma_t \varepsilon_t^*, \\ h_t^* &= \mu + \phi(h_{t-1}^* - \mu) + \tau \eta_t^*. \end{aligned} \tag{7}$$

Proof of Theorem 1. As the theorem assumes that there is an equivalent probability measure \mathbb{Q} , we only need to prove for the second and third conditions of LRNVR. Using the transformation $\varepsilon_t^* = \varepsilon_t + \beta$ and $\eta_t^* = \eta_t$, the SV model in (6) becomes

$$\begin{aligned} y_t &= r - \ln M(\sigma_t) + \sigma_t \varepsilon_t^*, \\ h_t &= \mu - \tau \beta \rho + \phi(h_{t-1} - \mu) + \tau \eta_t^*. \end{aligned}$$

As $h_t \in \mathcal{F}_{t-1}$, it is easy to check that

$$\begin{aligned} \mathbb{E}^{\mathbb{Q}} [S_{t+1} | \mathcal{F}_t] &= \mathbb{E}^{\mathbb{Q}} [S_t e^{y_{t+1}} | \mathcal{F}_t] = S_t e^r \frac{\mathbb{E}^{\mathbb{Q}} [e^{\sigma_t \varepsilon_t^*} | \mathcal{F}_t]}{M(\sigma_t)} = S_t e^r; \\ \mathbb{V}^{\mathbb{Q}} [y_{t+1} | \mathcal{F}_t] &= \sigma_t^2 = \mathbb{V}^{\mathbb{P}} [y_{t+1} | \mathcal{F}_t]. \end{aligned} \tag{8}$$

□

The SV model in (6) has a plausible economic interpretation. A simple consideration may assume a constant drift term m so that

$$y_t = m + \sigma_t \varepsilon_t.$$

When a constant expected geometric return is assumed, it violates the economic belief on the return–risk trade off. In terms of the number of model parameters, both the constant drift SV model and that in (6) contain five basic parameters, and other parameters associated with the distributions of ε_t and η_t . Without increasing the number of parameters, the SV model shows that the higher the volatility, the higher the expected geometric return for a positive β . In fact, the β can be viewed as the Sharpe ratio: the excess return over volatility. Thus, our model can be interpreted as a constant Sharpe ratio model in a financial econometrics. Badescu et al. [19] reported that the constant drift model needs a subtle Esscher transform to obtain the \mathbb{Q} process under GARCH models, and the computation of option price could be less straightforward. We find that the constant drift SV model also admits implementation difficulties.

For (6) to be well defined, we must ensure that there is an equivalent probability measure \mathbb{Q} under which the \mathbb{Q} moment-generating function of ε_t^* is well defined. Although it is not necessarily true for arbitrary distributions of ε_t and η_t , many implementable SV models which are consistent with Theorem 1 can be developed via SMN distributions.

2.2. The SV Model with VG and t Error Distributions

The examples of SMN distributions include the Student- t , VG distributions, and many others. See [20]. Analyses of models with both Student- t innovations on returns and volatility processes have already been conducted, making it a relatively mature field. However, we do not model ε_t using a t distribution, because Theorem 1 requires a finite MGF. To allow heavy-tailed distributions for asset return and its volatility, we consider a standard VG distribution with the shape parameter α for ε_t , and a standard Student- t distribution with ν degrees of freedom for η_t . Using the MGF of the VG distribution, Theorem 1 brings us the following SV model under \mathbb{P} .

$$y_t = r + \beta \sigma_t + \alpha \ln \left(1 - \frac{\sigma_t^2}{2\alpha} \right) + \sigma_t \varepsilon_t, \tag{9}$$

$$h_t = \mu + \phi(h_{t-1} - \mu) + \tau \eta_t, \tag{10}$$

where

$$M(\sigma_t) = \left(1 - \frac{\sigma_t^2}{2\alpha}\right)^{-\alpha}.$$

The probability density function (PDF) of the VG distribution and its moments are given in Appendix A.

2.3. Bayesian Framework

Bayesian inference combines expert opinions with observational evidence, supplementing classical statistical inference or the frequentist approach once the likelihood is not trivial to obtain, and is performed via the joint posterior distribution of all model parameters. In this context, unlike the GARCH models, the SV models contain an additional random process in the volatility equation, resulting in an additional latent variable for each observation and an intractable likelihood function which involves very high dimensional integrals. Suppose that N asset returns, $\mathbf{y} \in \mathbb{R}^N$, are collected for statistical analysis. Then, there are $N + 1$ latent volatility variables, $\mathbf{h} = (h_0, h_1, \dots, h_N)$, in the model. Let $\boldsymbol{\theta} = (\mu, \phi, \tau^2, \alpha, \nu, \beta)$ be the vector of parameters of the VG- t SV model. The likelihood function is an analytically intractable N -dimensional integral of the form:

$$L(\boldsymbol{\theta}) = \int \prod_{t=1}^N f(y_t | h_1, \dots, h_N) f(h_0, \dots, h_N | \boldsymbol{\theta}) dh_0 \dots dh_N.$$

The parameters spaces of $\boldsymbol{\theta}$ and \mathbf{h} together can be viewed as an augmented parameter space. In the Bayesian paradigm, a full Bayesian approach for performing Bayesian inference is via the simulation-based MCMC algorithms, which iteratively sample posterior realisations from the joint posterior distribution

$$f(\boldsymbol{\theta}, \mathbf{h} | \mathbf{y}) \propto f(\mathbf{y} | \boldsymbol{\theta}, \mathbf{h}) f(\mathbf{h} | \boldsymbol{\theta}) f(\boldsymbol{\theta}),$$

where $f(\boldsymbol{\theta})$ is the PDF of the joint prior distributions of $\boldsymbol{\theta}$.

Since the VG and Student- t distributions belong to the class of SMN distributions, we can facilitate an efficient MCMC algorithm for Bayesian inference using a data augmentation technique. Choy and Chan [20] demonstrated Bayesian inference using univariate SMN distributions, including the VG and Student- t distributions. Therefore, our VG- t SV model under the \mathbb{P} measure can be expressed hierarchically as

$$y_t | h_t, \mu, \phi, \tau^2, \lambda_{y_t} \sim N(Q, \lambda_{y_t} \sigma_t), \tag{11}$$

$$h_t | h_{t-1}, y_t, \mu, \phi, \tau^2, \lambda_{h_t} \sim N\left(\mu + \phi(h_{t-1} - \mu), \lambda_{h_t}^{-1} \tau^2\right), \tag{12}$$

$$h_0 | \mu, \phi, \tau^2, \lambda_{h_0} \sim N\left(\mu, \frac{\tau^2}{\lambda_{h_0}(1 - \phi^2)}\right), \tag{13}$$

$$\lambda_{y_t} | \alpha \sim Ga\left(\frac{\alpha}{2}, \frac{\alpha}{2}\right), \quad t = 1, 2, \dots, T, \tag{14}$$

$$\lambda_{h_t} | \nu \sim Ga\left(\frac{\nu}{2}, \frac{\nu}{2}\right), \quad t = 0, 1, \dots, T, \tag{15}$$

where $Ga(c, d)$ is the gamma distribution with mean c/d . See [20] for Bayesian implementation of SMN distributions and [13] for the SV models. In this setup, λ_{h_t} and λ_{y_t} are scale mixture variables which can be used as a global diagnostic of potential outliers ([20]) in the return and volatility equations, respectively. The shape parameters ν and α capture the heavy-tailed features of the innovations in the return and volatility processes. This representation of the VG- t SV model enables a simple Gibbs sampler for posterior inference, because the full conditional distributions of the log-volatilities and the scale mixture variables will be the univariate Gaussian and gamma distributions.

To complete the Bayesian framework, we adopt the following prior distributions for model parameters. $\mu \sim N(a_\mu, b_\mu), \beta \sim N(a_\beta, b_\beta), \tau^2 \sim IG(a_\tau, b_\tau), \phi^* \sim Be(a_\phi, b_\phi), \nu \sim Exp(a_\nu), \alpha \sim Exp(a_\alpha),$

where $\phi^* = \frac{\phi+1}{2}$ and $N(\cdot, \cdot), IG(\cdot, \cdot), Be(\cdot, \cdot)$ and $Exp(\cdot)$ are the Gaussian, inverse gamma, beta, and exponential distributions. In the empirical study in Section 3, vague and non-informative prior distributions are assigned to μ, β, τ^2 , and α , respectively. For ϕ^* , an informative beta prior distribution is assumed. To ensure the existence of the mean and variance for the Student- t distribution, the exponential prior distribution of the degrees of freedom ν is restricted to the range (2.01, 40). To implement the Bayesian VG- t SV model, one can easily program the Gibbs sampling scheme to obtain posterior samples of the model parameters and the scale mixtures variables for Bayesian inference. Alternatively, we can use the user-friendly WinBUGS software, which just requires the specification of the model and the prior distributions.

2.4. Option Pricing

Our primary interest is option pricing. We use Figure 1 to explain the valuation procedure. Consider the option with maturity T days after today, when we are standing at time 0 in the perspective of pricing, but time N in the perspective of parameter estimation. Figure 1 shows 375 historical prices of the underlying asset, namely S_1, \dots, S_N , depicted as a random path on the left side. The estimation conducted using the Bayesian approach is carried out based on that path, and the estimated model is defined under \mathbb{P} . On the right, five sample paths are generated according to the model under \mathbb{Q} measure for valuing the option numerically. For each generated path, we get S_T^* at the end. As we consider the European call option in this paper, the payoff would be $\max(S_T^* - K, 0)$, where K is the strike price. With all the simulated sample paths, the option price at time-0 state for pricing (after discounted) can be obtained through the equation

$$C = e^{-rT} E^{\mathbb{Q}} \{ \max(S_T^* - K, 0) \}. \tag{16}$$

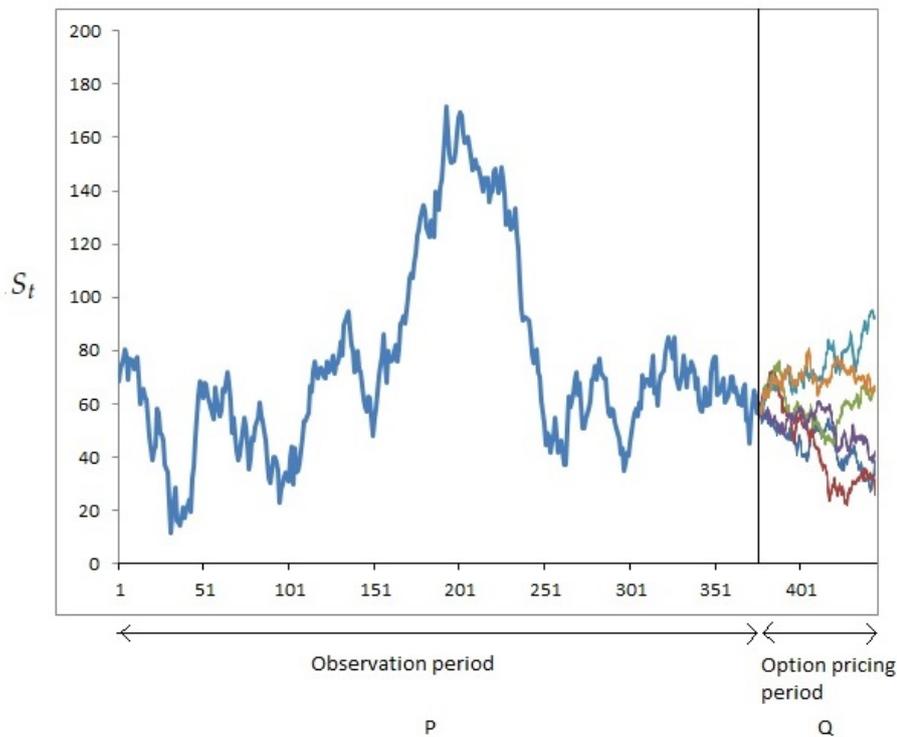


Figure 1. Historical asset price movement and simulated asset prices.

To simulate the \mathbb{Q} -process, the posterior mean of each parameter is regarded as the point estimate of the parameter value under \mathbb{P} . We transform the model into its risk-neutral equivalent using Theorem 1. The detailed algorithm is summarized in the algorithm below.

Posterior means of the model parameters $\theta = (\mu, \phi, \beta, \tau^2, \alpha, \nu)$, and log-volatilities $\mathbf{h} = (h_1, \dots, h_T)$ denoted by $\hat{\theta}$ and $\hat{\mathbf{h}}$, are obtained from WinBUGS output, and are used in a process to simulate realisations of the return and log-volatilities under the \mathbb{Q} -process. We perform the algorithm as follows.

- Step A: For $j = 1:J$ where J is the number of simulated stock price paths,
 - A1: Set $h_1^* = h_N$;
 - A2: Set $S_0^* = S_N$;
 - A3: Sample $\lambda_{h_1}^*, \dots, \lambda_{h_T}^*$ from $Ga(\frac{\hat{\nu}}{2}, \frac{\hat{\nu}}{2})$;
 - A4: Sample $\lambda_{y_1}^*, \dots, \lambda_{y_T}^*$ from $Ga(\frac{\hat{\alpha}}{2}, \frac{\hat{\alpha}}{2})$;
- Step B: For $i = 1:T$,
 - B1: Sample $y_i^*(j)$ from $N(r + \hat{\beta}e^{h_i^*(j)/2}, \lambda_{y_i(j)}^* e^{h_i^*(j)})$;
 - B2: Sample $h_i^*(j)$ from

$$N(\hat{\mu} + \hat{\phi}(h_{i-1}^*(j) - \hat{\mu}), (\lambda_{h_i(j)}^*)^{-1} \hat{\tau}^2);$$

$$B3: \text{Set } S_i^*(j) = S_{i-1}^*(j)e^{y_i^*(j)};$$

- Step C: By (16), set

$$C \cong \frac{e^{-rT}}{J} \sum_{j=1}^J \max(S_T^*(j) - K, 0). \tag{17}$$

As each simulated path starts the day that the observations end, h_{N+1} —the volatility right after the last point of the observations—is needed as the simulation’s initial volatility. More specifically, Step A shows that the posterior means of h_{N+1} is fixed for all simulated paths. S_0^* represents the initial price in the simulation part. The rest of the volatilities and mixing parameters are randomly drawn from their posterior distribution, with the hyperparameter values obtained from the estimation. Once h_{N+1} is fixed, the risk-neutral return at $(N + j)$, denoted as y_j^* , is generated, as shown in Step B. Note that h_1^* is used in the simulation, where h_1^* and h_{N+1} are related through Theorem 1. Then, option price is obtained through Step C.

3. Empirical Studies

This section examines the empirical performance of the proposed approach in FX option pricing. We collect the daily returns of an FX rate, estimate the $VG-t$ SV model parameters with the proposed Bayesian framework, and simulate option prices under the risk-neutral measure. The empirical performance is based on the out-of-sample comparison between the estimated option prices and the market-observed option prices. We pretend that the market option prices are unavailable during our implementation. The option prices evaluated by the Black-Scholes (BS) formula are also reported as a benchmark. Although it is known that GARCH option pricing models perform well in the FX market, they fail to capture the U-shaped volatility smile. By fitting our proposed $VG-t$ SV model to the FX market data, we empirically show that our model overcomes this problem by providing a set of implied volatilities within the bid-ask spread of the market-observed implied volatility smile.

Specifically, the estimation is implemented using the WinBUGS package (<http://www.mrc-bsu.cam.ac.uk/bugs/>) in R. We implement a Gibbs sampler by running a single Markov chain for 20,000 iterations and discarding the first 3000 iterations as the burn-in period to better guarantee the convergence of the estimation result. Sampling procedures and prior distributions are chosen according to Section 2.3. After estimation, simulation for option pricing starts at the last value of observation. From this starting point, we generate $J = 8000$ FX rates under the risk-neutral measure according to Section 2.4. Based on the simulated option prices across different strike prices, we calculate the implied volatility and make a comparison with BS formula and the market-observed volatility smile.

3.1. FX Market Data

The empirical study uses one year of data points on the YEN–USD exchange rate from 1 March 2012 to 28 February 2013. To confirm that the $VG-t$ SV model is the best model, the FX data is also fitted by the basic SV models ([13]) with $N-N$ (normal-normal), $t-t$, $VG-VG$ and $t-VG$ error distributions in the innovations terms. The Deviance Information Criterion (DIC) and Bayesian Information Criterion (BIC) values of the five models are displayed in Table 1. It is obvious that all models with heavy-tailed error distributions are superior to the model with Gaussian errors. Moreover, the models with VG innovation in the return equation outperform the models with Student- t innovation, and in particular, the $VG-t$ SV model is chosen to be the best amongst the five models, according to both DIC and BIC values.

Table 1. Deviance information criterion (DIC) and Bayesian information criterion (BIC) of stochastic volatility (SV) models with different error distributions. $N-N$: normal-normal; VG : variance gamma.

Model	$N-N$	$t-t$	$t-VG$	$VG-t$	$VG-VG$
DIC	383.3	379.9	378.6	353.3	360.2
BIC	386.6	376.9	379.7	370.9	371.0

European calls on the YEN–USD exchange rate with different strike prices are collected for out-of-sample comparison, with the strike price ranging from 87.94 to 99.79. Half of the options are in-the-money, while the others are out-of-the-money. We use the one-month Libor rates as the constant interest rate for both the estimation and option simulations. Table 2 presents the summary statistics for the source data on the YEN–USD exchange rate. Implied volatilities against different strike prices are given in Table 3, in which the market option price is calculated from these market volatilities through the BS formula.

Table 2. Summary statistics of return data on the foreign exchange (FX) rate.

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.	Sd.
−0.0188	−0.0027	−0.0003	0.0002	0.0031	0.0310	0.0052

Table 3. Implied volatilities of 1-month option data on the FX rate.

	$K = 87.94$	$K = 89.28$	$K = 90.15$	$K = 91.37$	$K = 92.32$	$K = 93.53$
Mid	0.1318	0.1281	0.1258	0.1229	0.1213	0.1205
Bid	0.1154	0.1184	0.1185	0.1176	0.1167	0.1163
Ask	0.1483	0.1377	0.1331	0.1283	0.1259	0.1247
	$K = 93.84$	$K = 95.88$	$K = 97.24$	$K = 98.23$	$K = 99.79$	
Mid	0.1219	0.1241	0.1274	0.1301	0.1343	
Bid	0.1175	0.1189	0.1203	0.1206	0.1182	
Ask	0.1264	0.1293	0.1345	0.1395	0.1503	

Table 4 summarizes the estimation results from WinBUGS using the historical FX rates. The small value of ν supports the adoption of a heavy-tailed innovation in the volatility process. The initial spot FX rate for the option simulation was 93.59, on 1 March 2013. The estimated option prices against different strike prices are presented in Table 5. It can also be seen from Figure 2 that our model greatly outperforms the BS model, because the model prices (crosses in the graph) are far over the straight line of the implied volatility from the BS model, and very close to the volatilities implied by market mid-prices. In addition, all implied volatilities from our model are between the market bid and ask volatilities. As actual transactions take place within the bid-ask prices rather than the mid-price, our model offers a reliable tool for practical option prices when the market lacks liquid option data.

Furthermore, our $VG-t$ SV model captures the volatility smile very well, and this is in consistent with the U-shaped smile of the market-implied volatility. On the contrary, the BS model only offers an unrealistic constant volatility. Although not presented in the figure, the standard GARCH option pricing models also fail to produce such a U-shaped volatility smile for FX options.

Table 4. Summary statistics of parameters: option on the YEN–USD FX rate from WinBUGS.

Par.	Mean	Sd	95% CI
β	0.078	0.064	(−0.046 , 0.203)
τ	0.225	0.067	(0.128 , 0.361)
μ	−10.420	0.351	(−10.950 , −9.585)
ϕ	0.873	0.088	(0.687 , 0.988)
ν	8.176	4.155	(2.531 , 18.114)
α	15.622	8.582	(4.290 , 34.868)

Table 5. Results of option prices on the YEN–USD FX rate from the Black-Scholes (BS) model, the SV model, and the market.

Model	K = 87.94	K = 89.28	K = 90.15	K = 91.37	K = 92.32	K = 93.53
Mkt	5.728	4.480	3.711	2.712	2.040	1.335
BS	5.658	4.339	3.507	2.410	1.676	0.934
SV	5.679	4.372	3.555	2.480	1.759	1.020

Model	K = 94.84	K = 95.88	K = 97.24	K = 98.23	K = 99.79
Mkt	0.049	0.509	0.272	0.169	0.077
BS	0.414	0.189	0.054	0.019	0.003
SV	0.497	0.259	0.104	0.055	0.022

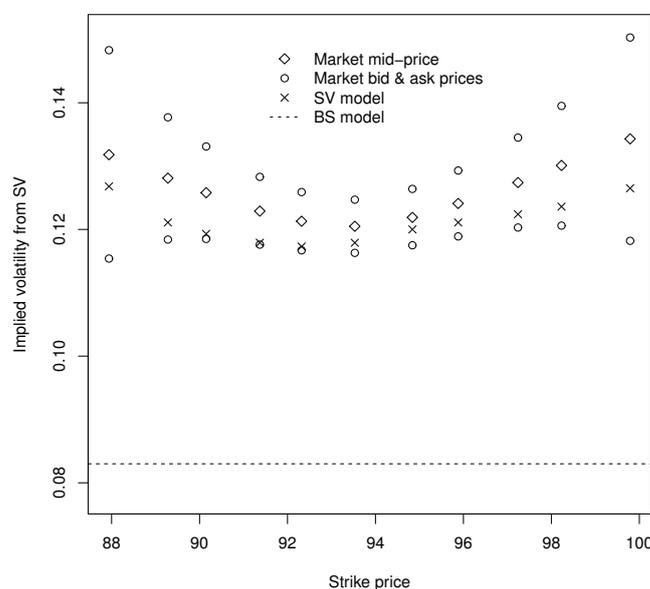


Figure 2. Comparison of volatility smile on the YEN–USD FX rate from the Black-Scholes (BS) model, $VG-t$ SV model, and the market.

Figure 3 plots volatility premium across strike price, which is the difference between the implied volatilities of the market mid-price and the model. The volatility premium is generally flat across strike prices, and the implied volatility from our model can be shifted up to approach the market price if a massive amount of option data is available.

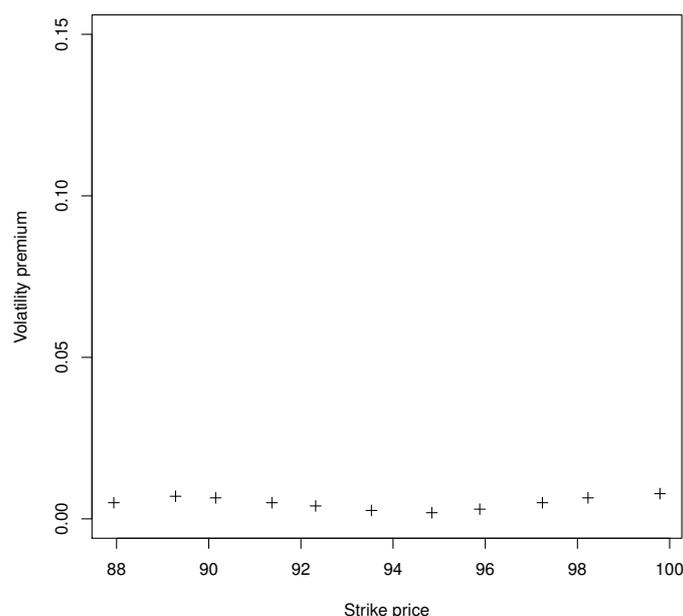


Figure 3. Volatility premium on the 1-month YEN–USD option

4. Conclusions

In this paper, we propose an SV option pricing model inspired by Duan’s GARCH option pricing model, and extend our model by allowing return and volatility to have VG and Student- t innovations, respectively. The result is an alternative way to price European options when calibration is not applicable. This model enjoys simplicity, reduces simulation inaccuracy, provides an easy and convenient way of transforming the model between physical and risk-neutral measures, and allows the Bayesian inference to apply in the estimation procedure with the benefit of the SMN representation in the MCMC method. In the empirical study, our SV pricing model outperforms its competitive models in the out-of-the-sample test using European call option data, as the model prices are close to the market option prices compared with the BS formula. Furthermore, the model implemented with a Bayesian framework effectively captures the volatility smile in the FX option market, whereas standard GARCH option pricing models only produce a volatility skew.

Future research can focus on alternative specifications of the drift term in the return process, such as the mean reversion process considered by [21]. The proposed Bayesian pricing framework can be generalized to other SV models such as the DSARV model by [7], and to incorporate information for the Volatility Index (VIX) data as suggested by [22].

Acknowledgments: The authors would like to thank the editor and the two anonymous reviewers for their constructive comments and suggestions, which helped improve the presentation of the paper.

Author Contributions: All authors contributed equally to the paper. H.Y.W. initiated the idea; S.T.B.C. and H.Y.W. established the theoretical results while Y.W. and S.T.B.C. worked on the numerical analysis; Y.W. and S.T.B.C. wrote the paper.

Conflicts of Interest: The author declares no conflict of interest.

Appendix A

Definition A1. A univariate random variable X is said to be distributed as a symmetric variance gamma $VG(\mu, \sigma^2, \alpha)$ distribution if its PDF is given by

$$f(x) = \frac{2\alpha^{\alpha+1/2} |z|^{\alpha-1/2}}{\sigma\pi^{1/2}\Gamma(\alpha)} K_{\alpha-1/2} \left((2\alpha)^{1/2} |z| \right), \quad x \in \mathbb{R},$$

where $\mu \in \mathbb{R}$ is the location parameter, $\sigma \in \mathbb{R}^+$ is the scale parameter, $\alpha \in \mathbb{R}^+$ is the shape parameter, $z = \sigma^{-1}(x - \mu)$ and $K_\alpha(\cdot)$ is a modified Bessel function of the second kind. The mean, variance and kurtosis are given by $E(X) = \mu$, $V(X) = \sigma^2$ and $Kur(X) = 3\alpha(\alpha - 1)$, $\alpha > 1$, and the moment generating function is

$$M_X(t) = \left(1 - \frac{\sigma^2 t^2}{2\alpha}\right)^{-\alpha} e^{\mu t}.$$

References

1. Engle, R.F. Autoregressive conditional heteroscedasticity with estimates of variance of United Kingdom inflation. *Econometrica* **1982**, *50*, 987–1008.
2. Bollerslev, T. Generalized autoregressive conditional heteroskedasticity. *J. Econ.* **1986**, *31*, 307–327.
3. Duan, J.C. The GARCH option pricing model. *Math. Financ.* **1995**, *5*, 13–32.
4. Nelson, D.B. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica* **1991**, *59*, 347–370.
5. Hull, J.C.; White, A. The pricing of options on assets with stochastic volatilities. *J. Financ.* **1987**, *42*, 281–300.
6. Heston, S.L. A closed solution for options with stochastic volatility, with application to bond and currency options. *Rev. Financ. Stud.* **1993**, *6*, 327–343.
7. Cordis, A.S.; Korby, C. Discrete stochastic autoregressive volatility. *J. Bank. Financ.* **2014**, *43*, 160–178.
8. Jacquier, E.; Polson, N.; Rossi, P. Bayesian analysis of stochastic volatility models (with discussion). *J. Bus. Econ. Stat.* **1994**, *12*, 371–417.
9. Shephard, N.; Pitt, M.K. Likelihood analysis of non-Gaussian measurement time series. *Biometrika* **1997**, *84*, 653–667.
10. Abanto-Valle, C.A.; Bandyopadhyay, D.; Lachos, V.H.; Enriquez, I. Robust Bayesian analysis of heavy-tailed stochastic volatility models using scale mixtures of normal distributions. *Comput. Stat. Data Anal.* **2010**, *54*, 2883–2898.
11. Meyer, R.; Yu, J. BUGS for a Bayesian analysis of stochastic volatility models. *Econ. J.* **2000**, *3*, 198–215.
12. Yu, J. On leverage in a stochastic volatility model. *J. Econ.* **2005**, *127*, 165–178.
13. Choy, S.T.B.; Wan, W.Y.; Chan, C.M. Bayesian Student-t stochastic volatility models via scale mixtures. *Adv. Econ.* **2008**, *23*, 595–618.
14. Wang, J.J.J.; Chan, S.K.J.; Choy, S.T.B. Stochastic volatility models with leverage and heavy-tailed distributions: A Bayesian approach using scale mixtures. *Comput. Stat. Data Anal.* **2011**, *55*, 852–862.
15. Abanto-Valle, C.A.; Migon, H.; Lachos, V.H. Bayesian analysis of heavy-tailed stochastic volatility in mean model using scale mixtures of normal distributions. *J. Stat. Plan. Inference* **2011**, *141*, 1875–1887.
16. Bates, D.S. Maximum likelihood estimation of latent affine processes. *Rev. Financ. Stud.* **2006**, *19*, 909–965.
17. Madan, D.B.; Seneta, E. The variance gamma (V.G.) model for share market returns. *J. Bus.* **1990**, *63*, 511–524.
18. Madan, D.B.; Carr, P.P.; Chang, E.C. The variance gamma process and option pricing. *Eur. Financ. Rev.* **1998**, *2*, 79–105.
19. Badescu, A.; Elliott, R.J.; Siu, T.K. Esscher transforms and consumption-based models. *Insur. Math. Econ.* **2009**, *45*, 337–347.
20. Choy, S.T.B.; Chan, J.S.K. Scale mixtures in statistics modelling. *Aust. N. Z. J. Stat.* **2008**, *50*, 135–146.
21. Chung, S.F.; Wong, H.Y. Analytical pricing of discrete arithmetic Asian options with mean reversion and jumps. *J. Bank. Financ.* **2014**, *44*, 130–140.
22. Kannianen, J.; Lin, B.; Yang, H. Estimating and using GARCH models with VIX data for option valuation. *J. Bank. Financ.* **2014**, *43*, 200–211.

