

Forecasting Latent Volatility through a Markov Chain Approximation Filter

CHIA CHUN LO,^{1*} KONSTANTINOS SKINDILIAS,^{2,3*} AND
ANDREAS KARATHANASOPOULOS⁴

¹ Department of Finance and Business Economics, Faculty of Business Administration, University of Macau

² Mathematical Sciences Department, University of Greenwich, London, UK

³ ABM-Analytics Ltd, London, UK

⁴ Suliman S. Olayan School of Business, American University of Beirut, Lebanon

ABSTRACT

We propose a new methodology for filtering and forecasting the latent variance in a two-factor diffusion process with jumps from a continuous-time perspective. For this purpose we use a continuous-time Markov chain approximation with a finite state space. Essentially, we extend Markov chain filters to processes of higher dimensions. We assess forecastability of the models under consideration by measuring forecast error of model expected realized variance, trading in variance swap contracts, producing value-at-risk estimates as well as examining sign forecastability. We provide empirical evidence using two sources, the S&P 500 index values and its corresponding cumulative risk-neutral expected variance (namely the VIX index). Joint estimation reveals the market prices of equity and variance risk implicit by the two probability measures. A further simulation study shows that the proposed methodology can filter the variance of virtually any type of diffusion process (coupled with a jump process) with a non-analytical density function. Copyright © 2015 John Wiley & Sons, Ltd.

KEY WORDS filtering variance; forecasting probability; continuous time Markov chains

INTRODUCTION

In this paper we propose a new methodology for estimating the latent variance process in a two-factor diffusion process with jumps. A by-product of our approach is the transition probability matrix of the return variance process which we use to forecast volatility from the discretely observed data generated by the continuous-time stochastic volatility jump diffusion models. The demand for volatility derivative contracts has grown substantially over the past years and admits the importance of modelling, and thus successfully forecasting, volatility, which constitutes an important factor when pricing derivative contracts and hedging volatility risk. Furthermore, the importance of appropriate volatility (variance) forecasting is further depicted in the forecastability of asset returns. Financial econometric theory suggests that, while asset returns are not easily forecast, asset return volatility exhibits dependence and is thus forecastable.¹ That is, given the evidence of volatility dependence one should expect sign (market direction) dependence and hence forecastability, a property of many profitable trading strategies (Christoffersen and Diebold, 2006).

The variance is not directly observable and thus direct inference is not possible. As a result a filtering algorithm must be used to model the unobserved (latent) variance process. Optimal filters have recently become popular for estimating the parameters of latent processes. Filtering distributions provide estimates of latent variables that are useful for forecasting volatility or price distributions. They also allow researchers to identify or disentangle transient from persistent shocks (Ait-Sahalia, 2004). The studies of Johannes (2004), Johannes *et al.* (2009) and Christoffersen *et al.* (2010) provide a seminal application of optimal filters for volatility forecasting, in which popular time discretization schemes (e.g. the Euler scheme) are combined with particle simulations methods such as those of Pitt and Shephard (1999) and Arulampalam *et al.* (2002). Their approach is quite generic, as it applies Monte Carlo simulations to the underlying multivariate jump diffusion models; thus this approach even allows for models whose probability density functions are not known in closed form that commonly arise in derivatives pricing applications. The fundamental characteristic of their approach is time discretization of the underlying problem. The accuracy of parameter estimates, one may argue, depends greatly on the number of simulations (particles) at each discrete time step. That is, the robustness of particle filters may come at a cost of computational time. In this respect, we introduce a filtering approach from a continuous-time perspective making use of the continuous-time Markov chain

* Correspondence to: Chia Chun Lo and Konstantinos Skindilias, Mathematical Sciences Department, University of Greenwich, London SE10 9LS, UK. E-mail: k.skindilias@gre.ac.uk

¹ The autocorrelation function of the volatility exhibits long-range dependence and is well described by a power law decay (Bollerslev and Mikkelsen, 1996; Andersen *et al.*, 2001a, 2001b).

approximation (MCA) of Kushner and DiMasi (1978), Kushner (1990) and Kushner and Dupuis (2001). Contrary to most time discretization schemes, continuous-time Markov chains enjoy probabilistic convergence proofs (see Kushner and DiMasi, 1978), which means that the proposed method in this paper does not require any time discretization between market observations.

Recent studies in the literature provide empirical evidence on the good approximating features of the MCA to continuous-time processes (Chourdakis, 2004; Skindilias and Lo, 2013; Lo and Skindilias, 2014b). Although for single-factor processes Markov chain filters have been already designed (see Chourdakis and Dotsis, 2011; Lo and Skindilias, 2014a), two-factor (or multi-factor) processes with correlated Brownian motions approximations do not exist in the literature. One viable reason is that no sound (asymptotic) convergence and empirical proofs exist for continuous-time multidimensional Markov chains as there are for the one-dimensional Markov chains (which rely on the local consistency condition of Kushner, 1990). This indicates that when we approximate models with correlated Brownian motions using an MCA the rate generator matrix of the Markov chain may not be well defined, leading to a negative transition probability issue. We will review the conditions of a well-defined rate generator matrix below ('The approximating chain to a variance process').

In this paper, for the proposed filtering algorithm we have taken advantage of the existing probabilistic convergence proofs that underlie the one-factor MCA, and exploit these by heuristically accommodating the two-factor problem in one dimension. As such we successfully design a continuous-time, discrete state space filtering algorithm for estimating the unobserved volatility process using a continuous-time Markov chain, which we thus call the Markov chain approximation filter (MCAF). The method is fast and accurate, as we discuss and show further in the text. Thus, in contrast to particle filters referenced above, the proposed MCAF does not require us to simulate an approximation of the distribution of the latent variables. Instead, we use continuous-time Markov chains to model the transition probabilities of the latent process. Additionally, we build the Markov chain in continuous time; this means there is no need for time discretization between two market observations. The approach primarily involves matrix operations and therefore naturally offers acceleration over simulation techniques. This is an important implication for real-time forecasting applications.

In order to illustrate the proposed MCAF we choose to work with the well-known affine (univariate) continuous-time stochastic volatility model of Heston (1993), a model that has enjoyed arguably significant popularity in the financial literature.² The fact that investors price equity and volatility risk has been thoroughly illustrated in the literature (see Pan, 2002; Carr and Wu, 2009; Duan and Yeh, 2010; Ait-Sahalia *et al.*, 2015; and references therein for further studies) and as such the risk premia required by investors are depicted in continuous-time models through the market prices of equity and volatility risk, respectively. Furthermore, the jump component is crucial for capturing the leptokurtic feature of the return distribution.³ But from a time series analysis perspective estimation of the jump amplitude parameters is ill posed,⁴ thus making the choice of appropriate parameters hard enough. Additionally, to address the leverage effect (between equity returns and volatility), often unrealistically high negative correlation between the stock index and volatility is required. Numerous authors have argued that it is very hard to obtain a concise value for this correlation coefficient (see Johannes *et al.*, 2009). Tankov and Cont (2003) have shown that one may set this correlation equal to zero.⁵ A volatility model must be able to forecast volatility; this is the central requirement in almost all financial applications. To this end, we examine model specification under three cases: (i) one in which we estimate the variance process including jumps in the return process while also estimating the correlation coefficient of the leverage effect; (ii) another, in which we fix the correlation to zero; and (iii) one in which we examine the case where no jump component is incorporated in the return process. We assess performance by a simple trading strategy in variance swaps. A variance swap is, in principle, a simple contract: it is consisted of a fixed leg agreed at contract inception which determines the fixed amount to be paid at maturity, the so-called variance swap rate, and a floating leg for which the counterpart receives in exchange a floating amount based on the sum of squared daily log-returns, i.e. the realized variance, over the lifespan of the variance swap contract. Volatility forecastability in this simple trading strategy is crucial as the fixed leg represents the expected integrated variance under the risk-neutral probability measure, and the floating leg represents the expected integrated variance under the objective (or statistical) probability measure over the contracted horizon. We find that the model-based signal from our variance forecasts provides significant returns when compared to a benchmark strategy. We note that the jump component is an important addition and, while the model with leverage effect parameter outperforms all other cases, one may avoid estimating the correlation coefficient as both cases give similar signals. We produce value-at-risk forecasts and note similar results for each model under consideration. A different conclusion is drawn when we examine sign

² The importance of introducing a continuous-time model relies on the fact that only a few discrete-time models with time-varying volatility enjoy a closed-form solution of the density function for increasing time horizons.

³ Andersen *et al.* (2002) conclude that allowing jump improves the performance of the model in terms of time series fitting, whereas Bakshi *et al.* (1997) and Pan (2002), among others, offer different conclusions in terms of option pricing.

⁴ There exist a substantial number of solutions that provide a good fit to the market observed data.

⁵ In which case one may still be able to capture the short-term skew implied in options.

forecastability. Our findings are in agreement with the observations of Christoffersen and Diebold (2006) but are not along the same lines as the above tests in terms of which model performs better.⁶

Furthermore, our simulation study shows that the proposed MCAF may accommodate virtually any type of diffusion process, particularly those with unknown density functions for which numerical approximations need to be employed. The implications of the simulation study are thus twofold. First, we show that the proposed methodology can accommodate more complicated functional forms. This mitigates the issue of parameter restrictions required in the Heston (1993) SV model that restricts the volatility of volatility to be small relative to the unconditional variance, which in turn restricts the amount of sign predictability (sign predictability is higher when the volatility of volatility is higher (Christoffersen and Diebold, 2006)). Secondly, we show that MCAF parameter estimates are free from spurious nonlinearity, making the proposed filtering algorithm a promising candidate for estimating such processes.

The study is performed on the S&P 500 index and volatility index (VIX) of the Chicago Board Options Exchange (CBOE). Joint estimation of the S&P 500 index and VIX volatilities allows one to obtain estimates on the market prices of risk and hence the market risk premia. This is because the S&P 500 admits the objective probability measure and VIX pricing admits the risk-neutral probability measure. The resulting latent variance (under the two different probability measures) naturally allows for computation of the expected integrated variances for the variance swap counterparts. Furthermore, the VIX index is the reference for the fixed leg of the variance swap contract with length 30 days, as the model-free construction of the VIX index provides a measure of the market's expectation of 30-day implied volatility. Henceforth, joint estimation under the two probability measures allows one to extract information based on the dynamics of the historical time series but also the market views of future expected variance thus naturally assisting in forecasting.

The rest of this paper expands in the following manner. In the next section we discuss the model specification of the Heston SV model with jumps and the counterparts of the joint estimation process. The MCA to a stochastic volatility process and the filtering algorithm (MCAF) for the latent variance are described. In the third section we provide results of our empirical analysis in which we test each model specification's capability in forecasting accurately the variance level and provide proofs of the strength of predictability. The fourth section consists of our simulation analysis and the fifth section concludes.

MODELLING AND FILTERING VARIANCE

Model specification

We examine a parametric model specification for which we choose the stochastic differential equation (SDE) to take a well-known structure that addresses stylized facts of observed asset return and its variance. Under the statistical probability measure \mathbb{P} the stochastic volatility model with affine drift (Heston, 1993) and jumps (Bates, 1996) in the return process is given by

$$\begin{aligned} d \ln S_t &= \mu_t dt + \sqrt{V_t} dW_t + J dN_t \\ dV_t &= \kappa(\theta - V_t) dt + v \sqrt{V_t} dB_t \end{aligned} \quad (1)$$

where μ_t is the expected return on the asset. The arrival of jumps is governed by a Poisson counting process N_t with jump intensity rate given by $\lambda(t, S_t) = \lambda$. Thus the probability of a jump in dt is given by $\Pr(dN_t = 1) = \lambda dt$. Following Merton (1976), we assume that the jump size J is normally distributed, $J \sim N(\mu_J, \sigma_J^2)$. W_t and B_t are two Brownian motions with a correlation parameter ρ , which addresses the so-called leverage effect; the Poisson process and Brownian motions are assumed to be independent. In the variance process κ is the coefficient of mean reversion speed over the long-term value of θ , and v represents the volatility of volatility parameter.

In order to compute variance swap rates, we use the standard approach of the risk-neutral probability measure \mathbb{Q} , which implies that the discounted asset price process is a martingale with respect to an equivalent probability measure. Therefore, the compensator parameter in the return equation must be defined by

$$m := E(e^J - 1) = \exp\left(\mu_J + \frac{\sigma_J^2}{2}\right) - 1$$

to ensure that under the risk-neutral probability measure the expected return in the asset is equal to the risk-free rate. When a jump occurs the asset price changes in value by $(S_t - S_{t-})/S_{t-} = (e^J - 1)$.⁷ Thus m represents the average jump size. The corresponding return process under the risk-neutral probability measure \mathbb{Q} is thus given as

$$d \ln S_t = \left(r - q - \frac{V_t}{2} - \lambda m\right) dt + \sqrt{V_t} dW_t^* + J dN_t$$

⁶This is an interesting observation, which is out of the scope of this paper.

⁷Note that $\log(S_t/S_{t-}) = J$.

where $dW_t^* = dW_t + \left(\frac{\mu-r}{\sqrt{V_t}}\right) dt$ and q is the dividend yield. Suppose the volatility risk premium $\eta(S, V; t)$ is equal to $\eta_V V_t$, so that the risk-neutral process for the variance is obtained by

$$\begin{aligned} dV_t &= (\kappa(\theta - V_t)dt - \eta(S, V; t)) dt + v\sqrt{V_t}dB_t^* \\ &= (k(\theta - V_t) - \eta_V V_t) dt + v\sqrt{V_t}dB_t^* \end{aligned} \tag{2}$$

where $dB_t^* = dB_t + \left(\frac{\eta(S, V, t)}{v\sqrt{V_t}}\right) dt$. W_t^* and B_t^* are the two Brownian motions under the risk-neutral measure maintaining the correlation ρ of the statistical measure and r is the risk-free rate. We assume that the jump dynamics remain the same under the change of probability measure. In order to interpret the equity risk premia, we rewrite the return parameter μ_t in equation (1) by $\mu_t = (r - q - V_t/2 + \eta_S V_t)$. This setting is also similar to the Garch-in-mean model in the discrete-time framework.

Our methodology consists of joint estimation of the return series of the the S&P 500 index and the observed variance swap rate (the fixed leg of the contract) for which we choose to use the VIX² index, the S&P500 time series corresponding (annualized) volatility index. The VIX index measures the market’s expectation of the 30-day forward S&P 500 index volatility implicit in the index option prices, and is thus the reference variance swap rate for a 30-day variance swap contact. Since VIX represents the volatility implicit in the index option prices, it constitutes the market’s risk-neutral expectation of the 30-day forward volatility, thereby requiring the use of the parametric SDE under the risk-neutral probability measure that we defined earlier. The market prices of risk, η_S and η_V , relate the two probability measures while allowing inference on the market risk premia. Given the stochastic volatility model above, the 30-day variance swap rate (VIX²) can be easily computed following the known result of Carr and Wu (2009), which relates variance swap rate computed from option prices to the time t risk-neutral expectation of the asset’s integrated variance over a period $[t, T]$ denoted by $IV_{t,T}$:

$$\frac{2}{T-t} \int_0^\infty \frac{O_t(K, T)}{P_t(T)K^2} dK + \epsilon_t = E^Q[IV_{t,T}] = \frac{1}{T-t} E^Q \left[\int_t^T V_s ds \right] + 2\phi \tag{3}$$

where $P_t(T)$ is the value of a bond at time t that pays one dollar at time T and K is the strike price. $O_t(K, T)$ denotes the time- t value of an out-of-the-money option with strike price $K > 0$ and maturity $T \geq t$ (a call option when $K > F_t$ and a put when $K \leq F_t$), where F_t is the time- t forward price. The term ϵ_t and ϕ determine the compensator of the discontinuous component, namely the jump process.

The right-hand side of equation (3) provides an alternative vehicle for making inference on the variance dynamics that contains rich information with respect to the stochastic processes followed by the unobserved latent variance and with respect to the variance risk premium. For the stochastic process in equation (2) the variance swap rate can be computed in closed form⁸ and is given by

$$VIX_t^2 = \frac{1}{T-t} \int_t^T E_t^Q(V_s) ds + 2\phi \tag{4}$$

with $\phi = \lambda(e^{\mu_J^2 + \sigma_J^2/2} - 1 - \mu_J)$ and where the risk-neutral expected cumulative variance is computed by

$$\int_t^T E^Q(V_s) ds = \frac{\kappa\theta}{\kappa^*} \left(\tau - \frac{1 - e^{-\kappa^*\tau}}{\kappa^*} \right) + \frac{1 - e^{-\kappa^*\tau}}{\kappa^*} V_t \tag{5}$$

where $\kappa^* = \kappa + \eta_V$ and $\tau = T - t$. In the following subsections we discuss the mechanics for approximating the variance process using a MCA, filtering the latent variance in the two-factor stochastic process, and the joint estimation setup of (S&P 500, VIX).

The approximating chain to a variance process

We let the variance process in equation (2) be approximated by an n -state Markov chain $\{V_t^h : t \geq 0\}$ taking positive values from a finite-element grid:

$$\mathbf{v} = \{v_1^h, v_2^h, \dots, v_n^h\}'$$

where h denotes the grid spacing or the distance between any two adjacent grid elements in set \mathbf{v} . We note that it is not necessary to assume the grid elements are equidistant (see Lo and Skindiliias, 2014b). Define an $n \times n$ rate

⁸ Where analytical solutions do not exist we show in the following sections how one can estimate this using the MCA

generator matrix by $\mathbf{Q} = (q_{i,j} : i \neq j)$, with the rate elements $q_{i,j}$ subject to the conditions: $q_{i,i} \leq 0$, $q_{i,j} \geq 0$ and $\sum_j q_{ij} = 0$. The transition probability of variance moves state v_i^h to v_j^h in time t is obtained by⁹

$$\mathbf{P}(t) = (p_{i,j}(t)) = e^{t\mathbf{Q}} = \sum_{j=0}^{\infty} \frac{(t\mathbf{Q})^j}{j!} = \mathbf{I} + \sum_{j=1}^{\infty} \frac{(t\mathbf{Q})^j}{j!} \tag{6}$$

where \mathbf{I} is the unit matrix. In this paper, we use the model proposed by Lo and Skindilias (2014b) to construct \mathbf{Q} . To save space, we present their formula in equation (28) in the Appendix. When pricing variance derivatives, we will need to approximate the risk-neutral variance process in equation (2) by constructing a risk-neutral rate generator matrix \mathbf{Q}^* , and the risk-neutral transition probability $\mathbf{P}^*(t)$ can also be obtained by equation (6).

MCAF and variance forecasting

We present our MCAF for calibrating the return variance system in equations (1) and (2). In contrast to the filter algorithm proposed by Chourdakis and Dotsis (2011), their methodology only applies to the univariate system. Our model is able to accommodate a multivariate system. Before we introduce our methodology, we define $a(V_t)$ and $b(V_t)$, the drift and diffusion functions of the variance process, in equation (2). In our return variance system, we have $a(V_t) = \kappa(\theta - V_t)$ and $b(V_t) = v\sqrt{V_t}$. From equation (2) we can express

$$dB_t = \frac{dV_t - a(V_t) dt}{b(V_t)} \tag{7}$$

Given that the correlation between B_t and W_t is ρ , we introduce another Brownian motion, Z_t , which is independent of B_t and W_t ; then we can write

$$dW_t = (\rho dB_t + \sqrt{1 - \rho^2} dZ_t) \tag{8}$$

Taking equations (7) and (8) in equation (1), we obtain

$$\begin{aligned} d \ln S_t &= \left(\mu_t - \frac{1}{2} V_t \right) dt + \sqrt{V_t} \rho \left(\frac{dV_t - a(V_t) dt}{b(V_t)} \right) + \sqrt{V_t} \sqrt{1 - \rho^2} dZ + dJ_t \\ &= \left(\mu_t - \frac{1}{2} V_t - \rho \sqrt{V_t} \frac{a(V_t)}{b(V_t)} \right) dt + \rho \sqrt{V_t} \frac{dV_t}{b(V_t)} + \sqrt{V_t} (1 - \rho^2) dZ_t + dJ_t \end{aligned} \tag{9}$$

Because financial data are observed in discrete time, in order to develop the econometrics inference for estimation we consider the discrete version of equation (9) and apply the Poisson approximation¹⁰ to evaluate the conditional density function of log-return $y_t = \Delta \ln S_t$. We obtain

$$\Phi_{y_t}(y_t|\Theta) = \sum_k^{\infty} \phi(y_t; \mu_y + k\mu_J, \sigma_y^2 + k\sigma_J^2) \frac{\lambda e^{-\lambda k}}{k!} \tag{10}$$

where Θ is the unknown parameters set, $\phi(\cdot; a, b)$ is the normal density function with mean a and variance b and

$$\mu_y = \left(\mu - q - \frac{1}{2} V_t - \rho \sqrt{V_t} \frac{a(V_t)}{b(V_t)} \right) \Delta t + \rho \sqrt{V_t} \frac{V_t - V_{t-1}}{b(V_t)} \tag{11}$$

and

$$\sigma_y^2 = V_t (1 - \rho^2) \Delta t \tag{12}$$

where Θ is the unknown parameters set, and $\phi(\cdot|a, b)$ is the normal density function with mean a and variance b .

When y_t becomes available, we note that V_t in equation (11) has only n different values from an n -state Markov chain $\mathbf{v} = \{v_1^h, v_2^h, \dots, v_N^h\}'$. This is the advantage of MCAF over the particle filter, which requires simulations of at least 500–1000 different particles (variances) at each time step. The proposed model keeps the states of variances unchanged, while adjusting or updating the transition probability through time when observations become available.

Because the variance process is hidden, V_{t-1} in equation(11) is still unknown. To overcome this issue, we propose to use the filtered variance \hat{V}_{t-1} to replace V_{t-1} in equation (11), and we found that the performance of this approximation works very well in our simulation exercises. We will discuss how to obtain the filtered variance \hat{V}_t series by using our MCAF shortly, and the results from our simulation exercises will also be presented.

⁹ For a rigorous discussion of Markov chain theory, we refer readers to Rogers and Williams (2000) and Elliott (1995).

¹⁰ When jump frequency λ and time interval Δt are small, we found that using a Bernoulli approximation can also generate very precise results.

As discussed in the previous subsection, the variance process in equation (2) is approximated by a local consistency Markov chain with a transition matrix in equation (6), and conditional density of log-return y_t is also available by using equation (10). We now turn our attention to the model calibration by means of MCAF.

We define $\mathbf{e}(t)$ as a stochastic process and we assume $\mathbf{e}(t) = \mathbf{e}_i$ when the value of variance at t is the value of the i th element on the Markov chain, i.e. $\mathbf{V}_t^h = v_i^h$. The vector \mathbf{e}_i is an $N \times 1$ vector whose value is 1 on the i th entry and 0 elsewhere, and $\mathbf{e}'(t)$ represents the transposition of $\mathbf{e}(t)$. We can then apply the one-dimensional MCAF of Lo and Skindilias (2014a) and similar methodology of Chourdakis and Dotsis (2011) to calibrate our two-dimensional system. We define an $N \times 1$ filtered probability $\xi(\mathbf{V}_t^h|\mathcal{F}_t)$ vector, where \mathcal{F}_t collects all available information up to t . Given the transition probability matrix \mathbf{P} in equation (6), we compute the forecasting probability

$$\xi(\mathbf{V}_{t+\Delta t}^h|\mathcal{F}_t) = \mathbf{P}'(\Delta t) \xi(\mathbf{V}_t^h|\mathcal{F}_t) \tag{13}$$

where \mathbf{P}' denotes the transpose of matrix \mathbf{P} . When new observations become available, we update the filtered (posterior) probability by

$$\begin{aligned} \xi(\mathbf{V}_{t+\Delta t}^h|\mathcal{F}_{t+\Delta}) &= \left(\xi(\mathbf{V}_{t+\Delta t}^h|\mathcal{F}_t) \bullet \mathbf{F}_{t+1} \right) \\ &\div \left(\xi'(\mathbf{V}_{t+\Delta t}^h|\mathcal{F}_t) \mathbf{F}_{t+1} \right) \end{aligned} \tag{14}$$

where \bullet and \div , respectively, denote the element-by-element product and division, and \mathbf{F}_{t+1} is an $n \times 1$ probability density vector obtained from equation (10). For instance, the i th entry in \mathbf{F}_{t+1} is obtained from equation (10) with $V_{t+1} = v_i$. That is

$$\mathbf{F}_{t+1} = [\Phi_{y_t}(y_t|\Theta, v_1), \dots, \Phi_{y_t}(y_t|\Theta, v_i), \dots, \Phi_{y_t}(y_t|\Theta, v_n)]' \tag{15}$$

Suppose the initial variance is $V_0 = v_i$; the initial filtered probability is thus given by $\xi_0 = (0, \dots, 0, \xi_0^{i\text{-th}} = 1, 0, \dots, 0)'$. We run equations (13) and (14) recursively from $t = 1 \dots T$, where T is the total number of time series observations in our dataset. Applying the Markov switching technique of Hamilton (1989), we obtain two important by-products from the MCAF system: (i) the filtered variance; and (ii) the quasi-log-likelihood function, which can be used to estimate model parameters. The filtered variance \tilde{V} from the system is computed by

$$\tilde{V}_{t+1} = \xi'(\mathbf{V}_{t+\Delta t}^h|\mathcal{F}_{t+\Delta}) \mathbf{v} \tag{16}$$

and the quasi-log-likelihood function is defined by

$$\mathcal{L}(\Theta) = \sum_t \log \left(\xi'(\mathbf{V}_{t-\Delta t}^h|\mathcal{F}_{t-\Delta}) \mathbf{P}(\Delta t) \mathbf{F}_t \right) \tag{17}$$

To forecast future variance is very easy with MCAF. Suppose filtered variance at time t is $\tilde{V}_t = v_i$; the variance forecast V_T^f at a future time T is

$$V_T^f = \mathbf{e}'_i \mathbf{P}(T-t) \mathbf{v} \tag{18}$$

Filtering for a higher-dimensional system

The system we proposed in the previous subsection can easily accommodate a system with higher dimensions. Suppose O_t is the $k \times 1$ vector collecting market prices of k different variance-sensitive derivatives; for example, $O_t = [\text{Call}_t, \text{VIX}_t^2]'$. The objective now is to calibrate the hidden variance process by using these market observations. We define the measurement equation

$$O_t = G(V_t, \Theta) + \eta_t \tag{19}$$

where $G(V_t, \Theta)$ is the $k \times 1$ model price vector obtained by analytical solutions or numerical approximations, and η_t is the multivariate normal distribution with zero mean and a diagonal covariance matrix Σ . The density value for each variance v_i on the Markov chain is then computed by

$$\phi^i(O_t, G(v_i, \Theta), \Sigma) \tag{20}$$

To include the market observations of variance-sensitive derivatives in our MCAF model, users only need to change the probability density vector \mathbf{F}_{t+1} in equation (15). Owing to the diagonal covariance matrix Σ , the i th entry in \mathbf{F}_{t+1} in equation (15) is replaced by

$$\mathbf{F}_{t+1}^i = \Phi_{y_t}(y_t|\Theta, v_i) \sum_{o=1}^k \phi_o^i(v_i)$$

where $\phi_o^i(v_i)$ is the o th entry in equation (20).

Table I. Time series summary statistics over the period January 2000 to June 2007 (in-sample) and over the period July 2007 to July 2014

	In-sample		Out-of-sample	
	SP500	VIX	SP500	VIX
Mean	0.0044	19.664	0.0374	26.303
SD	0.1757	7.046	0.2333	11.284
Skewness	0.0751	0.783	-0.2975	1.834
Kurtosis	5.7798	3.088	11.6360	6.810
Max.	0.0557	80.860	0.1096	80.860
Min.	-0.0600	9.890	-0.0947	11.980

EMPIRICAL ANALYSIS

Our evaluation of variance forecasting involves three particular applications (forecasting tests) that we describe below. For each of the tests we conduct, the forecastability of variance is profoundly assessed in three subclasses of the two-factor model specification. The first class, hereafter SVJ, leaves the specification untouched as given in equation (1). Thus, allow for jumps in returns and allow for estimation of the correlation coefficient which represents the leverage effect. Next, model SVJ0 sets the correlation coefficient to zero and essentially does not include it in the estimation process. Lastly, model SVH is actually the Heston (1993) stochastic volatility model with a non-zero correlation and no jumps in the return process. Our in-sample estimation period stops just before the financial crisis of 2007, thus allowing us to examine thoroughly the contribution of the jump component and/or the leverage effect parameter in the forecastability of the variance over the turmoil period after June 2007. For all three models, we perform joint estimations of their objective (actual) and risk-neutral dynamics using the time series $\{S_t, VIX_t\}$ of the S&P 500 index and its corresponding volatility index (VIX).¹¹

The VIX index measures the market's expectation of the 30-day forward S&P 500 index volatility implicit in the index option prices. VIX² stands as a reference for the fixed-leg counterpart of a 30-day variance swap (VS) contract. On daily frequencies, our data sample starts in January 2000 and ends in July 2014, though we run a subsample analysis that covers the period January 2000 up to June 2007 and use part of the remaining period for our out-of-sample analysis. Table I provides some basic statistics of two series of log-return of the S&P 500 index and VIX. We emphasize two important observations for each of the time series over the in-sample and out-of-sample periods: the return process after the financial crisis exhibits even more leptokurtic features, while the VIX index, although maintaining the natural positive skewness, is profoundly characterized by an increased standard deviation and excess kurtosis.

Estimation results

Our in-sample parameter estimates are reported in Table II for three model classes over the period 3 January 2000 to 29 June 2007. The spot variance for the two similar class models with jumps (SVJ and SVJ0) are relatively fast mean reverting, implying that the time necessary to halve a unit shock¹² in numbers of days is 45 and 50 days, respectively, as opposed to the slow mean reversion of 180 days postulated by the model with no jumps (SVH). This is most likely due to the fact that the model with no jumps requires keeping volatility elevated to capture adverse price movements. The correlation for the two models with a non-zero correlation can confirm the leverage effect for which their values are not necessarily consistent due to the inclusion of a jump. The long-run average volatilities, $\sqrt{\theta}$, are consistent with the summary statistics of Table I. For all three model the parameter η_V for the prices of volatility risk is negative, implying negative instantaneous variance risk premia. The jump parameters, for SVJ and SVJ0, exhibit a slightly negative mean jump¹³ and remain very close for both cases, although for one we do omit the leverage effect parameter estimation.

Figure 1 plots the estimated latent variance for all three models. Notably, the SVH model has elevated variance levels due to the omission of jumps. For the two cases with jumps, over most of the period their levels are similar, except for where the SVJ model exhibits sharp changes in the volatility level, which may mostly be due to inclusion of the leverage effect. That is, a sharp decline/increase in the underlying returns drives a sharp increases/decrease in the level of volatility for the SVJ model.

¹¹ The CBOE launched the VIX index in 1993 and switched to the new VIX index in September 2003. The VIX index values used in this study are the new VIX index series obtained by the CBOE.

¹² The half-life is given by $-\log(0.5)/k \times 252$

¹³ We have decided, for ease in estimation, to keep the jump parameters equal under both probability measures.

Table II. Calibrated parameters over the subsample period January 2000 to June 2007 using the MCAF, which reveals parameter estimates under both probability measures along with model likelihood estimate \mathcal{L}

κ	θ	v	η_S	η_V	λ	μ_j	σ_j	ρ	\mathcal{L}
3.84490	0.06765	0.70674	0.00009	-1.87530	80.28000	-0.00306	0.00195	-0.85047	13578
3.40590	0.05000	0.54000	0.00010	-1.90000	76.90000	-0.00310	0.00190	0	12826
0.94780	0.03500	0.26000	0.07620	-0.3721	—	—	—	-0.54050	7548

Note: In agreement with studies in the literature, the market price of volatility risk parameters η_V are negative for all models, indicating a negative variance risk premium. All parameter estimates are significant at the 5% level of significance.




Figure 1. The filtered variance process (in-sample) for each of the three model representations for the daily S&P 500 return series

To further examine the difference of each model, in Figure 2 we display predictive volatility distributions, $p(\sqrt{V_{t+\tau}}|y_t, V_t)$, and predictive return distributions, $p(y_{t+\tau}|y_t, V_t)$. Because variance is unobservable, we use filtered variance \hat{V}_t in equation 16 to produce the predictive distributions of future volatilities and returns. We locate over our in-sample data series the largest negative daily return, move our estimation 1 week earlier to this event, and produce predictive densities for the week ‘ahead’ ($\tau = 7$ days). This large movement corresponds to about a -6% change in return. Prior to this move the predictive return densities were already different. We note that the density of the SVH, relative to SVJ and SVJ0, is shifted to the right, indicating that relatively large moves prior to this date were in part explained by the jump components. The return densities for the jump models are almost identical and substantially different for the SVH model. The return density of two models with jump components are significantly more negatively skewed than the SVH model, also exhibiting fatter tails. We further show the densities for the period just after the market move and we further provide an out-of-sample prediction for the largest negative return in the out-of-sample period. Although predictive volatility and return estimates across models are different, there is no formal indication as to which estimates are more accurate. Qualitatively, we may note that while return distributions can be significantly different (at least in the jump and zero-jump cases) the forecastability of the volatility distribution looks more concise.¹⁴ We thus address the robustness of our parameter estimates by forecasting total variance over daily and weekly horizons in the out-of-sample part of the data. For each model, expected future variance is given by

¹⁴This may be in agreement with the study of Christoffersen and Diebold (2006), who note that whereas volatility is forecastable returns are not.

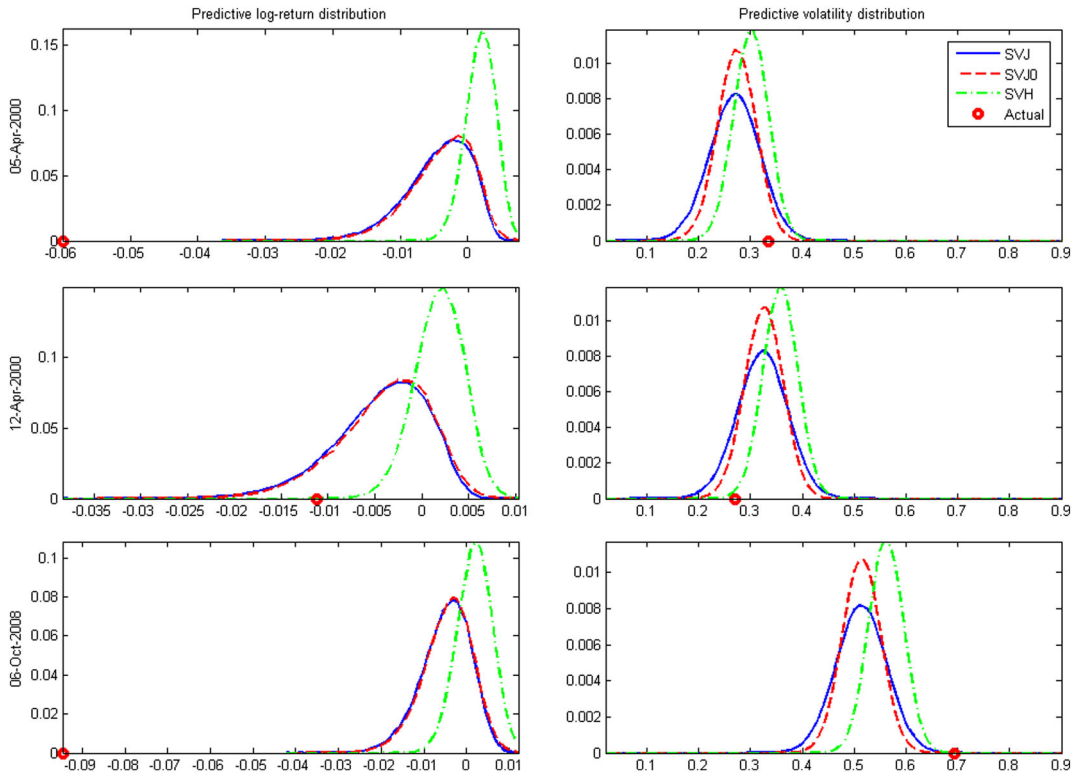


Figure 2. The predictive return $p(y_{t+\tau}|y_t)$ (left plots) and volatility $p(\sqrt{V_{t+\tau}}|y_t)$ (right plots) distribution around a date with the largest market move in the in-sample period (top and middle plots) and the out-of-sample predictive densities 1 week before the largest negative realized return in the out-of-sample period

Table III. Forecasting results over three time horizons for each model. Each value represents the RMSE of the forecast (expected) realized variance over the realized variance

	SVJ	SVJ0	SVH
Daily	0.0364	0.0386	0.0399
Weekly	0.0668	0.0721	0.0760
Monthly	0.1602	0.1687	0.1588

$E_t^P(\int_t^{t+\tau} V_s ds)$, which is known analytically, and the realized n -day variance is given by summing squared returns: $\sum_{i=1}^n (\log(S_{t_i}/S_{t_{i-1}}))^2$.

Table III reports the root mean square error (RMSE) of the variance forecast of each of the model specifications. It is clear that the SVJ model outperforms all other cases, and in some instances marginally the SVJ0 counterpart. Surprisingly, we note that the SVH model in the monthly forecast horizon outperforms (although marginally) all other models.

Variance swaps trading

Our first assessment of successful forecastability of the variance, given our observations in the preceding section, focuses in variance swap trading. This has a direct implication for the ability of the model to forecast variance effectively as the variance swap contract payoff is a function of the variance itself.

For variance swap contracts, as for any swap contract, no exchange of the principal amount takes place; the fixed leg of the VS agrees to pay at maturity the agreed amount fixed at time t : the VS rate ($SW_{t,T}$). At maturity T , the long position in the variance swap receives the difference between the realized variance $RV_{t,T}$ and the VS rate:¹⁵

$$(RV_{t,T} - SW_{t,T}) \times (\$ \text{ notional amount})$$

The studies of Carr and Wu (2009) and Ait-Sahalia *et al.* (2015) clearly indicate that the difference $VP_{t,T} = E_t^P[RV_{t,T}] - SW_{t,T}$, which defines the variance risk premia, is negative most of the time, thus implying that taking short positions in variance swap contracts will, on average, be more profitable. Our negative estimate of the market

¹⁵ The value of the contract at initiation is zero.

Table IV. Shorting variance swaps (short-and-hold): in-sample and out-of-sample trading

	SVJ	SVJ0	SVH
<i>In-sample</i>			
Mean payoff	0.009213	0.008297	0.006192
Sharpe ratio	0.121802	0.111622	0.092302
<i>Out-of-sample</i>			
Mean payoff	0.007731	0.007721	0.006726
Sharpe ratio	0.053468	0.051485	0.048327

Note: We present mean payoff as well as Sharpe ratios. Sharpe ratios are computed as the average return throughout the sample divided by its standard deviation. The signal for the strategy is defined as a positive expected payoff from shorting the VS contract, i.e. $VS_{t,T} - E_t^P[RV_{t,T}]$. The signal for each model is determined by the estimate of the cumulative expected realized variance $E_t^P[RV_{t,T}]$. We repeat the strategy for each t in our sample and hold the position until maturity.

price of variance risk η_V also proves this. Therefore the timing of whether a trader takes a short position in a variance swap contract comes as a signal of whether we expect the payoff of shorting the contract to be positive. That is, our forecast $RV_{t,T}$ can be used to construct this signal. Recent results from the theory of quadratic variation show that when the sampling frequency is high the realized (integrated) variance $IV_{t,T}$ converges to the unobserved true integrated variance:

$$RV_{t,T} = \frac{252}{n} \sum_{i=1}^n \left(\log \frac{S_{t_i}}{S_{t_{i-1}}} \right)^2 \rightarrow IV_{t,T} = \frac{1}{T-t} \int_t^T V_s ds + \frac{1}{T-t} \sum_{j=N_t}^{N_T} J_j^2 \tag{21}$$

which is the sum of two terms in our SDE: the continuous and discontinuous parts. In other words, to forecast the RV and assess our trading signal we need to compute the expectation of integrated variance under the objective probability measure P over the period of interest. This quantity is known in closed form, comes as a result by taking expectations in the continuous-time process under the objective probability measure.

$$E_t^P[IV_{t,T}] = \frac{1}{T-t} E_t^P \left[\int_t^T V_s ds \right] + 2\phi \tag{22}$$

The different parameter estimates of the three cases we examine (SVJ, SVJ0, SVH) are expected to give different signals. The strategy is to short variance swaps when the signal indicates a positive return, $E_t^P[IV_{t,T}] - SW_{t,T} > 0$. This is rather a simple trading strategy but profoundly assesses our variance forecastability among the three competing model representations. When the contract is shorted we hold the position until maturity, essentially undertaking a short-and-hold strategy. We repeat the strategy for each t in our in-sample period and also assess each forecast for each t of our out-of-sample period. That is, each day t , based on the trading signal we decide whether or not to short-and-hold.

Table IV provides a plot of each model performance over the in-sample and out-of-sample periods. We compare models by measuring the return in shorting VS contracts and their corresponding Sharpe ratios. Sharpe ratios are computed as the average return throughout our sample divided by its standard deviation. We note that SVJ and SVJ0, as opposed to SVH, give similar signals as depicted in their average returns and Sharpe ratios. In the preceding section we noted that the SVH model produces the least error amongst the other models for monthly forecasting. Essentially, as VIX represents the 30-day expected variance we are actually forecasting 30-day variance. Nonetheless, we note that here the SVH model is the worst performer. There are two underlying reasons for this. First, as we saw in the latent variance plots, the SVH model was the one with the most elevated variance level. It is then possible to expect the SVH model to give signals more frequently. Most importantly, the VS payoff is a function of the average variance over the life of the contract. Therefore, given the good performance of the other two models in the forecasting experiment in the preceding section, we should expect SVJ and SVJ0 to give better signals.¹⁶

¹⁶ We have also benchmarked the strategy to an S&P 500 buy-and-hold strategy that initiates whenever the VS signal kicks in. Results are given upon request. Ait-Sahalia *et al.* (2015), too, shows that on average shorting variance swaps outperforms the underlying index buy-and-hold strategies.

Table V. Evaluation of 1-day VaR forecast: in- and out-of-sample empirical shortfall probability $\hat{\alpha}$ and mean absolute percentage error (MAPE) based on 1-day VaR forecast are reported

α	1%	5%	10%	MAPE	T
<i>In-sample</i>					1882
SVJ	0.012*	0.053*	0.116	0.17	
SVJ0	0.015	0.06*	0.123	0.342	
SVH	0.006*	0.034	0.086	0.273	
<i>Out-of sample</i>					1779
SVJ	0.014*	0.057*	0.098*	0.211	
SVJ0	0.023	0.065	0.102*	0.562	
SVH	0.015	0.053*	0.085	0.247	

Note: In-sample period covers January 2000 to June 2007 and out-of-sample period covers July 2007 to July 2014. Values with an asterisk indicate the null hypothesis $H_0: \alpha = \hat{\alpha}$ is not rejected at the 5% level from the likelihood ratio test. The critical value at the 5% level from $\chi^2(1)$ is 3.84.

Value-at risk (VaR) forecast

The 1-day VaR forecast of S&P 500 daily log-returns R_t , given an associated probability threshold α , implied by a model \mathcal{M} is defined by

$$P r_{t-1}^{\mathcal{M}} (R_t < -\text{VaR}_t^\alpha) = \alpha \quad (23)$$

Following convention, we define VaR as a positive number. If the model is appropriately specified, we expect only $100 \times \alpha\%$ of the observed returns to exceed the the negative VaR forecast. We can access the performance of model \mathcal{M} by examining the percentage shortfall frequencies:

$$F^\alpha = 100 \times \frac{x}{T} = 100 \times \hat{\alpha} \quad (24)$$

where T is the number of observed returns and x is number of days that return is lower than negative VaR (shortfall frequency). Therefore, $\hat{\alpha}$ is the empirical shortfall probability. If the model can adequately assess (or quantify) the risk, the empirical probability of shortfall $\hat{\alpha}$ should not be statistically different from its theoretical counterpart α . To formally test the null hypothesis, $H_0: \alpha = \hat{\alpha}$, we use the standard likelihood ratio test:

$$\text{LRT} = -2 \left(\log \left(\frac{\alpha}{\hat{\alpha}} \right)^x + \log \left(\frac{1-\alpha}{1-\hat{\alpha}} \right)^{T-x} \right) \sim \chi^2(1) \quad (25)$$

In our analysis, we focus on $\alpha = 0.01, 0.05$ and 0.1 . For the performance summary, we also evaluate the mean absolute percentage error (MAPE) for each model, given by

$$\text{MAPE} = \frac{1}{3} \sum_{i=1}^3 \left| \frac{\hat{\alpha}_i - \alpha_i}{\alpha_i} \right| \quad (26)$$

where $\{\alpha_1, \alpha_2, \alpha_3\} = \{0.01, 0.05, 0.1\}$. The VaR_t^α forecast associated with the proposed system is simply the α -percentile of the log-return model given by equation 1, where the variance V_t is replaced by the variance forecast V_t^f from equation (18).

We report the empirical shortfall probability $\hat{\alpha}$ and MAPE of the 1-day VaR forecast of each model in Table V. The values with an asterisk indicate that the null hypothesis $H_0: \alpha = \hat{\alpha}$ is not rejected at the 5% level from the likelihood ratio test, given that the critical value at the 5% level from $\chi^2(1)$ is 3.84. As expected, the SVJ model performs best with the minimal MAPE, 0.17 and 0.211 respectively, in both in-sample and out-of-sample analyses. All VaR estimates based on SVJ models pass the likelihood ratio test, except for the in-sample $\alpha = 10\%$ case. We observe that both SVJ0 and SVH perform poorly in the likelihood ratio test and SVJ0 performs worst in terms of MAPE values.

Sign forecast

To further illustrate the power of forecastability over the three model structures, we incorporate a logic that relies on variance dependence and determine which of the three provides the optimal sign prediction and for which horizon. The idea is introduced in Christoffersen and Diebold (2006),¹⁷ who show that the ability of market direction forecastability

¹⁷ For an excellent treatment of the idea we refer the reader to the original manuscript from the authors.




Figure 3. Out-of sample forecast correlation $\text{corr}(P_{t+h|t}, I_{t+h})$. We examine the correlation between sign forecasts and realizations as a function of horizon. For each of the models under consideration we compute the conditional probability of a positive return as well as the return sign realization ranging from 1 to 500 days ahead. Some sign forecastability remains for the very long horizon but drops significantly for all models after 250 days

(return sign) is a case of the dependence and forecastability of asset return variance. In particular, the authors show that one should expect sign dependence, given the strong literature evidence of variance dependence. Further, building on an important principle of forecasting, it is shown that it is statistically possible for one to have sign and hence variance dependence when mean independence may exist. This phenomenon can be interpreted by the following return decomposition:

$$R_t = \text{sign}(R_t) \times |R_t| \tag{27}$$

It is well known in the empirical literature that both of the right-hand side components of the return decomposition equation are persistent and hence forecastable; however, returns themselves are unforecastable. Following Christoffersen and Diebold (2006), let $R_{t+1:t+h}$ be the h -day return, and define a positive return indicator variable $I_{t+h} = 1$ if $R_{t+1:t+h} > 0$, and $I_{t+h} = 0$ otherwise. Our objective here is to forecast sign I_{t+h} using the variance V_t . Because variance is unobservable, we use filtered variance \hat{V}_{t+1} in equation (16) to forecast the sign of returns. As suggested by Christoffersen and Diebold (2006), we use in-sample data to estimate the following logistic regression:

$$I_{t+h} = \Phi(V_t, \Theta) + e_{t+h}$$

where Θ is the parameter set, $\Phi(\cdot)$ is the logistic function

$$\Phi(x) = \frac{\exp(x)}{1 + \exp(x)}$$

and e is white noise. In our analysis, we focus on $h = \{20, 40, 60, 120, 250, 300, 500\}$. The out-of-sample correlation between sign forecasts $P_{t+h|h}(R_{t+1:t+h} > 0)$ obtained by logistic regression and realization I_{t+h} is presented in Figure 3. The results indicate that: (i) a hump shape is apparent, which means it is actually hard to forecast over high frequencies as well as low frequencies; and (ii) our findings are similar to those of the authors. The peak of correlation is on a 250-trading day horizon. Surprisingly, signs have higher forecastability for the SVH model. We were expecting the volatility of volatility restriction in the Heston model to limit the amount of sign predictability. This is not what we have observed. One explanation could be that the leverage effect parameter (ρ), which is non-zero, allows for sign forecastability even when expected returns are zero, as the authors have shown, thereby possibly increasing the forecastability in the Heston model.¹⁸ Further, we were hoping to see qualitatively similar results for SVJ and SVJ0, as the authors further showed that results are qualitatively similar if no leverage is incorporated. The underlying reason for examining this is beyond the scope of our paper but we will address this subject in further studies.

SIMULATION ANALYSIS

The performance of model calibration for a jump diffusion approximated by a continuous-time Markov chain has been documented in past literature from Lo and Skindilias (2014a,2014b) for a one-dimensional filtering problem. A Markov chain filter can accurately track the hidden state sequence.

¹⁸ The jump component may have some effect here too.

Table VI. Simulation results under the two settings

	True	Mean	SD	5%	95%
<i>Affine CEV</i>					
a_0	0.0179	0.0218	0.0039	0.0141	0.0297
a_1	-1.524	-1.7685	0.2445	-2.2599	-1.2697
b_1	4.7207	4.2123	0.5308	3.1629	5.2687
b_2	1.2563	1.1951	0.0611	1.0667	1.3192
μ	0.25	0.2503	0.0004	0.2496	0.2511
<i>Non-affine CEV</i>					
a_0	-0.0232	-0.0188	0.0089	-0.0378	-0.0011
a_1	3.4218	3.8601	0.6433	2.4448	5.2136
a_2	-57.8252	-59.7725	1.9407	-63.7081	-55.7092
a_3	0.0001	0.0001	0	0.0001	0.0001
b_1	13.7907	14.8	1.5588	11.6509	17.9645
b_2	1.5237	1.4688	0.118	1.2178	1.7242
μ	0.15	0.1529	0.0291	0.0885	0.2133

Note: An SV model with linear drift in the variance process (this is equivalent to the process used in the text for $a_0 = \kappa\theta$ and $a_1 = -\kappa$), and an SV model with nonlinear drift component and CEV-diffusion type for which the density function is not known in closed form.

In this section we perform a simulation study for calibrating model parameters in a two-dimensional system where the underlying process density is unknown. Our sole purpose is to show that the proposed filtering methodology may accommodate virtually any type of process.

We work with a two-factor stochastic volatility model. As before, we add jumps in the return process and, additionally, let the diffusion of the variance take a less restrictive form than that of the square root type. Furthermore, the drift of the variance process takes a nonlinear specification. As a result, the resulting process density function does not exist in closed form. In mathematical finance this process is known as the non-affine stochastic volatility model.

The return process specification is given by the SDE:

$$d \ln S = (\mu - V_t/2)dt + \sqrt{V_t}dW_t + dJ$$

with variance specification under the two different settings given by

Affine CEV:	$dV_t = (a_0 + a_1V_t)dt + b_1V^{b_2}dB_t$
Non-affine CEV:	$dV_t = (a_0 + a_1V_t + a_2V_t^2 + a_3/V_t)dt + b_1V^{b_2}dB_t$

We have considered an extra case in which, while the drift of the variance is linear, the diffusion is of constant elasticity of variance type (CEV).¹⁹ This process is denoted by 'Affine CEV', while the process with nonlinear drift and CEV-type diffusion is denoted by 'Non-affine CEV' in the table above. There is no consensus in the empirical literature as to which process better explains the dynamics of the latent variance. Certainly the Bates (2000) model used in our empirical section has enjoyed popularity owing to its analytical tractability and closed-form solutions for option prices. The affine process with CEV diffusion has no closed-form formula for computing option prices. This also applies to the variance process with nonlinear drift.

We simulate 5000 time steps of size $\Delta t = 1/252$ and run the MCAF 500 times to produce average sample parameter estimates. Estimation results are given in Table VI. Results clearly indicate that the MCAF is able to capture the 'true' parameters of the underlying pilot process.

Table VI displays the calibrated parameter estimates. The proposed MCAF can successfully track the latent variance process in either of the forms presented. These results allow us to extend our study and investigate the variance structure in a non-affine stochastic volatility framework for which no concise conclusions yet exist in the literature. Primarily, this is because of the limiting (or time-consuming) approaches that exist in the literature for estimating non-affine processes.

¹⁹ We include this specification, as in the main empirical analysis we used the square root type diffusion for which the density was known.

SUMMARY AND CONCLUSIONS

We have successfully designed and implemented a filtering algorithm to the latent variance process in a two-factor stochastic volatility model with jumps in the return process. The filter is built in continuous time and enjoys probabilistic convergence proofs whose roots rely on the well-established idea of local consistency in a Markov chain approximation setting as given by Kushner and Dupuis (2001). It shares similar features to the popular particle filters of Johannes *et al.* (2009) and Christoffersen *et al.* (2010), such as the ability to calibrate virtually any type of process whether the density is known or not. A major implication of the proposed filter is that this filter is built in continuous time.

We illustrate this through a study of the well-known Heston (1993) SV model, whose variance process has a linear drift function and square-root diffusion, and whose return process encompasses a jump component (Bates, 2000). This model enjoys the features of an analytical density function but also a closed-form formula for option prices. Empirical results are open to a number of interpretations. A major consideration of this type of model structure is the ill-posedness of jump parameter estimation and the hard-to-estimate leverage effect parameters. To this end, we have illustrated through the proposed filtering algorithm the implications of omitting from the estimation process the jump components and/or leverage effect parameter. We essentially dealt with three popular cases: (i) the SVJ model, which leaves intact the structure of the SV model with jumps in the return process; (ii) the SVJ0 model, which sets the leverage effect equal to zero but allows returns in the jump process; and (iii) the SVH model, which is a zero-jump model with non-zero leverage effect parameters. We assess each of the model through a series of forecasting experiments that rely primarily on accurate variance forecasting. We confine ourselves solely to forecasting variance as we note that financial econometric theory suggests that, while asset returns are not easily forecast, asset return volatility exhibits dependence and is thus forecastable. We find that the model that attempts to address all stylized facts of the return/variance relationship, namely the SVJ model, can outperform all other components in most of our experiments. Nonetheless, if one finds it hard to obtain a concise estimate of the correlation coefficient of the so-called leverage effect, results are not expected to differ substantially. Agreeing with most studies in the literature we note that, while the leverage effect parameter may (under some conditions) be omitted from the model, the addition of a jump component is necessary. Throughout the manuscript we have used a model that is analytically tractable. We choose to do so, first because the SV structure we work with is a model that has enjoyed great popularity, but also to ensure that our numerical approximation is consistent with the approximation of the closed-form solution, thereby relating all closed-form equations used in this paper with their semi-closed-form analogues of our proposed approach. Nonetheless, the methodology is not confined to processes with analytically tractable density functions. In a simulation study we show that the proposed methodology has applicability to a wider class of model with complicated structures for which the probability density function is not known analytically. We show that the MCAF parameter estimates are very close to the true parameters that have generated the ‘data’ process.

Our study allows us to extend our research and address a long-standing question in the literature regarding the use of nonlinear specifications in the drift of the variance process, for which closed-form solutions do not exist. For this class of model empirical results are not yet sufficient to yield a concise conclusion of whether inference using these non-affine models can better explain the dynamics of the latent unobserved variance process.

APPENDIX

Lo and Skindilias (2014b) propose a Markov chain approximation scheme for general jump diffusion. In this paper, we only consider the diffusion process for the variance process without any jump component. Therefore we use the continuous part of their formula to approximate the variance process $dv_t = a(v_t) dt + b(v_t) dW$. Lo and Skindilias (2014b) show that the proposed generalized formula for the \mathbf{Q} rate matrix of an n -state non-equidistant Markov chain with grid spacing $\mathbf{h} = \{h_1, \dots, h_{n-1}\}$ satisfies the local consistency condition of Kushner and DiMasi (1978):

$$\begin{aligned}
 q_{i,i-1} &= \frac{1}{h_{i-1}} a^-(v_i^h) + \frac{b^2(v_i^h) - (h_{i-1} \times a^-(v_i^h) + h_i \times a^+(v_i^h))}{h_{i-1}(h_{i-1} + h_i)} \\
 q_{i,i} &= -q_{i,i-1} - q_{i,i+1} \\
 q_{i,i+1} &= \frac{1}{h_i} a^+(v_i^h) + \frac{b^2(v_i^h) - (h_{i-1} \times a^-(v_i^h) + h_i \times a^+(v_i^h))}{h_i(h_{i-1} + h_i)} \\
 q_{i,j} &= 0 \quad \forall j \neq i, i-1, i+1
 \end{aligned}
 \tag{28}$$

where $a^+ = \max(a, 0)$ and $a^- = \max(-a, 0)$.

REFERENCES

- Ait-Sahalia Y. 2004. Disentangling diffusion from jumps. *Journal of Financial Economics* **74**(3): 487–528.
- Ait-Sahalia Y, Karaman M, Mancini L. 2015. The term structure of variance swaps and risk premia. DOI: 10.2139/ssrn.2136820.
- Andersen TG, Benzoni L, Lund J. 2002. An empirical investigation of continuous time equity return models. *The Journal of Finance* **57**(3): 1239–1284.
- Andersen TG, Bollerslev T, Diebold FX, Ebens H. 2001a. The distribution of realized stock return volatility. *Journal of Financial Economics* **61**: 43–76.
- Andersen TG, Bollerslev T, Diebold FX, Labys P. 2001b. The distribution of realized exchange rate volatility. *Journal of the American Statistical Association* **96**: 42–55.
- Arulampalam MS, Maskell S, Gordon N, Clapp T. 2002. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing* **50**(2): 174–188.
- Bakshi G, Cao C, Chen Z. 1997. Empirical performance of alternative option pricing models. *Journal of Finance* **52**(5): 2003–2047.
- Bates DS. 1996. Jumps and stochastic volatility: exchange rate processes implicit in Deutsche Mark options. *Review of Financial Studies* **9**(1): 69–107.
- Bates DS. 2000. Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics* **94**(1–2): 181–238.
- Bollerslev T, Mikkelsen HO. 1996. Modeling and pricing long memory in stock market volatility. *Journal of Econometrics* **73**: 151–184.
- Carr P, Wu L. 2009. Variance risk premiums. *Review of Financial Studies* **22**: 1311–1341.
- Chourdakis K. 2004. Non-affine option pricing. *Journal of Derivatives* **11**(3): 10–25.
- Chourdakis K, Dotsis G. 2011. Maximum likelihood estimation of non-affine volatility processes. *Journal of Empirical Finance* **18**(3): 533–545.
- Christoffersen PF, Diebold FX. 2006. Financial asset returns, direction- of-change forecasting, and volatility dynamics. *Management Science* **52**(8): 1273–1287.
- Christoffersen PF, Jacobs K, Mimouni K. 2010. Volatility dynamics for the S&P500: evidence from realized volatility, daily returns, and option prices. *Review of Financial Studies* **23**(8): 3141–3189.
- Duan JC, Yeh CY. 2010. Jump and volatility risk premiums implied by VIX. *Journal of Economic Dynamics and Control* **34**(11): 2232–2244.
- Elliott RJ. 1995. *Hidden Markov Models: Estimation and Control*, Vol. 29. Springer: New York.
- Hamilton JD. 1989. A new approach to the economic analysis of nonstationary time series and business cycle. *Econometrica* **57**(2): 357–384.
- Heston SL. 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* **6**(2): 327–343.
- Johannes M. 2004. The statistical and economic role of jumps in continuous-time interest rate models. *Journal of Finance* **59**(1): 227–260.
- Johannes MS, Polson NG, Stroud JR. 2009. Optimal filtering of jump diffusions: extracting latent states from asset prices. *Review of Financial Studies* **22**(7): 2559–2599.
- Kushner HJ. 1990. Numerical methods for stochastic control problems in continuous time. *SIAM Journal of Control and Optimization* **28**(5): 999–1048.
- Kushner HJ, DiMasi G. 1978. Approximations for functionals and optimal control problems on jump-diffusion processes. *Journal of Mathematical Analysis and Applications* **63**: 772–800.
- Kushner HJ, Dupuis PG. 2001. *Numerical Methods for Stochastic Control Problems in Continuous Time* (2nd edn), Stochastic Modelling and Applied Probability. Springer: London.
- Lo CC, Skindilias K. 2014a. Does the short rate revert to its mean in the risk-neutral world? Available: <http://ssrn.com/abstract=2446557>.
- Lo CC, Skindilias K. 2014b. An improved Markov chain approximation methodology: derivatives pricing and model calibration. *International Journal of Theoretical and Applied Finance* **17**.
- Merton RC. 1976. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* **3**(1–2): 125–144.
- 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 MK, Shephard N. 1999. Filtering via simulation: auxiliary particle filters. *Journal of the American Statistical Association* **94**(446): 590–599.
- Rogers LCG, Williams D. 2000. *Diffusions, Markov Processes and Martingales*, Ito Calculus, vol. 2. Cambridge University Press: Cambridge, UK.
- Skindilias K, Lo CC. 2013. Local volatility calibration during turbulent periods. *Review of Quantitative Finance and Accounting* **44**: 425–444.
- Tankov P, Cont R. 2003. *Financial Modelling with Jump Processes*. Chapman & Hall: London.

Authors' biographies:

Chia Chun Lo is an Assistant Professor of Finance at University of Macau. He was the external head of quantitative research and trading at CITICS Futures; a research associate at ABM Analytics; a quantitative ALM analyst at Cambridge Systems associates and Shin Kong Life; and the Asian Pacific operational manager of P&C Information Technology. He holds a PhD in Computational Finance from the CCFEA, University of Essex; MSc of actuarial science from Boston University, and MSc of financial engineering from National Chen Chi University. He is a CFA and FRM charter holder.

Konstantinos Skindilias is a senior research consultant at ABM Analytics as well as a Senior Lecturer at the Dept. of Mathematical Sciences, University of Greenwich. He holds a PhD in Computational Finance from the Centre of Computational Finance and Economic Agents (CCFEA), University of Essex, UK.

Andreas Karathanasopoulos is an Associate Professor at FAME, American University of Beirut, Lebanon. He has a series of publications in forecasting and has acted as a consultant for a number of fund management offices.

Authors' addresses:

Chia Chun Lo, Department of Finance and Business Economics, Faculty of Business Administration, University of Macau, Taipa, Macau, China.

Konstantinos Skindilias, Mathematical Sciences Department, University of Greenwich, London SE10 9LS, UK; ABM-Analytics Ltd, London, UK.

Andreas Karathanasopoulos, School of Business, American University of Beirut, Lebanon.