The Impact of Jump Dynamics on the Predictive Power of Option-implied Densities

Yaw-Huei Wang[*]

ABSTRACT

This study examines whether incorporating jumps with stochastic volatility can improve the predictive power of option-implied densities of the FTSE 100 index. A general double-jump model, as proposed by Duffie et al. (2000), is used to fit the market prices of options and to estimate 'risk-neutral' densities. 'Real-world' densities are then converted from their risk-neutral form by means of alternative statistical calibrations. Both the risk-neutral and real-world densities are evaluated, over five forecast horizons, using two different tests. Our empirical results indicate that adding jumps into the price and/or volatility processes not only substantially lowers the fitting errors of option prices, but also improves the predictive power of risk-neutral densities. Furthermore, satisfactory density prediction was consistently provided by the real-world densities, which were not dependent on the addition of jumps, the approach used to construct the densities, or the prediction horizon.

Keywords: Density prediction; Stochastic volatility; Jumps; Risk-neutral; Real-world.

JEL Classification: C53, G13

First version: 15 October 2006; **This version**: 6 January 2009

---

[*] Yaw-Huei Wang is an associate professor at the Department of Finance, National Taiwan University, No. 1, Sec. 4, Roosevelt Rd., Taipei 106, Taiwan. Tel: +886-2-33661092; E-mail: yhwang@management.ntu.edu.tw.

The author is indebted to the seminar participants at the European Financial Management Association Annual Meeting 2007, the National Taiwan University and Academia Sinica for their helpful comments and suggestions, and would also like to express his sincere gratitude for the financial support provided by the National Science Council of Taiwan.

1. INTRODUCTION

Density prediction has crucial implications from many perspectives, including policy making, risk management and derivatives pricing. Historical data has conventionally been used to estimate a price dynamic model, with the estimated model then being used to simulate the distribution of future prices. With forward-looking information, the derivatives markets are able to provide a much richer source for the prediction of the future price distribution of the underlying assets.

Many methods have been proposed as the means of inferring the risk-neutral density (RND) of the underlying asset price for the option maturity date.[1] A variety of methods were also surveyed in Cont (1997), Bahra (1997), Jackwerth (1999), Jondeau and Rockinger (2000) and Bliss and Panigirtzoglou (2002). It was demonstrated that most of these methods were likely to perform satisfactorily, providing that options were traded with sufficient strike prices, and that their range covered most of the distribution. There is, however, an inevitable limitation inherent within these methods, since the options used to estimate the RND must have the same time-to-maturity, which is of course, also the prediction horizon. As a result, the implied information within the option market cannot be used in its entirety, and we can only predict the distribution of the underlying asset price for a single horizon.[2]

However, as long as the corresponding option pricing formula can be derived either analytically or numerically, any risk-neutral specification of the price dynamic process potentially has the ability to conquer this limitation and yield RND estimates for any time horizon by simultaneously using all market prices of options with different maturities. Therefore, along with the progress being made in the advancement of option pricing models, density prediction methods have also been continually developing. In particular, continuous-time models have provided an ideal method of density prediction, essentially because there is often a multi-horizon analytical RND formula which corresponds to this type of option pricing model.

As the assumptions regarding price dynamics become more realistic, the implied volatility function can be more precisely captured. Those models that include jumps in the price and/or stochastic volatility dynamics have been examined in numerous studies including Bates (1996, 2000), Bakshi et al. (1997), Duffie et al. (2000), Eraker et al. (2003), Eraker (2004) and Carr and Wu (2004). The common finding in these studies is that incorporating jumps with stochastic volatility can substantially improve both the option pricing and the consistency of the parameters implied in option prices and estimated with relevant time-series data. It is therefore natural to question whether incorporating stochastic volatility with jumps in the asset price and/or jumps in volatility can also improve the forecasting power of option-implied densities. To the best of our knowledge, this issue has not yet been investigated; thus, the present paper aims to contribute to the literature by filling this gap.[3] In addition to presenting a more flexible framework for the estimation of multi-horizon RNDs, this paper offers important empirical evidence on the role played by jumps in density prediction.

In this study, we follow the affine-analysis framework of Duffie et al. (2000) to compare the performance, across five different forecast horizons, of three different density prediction models for the FTSE 100 index.[4] These models include the stochastic volatility (SV) model, the stochastic volatility model with jumps in price only (SVJ), and the stochastic volatility model with simultaneous and correlated jumps in both price and volatility (SVJJ). The RNDs and 'real world densities' (RWDs), transformed by the use of statistical calibrations, are evaluated based upon the Anderson-Darling and Berkowitz (2001) tests.[5]

Incorporating jumps with stochastic volatility essentially fattens the tails of the distribution of the underlying assets, thus making them more realistic. Therefore, in line with the prior studies, our initial finding is that the addition of jumps makes the parameter estimates to be much more reasonable and consistent with their theoretical expectations. Besides, as compared to the SV model, the SVJ model substantially lowers the fitting errors of option prices, whilst the average error for the SVJJ model is even less. Furthermore, adding jumps into the price and volatility dynamics can also improve the predictive power of RNDs.

By contrast, we find that remarkable improvements in the precision of the density predictions are achieved by RWDs over RNDs; furthermore, we also find that RWDs consistently provided satisfactory density predictions, which were not dependent on the approach used to construct the RWDs, the model specification, or the forecast horizon.

The remainder of the paper is organized as follows. Section 2 provides details of the option pricing model used in this study, followed in Section 3 by an explanation of the approaches used to derive the RNDs and then transform them into RWDs. The tests used to evaluate the density prediction performance are described in Section 4, followed by a description of the data used in this study in Section 5. Section 6 presents the analysis of the empirical results, with the conclusions drawn from this study finally being presented in Section 7.

2. THE OPTION PRICING MODEL

A general double-jump stochastic volatility model for option pricing was provided by Duffie et al. (2000) using the affine analysis method. Let S be the price process of a security which pays dividends at a constant proportional rate ξ, and $Y = \ln(S)$. The state process is $X = (Y, V)^{\mathrm{T}}$, where V is the variance process with a long-term level υ and volatility σ_V. Also, suppose that the short rate is a constant r, under an equivalent martingale measure Q.

$$d\begin{pmatrix} Y_t \\ V_t \end{pmatrix} = \begin{pmatrix} r-\xi-\lambda\mu-\frac{1}{2}V_t \\ \kappa_v(\upsilon - V_t) \end{pmatrix} dt + \sqrt{V_t}\begin{pmatrix} 1 & 0 \\ \rho\sigma_v & \sqrt{1-\rho^2}\sigma_v \end{pmatrix} dW_t^Q + dZ_t, \tag{1}$$

where W^Q is a standard Brownian motion under Q measure in $\mathbb{R}^2$; and Z is a pure jump process in $\mathbb{R}^2$ with constant mean arrival rate λ, whose bivariate jump-size distribution has the transform θ. Through the specification of θ, a flexible range of distributions of jumps can be explored. In order to satisfy the risk-neutral restriction, $\mu = \theta(1,0) - 1$.

Similar to the function proposed by Duffie et al. (2000), the jump transform function θ in this study is defined by:

$$\theta(c_1, c_2) = \lambda^{-1}(\lambda^y \theta^y(c_1) + \lambda^c \theta^c(c_1, c_2)) \tag{2}$$

where $\lambda = \lambda^y + \lambda^c$,

$$\theta^y(c_1) = \exp\left(\mu_y c_1 + \frac{1}{2}\sigma_y^2 c_1^{\ 2}\right) \tag{3}$$

and

$$\theta^c(c_1, c_2) = \frac{\exp\left(\mu_{cy} c_1 + \frac{1}{2}\sigma_{cy}^2 c_1^2\right)}{1 - \mu_{cv} c_2 - \rho_J \mu_{cv} c_1}. \tag{4}$$

The two functions referred to above correspond to two different types of jumps, as follows:

1. Jumps in price, Y, with arrival intensity λ^y and normally distributed jump size with mean μ_y and standard deviation σ_y.

2. Simultaneous correlated jumps in Y and V, with arrival intensity λ^c. The marginal distribution of the jump size in V is exponential, with mean μ_{cv}. Given

a realization of, say z_v, in the jump size in V, the conditional distribution of the jump size in Y is normal, with mean $\mu_{cy} + \rho_J z_v$ and standard deviation σ_{cy}.

Given the constraints on certain parameters, the following three different types of model are selected.

1. SV model: Stochastic volatility model, without jumps, obtained with $\lambda = 0$.

2. SVJ model: Stochastic volatility model, with jumps in price only, obtained with $\lambda^y > 0$ and $\lambda^c = 0$.

3. SVJJ model: Stochastic volatility model, with simultaneous and correlated jumps in both price and volatility, obtained with $\lambda^c > 0$ and $\lambda^y = 0$.

Under the affine transform framework, the price of a call option with the strike price K is given as:

$$C(K,T) = G_{1,-1}(-\ln K, Y_0, V_0, T) - KG_{0,-1}(-\ln K, Y_0, V_0, T), \tag{5}$$

where

$$G_{a,b}(-\ln K, Y_0, V_0, T) = \frac{\psi(a, Y_0, V_0, 0, T)}{2} - \frac{1}{\pi}\int_0^\infty \frac{\mathrm{Im}[\psi(a+ivb, Y_0, V_0, 0, T)e^{iv(\ln K)}]}{v}dv.$$

Here, $\psi(u, Y_0, V_0, 0, T) = \exp(\alpha(T,u) + uY_0 + \beta(T,u)V_0)$, and $\mathrm{Im}(c)$ denotes the imaginary part of c. By letting $b = \sigma_v \rho u - k_v, a = u(1-u)$ and $\gamma = \sqrt{b^2 + a\sigma_v^2}$, we have $\beta(T,u) = -\frac{a(1-e^{-\gamma T})}{2\gamma - (\gamma + b)(1-e^{-\gamma T})}$ and $\alpha(T,u) = \alpha_0(T,u) - \lambda T(1+\mu u) + \lambda\int_0^T \theta(u, \beta(s,u))ds$

with $\alpha_0(T,u) = -rT + (r-\xi)uT - \kappa_v \upsilon\left(\frac{\gamma+b}{\sigma_v^2}T + \frac{2}{\sigma_v^2}\ln\left[1-\frac{\lambda+b}{2\gamma}(1-e^{-\gamma T})\right]\right)$, and the term $\int_0^T \theta(u,\beta(s,u))ds$ being dependent upon the specification of the bivariate jump transform function $\theta(\cdot,\cdot)$.

Based upon the jump transform function θ detailed by (2)-(4), the term $\int_0^T \theta(u,\beta(s,u))ds$ is therefore given by:

$$\int_0^T \theta(u,\beta(s,u))ds = \lambda^{-1}(\lambda^y f^y(u,T) + \lambda^c f^c(u,T)) \tag{6}$$

where

$$f^y(u,T) = T\exp(\mu_y u + \frac{1}{2}\sigma_y^2 u^2) \tag{7}$$

and

$$f^c(u,T) = \exp\left(\mu_{cy}u + \sigma_{cy}^2\frac{u^2}{2}\right)d \tag{8}$$

with $a = u(1-u),\ b = \sigma_v \rho u - \kappa_v,\ c = 1-\rho_J \mu_{cv} u,$ and

$$d = \frac{\gamma-b}{(\gamma-b)c+\mu_{cv}a}T - \frac{2\mu_{cv}a}{(\gamma c)^2-(bc-\mu_{cv}a)^2}\ln\left(1-\frac{(\gamma+b)c-\mu_{cv}a}{2\gamma c}(1-e^{-\gamma T})\right).$$

We estimate the dynamic parameters by minimizing the following loss function:

$$SSE = \sum_{i=1}^{N}(C_i^*(K_i,T_i) - C_i(K_{i,}T_i))^2 \tag{9}$$

where C^* and C denote the respective real market and theoretical call prices.

3. RISK-NEUTRAL AND REAL-WORLD DENSITIES

3.1 Risk-neutral Densities (RNDs)

Rearranging the formula for the call price in Equation (5) to a Black-Scholes-type form of

$Se^{-\xi T}P_1$-$Ke^{-rT}P_2$, where P_2 is the probability of $S_T \geq K$ under a Q measure, we have:

$$P_2 = e^{rT} G_{0,-1}(-\ln K, Y_0, V_0, T). \tag{10}$$

The RND is therefore obtained by the derivative of 1–P_2 with respect to K evaluated at x, and given by:

$$f(x) = \frac{e^{rT}}{\pi x} \int_0^\infty \mathrm{Im}[\psi(-iv, Y_0, V_0, 0, T) e^{iv \ln x} \times i] dv. \tag{11}$$

3.2 Real-world Densities (RWDs)

The theory of asset pricing relates current asset prices to expectations of discounted future payoffs and in principle different measures result in the same asset price. Therefore, the transformation from RNDs to RWDs relies upon a pricing kernel describing the risk preference of the representative agent, i.e. $g = k \cdot f$, where g, f and k denote the RWD, RND and pricing kernel of a random variable, respectively. The pricing kernel can be specified by the assumptions of either an economic model or calibration theory. In order to minimize the computation load, our transformation is based upon the statistical calibration.

Given that $f(x)$ and $F(x)$ are the RND and cumulative distribution function (CDF) of the random variable X, we define U=$F(X)$, i.e. u=$F(x)$. Let the calibration function $C(u)$ be the real-world CDF of the random variable U. Then the real-world CDF of X is:

$$G(x) = \Pr(X \leq x) = \Pr(F(X) \leq F(x)) = \Pr(U \leq F(x)) = C(F(x)) \tag{12}$$

Accordingly, the RWD of X is given as:

$$g(x)=\frac{dC(F(x))}{dx}=\frac{dC(u)}{du}\frac{du}{dx}=c(u)f(x)\,, \tag{13}$$

where $c(u)$ is the probability density function (PDF) of U, i.e. a pricing kernel.

We use the CDF of the Beta distribution as the calibration function (as suggested by Fackler and King, 1990). The density transformation using this function was previously performed by Shackleton et al. (2006) and Liu et al. (2007). The CDF of the Beta distribution is defined as:

$$C(u)=\frac{1}{B(\alpha,\beta)}\int_0^u t^{\alpha-1}(1-t)^{\beta-1}\,dt\,, \tag{14}$$

where $B(\alpha,\beta)$ is the Beta function. In accordance with Equation (13), the calibrated RWD is:

$$g(x)=\frac{F(x)^{\alpha-1}(1-F(x))^{\beta-1}}{B(\alpha,\beta)}f(x)\,. \tag{15}$$

The estimates of the parameters α and β can be obtained by maximizing the likelihood assessed with the realized asset prices of the forecast horizon. For forecast horizon τ, the log-likelihood function with n RWDs and realized asset prices $S_{i,\tau}$ is:

$$\log(L(S_{1,\tau},S_{2,\tau},\cdots,S_{n,\tau}))=\sum_{i=1}^{n}\log(g_{i,\tau}(S_{i,\tau}))\,, \tag{16}$$

where $g_{i,\tau}$ denotes the i^{th} RWD with forecast horizon τ.

As suggested by Shackleton et al. (2006), we can also use an alternative non-parametric approach, i.e. an empirical calibration function, to implement the transformation. Let $\phi(.)$ and $\Phi(.)$ respectively denote the PDF and CDF of the standard normal distribution. Given a sample of n observations, we define $u_i = F(x_i)$ $(i = 1, 2, \ldots, n)$ and transform u_i to a new series $y_i = \Phi^{-1}(u_i)$. Under this series, the non-parametric kernels PDF and CDF, defined by normal kernels with bandwidth B, are respectively given as $h(y) = \frac{1}{nB}\sum_{i=1}^{n}\phi\left[\frac{y-y_i}{B}\right]$ and $H(y) = \frac{1}{n}\sum_{i=1}^{n}\Phi\left[\frac{y-y_i}{B}\right]$.[6] The empirical calibration function is then $C(u) = H(\Phi^{-1}(u))$ and the real-world CDF is:

$$G(x) = C(F(x)) = H(\Phi^{-1}(F(x))). \tag{17}$$

By definition, the RWD is then given as:

$$g(x) = \frac{dH(y)}{dx} = \frac{dH(y)}{dy}\frac{dy}{dx} = h(y)\frac{dy}{du}\frac{du}{dx} = \frac{h(y)}{\phi(y)}f(x). \tag{18}$$

Namely, $h(y)/\phi(y)$ is a pricing kernel.

4. DENSITY PREDICTION EVALUATION

On each density formation date, all of the option prices jointly produce a set of estimates for the model parameters; thus, we can construct the estimates of RNDs and RWDs for any particular forecast horizon. Using the Anderson-Darling test and the Berkowitz (2001) test, we aim to investigate whether the estimated densities are equal to the true densities.

Let $f_{i,\tau}(x)$ and $\hat{f}_{i,\tau}(x)$ respectively denote the true and estimated densities for forecast horizon τ. Under the null hypothesis that the realizations $S_{i,\tau}$ are independent and $f_{i,\tau}(x) = \hat{f}_{i,\tau}(x)$, the probability integral transformations (PITs) of the realizations,

$$u_{i,\tau} = \int_{-\infty}^{S_{i,\tau}} \hat{f}_{i,\tau}(x)dx \tag{19}$$

will be independently and uniformly distributed.[7]

Various statistical tests have been proposed and compared in the prior literature.[8] We adopt the Anderson-Darling A^2 test since it is particularly good at identifying mean errors, and extremely powerful when the forecasted distribution departs from the true distribution in the tails. The test statistic is defined as:

$$A^2 = -n - \sum_{i=1}^{n} \frac{(2i-1)}{N} [\ln U(u_i) + \ln(1 - U(u_{N+1-i}))] . \tag{20}$$

where U is the CDF of the uniform distribution. It should be noted that the u_i are the ordered data. The Anderson-Darling test is a one-side test; the null hypothesis that the PITs are uniformly distributed will be rejected if the statistic is higher than the critical value or the corresponding p-value is smaller than the significance level.

Arguing that most of the statistic tests based on the PITs were not sufficiently powerful for small samples, Berkowitz (2001) proposed a method for jointly testing uniformity and independence, further transforming $u_{i,\tau}$ with the inverse function of the standard normal distribution $\Phi^{-1}(\cdot)$, as follows:

$$z_{i,\tau} = \Phi^{-1}(u_{i,\tau}) = \Phi^{-1}(\int_{-\infty}^{S_{i,\tau}} \hat{f}_{i,\tau}(x) dx) . \tag{21}$$

Under the null hypothesis, $z_{i,\tau}$ will be i.i.d. and will follow a standard normal distribution. In order to test for both the independence and normality, Berkowitz first estimated the following model:

$$z_{i,\tau} - \mu = \rho(z_{i-1,\tau} - \mu) + \varepsilon_{i,\tau} . \tag{22}$$

This model was estimated using the maximum likelihood method with the aim being to determine the restrictions on the parameters using a likelihood ratio test. Under the null that $z_{i,\tau}$ is i.i.d., $N(0,1)$, $\mu=0$, $\rho=0$ and $\mathrm{Var}(\varepsilon_{i,\tau})=1$. If the log-likelihood function is denoted as $L(\mu,\sigma^2,\rho)$, the statistic for the likelihood ratio test is:

$$LR3 = -2[L(0,1,0) - L(\hat{\mu},\hat{\sigma}^2,\hat{\rho})], \tag{23}$$

which follows a $\chi^2(3)$ distribution under the null hypothesis.

However, if the data are serially correlated, which could arise from such overlapping forecasts, the rejection of the above test may arise from the autocorrelation of the data series. Berkowitz (2001) therefore tested for independence separately by examining the following likelihood ratio statistic:

$$LR1 = -2[L(\hat{\mu},\hat{\sigma}^2,0) - L(\hat{\mu},\hat{\sigma}^2,\hat{\rho})], \tag{24}$$

which follows a $\chi^2(1)$ distribution under the null hypothesis.

Thus, we will support the predictive power of a density only where we fail to reject the null hypotheses of both $LR3$ and $LR1$. If the null of $LR3$ is rejected, but the null of $LR1$ is not, we then have evidence indicating that the estimated densities do not provide accurate forecasts. Finally, if we reject the nulls of both $LR3$ and $LR1$, we will be unable to determine whether the rejection of $LR3$ is caused by serially correlated data.

5. DATA

The underlying asset for the density prediction is the FTSE 100 index; thus, the primary data used in this study is the daily settlement prices of the FTSE 100 index options, which are European style options traded on the London International Financial Futures Exchange (LIFFE). The dataset, which covers a sample period from 2 January 2000 to 31 December 2005, was obtained from Euronext.

Since the quarterly FTSE 100 futures contracts have the same expiration dates as the index options, if the mark-to-market effect is insignificant, the index options may then be regarded as the index futures options. The LIFFE reports the futures prices as the underlying asset prices in the dataset. For the serial-month options, the LIFFE calculates the theoretical futures prices based on the term structure of the quarterly futures prices. As serial-month options are less actively traded, with their underlying asset prices being constructed synthetically, only quarterly options are included in this study.

The option prices are filtered using additional criterion. Firstly, we exclude those option settlement prices that violate any arbitrage-free bounds, as well as those with a zero trading volume. Call and put option prices for the same strike and maturity should essentially contain the same information; however, given that the out-of-the-money (OTM) options are usually more heavily traded, we discard the in-the-money (ITM) options. Furthermore, in order to avoid the liquidity issue being imposed on the short-maturity options, we exclude those options with a time-to-maturity which is less than seven calendar days. We also eliminate those options with prices that are less than £1.00 since the minimum quote and change in the option prices is 50 pence, and as such, these option prices could be insensitive to the information contained therein.

In addition to the prices of futures and options, the risk-free rates adopted in this study are proxied by the three-month Euro-currency interest rates for Pounds Sterling, given that the three-month Euro-currency market has the best liquidity (these interest rates are obtained from Datastream). All of the put prices are converted to equivalent call prices using put-call parity. A summary of the option prices is provided in Exhibit 1, which is a cross-sectional table comprising of six sections of moneyness divided by 0.94, 0.97, 1.0, 1.03 and 1.06, and three sections of time-to-maturity divided by 90 and 180 calendar days.[9]

<Exhibit 1 is inserted about here>

Our study sample ultimately comprised of a total of 72,494 option prices with an average call (equivalent) price of £427.40. As expected, the call price has a positive association with time-to-maturity and a negative association with moneyness. The number of OTM put contracts (moneyness <0.97) is generally much higher than the number of OTM call option contracts (moneyness > 1.03). The average index values across moneyness are also given in the last column of Exhibit 1. During our sample period, the average level of the FTSE 100 index is 5009.98 and does not vary much across moneyness.

6. EMPIRICAL RESULTS

Given that in virtually all of the prior empirical studies on the role of jumps in the specification of price dynamics and option pricing the focus was on either the S&P 500 or NASDAQ 100 index, the investigation of other assets is quite sparse. Therefore, prior to exploring the impact of jumps on the predictive power of option-implied densities, we must first estimate the parameters of the SV, SVJ and SVJJ models to determine whether the parameters describing the price dynamics of the FTSE 100 index are consistent with the general expectation, and then use the estimated parameters to examine the way in which the jumps affect the option pricing errors of the FTSE 100 index. We then go on to investigate the impact of the jump dynamics on density predictions for the FTSE 100 index. Finally, we briefly look at the risk preference implied in the option market.

6.1 Estimation of Model Parameters

We begin with the estimation of the parameters of the SV, SVJ and SVJJ models by minimizing the sum of the squared differences between the market and theoretical prices of the FTSE 100 index options across different strike prices and different periods of time-to-maturity, once per day. The summary statistics of the parameter estimates of the three models are provided in Exhibit 2.

<Exhibit 2 is inserted about here>

As suggested in several of the prior studies, the addition of jumps lowers the average annualized level of the stochastic volatility ($\sqrt{\upsilon}$) from 25.96 per cent (SV model) to 21.10 per cent (SVJJ model).[10] The estimation of the initial stochastic volatility ($\sqrt{V_0}$) is also lower when the jump factor is considered. Although a higher value of the volatility of volatility (σ_v) implies fatter tails of the index level distribution, the value for the SV model is unreasonably high (0.73); however, adding jumps reduces this level substantially, to 0.47 for the SVJ model and 0.41 for the SVJJ model.[11] In the same vein, the speed of the mean reversion of volatility (κ_v) is also reduced from 2.06 to about 1.6.

Adding a jump factor into the price process (SVJ model) essentially fattens the tails of the returns distribution, with the volatility of the jump size fattening both tails, whereas a negative mean ($\mu_y = -0.41$) implies that the right tail is relatively less fat, which is consistent with the general findings for equity returns. The arrival intensity of the jump in the price process of the SVJ model is estimated to be an average of 0.32.

In the SVJ model, the significant movement in prices resulting from jumps has no impact on volatility. Correcting the shortcoming in the SVJ model, the SVJJ specification allows jumps to affect both prices and volatility. The addition of the jump factor in volatility makes the distribution of volatility more positively skewed, and hence, fattens the tails of the returns distribution.

The intensities of jumps in both prices and volatility in the SVJJ model follow the same Poisson process, with the sizes of the jumps all being correlated. The correlation (ρ_j)

is found to be negative (-0.47), which indicates that the larger the market price crash, the greater the increase in volatility. The patterns of the parameters associated with the jump process in prices (λ_c, μ_{cy} and σ_{cy}) are similar to those of the SVJ model, although the magnitudes of the former are smaller.

In summary, the estimates of all of the model parameters for the FTSE 100 index are consistent with the theoretical expectations, as well as the findings of the prior studies with regard to other assets; namely, a more general model can make the distribution of both prices and volatility more realistic, and explain more of the stylized facts observed in the real world stock market.

6.2 Fitting of Option Prices

Given the estimates of the model parameters from daily option market prices, we then go on to calculate the theoretical prices under the three price dynamic assumptions. In this section, we investigate whether adding jumps can improve the model price fitting for the FTSE 100 index options. A summary of the fitting errors across different levels of moneyness and periods of time-to-maturity, defined by the difference between the theoretical and market prices, is provided in Exhibit 3.

<Exhibit 3 is inserted about here>

As compared with the Black-Scholes model, the SV model offers a flexible price distribution with substantially improved option pricing. Following the lead of Bakshi et al.

(1997), in many subsequent studies on the role of jumps in option pricing it has been suggested that the incorporation of jumps into the price process is very important; however, Bakshi et al. (1997) found that the SVJ model still systematically overpriced the OTM call options. Nevertheless, Bates (2000), Duffie et al. (2000), Pan (2002) and Eraker et al. (2003) each demonstrated that adding jumps into volatility could attenuate the overpricing problem in the SVJ model.

In general, our results for the FTSE 100 index are consistent with those of the prior studies. The addition of jumps into the price process (SVJ model) lowers the fitting errors considerably. For short-maturity near-the-money calls, as compared to the SV model, the reduction in the average fitting error for the SVJ model was found to be as high as 46.03 per cent.

In addition to the significant improvement in short-maturity (<90 days) calls, the improvement in longer-maturity (>90 days) calls is also very impressive. When we allow simultaneous and correlated jumps in the volatility process (SVJJ model), the fitting errors are reduced still further. The average magnitude of the SVJJ model over the SVJ model is not as significant as that of the SVJ model over the SV model; nevertheless, the improvement is, on average, in excess of 10 per cent.

In summary, our results provide strong support for the argument that the incorporation of jumps into option pricing is an important development, since we find that the model with jumps incorporated into the price process can substantially improve the precision of option pricing. Adding jumps into the volatility process can further lower the pricing error, and indeed, the presence of jumps in volatility does not eliminate the requirement for jumps in prices. The contribution of the jumps comes from the more realistic distributions of prices and volatility.

6.3 Density Prediction

Although a more complex model can generally result in better in-sample performance, it may not necessarily have better out-of-sample performance since such performance may be penalized by overfitting. Therefore, in this section we investigate whether the addition of jumps can improve the density predictions of the future levels of the FTSE 100 index.

The density predictions are executed for five horizons, 1 day, 1 week, 2 weeks, 3 weeks and 4 weeks, evaluated under the Anderson-Darling and Berkowitz (2001) tests. We form the forecasts once per unit of horizon so as to avoid overlapping forecasts; for example, for the 2-week-ahead forecast, the forecast is formed once every two weeks so that there is no forecast between a formation date and the corresponding realization date. As a result, we have 1,486 observations for the 1-day-ahead predictions, and 38 observations for the 4-week-ahead predictions.

Given the estimates of the model parameters, the RNDs for the five forecast horizons can be constructed for the three different models using the closed-form solution of

Equation (11). We transform the RNDs to RWDs using either the Beta-distribution or the non-parametric calibration method. The RNDs and RWDs under the two different calibrations are then evaluated against the realized levels of the FTSE 100 index for every forecast horizon.

The results of the Anderson-Darling tests, including the statistics and the *p*-values, are presented in Exhibit 4. Basically, the lower (higher) the statistic (*p*-value), the more satisfactory the density prediction. The test results show that with the exception of the 3-week-ahead forecast horizon, the RNDs constructed for all other forecast horizons using the SV model provide poor prediction of the density of the future levels of the FTSE 100 index at the 5 per cent significance level. By contrast, with the exception of the 1-week-ahead forecasts, the addition of jumps in volatility results in acceptable forecasts for all other horizons at the 5 per cent significance level.

<Exhibit 4 is inserted about here>

Adding jumps into the volatility process further raises the significance levels of most of the forecast horizons above 10 per cent. In general, therefore, the incorporation of jumps with stochastic volatility can improve the overall precision of the density predictions made by the RNDs, with the SVJJ model demonstrating the most satisfactory performance.

Consistent with the prior studies, including Bliss and Panigirtzoglou (2004), Anagnou-Basioudis et al. (2005), Shackleton et al. (2006) and Liu et al. (2007), the RWDs provide satisfactory forecasts for the densities of future index levels for all prediction horizons, and indeed, these are not dependent on either the model specifications or the approaches used for the construction of the RWDs. All of the Anderson-Darling statistics are insignificant at a very high level, with the highest being those for the RWDs constructed under the SVJJ model. Most of the *p*-values used to evaluate the RWDs are greater than 0.8.

Since most of the statistical tests based on the PITs, such as the Anderson-Darling test, may not be sufficiently powerful for small samples, we also employ the test proposed by Berkowitz (2001) to evaluate our density prediction; the Berkowitz *LR*3 and *LR*1 test statistics and *p*-values are presented in Exhibit 5. From the results in Panel 1, we find that almost all of the RNDs constructed using the SV model provide poor density prediction for the FTSE 100 index, since all of the *LR*3 tests are rejected and all of the *LR*1 tests, with the exception of the 3-week forecast horizon, are accepted.

<Exhibit 5 is inserted about here>

In contrast to the results obtained using the Anderson-Darling test, adding jumps into prices and/or volatility results in satisfactory density predictions for only the 1-day-ahead forecast at the 10 per cent significance level, since both the *LR*3 and *LR*1 tests are accepted. According to the test statistics and *p*-values, although the incorporation of jumps with stochastic volatility does improve the prediction performance of the RNDs, the improvement may not be sufficiently significant to render such predictions acceptable under a reasonable level of significance. If we adopt a 1 per cent significance level, then the 1-day, 2-week and 3-week-ahead forecasts of the RNDs under the SVJJ model are acceptable.

From Panels 2 and 3 of Exhibit 5, we find that the RWD forecasts are satisfactory for all prediction horizons at the 5 per cent significance level, with some of the *p*-values being even greater than 0.8. In similar vein to our findings using the Anderson-Darling test, the improved performance provided by the RWDs is dependent on neither the model specifications nor the approaches used to transform the densities. However, given that the power of the Berkowitz test may be higher than that of the Anderson-Darling test, the *p*-values are, on average, lower.

To further visualize our findings, we use the realized index levels to calculate the values of their CDFs assessed by both the RNDs and RWDs constructed under the three different models for the five different forecast horizons and then go on to plot the differences between the predicted CDF values and their empirical CDF values against the empirical CDF values.[12] If a density prediction is perfect, then the differences are equal to zero.

Since the patterns for all forecast horizons are similar, in Exhibit 7 the plots are produced for the 1-day horizon only and those for other horizons are available upon request. As illustrated by the plots, the RWDs provide remarkable improvements over the prediction power of the RNDs, across all model specifications and forecast horizons, since the differences for the RWDs are much smaller.

<Exhibit 7 is inserted about here>

Following the same assessment rules, the addition of jumps also improves the density prediction of the SV RND, with the most significant improvement being found in the 1-day-ahead forecast. Moreover, as the differences for RNDs are highly positive (negative) in the left (right) tail, the risk-neutral distributions implied in the option prices are more negatively skewed than those of the spot index levels, with the calibrated RWDs mitigating this problem substantially.

In summary, based upon the results of this study, using two different tests, we find that the addition of jumps into the price and volatility dynamics does improve the precision of the density predictions made by RNDs; however, the magnitude of such improvement may not always be sufficiently significant. By contrast, regardless of whether we used the Beta-distribution or non-parametric calibration method, or whether or not jumps were incorporated, the RWDs provided consistently satisfactory density predictions for future FTSE 100 index levels.

6.4 Risk Preference Implied in the Option Market

Shackleton et al. (2006) derive the necessary and sufficient condition for the Beta-distribution calibrated density to imply a risk-averse utility function, $\alpha \leq 1 \leq \beta$ (with $\alpha \neq \beta$). If either α>1 or β<1, then the representative agent is a risk-seeker for an interval of wealth, with respect to which the second derivative of the utility function is positive. Exhibit 6 lists our estimates of α and β across models and forecast horizons. Most of the estimates of α are larger than 1, indicating that the risk preference implied in the option prices of the FTSE 100 index is different from the rational assumption imposed in most asset pricing models. If one does not constrain α to be 1, then Shackleton et al. (2006) also find that the implied risk preference is not risk-aversion. Using different approaches, Carr et al. (2002) and Constantinides et al. (2006) also find empirical support for the irrational risk preference.

In our study, the risk-preference implied by the 1-day density prediction of the SVJJ model is the only one case that fits the rational assumption. This finding leaves some questions for future studies, including whether option investors are really risk-seekers, whether there is any space to improve the SVJJ model for option pricing, and whether there is a better density for the pricing kernel.

7. CONCLUSIONS

We have used the Anderson-Darling and Berkowitz (2001) tests in this study to evaluate the density predictions of multi-horizon RNDs and RWDs implied under the three different continuous-time option pricing (SV, SVJ and SVJJ) models, nested by the general double-jump model of Duffie et al (2000). By so doing, we have been able to investigate whether the addition of jumps into the price and/or volatility dynamics can lead to improvements in density prediction. Our empirical examination was conducted for the FTSE 100 index at five different forecast horizons, with the RWDs being transformed from RNDs under either the Beta-distribution or the non-parametric calibration method.

Our findings indicate that (i) all estimates of the model parameters are consistent with their theoretical expectations, (ii) adding jumps into the price and volatility processes substantially lowers the fitting errors of the option prices, (iii) incorporating jumps with stochastic volatility can improve the predictive power of RNDs, and (iv) RWDs provide consistently satisfactory density predictions, which are dependent on neither the model selection nor the approaches used to construct the RWDs.

In conclusion, apart from presenting a more flexible framework for estimating multi-horizon RNDs, this paper also offers important empirical evidence on the role of jumps in density prediction. In addition, our empirical results regarding the option-implied risk preference leave some interesting questions for future studies.

REFERENCES

Anagnou-Basioudis, I., M. Bedendo, S. Hodges and R. Tompkins (2005), 'Forecasting the Accuracy of Implied and GARCH-based Probability Density Functions', *Review of Futures Markets*, **11**: 41-66.

Bakshi, G., C.G. Cao and Z. Chen (1997), 'Empirical Performance of Alternative Option Pricing Models', *Journal of Finance*, **52**: 2003-49.

Bahra, B. (1997), 'Implied Risk-neutral Probability Density Functions from Option Prices: Theory and Application', *Working Paper*, Bank of England.

Bates, D.S. (1996), 'Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutschemark Options', *Review of Financial Studies*, **9**: 69-107.

Bates, D.S. (2000), 'Post-'87 Crash Fears in the S&P 500 Futures Option Market', *Journal of Econometrics*, **94**: 181-238.

Berkowitz, J. (2001), 'Testing Density Forecasts with Applications to Risk Management', *Journal of Business and Economic Statistics*, **19**: 465-74.

Bliss R.R. and N. Panigirtzoglou (2002), 'Testing the Stability of Implied Probability Density Functions', *Journal of Banking and Finance*, **26**: 381-422.

Bookstaber, R. and J. McDonald (1987), 'A General Distribution for Describing Security Price Returns', *Journal of Business*, **60**: 401-24.

Carr, P., H. Geman, D. B. Madan and M. Yor (2002), 'The Fine Structure of Asset Returns: An Empirical Investigation', *Journal of Business*, **75**: 305-332.

Carr, P. and L. Wu (2004), 'Time-changed Levy Process and Option Pricing', *Journal of Financial Economics*, **71**: 113-41.

Constantinides, G. M., J. C. Jackwerth and S. Perrakis (2006), 'Mispricing of S&P 500 Index Options', *Working paper*, University of Chicago and NBER.

Cont, R. (1997), 'Beyond Implied Volatility', in J. Kertesz and I. Kondor (eds.), *Econophysics*, Dordrecht: Kluwer.

D'Agostino, R.B. and M.A. Stephens (1986), *Goodness-of-Fit Techniques*, New York: Marcel Dekker Inc.

Diebold, F.X., T.A. Gunter and A.S. Tsay (1998), 'Evaluating Density Forecasts', *International Economic Review*, **39**: 863-83.

Duffie, D., J. Pan and K.J. Singleton (2000), 'Transform Analysis and Asset Pricing for Affine Jump-diffusions', *Econometrica*, **68**: 1343-76.

Eraker, B. (2004), 'Do Equity Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices', *Journal of Finance*, **59**: 1369-403.

Eraker, B., M.S. Johannes and N.G. Polson (2003), 'The Impact of Jumps in Returns and Volatility', *Journal of Finance*, **53**: 1269-300.

Figlewski, S. (2008), 'Estimating the Implied Risk Neutral Density for the U.S. Market Portfolio', forthcoming in *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, UK: Oxford University Press.

Frachler, P.L. and R. King (1990), 'Calibration of Option-based Probability Assessments in Agricultural Commodity Markets', *American Journal of Agricultural Economics*, **72**: 73-83.

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

Jackwerth, J.C. (1999), 'Option-implied Risk-neutral Distributions and Implied Binomial Trees: A Literature Review', *Journal of Derivatives*, **7**: 66-82.

Jondeau, E. and M. Rockinger (2000), 'Reading the Smile: The Message Conveyed by Methods which Infer Risk-neutral Densities', *Journal of International Money and*

Finance, **19**: 885-915.

Liu, X., M. Shackleton, S.J. Taylor and X. Xu (2007), ‘Closed-form Transformations from Risk-neutral to Real-world Distributions’, *Journal of Banking and Finance*, **31**: 1501-20.

Madan, D.B. and F. Milne (1994), ‘Contingent Claims Valued and Hedged by Pricing and Investing in a Basis’, *Mathematical Finance*, **4**: 223-45.

Malz, A. (1996), ‘Using Option Prices to Estimate Realignment Probabilities in the European Monetary System: The Case of Sterling-Mark’, *Journal of International Money and Finance*, **15**: 717-48.

Malz, A. (1997), ‘Estimating the Probability Distribution of the Future Exchange Rate from Option Prices’, *Journal of Derivatives*, **5**: 18-36.

Noceti, P., J. Smith and S.D. Hodges (2003), ‘An Evaluation of Tests of Distributional Forecasts’, *Journal of Forecasting*, **22**: 447-55.

Pan, J. (2002), ‘The Jump Risk Premium Implicit in Options: Evidence from an Integrated Time-series Study’, *Journal of Financial Economics*, **63**: 3-50.

Ritchey, R. (1990), ‘Call Option Valuation for Discrete Normal Mixtures’, *Journal of Financial Research*, **13**: 285-95.

Shackleton, M.B., S.J. Taylor and P. Yu (2006), ‘A Multi-horizon Comparison of Density Forecasts for the S&P 500 using Index Returns and Option Prices’, *Working Paper*, Lancaster University.

Shimko, D. (1993), ‘Bounds of Probability’, *Risk*, **6**: 33-37.

Exhibit 1 Summary of option prices

Moneyness	Maturity				Average
	<90	90-180	>180	Total	Index Level
<0.94	711.54 (10,851)	830.50 (8,272)	958.06 (10,493)	832.10 (29,616)	5001.91
0.94-0.97	287.34 (2,970)	347.93 (1,663)	450.31 (1,630)	345.80 (6,263)	5084.65
0.97-1.00	174.61 (3,032)	246.48 (1,754)	350.63 (1,873)	243.10 (6,659)	5100.44
1.00-1.03	91.03 (2,996)	163.65 (1,721)	260.13 (1,692)	155.20 (6,409)	5061.87
1.03-1.06	46.45 (2,923)	102.31 (1,707)	182.96 (1,664)	97.69 (6,294)	5071.01
>1.06	22.67 (4,826)	46.79 (5,077)	81.73 (7,350)	54.93 (17,253)	4920.38
Total	348.63 (27,598)	424.61 (20,194)	517.73 (24,702)	427.40 (72,494)	5009.98

Both the average and total numbers of settlement prices are provided for call options on the FTSE100 index across 'time-to-maturity' and 'moneyness' for the years 2000-2005. All prices of put options are converted to their equivalent call prices using put-call parity; only out-of-the-money options are selected. Moneyness is defined by the ratio of the strike price to the forward price. Time-to-maturity is defined by the number of calendar days between trading and expiry dates. Figures in parentheses refer to the total number of option prices within each of the classifications. The average index levels across moneyness are also provided in the last column.

Exhibit 2 Parameter estimates for alternative models

Parameters	κ_v	υ	σ_v	ρ	V_0	λ_c	μ_{cy}	σ_{cy}	μ_{cv}	ρ_j
Panel A: SV Model										
Mean	2.0613	0.0674	0.7273	−0.6618	0.0476	–	–	–	–	–
Median	1.9483	0.0595	0.6796	−0.6588	0.0325	–	–	–	–	–
Std. Dev.	1.1439	0.0336	0.2594	0.0721	0.0468	–	–	–	–	–
Panel B: SVJ Model										
Mean	1.5492	0.0541	0.4713	−0.6475	0.0410	0.3411	−0.4102	0.2155	–	–
Median	1.1623	0.0439	0.3968	−0.6323	0.0284	0.0696	−0.3758	0.2142	–	–
Std. Dev.	1.5407	0.0419	0.2428	0.1087	0.0407	0.8465	0.2852	0.1351	–	–
Panel C: SVJJ Model										
Mean	1.6307	0.0445	0.4145	−0.6525	0.0397	0.3182	−0.2664	0.1305	−0.0483	−0.4673
Median	1.1669	0.0307	0.3673	−0.6378	0.0275	0.2516	−0.1330	0.0219	0.0679	−0.4712
Std. Dev.	1.7190	0.0497	0.2983	0.1585	0.0400	0.3457	0.3273	0.1836	0.4382	0.2168

Summary statistics of the parameter estimates of SV, SVJ and SVJJ models are provided for the years 2000-2005. The parameters are estimated by minimizing the loss function, defined as the sum of the squared errors of option price fitting.

Exhibit 3 Price fitting errors under alternative models

Moneyness	Maturity								
	<90			90-180			>180		
	SV	SVJ	SVJJ	SV	SVJ	SVJJ	SV	SVJ	SVJJ
<0.94	1.99	1.34	1.16	1.86	1.51	1.35	2.10	1.67	1.58
		(32.76)	(12.79)		(18.81)	(11.08)		(20.82)	(5.10)
0.94-0.97	2.38	1.38	1.20	2.28	1.58	1.39	2.18	1.92	1.77
		(42.28)	(12.45)		(30.86)	(11.66)		(11.96)	(7.60)
0.97-1.00	2.69	1.45	1.26	1.93	1.52	1.33	2.19	1.92	1.74
		(46.03)	(13.09)		(21.37)	(12.10)		(12.38)	(9.41)
1.00-1.03	2.15	1.42	1.17	2.47	1.52	1.42	2.27	1.86	1.67
		(34.01)	(17.90)		(38.49)	(6.65)		(18.27)	(9.74)
1.03-1.06	1.81	1.37	1.19	2.99	1.69	1.54	2.43	1.93	1.65
		(24.42)	(13.00)		(43.52)	(8.86)		(20.92)	(14.50)
>1.06	1.32	1.27	1.19	2.26	1.58	1.49	2.51	1.75	1.58
		(4.04)	(6.12)		(30.30)	(5.29)		(30.19)	(9.87)

Average price fitting errors are provided for call options on the FTSE100 index across 'time-to-maturity' and 'moneyness' using SV, SVJ and SVJJ option pricing models. The fitting error is defined as the absolute difference between the model-generating price and the market price of the option. Time-to-maturity is defined by the number of calendar days between trading and expiry dates. Moneyness is defined by the ratio of the strike price to the forward price. Figures in parentheses under the SVJ (SVJJ) pricing errors refer to the improvements in option price fitting achieved by the model over the SV (SVJ) model, defined as the difference in price fitting errors ÷ SV (SVJ) errors x 100.

Exhibit 4   Evaluation of density predictions under the Anderson-Darling test

Forecast Horizons	Model								
	RND			RWD (Beta distribution)			RWD (Non-parametric distribution)		
	SV	SVJ	SVJJ	SV	SVJ	SVJJ	SV	SVJ	SVJJ
1 day	6.2785 (0.0007)	2.2263 (0.0692)	1.5821 (0.1580)	1.1579 (0.2841)	1.1474 (0.2884)	1.0179 (0.3480)	0.4419 (0.8064)	0.4423 (0.8060)	0.4212 (0.8274)
1 week	6.3438 (0.0007)	4.4573 (0.0052)	3.8598 (0.0102)	0.2724 (0.9573)	0.1919 (0.9925)	0.1872 (0.9935)	0.2900 (0.9455)	0.2901 (0.9454)	0.2975 (0.9400)
2 weeks	2.5128 (0.0488)	1.9818 (0.0940)	1.8262 (0.1147)	0.3836 (0.8648)	0.4796 (0.7677)	0.5275 (0.7189)	0.3377 (0.9074)	0.3740 (0.8740)	0.3950 (0.8537)
3 weeks	1.5356 (0.1682)	1.2827 (0.2380)	1.1640 (0.2816)	0.3763 (0.8719)	0.4098 (0.8389)	0.3918 (0.8568)	0.3770 (0.8712)	0.3828 (0.8656)	0.3798 (0.8685)
4 weeks	2.6321 (0.0423)	2.3635 (0.0585)	2.2246 (0.0693)	0.4271 (0.8215)	0.5693 (0.6775)	0.4938 (0.7532)	0.3996 (0.8491)	0.3835 (0.8649)	0.3512 (0.8953)

The Anderson-Darling statistics are used to evaluate the predictive power of the ‘risk-neutral densities’ (RNDs) and the ‘real-world densities’ (RWDs) estimated with alternative models from the FTSE 100 option prices for the years 2000-2005. The forecast horizons include 1 day, 1 week, 2 weeks, 3 weeks and 4 weeks, with the periods for all horizons being non-overlapping. The RWDs are transformed from the RNDs using either parametric (Beta distribution function) or non-parametric statistical calibration. Figures in parentheses refer to *p*-values.

Exhibit 5 Evaluation of density predictions under the Berkowitz (2001) test

Forecast Horizons	Model					
	SV		SVJ		SVJJ	
	LR3	LR1	LR3	LR1	LR3	LR1
Panel 1: RND						
1 day	23.9067 (0.0000)	2.2189 (0.1363)	5.1888 (0.1575)	2.3522 (0.1251)	3.8483 (0.2783)	2.3960 (0.1216)
1 week	35.8530 (0.0000)	0.0172 (0.8958)	25.2688 (0.0000)	0.0538 (0.8166)	21.3113 (0.0001)	0.0460 (0.8301)
2 weeks	15.1981 (0.0017)	2.5071 (0.1133)	12.5399 (0.0057)	2.2822 (0.1309)	11.1067 (0.0112)	2.1954 (0.1384)
3 weeks	11.7437 (0.0083)	3.6859 (0.0549)	9.9599 (0.0189)	3.2151 (0.0730)	9.4303 (0.0241)	3.0066 (0.0829)
4 weeks	16.0759 (0.0011)	1.2821 (0.2575)	13.8448 (0.0031)	1.1404 (0.2856)	13.3838 (0.0039)	1.1302 (0.2877)
Panel 2: RWD (Beta distribution)						
1 day	2.2294 (0.5262)	2.2219 (0.1361)	2.3787 (0.4976)	2.3589 (0.1246)	2.4144 (0.4910)	2.4024 (0.1211)
1 week	0.0229 (0.9991)	0.0151 (0.9021)	0.0510 (0.9970)	0.0497 (0.8235)	0.0440 (0.9976)	0.0420 (0.8376)
2 weeks	2.5106 (0.4734)	2.5063 (0.1134)	2.2870 (0.5150)	2.2821 (0.1319)	2.2025 (0.5315)	2.1959 (0.1384)
3 weeks	3.6875 (0.2972)	3.6857 (0.0549)	3.2252 (0.3582)	3.2196 (0.0728)	3.0156 (0.3892)	3.0120 (0.0826)
4 weeks	1.2758 (0.7349)	1.2756 (0.2587)	1.1357 (0.7685)	1.1341 (0.2869)	1.1251 (0.7710)	1.1242 (0.2890)
Panel 3: RWD (Non-parametric distribution)						
1 day	4.0345 (0.2578)	2.1506 (0.1425)	4.2503 (0.2357)	2.3266 (0.1272)	4.4241 (0.2192)	2.4691 (0.1161)
1 week	1.4255 (0.6996)	0.0578 (0.8100)	1.4688 (0.6895)	0.0626 (0.8025)	1.4695 (0.6892)	0.0619 (0.8035)
2 weeks	3.8209 (0.2815)	2.4467 (0.1178)	3.7585 (0.2887)	2.3377 (0.1263)	3.6817 (0.2979)	2.2365 (0.1348)
3 weeks	4.7616 (0.1901)	3.5474 (0.0596)	4.7327 (0.1925)	3.4100 (0.0648)	4.6497 (0.1993)	3.2929 (0.0696)
4 weeks	2.2472 (0.5227)	1.2074 (0.2719)	2.2991 (0.5127)	1.1373 (0.2862)	2.2614 (0.5199)	1.1278 (0.2882)

The Berkowitz (2001) LR1 and LR3 statistics are used to evaluate the predictive power of 'risk-neutral densities' (RNDs) and 'real-world densities' (RWDs) estimated under alternative models using the FTSE 100 option prices for the years 2000-2005. The forecast horizons include 1 day, 1 week, 2 weeks, 3 weeks and 4 weeks, with the periods for all horizons being non-overlapping. Figures in parentheses refer to *p*-values. The RWDs are transformed from the RNDs using either parametric (Beta distribution function) or non-parametric statistical calibration.

Exhibit 6　Parameter estimates of the Beta function

Model / Horizon	SV α	SV β	SVJ α	SVJ β	SVJJ α	SVJJ β
1 day	1.1403	1.1784	1.0255	1.0623	0.9870	1.0235
1 week	1.5131	1.6150	1.3983	1.4899	1.3490	1.4409
2 weeks	1.3994	1.4997	1.3451	1.4340	1.3076	1.4016
3 weeks	1.3823	1.4794	1.3373	1.4306	1.3256	1.4151
4 weeks	1.7389	1.8784	1.6648	1.7834	1.6512	1.7589

α and β are the parameters of the Beta distribution function used to transform the option-implied risk-neutral densities to the real-world densities.

Exhibit 7   Variations in predicted (option implied) and empirical CDF values

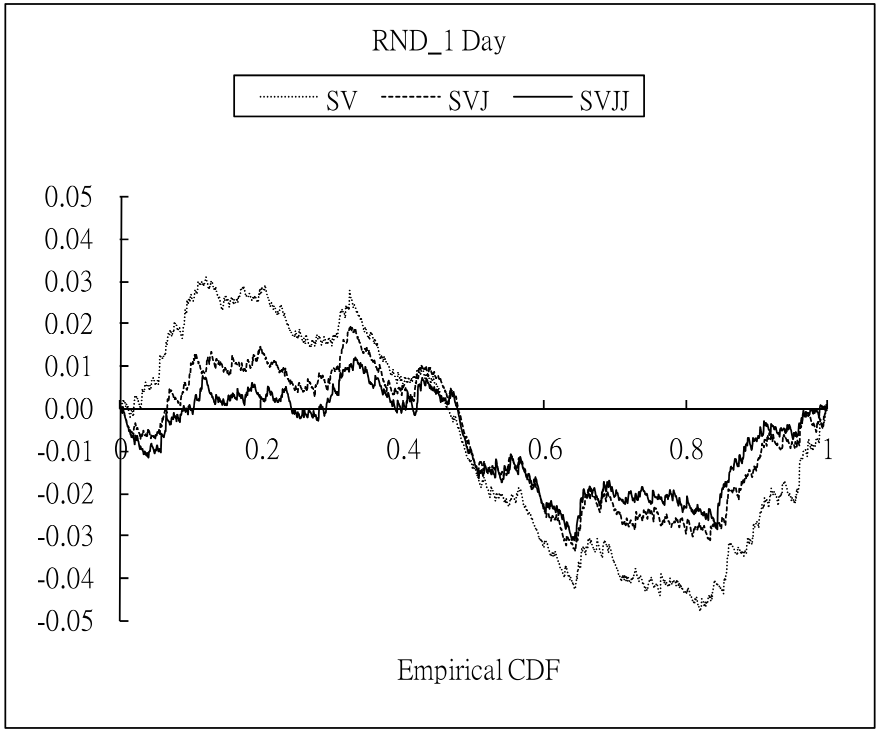

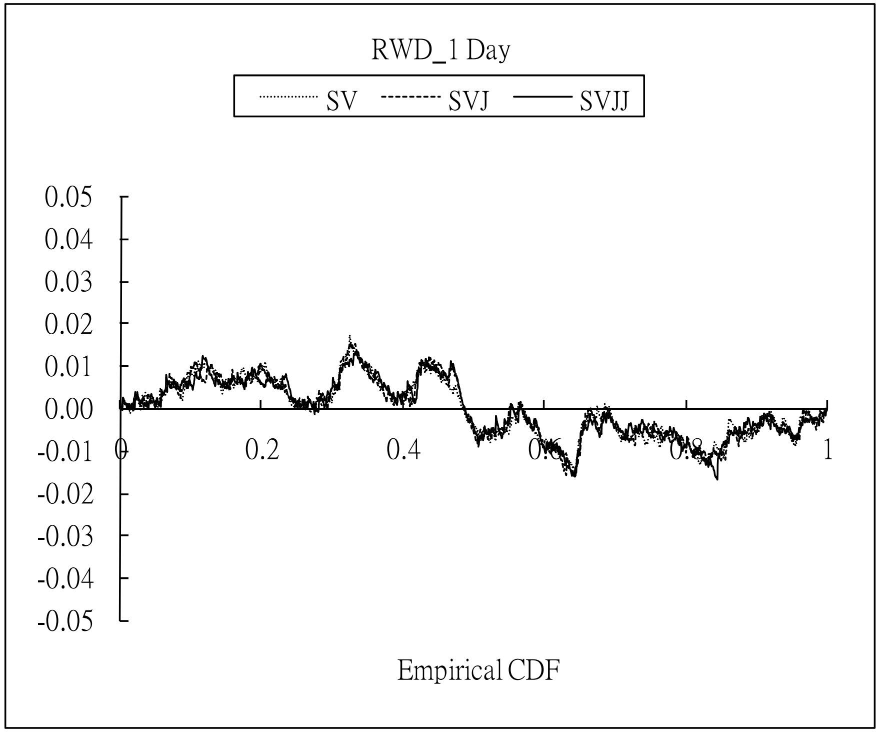

The Figures plot the differences between the predicted CDF values and the empirical CDF values across each of the models, both of which are evaluated with RNDs and RWDs (plotted individually against the empirical CDFs) for the purpose of comparison. As the patterns for all forecast horizons are similar, the plots are produced for the 1-day horizon only. Since the transformed RWDs (using the two different calibration methods) are quite similar, we report only the results for the RWDs transformed under non-parametric calibration.

ENDNOTES

[1] In some studies, a particular type of distribution is assumed for the underlying asset price at the option maturity date; examples include lognormal mixtures (Ritchey, 1990), generalized beta distribution of the second kind (GB2) (Bookstaber and McDonald, 1987) and the Hermite lognormal-polynomial method (Madan and Milne, 1994). As opposed to directly specifying the density, some specify the implied volatility function of the Black-Scholes formula, usually as a polynomial (Shimko, 1993; Malz, 1996, 1997). Furthermore, in order to achieve greater flexibility, some studies have used real data with a density smoothing technique to form a non-parametric distribution. A more recent study, Figlewski (2008), proposes a comprehensive approach that takes into account of the market's bid-ask spread for smoothing and completes the density beyond the available range of strike prices with a Generalized Extreme Value distribution.

[2] In particular, the prices for short-maturity options (for example, 5 days) are usually excluded due to liquidity concerns, which results in difficulties in predicting the distribution over such a short horizon.

[3] Jondeau and Rockinger (2000) compared the performance of the stochastic volatility model with certain parametric methods using only the market prices of options with the same maturity as the forecast horizon. Shackleton et al. (2006) used the stochastic volatility model of Heston (1993) to execute a multi-horizon comparison of density forecasts for the S&P 500 index.

[4] In almost all of the prior empirical studies on the role of jumps in the specification of price dynamics and option pricing, the focus was on either the S&P 500 or NASDAQ 100 index; the investigation of other assets is quite sparse.

[5] Numerous studies over the past decade have placed considerable effort into determining ways of transforming an RND into an RWD; these include Bliss and Panigirtzoglou (2004), Anagnou-Basioudis et al. (2005) and Liu et al. (2007). Basically, two methods have been adopted for such transformation, the first method involving applying the relationship between a utility function and the pricing kernel, and the second involving the use of a statistical calibration. In our analysis in the present study, the computation load involved in the first of these methods is regarded as being too heavy.

[6] Since the standardized values have unit variance, it is acceptable to use $B = n^{-0.2}$.

[7] See Diebold et al. (1998).

[8] See D'Agostino and Stephens (1986) and Noceti et al. (2003).

[9] The moneyness level is defined by the ratio of the strike price to the forward price.

[10] This is consistent with the findings of Eraker et al. (2003).

[11] Bates (1996, 2000) and Bakshi et al. (1997) found that the volatility level in the SV model appeared to be too high.

[12] According to the evaluation results of the density prediction, the performance of RWD is not dependent on the approach used to transform the densities. Here we report only the results for the RWDs transformed under the non-parametric calibration method.