Vix option pricing in a jump-diffusion model

Artur Sepp discusses Vix futures and options and shows that their market prices exhibit positive volatility skew. To better model the market behaviour of the S&P 500 index and its associated volatility skew, he introduces the stochastic dynamics of the volatility of the S&P 500 index with volatility jumps. Then he develops closed-form solutions for unified pricing of options on the S&P 500 index and its volatility

The Chicago Board Options Exchange (CBOE) Volatility Index (Vix) measures the implied volatility of S&P 500 stock index options with a maturity of 30 days. In a broad sense, the Vix represents the market expectation of the annualised at-the-money (ATM) implied volatility over the next 30-day period. The Vix spot value is calculated by the CBOE minute-to-minute using real-time bid/ask market quotes of S&P 500 index (SPX) options with nearby and second nearby maturities and applying the multiplier of $100.

The exchange-listed Vix-based derivatives include futures contracts, which began trading in 2004, and call and put options on the Vix, which began trading in February 2006. The final settlement date of the Vix futures contract is the third Wednesday of each month. Typically, there are listed futures contracts with a settlement date up to six near-term months and a few longer-term contracts. The underlying of the Vix call and put options is the Vix spot value observed on the option expiry date, which is specified in the same way as the settlement day for futures contracts. The Vix options are of European exercise style.

According to the CBOE Futures Exchange press release on July 11, 2007, in June 2007 the average daily volume of Vix options was 95,283 contracts, making the Vix the second most actively traded index and the fifth most actively traded product on the CBOE. On July 11, open interest in Vix options stood at 1,845,820 contracts (1,324,775 calls and 521,045 puts). In the same month, the Vix futures totalled 78,578 contracts traded, with open interest at 49,894 contracts at the end of June.

There are a number of reasons for trading in the volatility futures. The main one is that since options on the Vix derive their values from the implied volatility of the S&P 500 index, they are attractive for investors, who want to get exposure to the volatility of the S&P 500 without taking direct positions in the index and without the need to delta-hedge their portfolios. In addition, for equity portfolio insurance it might be less expensive to hedge exposure to the S&P 500 index by taking a position in out-of-the-money call options on the Vix rather than buying out-of-the-money puts on the S&P 500 index.

Given the growing popularity of contracts deriving their values from the implied volatilities of major indexes, including S&P 500, Dax, Eurostoxx 50 and Nasdaq, it is important to develop a dynamic model for these types of product, and analyse the model implied distributions and hedging strategies.

Empirical analysis of the Vix within an econometric framework, as well as general discussion of Vix futures, can be found in recent papers by Psychoyios & Skiadopoulos (2006), Dotsis, Psychoyios & Skiadopoulos (2007), Zhang & Zhu (2006) and Zhu & Zhang (2007). Although the first two of these papers do consider jumps in the volatility index, which are extremely important to fit the skew observed in Vix options, they treat the Vix spot value as a stand-alone process, which makes it difficult to deal with Vix and other volatility products on the S&P 500 index in a consistent way.

In this article, we consider the unified pricing of volatility products and options on Vix products and assume that the variance of returns on the S&P 500 is driven by square-root diffusion (Heston, 1993) with variance jumps and time-dependent parameters. Since model parameters can be calibrated using market data on Vix products and/or S&P 500 products, including vanilla options and variance swaps, this model can also serve for the relative-value analysis and cross-hedging of different volatility products.

Vix options implied volatility

The first model to price options on an implied volatility index was originated by Whaley (1993), who applied Black's formula (1976) to value the call option on the futures contract. Grünblicher & Longstaff (1996) applied the mean-reverting square-root process for the implied volatility index. In this section, we follow Whaley and apply Black's formula to price the call option on the Vix futures:

$$C(t,T,F,K) = DF(t,T)\big(F(t,T)\mathcal{N}(d_+) - K\mathcal{N}(d_-)\big),$$
$$d_\pm = \frac{\ln\frac{F(t,T)}{K} \pm \frac{1}{2}(T-t)\sigma^2}{\sqrt{(T-t)}\sigma}, \mathcal{N}(x) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{x} e^{-\frac{y^2}{2}}dy \quad (1)$$

where t is the current valuation time, $C(t, T, F, K)$ stands for the value of the call option with expiry time T and strike K, $F(t, T)$ is the futures price with expiry time T, $DF(t, T)$ is the discount factor applied to time T, and σ is the volatility of the futures price, which can be used for quoting or parametrisation of the market price of the call option.

Thus, given the market prices of the Vix options along with the prices of the Vix futures, we can infer their implied volatilities by means of formula (1). Since losses in the S&P 500 index are typically followed by high levels of the Vix, out-of-the-money call options on the Vix provide protection against a market crash. As a result, the call option writer takes on risk and charges compensation for taking this risk, in a similar way to when a writer of an out-of-the-money put option on the S&P 500 index charges extra compensation for their risk. Thus, in contrast to the downward sloping implied volatility skew (negative skew) observed in the S&P 500 put options, the skew observed in Vix call options has an upward sloping skew (positive skew). The term structure of the corresponding implied skews in Vix call options observed on July 25, 2007 is shown in figure 1.

The weakness of the Black and Grünblicher & Longstaff approaches to Vix options pricing is that these models are separated from the actual evolution of the S&P 500 index volatility and, as a result, they can mis-specify risks of Vix futures and calls, especially the volatility-of-volatility risk. Since Vix futures are a non-linear function of the expected future realised variance, they have their own time-decay and volatility-of-volatility risk that needs to be taken into account by hedging options on the Vix.

1 Term structure of the implied volatility skew backed out from Vix call options by strike and option expiry

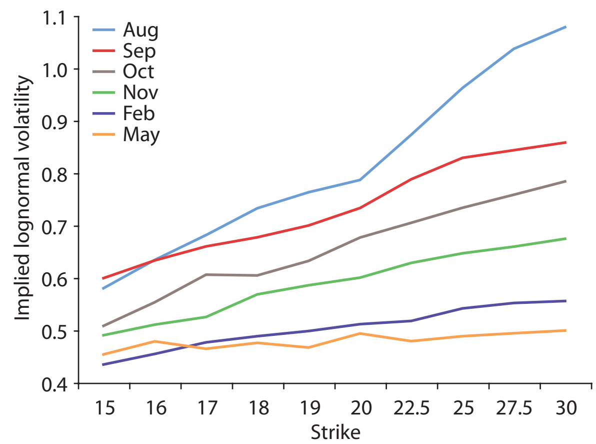

The dynamic model

To model the positive skew observed in implied volatilities of Vix options, we can follow two routes. First, we can assume that the volatility of the Vix dynamics is stochastic and positively correlated with the Vix dynamics. This is equivalent to introducing a model for the S&P 500 index dynamics with stochastic volatility and stochastic volatility of volatility that is positively correlated to the S&P 500 volatility dynamics and negatively correlated to the S&P 500 index dynamics. This model will produce the Vix skew by implying that low values of the S&P 500 index are followed by high values of its volatility and volatility of volatility. Ren, Madan & Qian Qian (2007) suggested a similar approach based on specifying the volatility-of-volatility parameter to be a local function of time, the S&P 500 spot value and its stochastic variance.

Second, we can introduce a jump process in the dynamics of the stochastic volatility of the S&P 500 index. This model will imply the Vix options skew by assigning higher probabilities to larger values of the Vix in the short term as a consequence of anticipating big jumps in the dynamics of the S&P 500 index volatility.

We choose the second option because, in our opinion, it is more financially justifiable and empirically observable, and, since this model results in a two-dimensional pricing problem, it can be handled by both analytical and numerical methods.

To solve the pricing problem for a variety of volatility products, we will consider the joint dynamics of the asset price $S(t)$, its variance $V(t)$ and its realised variance $I(t)$. To model the positive skew observed in the Vix options, we augment the variance dynamics with a jump process. The asset-price jumps, or more general Duffie, Pan & Singleton (2000) model, can easily be accommodated in our framework, but for the sake of brevity we omit price jumps by noting that jumps in the S&P 500 index alone cannot explain the variance skew observed in the Vix options. This is because, as can be seen from formula (10), the Vix futures dynamics is driven by the asset variance and not by the dynamics of its realised variance.

We finally note that, since hedging of Vix options is typically done with trading in Vix futures contracts, the pricing model needs to be consistent with the term structure of Vix futures. As a result, we assume that some of the model parameters are time-dependent in order to reproduce term-structure effects observed in market prices of Vix futures.

Specifically, we model the dynamics of these variables under the pricing measure $\mathbb{Q}$ using the square-root stochastic variance model:

$$\begin{cases} dS(t) = (r(t) - d(t))S(t)dt + \sigma(t)\sqrt{V(t)}S(t)dW^s(t), \\ \quad S(0) = S \\ dV(t) = \kappa(1 - V(t))dt + \varepsilon\sqrt{V(t)}dW^v(t) + JdN(t), \\ \quad V(0) = 1 \\ dI(t) = \sigma^2(t)V(t)dt, \\ \quad I(0) = I \end{cases} \tag{2}$$

where $r(t)$ and $d(t)$ are the deterministic risk-free interest and dividend rates, respectively, $\sigma(t)$ is the deterministic at-the-money volatility, κ is the mean-reverting rate, ε is volatility of volatility, $W^s(t)$ and $W^v(t)$ are correlated Wiener processes with constant correlation ρ, $N(t)$ is the Poisson process with intensity γ, and J is an exponentially distributed random jump with mean size η and probability density function:

$$\varpi(J) = \frac{1}{\eta}e^{-\frac{1}{\eta}J}$$

The dynamics for the asset realised variance $I(t)$ is derived by the augmentation principle developed by Lipton (2001).

Here, $\sigma(t)$ is the deterministic level of the at-the-money volatilities and $V(t)$ is the variance process scaled to unity. We assume that $\sigma(t)$ is piece-wise constant with local parameters chosen to match the term structure of Vix futures.

Now we introduce the expected values of the realised variance, $\bar{I}(t, T)$, at time T as follows:

$$\bar{I}(t,T) = \frac{1}{T-t}\mathbb{E}^{\mathbb{Q}}\left[\int_t^T \sigma^2(t')V(t')dt' \middle| \mathcal{F}(t)\right] \tag{3}$$

where $\mathcal{F}(t)$ is the information set available at time t, and at time $t = 0$ we obtain:

$$\bar{I}(0,T) = \frac{1}{T}\int_0^T \sigma^2(t')\left(1 + \frac{\gamma\eta}{\kappa}\left(1 - e^{-t'\kappa}\right)\right)dt' \tag{4}$$

Thus, given the values of mean-reversion parameters κ and jump parameters η and γ, we fit the term structure of $\sigma^2(t)$ to reproduce the term structure of the Vix futures. If, in addition, we want to fit the term structures of the at-the-money volatilities of SPX or Vix options, we introduce the term structure of parameters γ and η.

To build efficient analytics, we derive the joint transition density (Green) function, denoted by $G(t, T, X, X', V, V', I, I')$, of the joint evolution of the log of the S&P 500 index, $X(t) = \ln S(t)$, its variance, $V(t)$, and its realised variance, $I(t)$, and use it to price and calibrate vanilla options on the S&P 500 index, products on its implied and realised variance, along with contracts on the Vix.

We denote the Fourier transform of G with respect to X', V', and I' by $\hat{G}(t, T, X, \Phi, V, \Theta, I, \psi)$ with respective transform variables $\Phi = \Phi_R + i\Phi_I$, $\Theta = \Theta_R + i\Theta_I$ and $\psi = \psi_R + i\psi_I$ ($i = \sqrt{-1}$ and $\Phi_R, \Phi_I, \Theta_R, \Theta_I, \psi_R, \psi_I, \in \mathbb{R}$). Following the derivation outlined in Sepp (2007), we calculate the value function $U(t, X, V, I)$ of the contract with payout function $u(X, V, I)$ at maturity time T by inversion:

$$\begin{aligned} U(t,X,V,I) &= \frac{DF(t,T)}{8\pi^3}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \\ &\Re\left[\hat{G}(t,T,X,\Phi,V,\Theta,I,\Psi)\hat{u}(\Phi,\Theta,\Psi)\right]d\Phi_I d\Theta_I d\Psi_I \end{aligned} \tag{5}$$

where:

$$\hat{u}(\Phi,\Theta,\Psi) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{\Phi X' + \Theta V' + \Psi I'}u(X',V',I')dX'dV'dI' \tag{6}$$

and the solution to $\hat{G}$ is given by formula (17) in Appendix I.

Formula (5) along with (17) is one of our key results, which generalises the methodology proposed by Lipton (2002) for pricing equity options, and allows us to price European-style vanilla and volatility options jointly. In the one- (two-) dimensional case, formula (5) reduces to one-(two-) dimensional integrals. It is also possible to extend formula (5) to cover forward-start vanilla options, which are in fact volatility products, and forward-start volatility options. Numerical inversion of pricing formula (5) is achieved with standard fast Fourier transform and quadrature methods.

The dynamics of the Vix

We denote by $F(t, T)$ the value of the Vix futures at time t with settlement at time T and by $F(t)$ the spot value of the Vix, respectively. The square of $F(t)$ measures the expected annualised realised variance for options with maturity time Δ_T, while the square of $F(t, T)$ measures the expected annualised realised variance at future time T for options with maturity time $T + \Delta_T$, where Δ_T corresponds to the year fraction of 30 days ($\Delta_T = 30/365$). Accordingly, at future time $t = T$, $F(t, T)$ is the square root of the expected variance realised over time period $[T, T + \Delta_T]$:

$$\begin{aligned} F(T,T) &= \sqrt{\mathbb{E}^{\mathbb{Q}}\left[\frac{1}{\Delta_T}\int_T^{T+\Delta_T}\sigma^2(t')V(t')dt' \middle| \mathcal{F}(T)\right]} \\ &= \sqrt{\bar{I}(T, T+\Delta_T)} \end{aligned} \tag{7}$$

Now, using stochastic differential equation (2), we obtain:

$$\bar{I}(T, T+\Delta_T) = m_1(T) + m_2(T)V(T) \tag{8}$$

where:

$$m_1(T) = \frac{1 + \frac{\gamma\eta}{\kappa}}{\Delta_T}\int_T^{T+\Delta_T}\sigma^2(t')\left(1 - e^{-(T+\Delta_T-t')\kappa}\right)dt'$$

$$m_2(T) = \frac{1}{\Delta_T}\int_T^{T+\Delta_T}\sigma^2(t')e^{-(T+\Delta_T-t')\kappa}dt'$$

Accordingly, the Vix spot values $F(t)$ can be represented as a function of the variance at time t:

$$F(t) = \sqrt{m_1(t) + m_2(t)V(t)} \tag{9}$$

while the Vix futures $F(t, T)$ with settlement time T can be represented as time-t expectation of the variance at time T:

$$F(t,T) = \mathbb{E}^{\mathbb{Q}}\left[\sqrt{m_1(T) + m_2(T)V(T)} \middle| \mathcal{F}(t)\right] \tag{10}$$

From equations (9) and (10), we see that a derivative on Vix is essentially a bet on the future implied variance. Invoking Itô's lemma for equation (9), we can show that the dynamics of $F(t)$ is mean-reverting with reversion speed $\kappa/2$ and a stochastic mean level being the function of $F(t)$.

Pricing derivatives on the Vix

Now we consider the problem of pricing Vix futures and call options by using our general pricing formula (5). For the value function $U^F(t, T, V)$ of the Vix futures, the payout function to be used in formula (6) is:

$$u^F(X,V,I) = \sqrt{m_1(T) + m_2(T)V(T)}$$

Performing integration in (6), we obtain:

$$\hat{u}^F(\Phi,\Theta,\Psi) = \frac{2\pi^2\sqrt{\pi}e^{-\frac{m_1(T)\Theta}{m_2}}}{m_2(T)\left(-\frac{\Theta}{m_2(T)}\right)^{3/2}}\delta_0(\Psi)\delta_0(\Phi) \qquad (11)$$

where $\Theta_R < 0$ and $\delta_a(C)$ is the delta function of the complex-valued argument defined by formula (27) in Sepp (2007).

Simplifying formula (5), we obtain (the discount is not used for Vix futures):

$$U^F(t,T,V) = \frac{1}{\pi}\int_0^\infty \Re\left[\hat{G}^V(t,T,V,\Theta)\hat{\mathcal{P}}(\Theta)\right]d\Theta_I \qquad (12)$$

where $\hat{G}^V$ is the Fourier transform of the marginal density of the variance:

$$\hat{G}^V(t,T,V,\Theta) = \hat{G}(t,T,\cdot,0,V,\Theta,\cdot,0) \qquad (13)$$

and:

$$\hat{\mathcal{P}}(\Theta) = \frac{\sqrt{\pi}e^{-\frac{m_1(T)\Theta}{m_2(T)}}}{2m_2(T)\left(-\frac{\Theta}{m_2(T)}\right)^{3/2}}$$

The payout functions of the Vix call option are given respectively by:

$$u^C(X,V,I) = \max\left\{\sqrt{m_1(T) + m_2(T)V(T)} - K, 0\right\}$$

where K is the strike measured per Vix point.

Calculating transform (6) and simplifying formula (5), we obtain the value of the Vix call by formula (12) discounted with factor $DF(t, T)$ and:

$$\hat{\mathcal{P}}(\Theta) = \frac{\sqrt{\pi}\left(1 - erf\left(K\sqrt{-\frac{\Theta}{m_2(T)}}\right)\right)e^{-\frac{m_1(T)\Theta}{m_2(T)}}}{2m_2(T)\left(-\frac{\Theta}{m_2(T)}\right)^{3/2}}$$

where $\Theta_R < 0$ and:

$$erf(z) = \frac{2}{\sqrt{\pi}}\int_0^z e^{-s^2}ds$$

is the error function of the complex-valued argument, which can be evaluated using series representation (7.1.29) in Abramowitz & Stegun (1972).

Illustration

To illustrate our model, we calibrate it to the Vix options data observed on July 25, 2007. By changing the term structure of the piece-wise at-the-money volatility $\sigma(t)$, we ensure that the Vix futures prices are reproduced by the model and other model parameters are obtained by the global fitting to the call option prices. The estimates of these parameters are as follows: $\kappa = 2.26$, $\varepsilon = 1.66$, $\eta = 2.54$, $\gamma = 0.31$, the time-average of the $\sigma(t)$ is 0.18. The calibrated model adequately fits call option prices within the bid-ask spreads.

2 Model implied Vix density at option expiry month

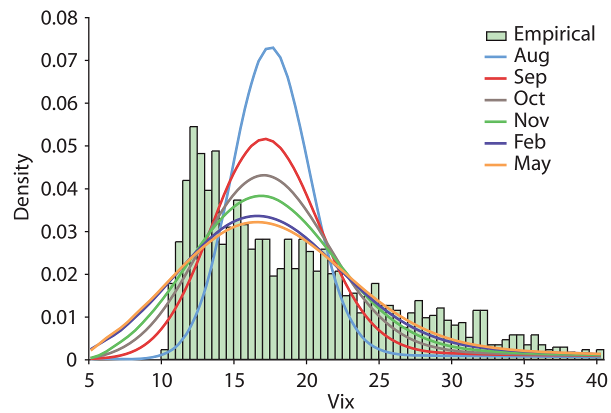

We see that the model implies a big jump (almost 100%) in the Vix with high probability. In figure 2, we show the model implied density for the Vix and the empirical frequency of the Vix calculated from Vix closing levels observed from July 25, 2000 to July 24, 2007 (the CBOE provides historical records for the Vix, calculated by the current approach, dating back to 1986). We note that the right tail of the empirical Vix density is based on Vix values that occurred before 2003, when the new methodology for the Vix calculation was implemented and the implied volatility dropped from its historical highs in 2001 and 2002. However, the model implied density covers the high values of the Vix experienced in the second half of 2007.

Hedging

We now consider some hedging strategies implied by this calibrated model. In general, the hedging of a Vix futures contract can be done by trading in Vix futures contracts with different maturities, while a Vix option can be hedged by trading in Vix futures contracts. The corresponding hedge ratios can be calculated by applying implicit differentiation to pricing formula (12). For example, the futures delta of the call option, $U^{call}(t, T_2, V)$, with respect to the Vix futures, $F(t, T_1, V)$, is calculated by:

$$\frac{\partial U^{call}(t,T_2,V)}{\partial F(t,T_1,V)} = \frac{\partial U^{call}(t,T_2,V)}{\partial V}\left(\frac{\partial F(t,T_1,V)}{\partial V}\right)^{-1} \qquad (14)$$

In figure 3 (left-hand graph), we examine the model implied futures deltas and gammas for: 1) the strategy involving hedging the futures with expiry in December ($T = 0.4$) by trading in Vix futures with the shortest maturity ($T = 0.08$) (denoted by FuturesDelta1 and FuturesGamma1, respectively); 2) the hedging strategy for the call option with strike $K = 19$ and expiry in December ($T = 0.4$) by trading in Vix futures with the shortest maturity (CallDelta1 and CallGamma1); and 3) the hedging strategy for the call struck at $K = 19$ with both the futures and the call expiring in December (CallDelta2 and CallGamma2). We see that for high values of the Vix, the futures delta in the second strategy converges to the same delta as in the first strategy because the underlying for the Vix option is essentially the Vix forward value.

3 Futures delta and futures gamma as functions of the Vix value (left) and volatility of volatility risk as function of the Vix value (right)

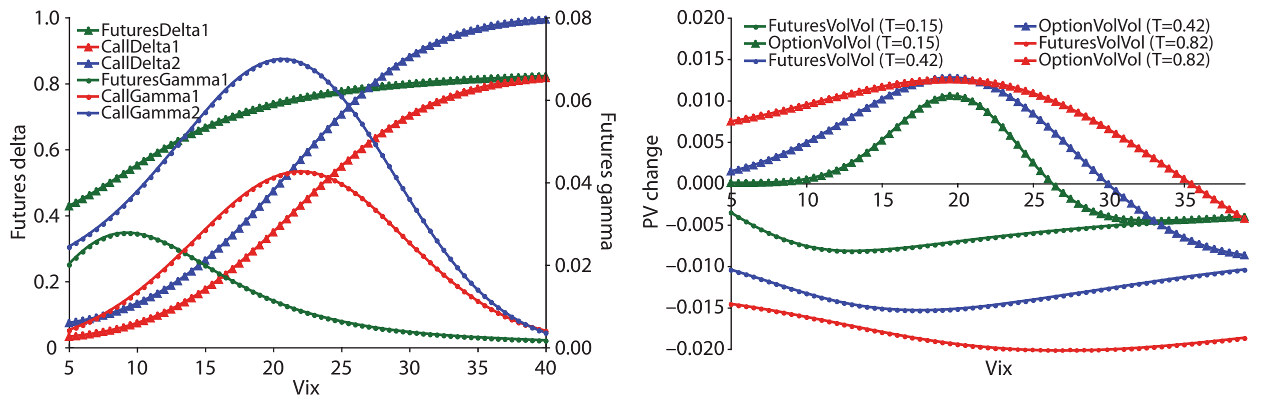

4 Profit and loss histogram under the variance delta-hedging strategy (left) and under the variance delta-jump-hedging strategy (right)

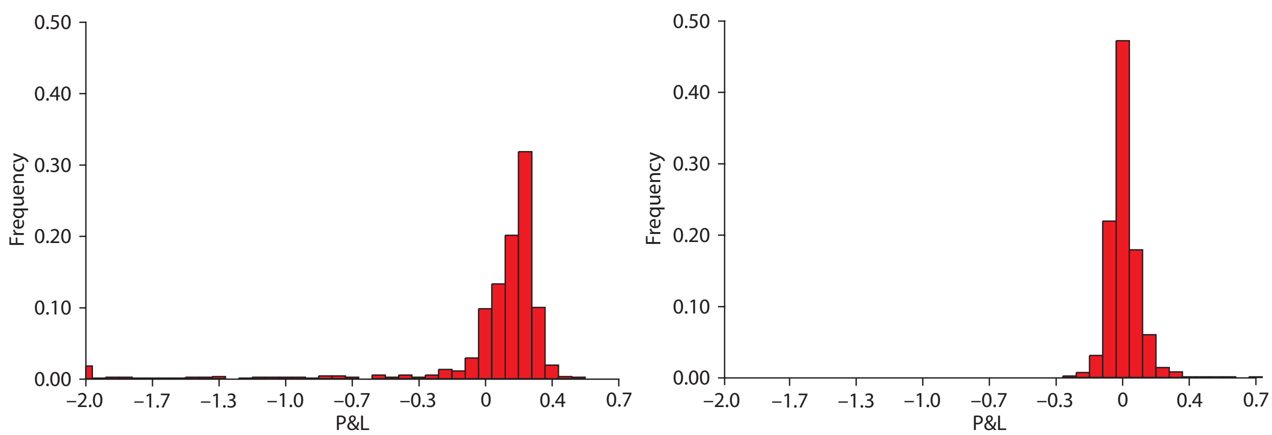

In figure 3 (right-hand graph), we show the volatility-of-volatility risk, obtained by increasing the value of the volatility-of-volatility parameter ε by 1%, for Vix futures and call option contracts with different maturities. We see that the Vix futures contract has negative volatility-of-volatility risk (because the square root is a concave function). The call option can be thought of as a composite function of the Vix (as the underlying) and the Vix futures (as the limiting call value for large values of the Vix), so that the positive call convexity to the Vix results in the positive volatility-of-volatility risk. For large values of the Vix, the positive call delta with respect to the Vix futures and the negative volatility-of-volatility risk of the latest results in the negative volatility-of-volatility risk of the call option.

We see that, in general, the risk measures of Vix futures and options behave differently to those of vanilla options. It is important to note that the standard delta-gamma hedging is not appropriate for Vix options hedging because of the infrequent and large jumps observed in Vix dynamics. There are a few potential ways to hedge jumps, and we follow the one suggested by Andersen & Andreasen (2000), which is based on elimination of the expected jump impact. If we hedge the Vix call, $U^{call}(t, T, V)$, with two futures contracts, $F(t, T_1, V)$ and $F(t, T_2, V)$, where $T_1 \neq T_2$, then we calculate the expected jump risk of the hedging position as follows:

$$\bar{\Pi}(t) = \Delta_1 \bar{F}(t,T_1,V) + \Delta_2 \bar{F}(t,T_2,V) - \bar{U}^{call}(t,T,V) \quad (15)$$

with:

$$\bar{F}(t,T_1,V) = \int_0^\infty \left(F(t,T_1,V+J) - F(t,T_1,V)\right)\varpi(J)\,dJ$$
$$= \frac{1}{\pi}\int_0^\infty \Re\left[\left(\frac{1}{1-\eta B(t,T_1)} - 1\right)\hat{G}^V(t,T_1,V,\Theta)\hat{\mathcal{P}}(\Theta)\right]d\Theta_I \quad (16)$$

where to explicitly compute $\bar{F}(t, T_1, V)$ we use formula (12) along with (17). Hedge ratios Δ_1 and Δ_2 are calculated by setting (15) and the variance delta of the hedging position $\Pi(t)$ to zero.

To test the delta-jump-hedging strategy in our model, we apply Monte Carlo simulation of our calibrated model to hedge the short position in the Vix call with maturity $T = 0.4$ and strike $K = 19$ by using two strategies: 1) the variance delta-hedging by trading in Vix futures with the same expiry date; and 2) the variance delta and jump risk hedging by trading in the two futures contracts with maturities $T = 0.4$ and $T = 0.5$. For both strategies

A. Statistics of the variance delta-hedged portfolio (Hedge 1) and the variance delta-jump-hedged portfolio (Hedge 2)

	Hedge 1	Hedge 2
Minimum	–10.3492	–0.3037
Maximum	0.4838	1.7720
Average	–0.0137	0.0018
Standard deviation	0.7470	0.1048
Skew	–7.6472	7.2754
Kurtosis	72.8494	102.6192

we use the same set of random numbers and generate 1,000 scenarios. Within each scenario, we simulate the Vix, applying formula (9) at each business day during the option life and rehedge the position daily with the number of rebalancing trades totalling 100. We normalise the final outcome of each scenario by the initial option price ($2.4311) and, for brevity, we ignore transaction and financing costs. In table A, we show the statistics of both strategies and in figure 4 we show the histograms of the final profit-and-loss distribution.

We see that if the writer of the Vix call option follows the plain variance delta-hedging strategy, they are short variance gamma, which means that their frequent and small gains are compensated by infrequent but rather huge losses when the Vix jumps. In contrast, if the hedger follows the delta-jump-hedging strategy, they are practically both variance delta- and gamma-neutral. As a result, the profit-and-loss distribution under the delta-jump-hedging strategy peaks at zero with little variation.

Conclusions

We have developed a dynamic model for the joint evolution of the Vix spot value and the S&P 500 index that can be made consistent with both the market prices of Vix futures and options, and options on the S&P 500 index. Utilising the generalised Fourier transform, we have obtained closed-form solutions for values of volatility products, including futures and options on the Vix. Finally, we have examined some hedging strategies for this model. ■

Appendix I: the transformed Green function

The Fourier transform of the Green function corresponding to stochastic differential equation (2) is given by:

$$\hat{G}(t,T,X,\Phi,V,\Theta,I,\Psi) = e^{-\Phi Y - I\Psi + A(t,T) + B(t,T) + \Gamma(t,T)} \quad (17)$$

where $Y = X + \int_t^T (r(t') - d(t'))dt'$ and the functions $A(t, T)$, $B(t, T)$ and $\Gamma(t, T)$ are defined by:

$$A(t,T) = \sum_{n=1}^{N} A_n, \quad B(t,T) = B_1, \quad \Gamma(t,T) = \sum_{n=1}^{N} \Gamma_n$$

with $N = \min\{n : T_n^J \geq T, n = 1, \ldots, N^J\}$ and B_n, A_n, Γ_n computed by recursion:

$$B_n = B(T_{n-1}, T_n, \bar{B}_n), \quad n = N, \ldots, 1$$

$$\bar{B}_N = -\Theta, \quad \bar{B}_n = B_{n+1}, \quad n = N-1, \ldots, 1$$

$$A_n = A(T_{n-1}, T_n, \bar{B}_n), \quad \Gamma_n = \Gamma(T_{n-1}, T_n, \bar{B}_n)$$

$$n = N, \ldots, 1, \quad T_0^J = t, \ldots, T_N^J = T$$

$$B(t,T,\bar{B}) = -\frac{-\psi_- C_+ e^{-\zeta\tau} + \psi_+ C_-}{\varepsilon^2 (C_+ e^{-\zeta\tau} + C_-)}$$

$$A(t,T,\bar{B}) = -\frac{2\kappa}{\varepsilon^2}\left(\frac{1}{2}\psi_+ \tau + \ln(C_+ e^{-\zeta\tau} + C_-)\right)$$

$$\Gamma(t,T,\bar{B}) = -\frac{2\varepsilon^2\gamma}{\bar{\psi}_+ \bar{\psi}_-}\left(\eta \ln\left(\frac{\bar{\psi}_- C_+ e^{-\zeta\tau} + \bar{\psi}_+ C_-}{\bar{\psi}_- C_+ + \bar{\psi}_+ C_-}\right) - \frac{1}{2}\tau\bar{\psi}_-\right) - \tau\gamma$$

$$C_\pm = \frac{1}{\zeta}\left(\frac{\psi_\pm}{2} \pm \frac{\varepsilon^2}{2}\bar{B}\right)$$

$$b = -\kappa - \rho\varepsilon\sigma(T)\Phi, \quad \zeta = \sqrt{b^2 - \varepsilon^2\sigma^2(T)(\Phi^2 + \Phi - 2\Psi)}$$

$$\psi_\pm = \pm b + \zeta, \quad \bar{\psi}_\pm = \varepsilon^2 \pm \eta\psi_\pm, \quad \tau = T - t$$

Artur Sepp is an associate at Merrill Lynch in New York. He is grateful to his colleagues and two anonymous referees for helpful and instructive comments. The views expressed in this article are those of the author alone.
E-mail: artur_sepp@ml.com

References

Abramowitz M and I Stegun, 1972
Handbook of mathematical functions
Dover, New York

Andersen L and J Andreasen, 2000
Jump-diffusion processes: volatility smile fitting and numerical methods for option pricing
Review of Derivatives Research 4, pages 231–262

Black F, 1976
The pricing of commodity contracts
Journal of Financial Economics 3, pages 167–179

Dotsis G, D Psychoyios and G Skiadopoulos, 2007
An empirical comparison of continuous-time models of implied volatility indices
Journal of Banking and Finance 31(12), pages 3,584–3,603

Duffie D, J Pan and K Singleton, 2000
Transform analysis and asset pricing for affine jump-diffusion
Econometrica 68(6), pages 1,343–1,376

Grünbichler A and F Longstaff, 1996
Valuing futures and options on volatility
Journal of Banking and Finance 20, pages 985–1,001

Heston S, 1993
A closed-form solution for options with stochastic volatility with applications to bond and currency options
Review of Financial Studies 6, pages 327–343

Lipton A, 2001
Mathematical methods for foreign exchange: a financial engineer's approach
World Scientific, Singapore

Lipton A, 2002
The vol smile problem
Risk February, pages 81–85

Psychoyios D and G Skiadopoulos, 2006
Volatility options: hedging effectiveness, pricing, and model error
Journal of Futures Markets 26, pages 1–31

Ren Y, D Madan and M Qian Qian, 2007
Calibrating and pricing with embedded local volatility models
Risk September, pages 138–143

Sepp A, 2007
Variance swaps under no conditions
Risk March, pages 82–87

Whaley R, 1993
Derivatives on market volatility: hedging tools long overdue
Journal of Derivatives 1, pages 71–84

Zhang J and Y Zhu, 2006
Vix futures
Journal of Futures Markets 26, pages 521–531

Zhu Y and J Zhang, 2007
Variance term structure and Vix futures pricing
International Journal of Theoretical and Applied Finance 10(1), pages 1–18