# Realised volatility and variance: options via swaps

*Peter Carr and Roger Lee present explicit and readily applicable formulas for valuing options on realised variance and volatility. They use variance and volatility swaps – or alternatively vanilla options – as pricing benchmarks and hedging instruments. They also cover Vix options*

In this article, we develop strategies for pricing and hedging options on realised variance and volatility. Our strategies have the following features.

- *Readily available inputs.* We can use vanilla options as pricing benchmarks and as hedging instruments. If variance or volatility swaps are available, then we use them as well. We do not need other inputs (such as parameters of the instantaneous volatility dynamics).
- *Comprehensive and readily computable outputs.* We derive explicit and readily computable formulas for prices and hedge ratios for variance and volatility options, applicable at all times in the term of the option (not just inception).
- *Accuracy and robustness.* We test our pricing and hedging strategies under skew-generating volatility dynamics. Our discrete hedging simulations at a one-year horizon show mean absolute hedging errors under 10%, and in some cases under 5%.
- Easy modification to price and hedge options on implied volatility (Vix).

Specifically, we price and hedge realised variance and volatility options using variance and volatility swaps. When necessary, we in turn synthesise volatility swaps from vanilla options using the Carr & Lee (2006) methodology and variance swaps from vanilla options using the standard log-contract methodology.

### Notation

Let $S_t$ denote the value of a stock or stock index at time $t$. Given a variance/volatility option to be priced and hedged, let us designate as time zero the start of its averaging period and time $T$ the end. For $t \in [0, T]$ and $\tau \le t$, let $Q_{\tau,t}$ denote the realised variance of returns over the time interval $[\tau, t]$.

The mathematical results about the synthesis of volatility and variance swaps will hold exactly if $Q$ refers to the continuously monitored variance. This means the quadratic variation of $\log S$, times a constant conversion factor $u^2$ that includes any desired annualisation and/or rescaling:

$$u^2 \lim \sum_{\tau < t_n \le t} \left( \log \frac{S_{t_n}}{S_{t_{n-1}}} \right)^2 \tag{1}$$

where lim denotes limit in probability as the mesh of the partition $\{t_1 < t_2 < \ldots\}$ tends to zero. For example, choosing $u = 100 \times \sqrt{1/T}$ expresses $Q_{0,T}$ in basis points per unit time.

For practical implementation with daily monitoring, however, let $N$ be the number of days in the period $[0, T]$ between the daily closing times zero and $T$, and let:

$$Q_{\tau,t} := u^2 \sum_{\tau < t_n \le t} \left( \log \frac{S_{t_n}}{S_{t_{n-1}}} \right)^2 \tag{2}$$

where $t_0 := \tau$ and the $t_1 < t_2 < \ldots$ are the successive daily closing times in $(\tau, t]$, together with $t$ itself, and $u$ is a constant annualisation/rescaling factor. For example, choosing:

$$u := 100 \times \sqrt{252 / N} \tag{3}$$

expresses $Q_{0,T}$ in units of annual basis points.

The $t$ will denote the valuation date. We allow arbitrary $t \in [0, T]$, because we will solve for prices and hedges not just at inception, but throughout the term of the option. The $\tau$ will denote the start date of the variance and volatility swaps that will serve as pricing benchmarks and hedging instruments; by not constraining $\tau = 0$, we have the freedom to use swaps whose averaging periods $[\tau, T]$ do not necessarily coincide with the option's period $[0, T]$.

For a given variance or volatility option, the $u$ is fixed. It may depend on $T$, but $T$ is fixed. It does not depend on $t$. Thus the $Q_{\tau,t}$ is a scale factor times a running 'cumulative' variance – not a running 'average' variance, because the scale factor is designed to give a proper average only for the full interval $[0, T]$.

Now consider swaps and options written on $Q$ and $Q^{1/2}$. A genuine $[\tau, T]$ variance swap with fixed leg $f^2$ pays the holder some notional amount times:

$$Q_{\tau,T} - f^2 \tag{4}$$

A genuine $[\tau, T]$ volatility swap with fixed leg $f$ pays the holder some notional amount times:

$$Q_{\tau,T}^{1/2} - f \tag{5}$$

In this article, all swaps have a notional of one and a fixed leg of zero, unless otherwise stated.

Our options treatment focuses on calls; results for puts can be obtained from parity relations. A $[0, T]$ variance call with strike $K_{vol}^2$ pays the holder at time $T$ some notional amount times:

$$\left(Q_{0,T} - K_{vol}^2\right)^+ \quad (6)$$

A $[0, T]$ volatility call with strike $K_{vol}$ pays the holder at time $T$ some notional amount times:

$$\left(Q_{0,T}^{1/2} - K_{vol}\right)^+ \quad (7)$$

In this article, all options have a notional of one, unless otherwise stated. The strike $K_{vol}$ may be any non-negative number. We say that a $[0, T]$ variance option struck at $K_{vol}^2$ is at-the-money at time $t$ if a $[0, T]$ variance swap (on the same underlier) with fixed leg $K_{vol}^2$ has time-$t$ value zero.

This article proposes a methodology to price and hedge options on realised variance and volatility. Specifically, we price variance options using an explicit formula that takes as inputs the prices of variance and volatility swaps. We hedge variance options by trading variance and volatility swaps. We do likewise for volatility options.

If variance and volatility swaps are unavailable to trade, then we propose to synthesise them using vanilla options. So we begin with swaps, and build towards options.

### Variance and volatility swaps

Variance option prices depend on the expectation and volatility of variance. The expectation is revealed by variance swap prices, and the volatility can be inferred from variance and volatility swap prices together. Specifically:

- Let $A_t$ be the time-$t$ value of the variance swap that pays $Q_{0,T}$
- Let $B_t$ be the time-$t$ value of the volatility swap that pays $Q_{0,T}^{1/2}$
- Let $r$ be the assumed constant interest rate, and let $A_t^* := A_t e^{r(T-t)}$ and $B_t^* := B_t e^{r(T-t)}$ be the time-$t$ variance swap rate and volatility swap rate respectively; by definition this is the unique fixed-leg payment that would make the variance swap (respectively volatility swap) have time-$t$ value zero.

As shown in figure 1, the volatility swap's concave square-root payout is dominated by the linear payout consisting of $\sqrt{A_t^*}$ in cash, plus $1/(2\sqrt{A_t^*})$ variance swaps with fixed leg $A_t^*$. The dominating payout has forward value $\sqrt{A_t^*}$, because the variance swaps have value zero. Thus we have enforced Jensen's inequality $\sqrt{A_t^*} \geq B_t^*$ by super-replication.

This concavity's price impact – as measured by how much $\sqrt{A_t^*}$ exceeds $B_t^*$ depends on the volatility of volatility. More precisely, letting $\mathbb{E}$ denote the risk-neutral expectation:

$$A_t^* - \left(B_t^*\right)^2 = \mathbb{E}_t Q_{0,T} - \left(\mathbb{E}_t Q_{0,T}^{1/2}\right)^2 = Var_t Q_{0,T}^{1/2} \quad (8)$$

So if we can obtain the swap values $A$ and $B$, we can back out the volatility of volatility, and use it to price options on $Q_{0,T}$ and $Q_{0,T}^{1/2}$. If we can moreover trade the variance and volatility swaps, we can hedge the volatility-of-volatility risk. Similar reasoning holds for volatility options.

In the absence of genuine variance or volatility swaps, we obtain $A$ and $B$ from synthetic swaps.

**Synthetic variance swap.** Using the theory developed in Neuberger (1990), Dupire (1992), Carr & Madan (1998) and Derman *et al* (1999), who assume essentially only the positivity and continuity of price paths, the following self-financing trading strategy replicates the continuously monitored $Q_{\tau,t}$ for a non-dividend-paying asset.

Write $F_t := S_t e^{r(T-t)}$ for the forward price, and choose an arbitrary put-call separator $\kappa > 0$.

At each time $t \geq \tau$, hold the following static position in options, and dynamic position in shares:

$$\begin{aligned}
&u^2 \frac{2}{K^2} dK \ \text{ calls at strikes } K > \kappa, \text{ puts at strikes } K < \kappa \\
&u^2 \left(\frac{2}{F_t} - \frac{2}{\kappa}\right) \ \text{ shares} \\
&e^{-r(T-t)}\left[Q_{\tau,t} + 2u^2 \log\left(F_t / \kappa\right)\right] \ \text{ cash}
\end{aligned} \quad (9)$$

where all options have expiry $T$. We call this portfolio a synthetic variance swap. Its initial payout profile appears in figure 2. With continuous trading and a continuum of strikes, the final portfolio value will match the continuously monitored variance.

**Volatility swap: valuation under independent volatility.** If one desires only to know the $[\tau, T]$ volatility swap's initial $t = \tau$ value (not the full replicating strategy), and if one assumes that instantaneous volatility evolves independently of the risk that drives price moves, then the at-the-money-forward implied volatility $\sigma_{imp}(F_t)$ approximates the desired volatility swap rate $B_t^*$, as the following argument shows.

Brenner & Subrahmanyam (1988) found, by a Taylor expansion of the normal cumulative distribution function about zero, that the at-the-money-forward Black-Scholes formula with vola-

1 Payout comparison

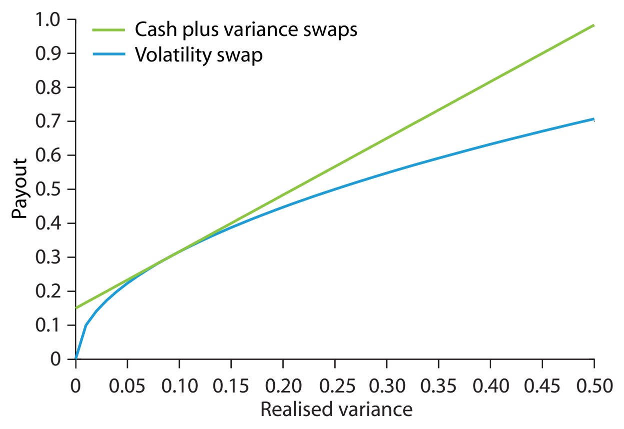


2 Initial payout profile of a synthetic variance swap

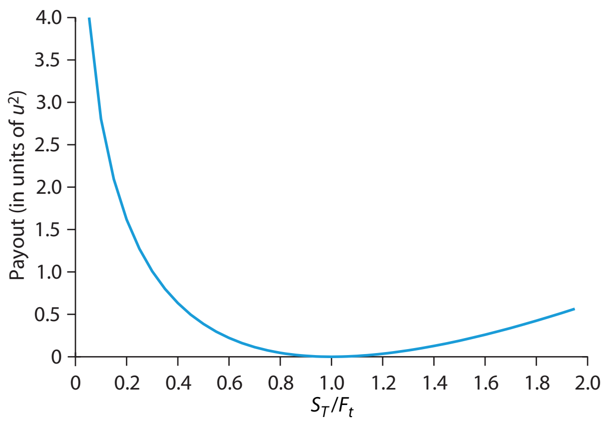

3 Initial payout profile of a synthetic volatility swap: $Q_{\tau,t} = 0$

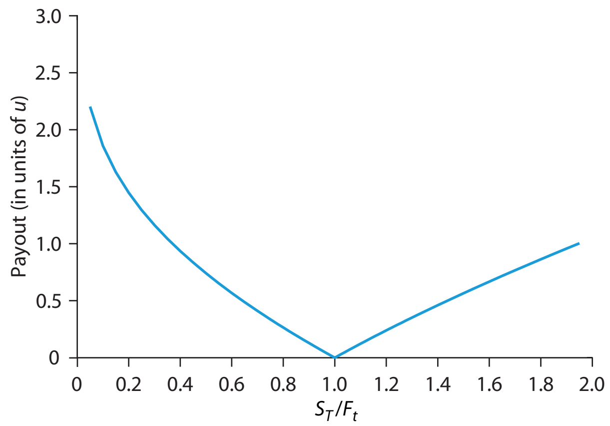


4 Volatility swap values: genuine, synthetic and the at-the-money forward implied approximation

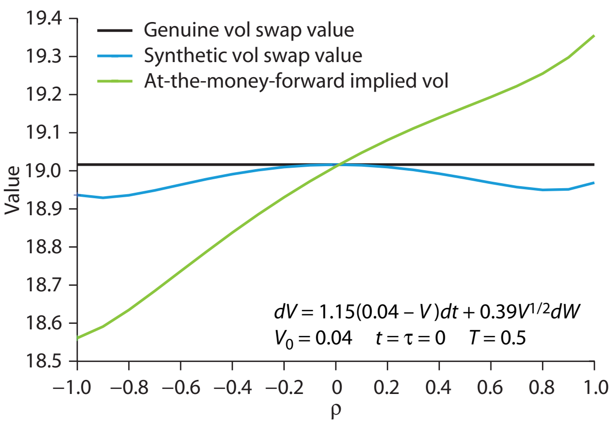


tility parameter $\sigma$ satisfies:

$$C^{BS}(\sigma) \approx \frac{S_t \sigma \sqrt{T-t}}{\sqrt{2\pi}} \tag{10}$$

Applying this twice:

$$\begin{aligned} \frac{S_t \sigma_{imp}(F_t)\sqrt{T-t}}{\sqrt{2\pi}} &\approx C^{BS}\left(\sigma_{imp}(F_t)\right) \\ = \mathbb{E}_t C^{BS}\left(Q_{t,T}^{1/2}\right) &\approx \mathbb{E}_t \frac{S_t Q_{t,T}^{1/2}\sqrt{T-t}}{\sqrt{2\pi}} = \frac{S_t B_t^* \sqrt{T-t}}{\sqrt{2\pi}} \end{aligned} \tag{11}$$

where the first equality holds due to the independence assumption. Therefore:

$$B_t^* \approx \sigma_{imp}(F_t) \tag{12}$$

as Feinstein (1989) observed.

However, this estimate does not establish a replication strategy, does not apply at times $t > \tau$ after inception, and does not suggest how to handle the important case of correlated volatility. Our approach below addresses all these issues.

■ **Synthetic volatility swap: the Carr-Lee approach.** The conventional wisdom holds that the pricing and hedging of a volatility swap is, unlike variance swaps, highly model-dependent.

Carr & Lee's 2006 paper challenges this notion. Without imposing a model on the dynamics of volatility, it shows how to replicate volatility swaps by trading vanilla options. It begins by making the same independence assumption as in the previous section, but then relaxes that assumption to produce correlation-robust strategies. It also assumes frictionless trading in vanilla options, ignoring transactions costs, but these may be mitigated by the fact that only the net exposures in a portfolio need to be hedged via trades.

To specify the correlation-robust replication strategy, let $I_\nu$ denote the modified Bessel function of order $\nu$ – for which numerical implementations are readily available (for example, see Matlab's `besseli` or Mathematica's `BesselI`).

Initially, starting at time $t = \tau$, when $Q_{\tau,t} = 0$, the replicating portfolio holds:

$u \times \sqrt{\pi/2}\,/\,F_\tau$ straddles at strike $K = F_\tau$

$$u \times \sqrt{\frac{\pi}{8K^3F_\tau}}\left[I_1\left(\log\sqrt{K/F_\tau}\right) - I_0\left(\log\sqrt{K/F_\tau}\right)\right]dK$$

calls at strikes $K > F_\tau$ (13)

$$u \times \sqrt{\frac{\pi}{8K^3F_\tau}}\left[I_0\left(\log\sqrt{K/F_\tau}\right) - I_1\left(\log\sqrt{K/F_\tau}\right)\right]dK$$

puts at strikes $K < F_\tau$

together with a zero-cost delta-hedge. Afterwards, at times $t > \tau$ when $Q_{\tau,t} > 0$, it holds:

$$u \times \frac{dK}{\sqrt{\pi}}\int_0^\infty \frac{e^{-zQ_{\tau,t}/u^2}}{K^2 z^{1/2}}\left[\theta_+ (K/F_t)^{p_+} + \theta_- (K/F_t)^{p_-}\right]dz$$

calls at strikes $K > F_t$ or puts at strikes $K < F_t$ and (14)

$e^{-r(T-t)}Q_{\tau,t}^{1/2}$ cash

together with a zero-cost delta-hedge, where:

$$\theta_\pm := \frac{1}{2} \mp \frac{1}{2\sqrt{1-8z}} \qquad p_\pm := \frac{1}{2} \pm \sqrt{1/4 - 2z} \tag{15}$$

and all options have expiry $T$. The portfolio holdings at strike $K$ depend only on the observables $K$, $F$, $Q_{\tau,t}$ and $u$.

Under the independence assumption, with continuous trading in a continuum of strikes, we show in Carr & Lee (2006) that the portfolio self-finances and has time-$T$ value equal to the continuously monitored volatility swap payout $Q_{\tau,T}^{1/2}$. We call this portfolio a synthetic volatility swap.

Figure 3 plots the initial payout profile of the synthetic volatility swap, and the appendix of Carr & Lee (2007) implements this strategy using a discrete set of strikes, and gives a numerical example.

Most of the synthetic volatility swap's value resides in the at-the-money-forward straddles in (13). The short out-of-the-money call positions and long out-of-the-money put positions in (13) are precisely chosen to gain robustness to violations of the independence condition, by neutralising the first-order effect of price/volatility correlation on the synthetic swap's value. This robustness to correlation is important in typical equity markets, where downward-sloping volatility skews indicate the presence of negative correlation.

Let us test how accurately the synthetic volatility swap value matches the genuine volatility swap value, under Heston (1993) dynamics, with various correlation parameters $\rho \in [-1, 1]$. For comparison, we also include the at-the-money Black-Scholes implied volatility, as motivated by (12). Figure 4 shows the results for $t = \tau = 0$ and $T = 0.5$ and $r = 0$, using parameters estimated by Bakshi, Cao & Chen (1997), and $u = 100\sqrt{1/T}$. We calculate the genuine volatility swap value using the identity $2\sqrt{\pi}\mathbb{E}Q^{1/2}_{0,T} = \int_0^\infty (1 - \mathbb{E}e^{-zQ_{0,T}})z^{-3/2}dz$ and the known Laplace transform of $Q_{0,T}$.

Our synthetic volatility swap clearly outperforms implied volatility as an approximation of the genuine volatility swap value, across essentially all correlation assumptions. In the case $\rho = 0$, our method is (as promised) exact and the implied volatility approximation is nearly exact, but, more importantly, in the empirically relevant case of $\rho \neq 0$, our synthetic volatility swap's relative 'flatness' with respect to correlation results in its greater accuracy.

Taking $r \neq 0$ would change nothing, except to scale the genuine and synthetic volatility swap values by $e^{-rT}$. Taking $T$ smaller (larger) tends to improve (worsen) the accuracy of both approximations: our synthetic volatility swap/the naive at-the-money implied rule. For example, with $\rho = -0.64$, the respective –6/–30bp discrepancies at $T = 0.5$ shown in figure 4 would become –1/–13bp at $T = 0.25$, and –18/–57bp at $T = 1.0$. At each maturity our synthetic volatility swap still has clearly greater accuracy than the naive at-the-money implied volatility rule.

■ *Remark 1.* As variance accumulates during the life of the volatility swap, the hedge evolves. Expressions (13) and (14) make this precise, but here let us give some intuition. The initial hedge resembles an at-the-money straddle. The dynamics of the hedge depend on two factors. First, as the spot moves, the 'strike' of the 'straddle' floats to stay at-the-money. Second, as the running variance $Q_{\tau,t}$ accumulates, the 'straddle' smooths out, losing its kink; indeed, only when $Q_{\tau,t} = 0$ does the kink literally exist.

Eventually, the synthetic volatility swap evolves towards a position in cash plus synthetic variance swaps, as shown in figures 5 and 6. Intuitively, as variance accumulates, we progress rightwards in figure 1. In that direction, the square-root function loses convexity, and becomes more linear. Thus the cash-plus-variance-swaps portfolio becomes not merely an upper bound, but indeed improves as an approximation to the volatility swap.

### Variance and volatility options

■ **Pricing of variance and volatility options.** We have the following pricing problem. At time $t$, where $0 \leq t < T$, we observe the following:

- $Q_{0,t}$, the variance from time zero to $t$, where $0 \leq t < T$.
- $Q_{\tau,t}$, the variance from some time $\tau$ to $t$, where $\tau \leq t$.
- $A_t$, the time-$t$ price of a (genuine or synthetic) variance swap, which pays $Q_{\tau,T}$.
- $B_t$, the time-$t$ price of a (genuine or synthetic) volatility swap, which pays $Q^{1/2}_{\tau,T}$.
- $K_{vol}$, the strike (quoted as a volatility).
- $r$, the interest rate. Let $G_t := e^{r(T-t)}$ denote the associated discount factor's reciprocal.

We have allowed the reference swaps' start date $\tau$, the option's

5 Evolution of the synthetic volatility swap: $Q_{\tau,t} = 0.25u^2$

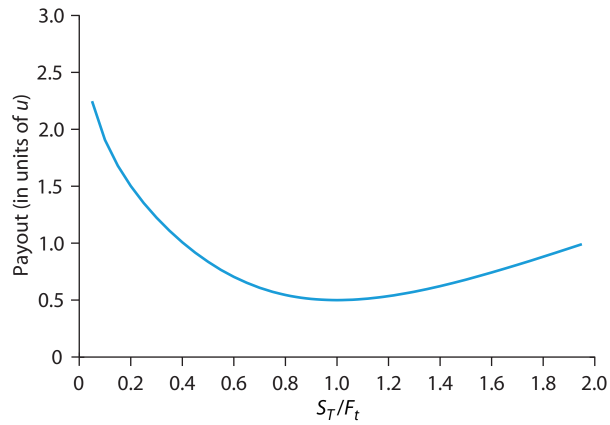


6 Evolution of the synthetic volatility swap: $Q_{\tau,t} = 1.0u^2$

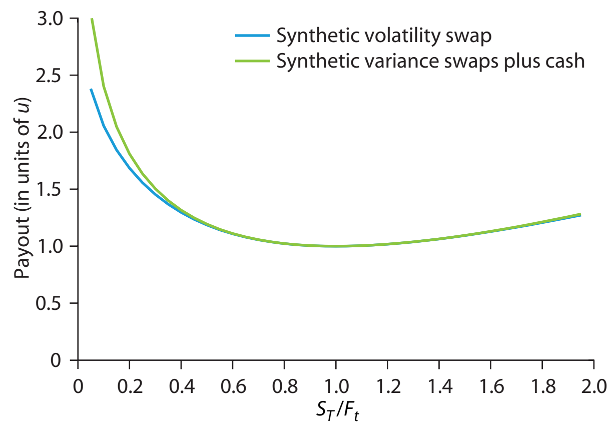


start date zero, and the valuation date $t$, to be distinct or identical. The condition $\tau \leq t$ allows either spot-starting or seasoned volatility swaps.

We intend to solve for:

- Variance call: the time-$t$ value of a claim on $(Q_{0,T} - K^2_{vol})^+$.
- Volatility call: the time-$t$ value of a claim on $(Q^{1/2}_{0,T} - K_{vol})^+$.

Our solution will approximate the time-$t$ conditional distribution of $Q^{1/2}_{\tau,T}$ as a displaced lognormal distribution. This extends the use of a lognormal distribution by Friz & Gatheral (2005), who assume that the valuation date, swap start date and option start date coincide. Moreover, they do not address hedging, and for estimation of the input $B$, they do not use *correlation-robust* volatility swaps.

Specifically, let us approximate as lognormal the time-$t$ conditional distribution of 'remaining volatility' $Q^{1/2}_{\tau,T} - Q^{1/2}_{\tau,t}$. The lognormal distribution has two parameters (mean and variance), which we calibrate to the given variance swap and volatility swap prices. Using the calibrated lognormal distribution, we solve for prices of volatility and variance options.

Thus we obtain the following explicit pricing formulas that, like the Black-Scholes formula, involve $N$, the normal cumulative density function.

The displaced lognormal price of a $K^2_{vol}$-strike variance call is:

7 Call price approximations, against strike

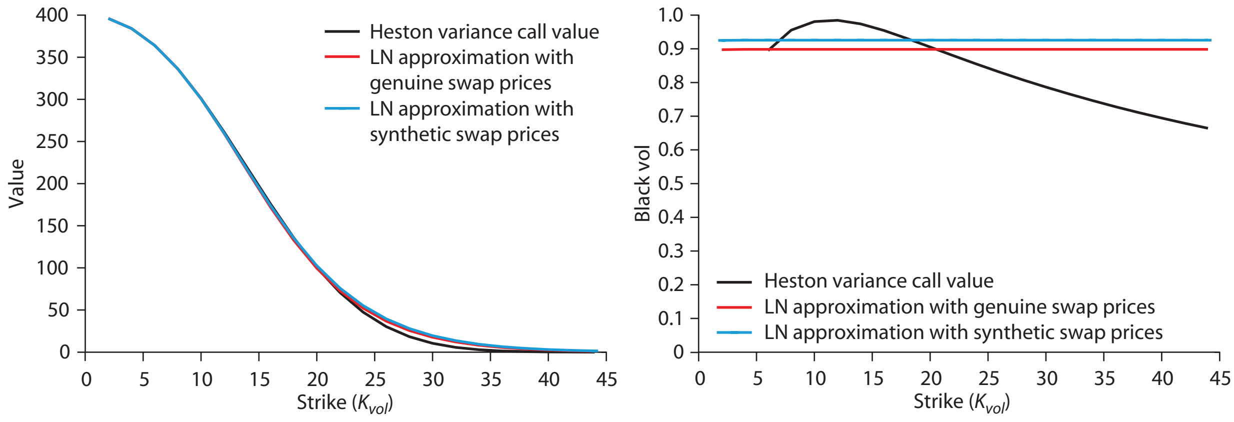


Note: correlation ρ is fixed at –0.64. The at-the-money strike (the variance swap rate, quoted as a volatility) is 20, and expiry is 0.5 years. In the left-hand plot, the vertical axis is the call's dollar value; in the right-hand plot, the vertical axis is the call's Black implied volatility. Heston dynamics are as in figure 4

$$C_{LN}^{var} = G_t^{-1} \times \begin{cases} A_t G_t + Q_{0,t} - Q_{\tau,t} - K_{vol}^2 & \text{if } K_{vol}^2 \le Q_{0,t} \\ \mu_2 N(d_0) + 2Q_{\tau,t}^{1/2} \mu_1 N(d_1) & \\ -\left(K_{vol}^2 - Q_{0,t}\right) N(d_2) & \text{if } K_{vol}^2 > Q_{0,t} \end{cases} \tag{16}$$

where:

$$\mu_1 := B_t G_t - Q_{\tau,t}^{1/2} \tag{17}$$

$$\mu_2 := A_t G_t + Q_{\tau,t} - 2B_t G_t Q_{\tau,t}^{1/2} \tag{18}$$

and:

$$d_j := \frac{m_t - \log\left(\left(K_{vol}^2 + Q_{\tau,t} - Q_{0,t}\right)^{1/2} - Q_{\tau,t}^{1/2}\right)}{s_t} + (2-j)s_t \tag{19}$$

$j = 0, 1, 2$

and where the log remaining volatility $\log(Q_{\tau,T}^{1/2} - Q_{\tau,t}^{1/2})$ has time-$t$ conditional mean $m_t$ computable via:

$$m_t = 2\log\mu_1 - \frac{1}{2}\log\mu_2 \tag{20}$$

and has time-$t$ conditional variance $s_t^2$ (or, loosely speaking, the squared 'vol of vol') computable, from the relative sizes of the variance swap rate and volatility swap rate, via:

$$s_t^2 = \log\mu_2 - 2\log\mu_1 \tag{21}$$

The displaced lognormal price of a $K_{vol}$-strike volatility call is:

$$C_{LN}^{vol} = G_t^{-1} \int_{-\infty}^{\infty} \frac{e^{-z^2/2}}{\sqrt{2\pi}} \left( \left(Q_{0,t} + 2Q_{\tau,t}^{1/2} e^{m_t + s_t z} + e^{2m_t + 2s_t z}\right)^{1/2} - K_{vol} \right)^{+} dz \tag{22}$$

In the case that $\tau = 0$, this simplifies to:

$$C_{LN}^{vol} = G_t^{-1} \times \begin{cases} B_t G_t - K_{vol} & \text{if } K_{vol} \le Q_{0,t}^{1/2} \\ \mu_1 N(d_1) - \left(K_{vol} - Q_{0,t}^{1/2}\right) N(d_2) & \text{if } K_{vol} > Q_{0,t}^{1/2} \end{cases} \tag{23}$$

where $\mu_j$ and $d_j$ are obtained by taking $\tau = 0$ in (17), (18) and (19).

The pricing formulas are easily computable functions of observable quantities. Moreover, the pricing distribution prices exactly correctly any variance/volatility option known to finish in-the-money, including the benchmark swaps (which are zero-strike calls), as well as any option whose strike is smaller than the running variance/volatility. Using an undisplaced lognormal distribution would violate this latter property.

■ *Remark 2.* Consider a volatility option in the case $t = \tau = 0$ and $\mu_1 = K_{vol}$. Thus the valuation date is zero, the pricing benchmarks are $[0, T]$ swaps, and the volatility option is at-the-money. Then, by (23):

$$C_0^* := e^{rT} C_{LN}^{vol} = \mu_1\left[N(s_0/2) - N(-s_0/2)\right]$$
$$\approx \frac{B_0^* s_0}{\sqrt{2\pi}} = B_0^*\left[\frac{1}{\pi}\log\left(\sqrt{A_0^*}/B_0^*\right)\right]^{1/2} \tag{24}$$

so the forward value of the at-the-money volatility option has a particularly simple expression in terms of the volatility swap rate $B^*$ and the variance swap rate $A^*$. Moreover, rearranging (24) gives:

$$\sqrt{A_0^*}/B_0^* = e^{\pi\left(C_0^*/B_0^*\right)^2} \tag{25}$$

so the 'convexity correction' $\sqrt{A_t^*}/B_0^*$ has a simple explicit monotonic relationship with the at-the-money volatility option value $C_0^*/B_0^*$ (expressed relative to the volatility swap rate).

■ *Tests.* The lognormal distribution is empirically a plausible approximation, but to test robustness, suppose that instantaneous volatility follows Heston dynamics (which do not generate lognormal distributions for realised variance), with the same parameters as in figure 4. Hence, the genuine and synthetic volatility swap values are as plotted in that figure, and the genuine and synthetic variance swap values are both $20^2$.

In figures 7 and 8, we plot the lognormal approximations (using genuine and synthetic volatility swap prices) to a variance call price, in comparison with the true Heston variance call price. Figure 7 varies the strike $K_{vol}$, whereas figure 8 varies the price-volatility correlation ρ. In each figure, the conversion factor is $u = 100\sqrt{1/T}$.

**8 Call price approximations, against correlation**

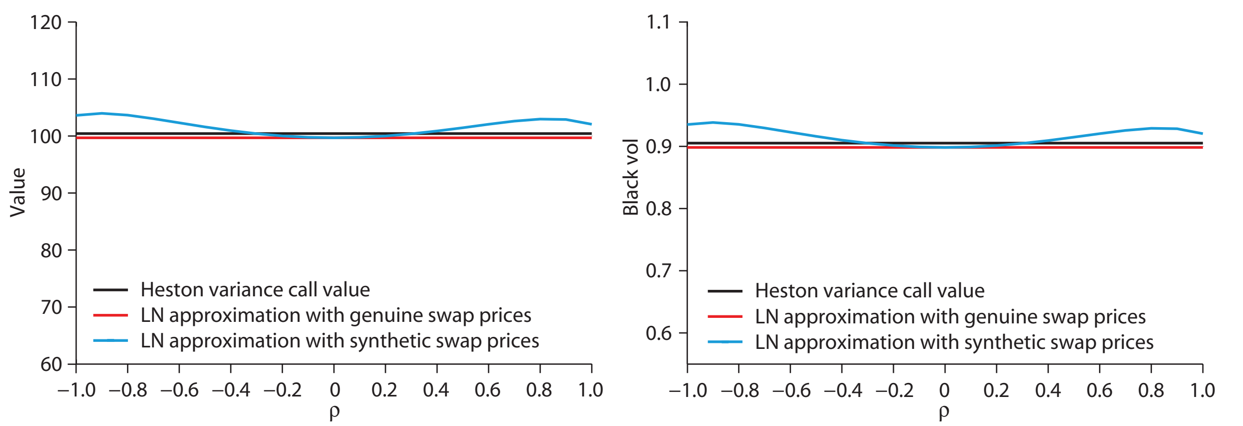


Note: the strike is fixed at-the-money, and expiry is 0.5 years. In the left-hand plot, the vertical axis is the call's dollar value; in the right-hand plot, the vertical axis is the call's Black implied volatility. Heston dynamics are as in figure 4

Even with non-lognormal Heston dynamics, and even when using synthetic volatility swaps, the figures reveal errors of less than 4bp (relative to a true value of about 100.4bp), across all correlation assumptions, in the important at-the-money 20 strike.

At out-of-the-money strikes such as 30 and 35, our approximation has significant error (if the true dynamics are Heston), but the error's sign is consistent with the relative thinness of the Heston variance distribution's tail, compared with the lognormal tail. Indeed, if the true distribution of $Q$ has the right-hand tail equal or thinner than lognormal, then the displaced lognormal approximation can be seen as a conservative estimate, prudent from the standpoint of a dealer selling an out-of-the-money variance call.

**Hedging of variance and volatility options.** We propose to hedge variance and volatility options by trading variance swaps and volatility swaps – either synthetic or genuine. Specifically, to hedge against a short position (or replicate a long position) in one variance call, hold at each time $t$:

$$\frac{\partial C_{LN}^{\mathrm{var}}}{\partial A} \text{ variance swaps and } \frac{\partial C_{LN}^{\mathrm{var}}}{\partial B} \text{ volatility swaps}$$

and to hedge against a short position (or replicate a long position) in one volatility call, hold at each time $t$:

$$\frac{\partial C_{LN}^{\mathrm{vol}}}{\partial A} \text{ variance swaps and } \frac{\partial C_{LN}^{\mathrm{vol}}}{\partial B} \text{ volatility swaps}$$

where all partial derivatives are evaluated at ($A_t$, $B_t$, $Q_{0,\,t}$, $Q_{\tau,\,t}$, $K_{vol}$).

The hedge ratios have the following explicit formulas. In the case of variance calls:

$$\frac{\partial C_{LN}^{\mathrm{var}}}{\partial A} = \begin{cases} 1 & \text{if } K_{vol}^2 \le Q_{0,t} \\ N(d_0) + \mu_2 N'(d_0)\frac{\partial D_0}{\partial A} & \\ \quad + 2Q_{\tau,t}^{1/2}\mu_1 N'(d_1)\frac{\partial D_1}{\partial A} - \chi N'(d_2)\frac{\partial D_2}{\partial A} & \text{if } K_{vol}^2 > Q_{0,t} \end{cases} \tag{26}$$

$$\frac{\partial C_{LN}^{\mathrm{var}}}{\partial B} = \begin{cases} 0 & \text{if } K_{vol}^2 \le Q_{0,t} \\ 2Q_{\tau,t}^{1/2}\left[N(d_1) - N(d_0)\right] + \mu_2 N'(d_0)\frac{\partial D_0}{\partial B} & \\ \quad + 2Q_{\tau,t}^{1/2}\mu_1 N'(d_1)\frac{\partial D_1}{\partial B} - \chi N'(d_2)\frac{\partial D_2}{\partial B} & \text{if } K_{vol}^2 > Q_{0,t} \end{cases} \tag{27}$$

where $\chi := K_{vol}^2 - Q_{0,\,t}$ and:

$$\frac{\partial D_j}{\partial B} := G_t^{-1}\frac{\partial d_j}{\partial B} = \frac{1}{s_t\mu_1} - \left(1 - j - \frac{m_t - \log\left(\left(K_{vol}^2 + Q_{\tau,t} - Q_{0,t}\right)^{1/2} - Q_{\tau,t}^{1/2}\right)}{s_t^2}\right)\left(\frac{1}{s_t\mu_1} + \frac{Q_{\tau,t}^{1/2}}{s_t\mu_2}\right) \tag{28}$$

$$\frac{\partial D_j}{\partial A} := G_t^{-1}\frac{\partial d_j}{\partial A} = \frac{1}{2s_t\mu_2} - \left(1 - j - \frac{m_t - \log\left(\left(K_{vol}^2 + Q_{\tau,t} - Q_{0,t}\right)^{1/2} - Q_{\tau,t}^{1/2}\right)}{s_t^2}\right) \tag{29}$$

for $j = 0, 1, 2$.

In the case of volatility calls with $\tau = 0$:

$$\frac{\partial C_{LN}^{vol}}{\partial A} = \begin{cases} 0 & \text{if } K_{vol} \le Q_{0,t}^{1/2} \\ \mu_1 N'(d_1)\frac{\partial D_1}{\partial A} - \left(K_{vol} - Q_{0,t}^{1/2}\right)N'(d_2)\frac{\partial D_2}{\partial A} & \text{if } K_{vol} > Q_{0,t}^{1/2} \end{cases} \tag{30}$$

$$\frac{\partial C_{LN}^{vol}}{\partial B} = \begin{cases} 1 & \text{if } K_{vol} \le Q_{0,t}^{1/2} \\ N(d_1) + \mu_1 N'(d_1)\frac{\partial D_1}{\partial B} & \\ \quad - \left(K_{vol} - Q_{0,t}^{1/2}\right)N'(d_2)\frac{\partial D_2}{\partial B} & \text{if } K_{vol} > Q_{0,t}^{1/2} \end{cases} \tag{31}$$

For general $\tau$, see equations (C.1) and (C.2) in Carr & Lee (2007).

The formulas of this section and the previous section allow arbitrary $u$, but assume a notional of one. To apply the formulas to general notionals, just ensure that $A$ and $B$ are swap prices per unit notional, then multiply the call price formulas by the call notional, and multiply the hedging formulas by the ratio of call notional to swap notional.

**Implementation choices.** To implement this strategy, the hedger must choose what type of variance and volatility swaps (genuine, synthetic or a hybrid), and what start dates (fixed or rolling) to use for the swaps. Specifically, consider the following types of variance/volatility swaps:

- 'Genuine': use genuine swap quotes to infer the option's initial price and calculate hedge ratios. Trade genuine swaps to implement the hedge.
- 'Synthetic': use synthetic swap quotes to infer the option's initial price and calculate hedge ratios. Trade synthetic swaps to implement the hedge.
- 'Hybrid': use genuine swap quotes to set the option's initial price and calculate hedge ratios. Trade synthetic swaps to implement the hedge.

Genuine quotes and genuine swap trades are desirable, but may not always be liquidly available, hence the importance of the synthetic alternatives.

All variance/volatility swaps will have terminal date $T$, but we have a choice of start dates $\tau$.

- 'Fixed': use swaps that start at a fixed time $\tau$. A natural choice is $\tau = 0$, which coincides with the start date of the option to be hedged.
- 'Rolling': at each discrete time $t_k$ when the hedge is rebalanced, use fresh swaps. Thus $\tau = t_k$ for the swaps that are held from time $t_k$ until the next rebalancing time $t_{k+1}$.

For variance swaps, the distinction between fixed and rolling is immaterial, by the additivity of variance. The dynamics of a variance swap that starts today differ only by a constant from one that started yesterday. For volatility swaps, however, the distinction matters, due to the square root.

In trading rolling-start synthetic volatility swaps, the bulk of the constituent transactions occur mainly among the most liquid vanilla options – opening a position in at-the-money-forward straddles, and closing a position in nearly-at-the-money-forward straddles (see figure 3 and table B.1(c) in the appendix of Carr & Lee, 2007).

**Simulation tests.** To test these hedging strategies, we run the following simulation. Assume Heston dynamics with $r = 0$ and Bakshi, Cao & Chen's (1997) estimated parameters $\rho = -0.64$ and:

$$dV_t = 1.15\left(0.20^2 - V_t\right)dt + 0.39\sqrt{V_t}\,dW_t, \quad V_0 = 0.20^2 \qquad (32)$$

under the risk-neutral measure. Under the physical measure, assume that $S$ has drift coefficient 0.06 and $V$ has Heston dynamics with the same parameters as (32), except a long-run mean of $0.18^2$. As the Heston model does not capture all the observed features of equity markets, the future research agenda includes supplementing these simulations with more sophisticated tests.

Let $T = 1$ year, and consider a variance call struck at-the-money at $K^2_{vol} = 0.20^2$. Although we defer to future research any theoretical analysis of the possibly material effect of discrete sampling on variance option valuations, we do use in our simulation analysis the discretely (daily) sampled returns typical in practice.

Suppose we sell the option at the initial price inferred from (genuine or synthetic) volatility swap quotes. Then we hedge by trading (genuine or synthetic) variance and volatility swaps once a day, allowing synthetic volatility swaps to use vanilla options of all strikes. Define the hedging error to be the initial sale proceeds, plus the hedging profit and loss, minus the contractual option payout, defined discretely using daily returns.

We simulate 400 paths, and report in table A the mean absolute hedging error for the variance call, under the strategy variations listed in the previous section. The errors are expressed as a percentage of the true initial option price. Figure 9 plots the distribution of the simulated error in hedging the variance call using rolling-start synthetic swaps.

The rolling-start strategy outperforms the fixed-start strategy in these simulations. Intuitively, greater sensitivity to volatility of variance can be captured in a newly issued volatility swap than in a seasoned volatility swap, making the former better suited to hedge volatility of variance.

The mean absolute hedging errors of under 10%, and under 5% for some strategies, indicate the robustness of our methodology in unfavourable conditions, including the non-lognormal dynamics and non-zero price correlation of the simulation's volatility process.

### Hedging Vix options

A special case of our set-up, with modified inputs, produces prices and hedges of Vix options using variance swaps and Vix futures. Define $V^2_{0,T} := \mathbb{E}_0 Q_{0,T}$ to be the time-zero implied variance given by the time-zero variance swap rate. Define the (idealised) time-zero Vix to be $V_{0,T} = \sqrt{V^2_{0,T}}$, where $T = 1$ month.

Consider a Vix call paying $(V_{0,T} - K_{vol})^+$ at time zero, which we will price and hedge at times $t \le 0$. Redefine the discount factor $G_t := e^{r(0-t)}$ to reflect the pay date zero, not $T$. Redefine $A_t$ as the time-$t$ price of the payout $V^2_{0,T}$, which admits replication by holding (when $t < 0$) variance swaps paying $Q_{0,T}$ (and closing the position when $t = 0$), or replication via futures on $Q_{0,T}$. Redefine $B_t$ as the time-$t$ price of the payout $V_{0,T}$, which admits replication via Vix futures.

Define $Q_{0,t} := Q_{\tau,t} := 0$. Assuming the lognormality of $V_{0,T}$, the formula (23) prices the Vix option – in agreement with Carr & Wu (2006), who, however, do not propose any hedge. Our formulas (30) and (31) propose how many units of $A$ and $B$ to hold, in a dynamic hedge of the Vix option.

More generally, assume the time-$t$ conditional lognormality of $V_{0,T} - \beta$, for some non-negative displacement parameter $\beta < B_t$. Then, with $Q^{1/2}_{0,t} := Q^{1/2}_{\tau,t} := \beta$, our formulas (23), (30) and (31) price and hedge Vix options. For $\beta > 0$, the resulting Black-implied volatility-of-Vix exhibits an upward skew at-the-money. We leave analysis of its accuracy to future research. ■

A. Mean absolute errors of six variants of the hedging strategy for at-the-money one-year variance calls, simulated under the Heston dynamics

| | Genuine swaps | Hybrid swaps | Synthetic swaps |
|---|---|---|---|
| Fixed-start swaps | 4.7% | 6.7% | 8.1% |
| Rolling-start swaps | 3.5% | 5.5% | 6.2% |

9 Distribution of the hedging errors using rolling-start synthetic swaps in table A

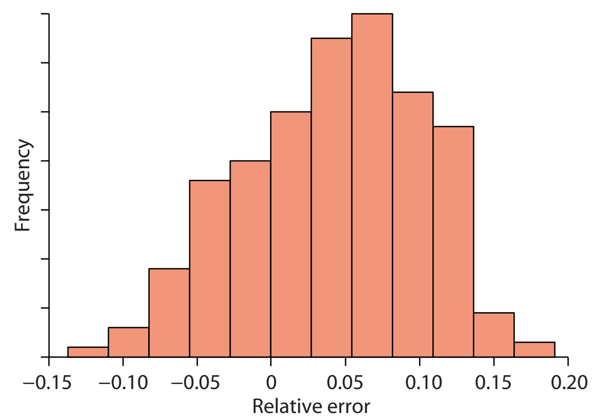


Peter Carr is head of quantitative financial research at Bloomberg and director of the Master's in Mathematical Finance Program at Courant Institute NYU. Roger Lee is assistant professor of mathematics at University of Chicago. They would like to thank the International Securities Exchange, Kris Monaco and two anonymous referees for helpful advice and support. Email: pcarr@nyc.rr.com, RL@math.uchicago.edu

## References

**Bakshi G, C Cao and Z Chen, 1997**
*Empirical performance of alternative option pricing models*
Journal of Finance 52, pages 2,003–2,049

**Brenner M and M Subrahmanyam, 1988**
*A simple formula to compute the implied standard deviation*
Financial Analysts Journal 44(5), pages 80–83

**Carr P and R Lee, 2006**
*Robust replication of volatility derivatives*
Bloomberg LP and University of Chicago. Available at *http://math.uchicago.edu/~rl/rrvd.pdf*

**Carr P and R Lee, 2007**
*Realized volatility and variance: options via swaps. With appendices*
Bloomberg LP and University of Chicago. Available at *http://math.uchicago.edu/~rl/OVSwithAppendices.pdf*

**Carr P and D Madan, 1998**
*Towards a theory of volatility trading*
In Volatility, pages 417–427, edited by Robert Jarrow, Risk Publications

**Carr P and L Wu, 2006**
*A tale of two indices*
Journal of Derivatives 13(3), pages 13–29

**Derman E, K Demeterfi, M Kamal and J Zou, 1999**
*More than you ever wanted to know about volatility swaps*
Journal of Derivatives 6(4), pages 9–32

**Dupire B, 1992**
*Arbitrage pricing with stochastic volatility*
Société Générale

**Feinstein S, 1989**
*The Black-Scholes formula is nearly linear in sigma for at-the-money options: therefore implied volatilities from at-the-money options are virtually unbiased*
Federal Reserve Bank of Atlanta

**Friz P and J Gatheral, 2005**
*Valuation of volatility derivatives as an inverse problem*
Quantitative Finance 5(6), pages 531–542

**Heston S, 1993**
*A closed-form solution for options with stochastic volatility and applications to bond and currency options*
Review of Financial Studies 6(2), pages 327–343

**Neuberger A, 1990**
*Volatility trading*
Working paper, London Business School