Hedging Options on Realized Variance

Fields Institute

February 26, 2008

Roger Lee

University of Chicago

RL@math.uchicago.edu

Joint with Peter Carr (Bloomberg LP and NYU Courant)

Some uses of volatility derivatives

- Directional: The user has a forecast/view (perhaps based on statistical time-series or macroeconomic or event-driven analysis) that volatility will be higher or lower than the market expects.
- Relative value: The user believes that the market is overpricing the volatility of one index relative to another index. Or the user believes that the market is over/under-pricing the volatility of an index relative to the volatility of the index's constituent stocks.
- Hedging: The user wants to hedge volatility exposure.

Variance swap pays realized variance less a fixed leg

Variance swaps (RISK, August 2006):

The market for variance swaps has grown sharply over the past two years. Deutsche Bank estimates that volumes grew fivefold last year and are up 50% year-on-year for the first half of 2006. According to BNP Paribas estimates, daily trading volumes for variance swaps on indexes reached $4 million – 5 million in vega in the first half of 2006. Credit Suisse, meanwhile, reckons daily trading volumes are around EUR 2 million – 3 million in vega, compared with EUR 250,000 – 500,000 three years ago.

Variance swaps (Financial Times, May 2006):

Volatility is becoming an asset class in its own right. A range of structured derivative products, particularly those known as variance swaps, are now the preferred route for many hedge fund managers and proprietary traders to make bets on market volatility.

Todd Steinberg, head of equities and derivatives at BNP Paribas in New York, said: "Variance swaps isolate volatility and take away all of the other attributes that you would ordinarily have to factor in, such as . . . movements in the price of the underlying asset." He pointed out that . . . the contracts are . . . almost as liquid as S&P 500 listed options, with spreads that are just as tight and with smaller capital requirements. Of course, losses have the potential to build up at a greater rate for those who turn out to take the wrong side of the swap.

Variance options (Financial Times, May 2007):

Mr Fields [head of flow and structured products, SocGen] remembers a call on variance on the Eurostoxx in 2005 ... was unexpectedly crossable between clients and became very liquid within months. At Deutsche they are not so sure ... head of equity derivatives for Europe Nino Kjellman ... admits to being "disappointed" at the flow in options on variance ... He blames a combination of bid-ask margins refusing to contract and the fact that risk management systems still need improvements to book the trades. "I wouldn't be happy to see another tranche on top of what we have already," says Mr Ankaoua. "The industry is taking a big risk writing such products and at some point that will be a risk that you can't assess. This industry has to fulfill investors' needs, but ... I don't want to write a ticking bomb."

This talk

- Objective: create/replicate/synthesize contracts on realized variance using vanilla options and the underlier.
- We want the replication strategy to be *model-free*, not dependent on a specific model of the underlying dynamics.
- For variance swaps, we review the known replication strategy.
- For variance options, it's impossible to replicate perfectly in a model-independent way, so we will *superreplicate* using vanilla options and the underlying. (Subreplication: Dupire 2005)
- Carr-Lee's "Robust Replication of Volatility Derivatives" made an independence assumption to find [infinitely many] perfect replication strategies for a given variance contract.

  Here we drop the independence assumption entirely.

Assumptions and notation

- Frictionless arbitrage-free markets
- Let $Y > 0$ denote the share price of an underlying asset.

  Assume Y is a continuous semimartingale (local martingale + finite variation "drift"). For example, an Itô process

  $$\mathrm{d}Y_t = \mu_t Y_t \mathrm{d}t + \sigma_t Y_t \mathrm{d}W_t,$$

  where W is a Brownian motion, and μ and σ may be stochastic.
- Let $X_t := \log(Y_t/Y_0)$ be the log returns process.
- No dividends.

  Zero interest rates: "cash"/"bond" has price 1 at all times.

Assumptions and notation

- Let $[X]$ and $[Y]$ denote the *quadratic variation*, which is the "continuously sampled cumulative realized variance" of respectively X and Y. The definition can be expressed as:
$$[X]_t = \int_0^t \mathrm{d}[X]_s = \int_0^t (\mathrm{d}X_s)^2.$$
So $X_t = \log(Y_t/Y_0)$ implies
$$\mathrm{d}X_t = \frac{1}{Y_t}\mathrm{d}Y_t + (\text{FV term}) \quad \Rightarrow \quad \mathrm{d}[X]_t = \frac{1}{Y_t^2}\mathrm{d}[Y]_t$$
In the Itô process case, $[X]_T = \int_0^T \sigma_t^2 \mathrm{d}t$ becuase
$$\mathrm{d}X_t = \sigma_t \mathrm{d}W_t + (\mathrm{d}t \text{ term}) \quad \Rightarrow \quad \mathrm{d}[X]_t = \sigma_t^2 \mathrm{d}t.$$
- A *variance swap* pays $[X]_T$. (So we're assuming zero fixed leg, unit notional, no annualization, continuous sampling.)
- A *variance call* with strike Q and expiry T pays $([X]_T - Q)^+$.

Variance swaps

Variance options

- Superreplicating $[X]_T - [X]_{\tau_b}$
- Superreplicating $([X]_{\tau_b} - Q)^+$
- The combined superhedge

Log contract is a “synthetic variance swap”

The payoff $[X]_T$ admits replication by

the European-style payoff $-2\log(Y_T/Y_0)$

plus the P&L from holding $2/Y_t$ shares at each time t

Therefore, a variance swap has time-0 value equal to a claim on

$$-2\log(Y_T/Y_0).$$

Model-independent! Essentially no assumptions on σ_t.

(Dupire 92, Neuberger 94, Carr-Madan 98, Derman et al 99)

Synthetic variance swap

Log contract

$$-2\log(Y_T/Y_0)$$

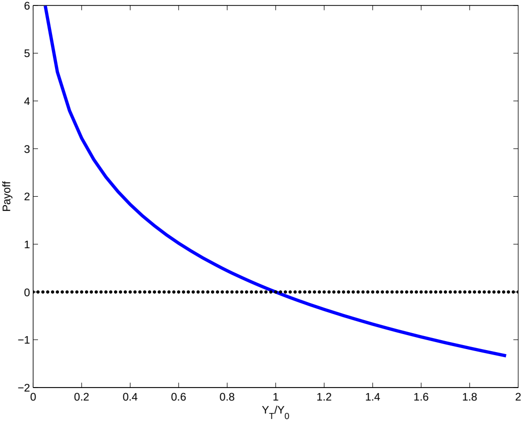

Synthetic variance swap

Log contract including also the financed share position:

$$-2\log(Y_T/Y_0) + 2(Y_T/Y_0 - 1)$$

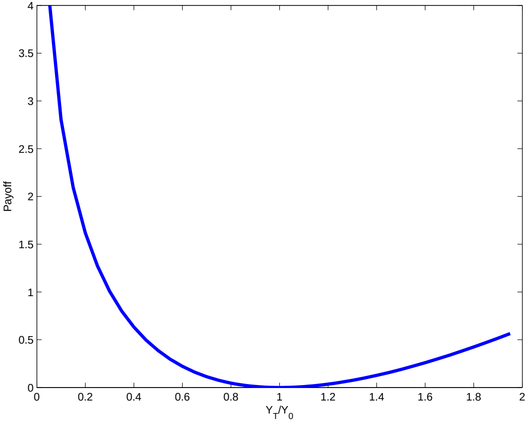

Synthetic variance swap

- Equivalently, it's the payoff of

$$(2/K^2)\mathrm{d}K$$

  OTM calls or puts at each strike K.
- The value of this option portfolio is the model-free risk-neutral expectation of realized variance on $[0, T]$, which can be described as the "model-free implied variance".
- Since 2003, the CBOE's VIX index has been defined as the square root of this value, using SPX options, for $T = 30$ days, after interpolation.

Derivation

By Itô's rule, $X_t = \log(Y_t/Y_0)$ implies

$$\mathrm{d}X_t = \frac{1}{Y_t}\mathrm{d}Y_t + \frac{1}{2}\left(\frac{-1}{Y_t^2}\right)\mathrm{d}[Y]_t = \frac{1}{Y_t}\mathrm{d}Y_t - \frac{1}{2}\mathrm{d}[X]_t.$$

Rewrite as

$$X_T = \int_0^T \frac{1}{Y_t}\mathrm{d}Y_t - \frac{1}{2}[X]_T.$$

Rearranging,

$$[X]_T = -2X_T + \int_0^T \frac{2}{Y_t}\mathrm{d}Y_t.$$

This is the sum of the European-style payoff $-2\log(Y_T/Y_0)$ and the gains from dynamically holding $2/Y_t$ shares.

Synthetic variance swaps: adding linear functions

Let λ, θ be arbitrary constants. Let

$$L(y) := -2\log y + \lambda y + \theta$$

The payoff $[X]_T$ admits replication by

the European-style payoff $L(Y_T) - L(Y_0)$

plus the P&L from holding $2/Y_t - \lambda$ shares at each time t

Therefore, a variance swap has time-0 value equal to a claim on

$$L(Y_T) - L(Y_0)$$

Or simply

$$L(Y_T)$$

if L is "at the money" at time 0. Model-independent!

Synthetic variance swaps: adding linear functions

At time 0 each of these price payoffs has same value as $[X]_T$ payoff

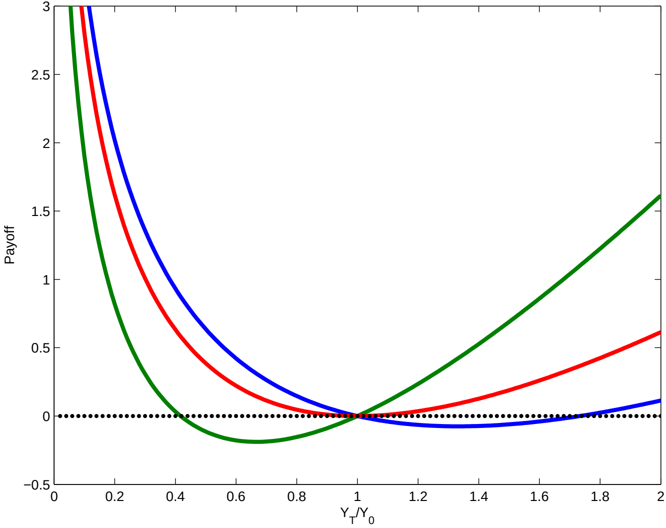

What about discrete monitoring and hedging

Discretely monitor and hedge, at times t_n. The variance swap pays

$$\sum_n (\log(Y_{t_{n+1}}/Y_{t_n}))^2$$

The hedge pays $-2\log(Y_T/Y_0)$ plus the share-trading P&L:

$$\sum_n \Big(-2\log(Y_{t_{n+1}}/Y_{t_n}) + \frac{2}{Y_{t_n}}(Y_{t_{n+1}} - Y_{t_n})\Big)$$

As a function of the one-period simple return $R_{n+1} := Y_{t_{n+1}}/Y_{t_n} - 1$, the one-period variance swap P&L is

$$[\log(1+R)]^2$$

and the one-period synthetic variance swap P&L is

$$-2\log(1+R) + 2R.$$

For small R, the P&L error (synthetic $-$ genuine) $\approx R^3/3$ is small.

What about discrete monitoring and hedging

P&L of synthetic variance swap ≈ P&L of variance swap because

$$\sum_n \Big(-2\log(Y_{t_{n+1}}/Y_{t_n}) + \frac{2}{Y_{t_n}}(Y_{t_{n+1}} - Y_{t_n})\Big) \approx \sum_n (\log(Y_{t_{n+1}}/Y_{t_n}))^2$$

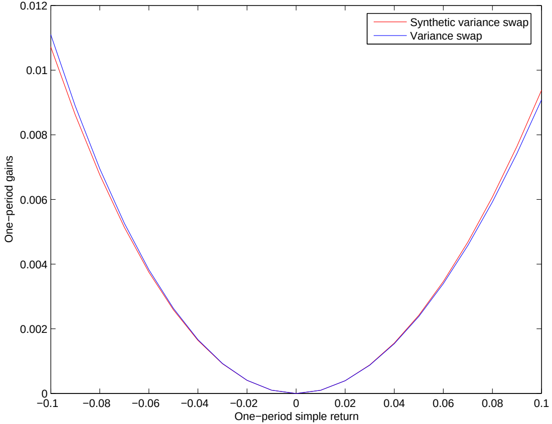

Variance swaps

Variance options

Superreplicating $[X]_T - [X]_{\tau_b}$

Superreplicating $([X]_{\tau_b} - Q)^+$

The combined superhedge

Variance call: *Option* to receive variance less strike

Synthesizing variance *options*?

- Variance *options* give portfolio managers more control over risk exposure. But they pose greater hedging problems to the dealer.
- The difficulties arise from the absence of a model-independent strategy to replicate variance options using Europeans and Y.
- But we will find model-independent strategies to *superreplicate* variance options using Europeans and Y.
- If a variance option is bid above our upper bound, then short it and go long our superreplicating strategy. Model-free arbitrage.

Variance call upper bound: general strategy

- Introduce a double barrier $\{b_d, b_u\}$. Let τ_b be Y's passage time.
- Decompose

$$([X]_T - Q)^+ \leq ([X]_{\tau_b \wedge T} - Q)^+ + ([X]_T - [X]_{\tau_b \wedge T}),$$

  the sum of a pre-barrier-variance call and post-barrier-variance.
- We will use cash and shares (dynamic) to replicate $([X]_{\tau_b} - Q)^+$,
  hence superreplicate $([X]_{\tau_b \wedge T} - Q)^+$.
- We will use options (static) and shares (dynamic) to:
  superreplicate $[X]_T - [X]_{\tau_b}$ in the event that $\tau_b \leq T$
  and produce 0 in the event that $\tau_b > T$.
  Hence superreplicate $([X]_T - [X]_{\tau_b \wedge T})$.

The basic identity

Both superhedges come from a basic identity. Recall

$$
\begin{aligned}
\mathrm{d}X_t &= \frac{1}{Y_t}\mathrm{d}Y_t + (\mathrm{d}t \text{ term}) \\
\mathrm{d}[X]_t &= \frac{1}{Y_t^2}\mathrm{d}[Y]_t.
\end{aligned}
$$

Let $h(x,q)$ be smooth. Again by Itô

$$
\begin{aligned}
h(Y_T,[X]_T) &= h(Y_0,0) + \int_0^T h_y \mathrm{d}Y_s + \int_0^T \frac{1}{2} h_{yy}\mathrm{d}[Y]_s + \int_0^T h_q \ \mathrm{d}[X]_s \\
&= h(Y_0,0) + \int_0^T h_y \mathrm{d}Y_s + \int_0^T \frac{1}{2} Y_s^2 h_{yy} + h_q \ \mathrm{d}[X]_s
\end{aligned}
$$

suppressing the arguments $(Y_s,[X]_s)$. See Bick (1995).

More generally we can replace T with a stopping time $\tau \wedge T$.

Variance swaps

Variance options

- Superreplicating $[X]_T - [X]_{\tau_b}$
- Superreplicating $([X]_{\tau_b} - Q)^+$
- The combined superhedge

Create $[X]$ using the $\mathrm{d}[X]$ integral

The basic identity says

$$h(Y_T, [X]_T) = h(Y_0, 0) + \int_0^T h_y \mathrm{d}Y_s + \int_0^T \frac{1}{2} Y_s^2 h_{yy} + h_q \ \mathrm{d}[X]_s$$

Making $h(y, q)$ depend only on y, let's have $[X]$ dependence in RHS.

- In particular, if $h(y, q) = L(y) = -2 \log y + \lambda y + \theta$ then

$$L(Y_T) = L(Y_0) + \int_0^T L_y \mathrm{d}Y_s + [X]_T$$

  This recovers the classic synthetic variance swap: hold a claim on $L(Y_T) - L(Y_0)$ hedged dynamically with $-L_y$ shares.

- If λ, θ are such that L is ATM, then replicating portfolio value is

$$\mathbb{E}_0 L(Y_T) - L(Y_0) = \mathbb{E}_0 L(Y_T)$$

Superreplicating $[X]_T - [X]_{\tau_b}$

Let

$$L(y) := -2\log y + \lambda y + \theta$$

where θ, λ are constants chosen such that $L(b_u) = L(b_d) = 0$.

At time τ_b we want to have a claim on $L(Y_T) - L(Y_{\tau_b}) = L(Y_T)$, or larger. Simply holding a $L(Y_T)$ claim suffices if $\tau_b \leq T$.

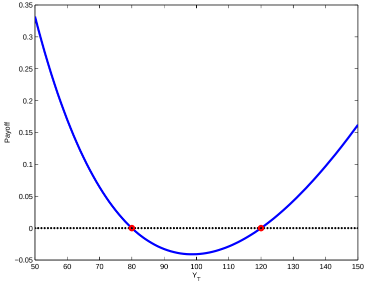

Superreplicating $[X]_T - [X]_{\tau_b}$

However, if $\tau_b > T$ then $L(Y_T) < 0$ and we don't want that liability.

So holding a claim on $L(Y_T)^+$ suffices in all cases.

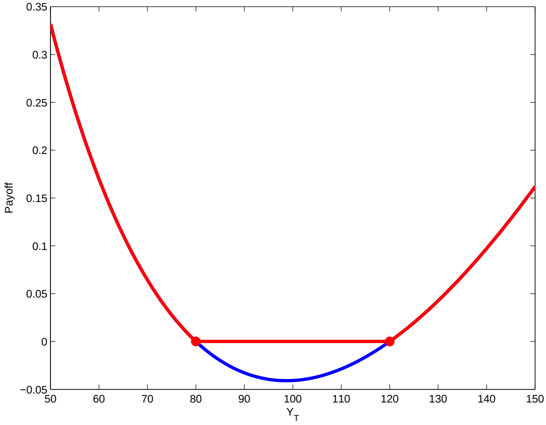

Variance swaps

Variance options

Superreplicating $[X]_T - [X]_{\tau_b}$

Superreplicating $([X]_{\tau_b} - Q)^+$

The combined superhedge

Create functions of $[X]$ using $h(Y_T, [X]_T)$

Recall the basic identity

$$h(Y_T, [X]_T) = h(Y_0, 0) + \int_0^T h_y \mathrm{d}Y_s + \int_0^T \frac{1}{2} Y_s^2 h_{yy} + h_q \; \mathrm{d}[X]_s$$

Now let's generate $[X]$ dependence in the LHS instead of the RHS.

- If $h(y, q)$ satisfies $\frac{1}{2} y^2 h_{yy} + h_q = 0$ then $h(Y_T, [X]_T)$ can be synthesized as $h(Y_0, 0)$ plus P&L from holding h_y shares. (Start with "intrinsic" in cash. Dynamically "delta hedge.")
- Imposing "terminal conditions" at $q = Q$ lets us create general payoff functions of price-when-variance-reaches-Q.
- Imposing "boundary conditions" at $y = b$ lets us create general payoff functions of variance-when-price-reaches-b.

Intrinsic+Delta replicates $h(Y, [X])$ if h solves PDE

- If h solves the PDE
$$\frac{1}{2}y^2 h_{yy} + h_q = 0$$
then the payoff $h(Y_{T\wedge\tau}, [X]_{T\wedge\tau})$ is replicated by holding
$$\begin{aligned} h_y(Y_t, [X]_t) &\quad \text{shares} \\ h(Y_t, [X]_t) - Y_t h_y(Y_t, [X]_t) &\quad \text{cash} \end{aligned}$$
at each $t \leq T \wedge \tau$. Initial replicating portfolio value is $h(Y_0, 0)$.
- The PDE can be understood as a "Black-Scholes" PDE (or a backward Kolmogorov equation) for driftless GBM with unit volatility, but with q in place of t.

The time-change point-of-view

This result, and what follows, can be understood via *time change*.

In particular recall the theorem of Dambis/Dubins-Schwarz:

If M is a continuous martingale with $M_0 = 0$ and $[M]_\infty = \infty$ then there exists a Brownian motion B such that for all t

$$M_t = B_{[M]_t}$$

The DDS Brownian motion $(B_u, \mathcal{F}_{A_u})$ is given by $B_u := M_{A_u}$ where $A_u := \{\inf u : [M]_t > u\}$.

- Intuition: Every continuous driftless process M is a Brownian motion run on a stochastic clock. The clock ("business time") is the realized variance of M. The clock runs faster when M is more volatile, slower when M is less volatile.

The time-change point-of-view

To apply DDS, we have

$$\mathrm{d}X_t = \frac{1}{Y_t}\mathrm{d}Y_t - \frac{1}{2Y_t^2}\mathrm{d}[Y]_t = \frac{1}{Y_t}\mathrm{d}Y_t - \frac{1}{2}\mathrm{d}[X]_t$$

so $M_t := X_t + \frac{1}{2}[X]_t$ is (under risk-neutral measure) a martingale.

By DDS, there exists Brownian motion B with $B_{[X]_t} = M_t$ for all $t \le T$. Hence

$$X_t = B_{[X]_t} - \frac{1}{2}[X]_t.$$

So X is drift $-1/2$ Brownian motion, run on a stochastic clock $[X]_t$.

So Y is driftless GBM, wrt this *business time* clock: $Y_t = \mathrm{GBM}_{[X]_t}$.

[Dupire (2005) uses DDS to relate volatility derivatives to the Skorohod problem. Forde (2005) notes the relevance of DDS to pre-barrier variance contracts.]

Intrinsic+Delta replicates general $g([X]_{\tau_b})$

Let τ_b be the first passage time of Y to b. For general boundary data $g : [0,\infty) \to \mathbb{R}$, let $h(y,q) := BP(y,q;b,g)$ where

$$BP(y,q;b,g) := \begin{cases} \int_0^\infty g(q+z)\ \frac{|\log(b/y)|}{\sqrt{2\pi z^3}} \exp\Big[-\frac{(\log(b/y)+z/2)^2}{2z}\Big]\, \mathrm{d}z & \text{if } q \neq b \\ g(q) & \text{if } q = b \end{cases}$$

Then the intrinsic+delta strategy creates $BP(X_{T\wedge\tau_b}, [X]_{T\wedge\tau_b}) =$

$$g([X]_{\tau_b})\mathbb{I}_{\tau_b \le T} + BP(X_T, [X]_T)\mathbb{I}_{\tau_b > T} \ge g([X]_{\tau_b\wedge T}) \text{ if } g \text{ is increasing}$$

Replicating portfolio has time-0 value $BP(Y_0, 0; b, g)$.

Proof.

Directly verify that $BP(y,q)$ satsifies the PDE. $\square$

Intuition of BP formula for pricing general $g([X]_{\tau_b})$

- $\mathbb{E}g(\text{variance until } Y \text{ hits } b) =$

  $\mathbb{E}g(\textit{time} \text{ until unit-vol driftless GBM hits } b)$.

  The $BP(y, q; b, g)$ computes this $\mathbb{E}g$(“Brownian passage” time).

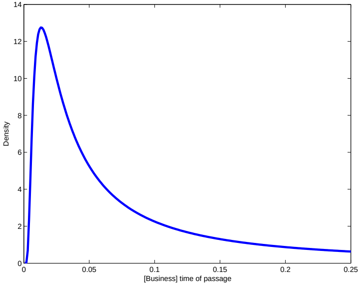

- The kernel $z \mapsto \frac{|\beta|}{\sqrt{2\pi z^3}} \exp\left[-\frac{(\beta+z/2)^2}{2z}\right]$ is the [defective if $\beta \geq 0$] density of passage time of drift $-1/2$ Brownian motion to level $\beta = \log(b/y)$ (or driftless GBM from $e^0 = 1$ to $e^\beta = b/y$).

The double barrier case

- Likewise, $g([X]_{\tau_b})$, where τ_b is first passage time to a *double* barrier $\{b_d, b_u\}$, admits replication by intrinsic+delta. In this case
$$BP(y,q;b_d,b_u,g) = \int_0^\infty g(q+z)\ p(\log(b_d/y), \log(b_u/y), z)\mathrm{d}z$$
involves a double-barrier density p.
- In the case $g(q) = (q-Q)^+$ we have the representation
$$BP(y,q;b_d,b_u,Q) = \int_{-\infty-\alpha i}^{\infty-\alpha i} \frac{\sqrt{y/b_u}\sinh(\log(b_d/y)\sqrt{1/4-2iz}) - \sqrt{y/b_d}\sinh(\log(b_u/y)\sqrt{1/4-2iz})}{2\pi z^2 e^{i(Q-q)z}\sinh(\log(b_u/b_d)\sqrt{1/4-2iz})}\mathrm{d}z$$
where $\alpha > 0$ is arbitrary.
Proof: Combine the well-known Laplace transform of p with the call pricing formula in Lee (2004).

Variance swaps

Variance options

Superreplicating $[X]_T - [X]_{\tau_b}$

Superreplicating $([X]_{\tau_b} - Q)^+$

The combined superhedge

Variance call upper bound: general strategy

- Introduce a double barrier $\{b_d, b_u\}$. Let τ_b be Y's passage time.
- Decompose

$$([X]_T - Q)^+ \leq ([X]_{\tau_b \wedge T} - Q)^+ + ([X]_T - [X]_{\tau_b \wedge T}),$$

  the sum of a pre-barrier-variance call and post-barrier-variance.
- We will use cash and shares (dynamic) to replicate $([X]_{\tau_b} - Q)^+$, hence superreplicate $([X]_{\tau_b \wedge T} - Q)^+$.
- We will use options (static) and shares (dynamic) to:
  superreplicate $[X]_T - [X]_{\tau_b}$ in the event that $\tau_b \leq T$
  and produce 0 in the event that $\tau_b > T$.
  Hence superreplicate $([X]_T - [X]_{\tau_b \wedge T})$.

Superreplication of variance call

Fix $b_d \leq Y_0$ and $b_u \geq Y_0$. Let

$$N_t := \begin{cases} BP_y(Y_t, [X]_t; b_d, b_u; Q), & t \leq \tau_b \\ -L_y(Y_t), & t > \tau_b. \end{cases}$$

The variance call is superreplicated by the strategy

$$\begin{array}{rl} 1 & \text{claim on } L(Y_T)^+ \\ N_t & \text{shares} \\ BP(Y_0, 0; Q) + \int_0^t N_s \mathrm{d}Y_s - N_t Y_t & \text{cash} \end{array}$$

Superreplicating portfolio has time-0 value $BP(Y_0, 0; Q) + \mathbb{E}_0 L(Y_T)^+$.

Proof.

The strategy clearly self-finances.

If $\tau_b \geq T$, then the portfolio has time-T value

$$L(Y_T)^+ + BP(Y_T, [X]_T; Q) \geq 0 + ([X]_T - Q)^+.$$

If $\tau_b < T$ then the portfolio has time-T value

$$\begin{aligned}([X]_{\tau_b} - Q)^+ + \int_{\tau_b}^{T} \frac{2}{Y_t}\mathrm{d}Y_t + L(Y_T)^+ &\geq ([X]_{\tau_b} - Q)^+ + [X]_T - [X]_{\tau_b} \\ &\geq ([X]_T - Q)^+\end{aligned}$$

where the first $\geq$ is because $L(Y_T)^+ \geq L(Y_T)$. $\square$

Variance call upper bound: a better decomposition

- Instead of

$$([X]_T - Q)^+ \le ([X]_{\tau_b \wedge T} - Q)^+ + ([X]_T - [X]_{\tau_b \wedge T}),$$

decompose

$$([X]_T - Q)^+ \le ([X]_{\tau_b} - Q)^+ - ([X]_{\tau_b} - [X]_T - Q)^+ + ([X]_T - [X]_{\tau_b \wedge T}),$$

long a pre-barrier-variance perpetual call,

short a post-barrier-variance perpetual call,

long post-barrier-variance

Variance call upper bound: Improvement

A better (cheaper) alternative to the payoff $L(Y_T)^+$ is $L^*(Y_T)$, where

$$L^*(y) := L^*(y; b_d, b_u, Q) := \begin{cases} L(y) & \text{if } y \notin (b_d, b_u) \\ -BP(y, 0; Q) & \text{if } y \in (b_d, b_u) \end{cases}$$

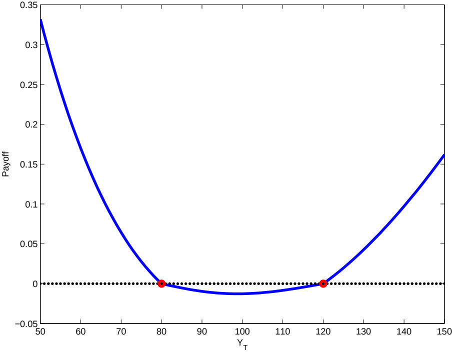

Variance call upper bound: Improvement

If $\tau_b \geq T$, then the portfolio still has time-T value

$$-BP(Y_T, 0; Q) + BP(Y_T, [X]_T; Q) \geq 0 + ([X]_T - Q)^+.$$

because

$$([X]_{\tau_b} - Q)^+ \geq ([X]_T - Q)^+ + ([X]_{\tau_b} - [X]_T - Q)^+$$

If $\tau_b < T$ then the portfolio still has time-T value

$$\begin{aligned}([X]_{\tau_b} - Q)^+ + \int_{\tau_b}^{T} \frac{2}{Y_t} \mathrm{d}Y_t + L^*(Y_T) &\geq ([X]_{\tau_b} - Q)^+ + [X]_T - [X]_{\tau_b} \\ &\geq ([X]_T - Q)^+\end{aligned}$$

where the first $\geq$ is because $L^*(Y_T) \geq L(Y_T) - L(Y_{\tau_b})$.

Variance call upper bound: Minimization

- The upper bound depends on the arbitrarily chosen $b_d \leq Y_0$ and $b_u \geq Y_0$. Optimized upper bound is

$$\inf_{(b_d, b_u)} \left[BP(Y_0, 0;\ b_d, b_u, Q) + \mathbb{E}_0 L^*(Y_T;\ b_d, b_u, Q) \right]$$

- The first term has an explicit formula (Fourier-style).

  The second term is observable from T-expiry European prices.

- Intuition:

  The bigger the strike Q, the wider the optimal interval (b_d, b_u).

  An interval too narrow includes in the post-barrier variance too much pre-Q variance.

  An interval too wide includes in the [call on] pre-barrier variance too much variance occurring after expiry.

Variance call: Universal superreplicating portfolios

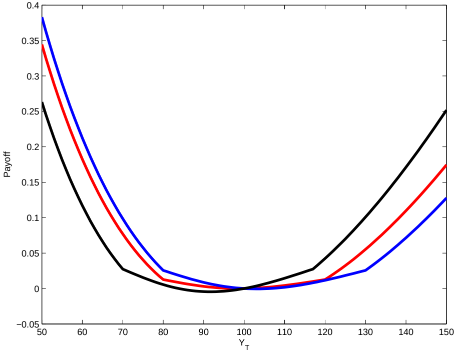

Let $Y_0 = 100$. These 3 time-T payoffs arise from 3 different choices of (b_d, b_u). A claim on any one of these payoffs, together with dynamic trading of shares, model-independently superreplicates a spot-starting variance call with strike 0.04 and expiry T.

Example of upper and lower bounds

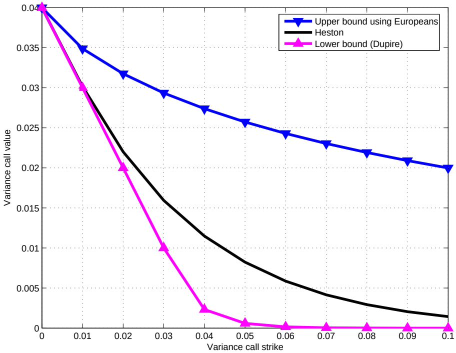

Bounds on variance calls with expiry $T = 1$, assuming the prices of T-expiry Europeans are consistent with Heston dynamics

$$dV_t = 1.15(0.04 - V_t) + 0.39\sqrt{V_t}dW_{2t}, \qquad V_0 = 0.04$$

where W_1 and W_2 are independent Brownian motions.

Conclusions

- Exactly replicate functions of pre-barrier variance (which, by DDS, has the same law as *time*-until-unit-GBM-hits-barrier).
- Superreplicate post-barrier variance.
- Hence superreplication of variance call.
- All results are model-independent, rigorous, and extendable (to forward-starting variance, arithmetic variance, corridor variance).
- Trading strategies are simple (static in options, "delta" in shares) and give explicit arbitrage if our bounds are violated.