

Variance Swaps and Volatility Derivatives

John Crosby

Glasgow University

My website is: <http://www.john-crosby.co.uk>

If you spot any typos or errors, please email me.

My email address is on my website

Lecture given 20th February 2009
for the M.Sc. and Ph.D. courses
in Quantitative Finance
in the Department of Economics
at Glasgow University

File date 21st June 2009 20.08

Motivation

- The market prices of stocks (or other assets such as foreign exchange rates or commodity prices) fluctuate randomly. Once we have observed a time-series of market prices, we can compute the realised variance. If we take the square root, we can compute the realised volatility. Suppose a trader wishes to take a view (via a trading position) today on the realised variance that will be observed over some given future time period. How can she do this?
- What sort of derivatives can be used for this and how are they priced and hedged?
- How is the realised variance over this given time period (which is unknown today but will be known at the end of the time period) related to the implied volatilities, observable today, of vanilla options which mature at the end of the time period?
- These are questions which we will try to answer today.

Why this is not easy

- It might be tempting to think that if a trader thinks that, for example, realised volatility over a given time period will be higher than the implied volatility of an option maturing at the end of the time period, then she should buy the vanilla option. However, what strike should the option have? Vanilla options which are struck at-the-forward forward have the largest vega (sensitivity to volatility). But options, which are at-the-forward forward at the time they are written, may be deep in or out of the money later (because the stock price moves) at which time they will have a much lower vega.
- It is clear that vanilla options are an imperfect vehicle for a trader to take a view on volatility or variance. This is because the price of the vanilla and the sensitivity of the price of the vanilla to variance (ie the partial derivatives of the vanilla price with respect to variance) depends on the stock price.
- What sort of instrument or derivative might be a better vehicle to take a view on variance.

A primer

- Let us introduce some notation. Suppose today, time t_0 , we write a European option which matures at time T on a stock whose price, at time t , is denoted by $S(t)$. We denote the price of the option, at time t , by $C(t)$. We assume that the stock price follows geometric Brownian motion with volatility σ . We'll assume at this stage, for simplicity, that interest-rates are zero and the stock pays no dividends. We delta-hedge our short position in the European option and rebalance our portfolio every Δt .
- Note that Δt is finite - not infinitesimal.
- The P+L (profit and loss) over the time interval from t to $t + \Delta t$ is:

$$\frac{\partial C}{\partial t} + \frac{1}{2} \frac{\partial^2 C}{\partial S^2} (\Delta S)^2,$$

where $\Delta S \equiv S(t + \Delta t) - S(t)$.

- In the last line, we have used a Taylor series expansion and cancelled out the delta terms.

A primer 2

- However, the Black and Scholes (1973) pde says:

$$\frac{\partial C}{\partial t} + \frac{1}{2}\sigma^2(S(t))^2\frac{\partial^2 C}{\partial S^2} = 0.$$

Hence, substituting, we get that the P+L over the time interval from t to $t + \Delta t$ is:

$$\frac{1}{2}S^2\frac{\partial^2 C}{\partial S^2} \left(\frac{(\Delta S)^2}{S^2} - \sigma^2 \right).$$

- Note that if we were to let Δt tend to zero, the P+L would tend to zero (this is simply the Merton (1973) hedging argument). However Δt is not infinitesimal.
- We can sum up the P+L over each time interval Δt . Then the P+L over the time interval from t_0 to T is:

$$\sum \frac{1}{2}S^2\frac{\partial^2 C}{\partial S^2} \left(\frac{(\Delta S)^2}{S^2} - \sigma^2 \right).$$

- Notice how there is a path-dependency in this P+L. If, for example, ΔS were to tend to be large, when $\frac{\partial^2 C}{\partial S^2}$ was large and positive, then the P+L would tend to be large and positive. If, for example, ΔS were to tend to be small relative to σS , and if $\frac{\partial^2 C}{\partial S^2}$ is positive (which it certainly is for a vanilla option), then the P+L would tend to be negative.

A primer 3

- So, in general, the P+L of the delta-hedging strategy is path-dependent. However, while we assumed the option was European, we never assumed it had a vanilla payoff. The option could have any payoff at time T .
- Suppose that the option is such that its gamma $\frac{\partial^2 C}{\partial S^2}$ is identically equal to $1/S^2$. Then the P+L over the time interval from t_0 to T is:
$$\sum \frac{1}{2} \left(\frac{(\Delta S)^2}{S^2} - \sigma^2 \right).$$
- Note that σ^2 is constant (we assumed this at the beginning). Furthermore, $\sum \frac{(\Delta S)^2}{S^2}$ is a possible definition for realised variance. In the market, variance swaps (which we will define and explain shortly - but which are essentially forward contracts on realised variance) have a payoff whose floating part is $\sum (\log(S(t + \Delta t)/S(t)))^2$. However, if Δt and ΔS are small then a Taylor's series expansion implies $\frac{(\Delta S)^2}{S^2}$ and $(\log(S(t + \Delta t)/S(t)))^2$ are approximately equal. So the P+L (upto a scaling factor) is approximately the same as that of a variance swap.

A primer 4

- What sort of derivative has a gamma equal to $1/S^2$?
- Integrating twice, we get $C(t) = a - \log(S(t)) + bS(t)$, where a and b are constants of integration.
- Notice how we can interpret a as cash (or equivalently a bond) and the term $bS(t)$ as a forward contract.
- The term $\log(S(t))$ represents a derivative whose payoff is \log of the stock price at maturity T . It is called a log contract (actually we often normalise by the initial stock price $S(t_0)$ so the payoff of the log contract is $\log(S(T)/S(t_0))$) and we will see that it plays a pivotal role in the pricing of variance swaps. (Note that a log contract can have a negative payoff).
- We will now consider the pricing and hedging of variance swaps.

Variance swaps (definition)

- A variance swap is a financial derivative whose payoff is defined as follows: It is written at time t_0 and matures at time T . The time interval $[t_0, T]$ is partitioned into N time periods $t_i, i = 1, 2, \dots, N$ where $t_N = T$. The time periods do not have to be equal although they are often approximately equal. The payoff of a (discretely monitored) variance swap at time T is:

$$\frac{1}{(T - t_0)} \sum_{i=1}^N ((\log(S(t_i)/S(t_{i-1})))^2 - K^2),$$

where K is a constant (called the fixed leg).

- K is often chosen (as for IR swaps) to make the initial (ie time t_0) price of the variance swap equal to zero.
- Note that, in practice in the markets, the floating leg does not subtract the square of the mean (so it is not really a variance).
- However, the mean squared is typically tiny so it doesn't make much difference. Furthermore, the definition means variances are additive in the sense that we can define a forward starting variance swap which starts in three months time and which is based on the computed realised variance for a further six months, say. Then if we own such a forward starting variance swap and a three month variance swap (starting today), then it is the same as owning a nine month variance swap (starting today).

Variance swap pricing methodologies

- In practice, all vanilla variance swaps have payoffs which are discretely monitored. However, from a theoretical standpoint, it is also relevant to consider continuously monitored variance swaps. We will consider pricing variance swaps from two different viewpoints.
- The first viewpoint is the classic "log-contract" replication approach. It has the benefit that it also shows how to hedge variance swaps. This approach requires some (fairly weak) assumptions and actually gives prices for continuously monitored variance swaps.
- The second approach prices discretely monitored variance swaps. It has the advantage that it is very generic because it works for almost all stochastic processes that might be used in mathematical finance. It has the disadvantage that it does not show how to hedge variance swaps.

Variance swap practicalities

- Before considering the pricing of variance swaps, we will mention a few practical issues.
- Variance swaps are now very, very actively traded on stock indices (and sometimes on individual stocks). They are also traded, but less commonly, in other asset classes such as fx.
- There are futures and options contracts on the CBOE VIX index which are now also very actively traded. The VIX index is the market price of a portfolio of vanilla options which (as we will show) replicates future realised variance. Specifically, the VIX index squared, at time t , is (essentially) the risk-neutral conditional time t expectation of the annualised realised variance between time t and time t plus 30 calendar days.
- The prices of vanilla options, variance swaps, VIX futures and VIX options are all closely linked - both practically and theoretically.
- Swaps on volatility are occasionally traded.

”Log-contract” replication approach

- We make the standard assumptions of a market with no-arbitrage as well as continuous and frictionless trading (no transactions costs).
- We assume that the stock price has continuous sample paths i.e. there are no jumps.
- We make no assumptions about the volatility of the stock - it could be constant, deterministic, stochastic with its own source of randomness (stochastic volatility) or, in principle, a function of the stock price (local volatility).
- We assume that the stock price is strictly positive at all times (this, in fact, rules out a Bachelier type arithmetic process with normal volatility so not all local volatility functions are possible - in addition, it, typically, rules out models with default).
- Hence, we write the dynamics of the stock price $S(t) \equiv S$ at time t under the risk-neutral equivalent martingale measure \mathbb{Q} in the form:

$$\frac{dS}{S} = (r - q)dt + \sigma(t, S, \dots)dz,$$

where dz denotes standard Brownian increments and r and q denote the interest-rate and the dividend yield (both assumed constant) respectively.

”Log-contract” replication approach 2

- We want to value a variance swap written at time t_0 , which matures at time T and which has a continuously monitored floating-leg payoff equal to:

$$\frac{1}{(T - t_0)} \int_{s=t_0}^T \sigma^2(s, S, \dots) ds.$$

- We know that the price $V(t_0)$, at time t_0 , of the floating leg of the variance swap is the expected discounted payoff i.e. it is:

$$V(t_0) = E_{t_0}^{\mathbb{Q}}[\exp(-r(T - t_0)) \frac{1}{(T-t_0)} \int_{s=t_0}^T \sigma^2(s, S, \dots) ds] = \exp(-r(T - t_0)) \frac{1}{(T-t_0)} E_{t_0}^{\mathbb{Q}}[\int_{s=t_0}^T \sigma^2(s, S, \dots) ds].$$

- If we apply Ito's lemma, we know:

$$d(\log S) = (r - q - \frac{1}{2}\sigma^2(t, S, \dots))dt + \sigma(t, S, \dots)dz.$$

Eliminating the term $\sigma(t, S, \dots)dz$, implies:

$$\frac{dS}{S} - d(\log S) = \frac{1}{2}\sigma^2(t, S, \dots)dt.$$

Hence, integrating from t_0 to T implies:

$$\frac{1}{2} \int_{s=t_0}^T \sigma^2(s, S, \dots) ds = \int_{t_0}^T \left(\frac{dS(s)}{S(s)} - d(\log S(s)) \right).$$

”Log-contract” replication approach 3

- Note that no expectations have been taken (yet). The last equation says that future realised variance can be captured no matter which path the stock price takes (assuming our assumptions hold - the assumption of no jumps in the stock price is crucial here). Simplifying, we can write:

$$\frac{1}{2} \int_{s=t_0}^T \sigma^2(s, S, \dots) ds = \int_{t_0}^T \frac{dS(s)}{S(s)} - \log(S(T)/S(t_0)).$$

- In the last equation, the term $\int_{t_0}^T \frac{dS(s)}{S(s)}$ is a stochastic integral. Or to put it another way, it is the gain (or loss) from a self-financing trading strategy. What strategy?

”Log-contract” replication approach 4

- It is the trading strategy of holding at all times between t_0 and T a position in $1/S$ units of stock. In other words, at any time t , $t \in [t_0, T]$, hold $1/S(t)$ units of stock. Since one unit of stock is worth $S(t)$, $1/S(t)$ units of stock are worth:

$$(1/S(t))S(t) = 1.$$

- To put it even more simply, the trading strategy is to dynamically trade the stock in such a way that at all times, the value of the position in the stock is worth one unit of account (one dollar, for example).
- Note that it is a dynamic trading strategy - as the stock price changes so does the position. In that respect, it is like delta-hedging where the delta equals $1/S(t)$. The value of the position is always one dollar.

"Log-contract" replication approach 5

- Note:

$$E_{t_0}^{\mathbb{Q}}\left[\int_{s=t_0}^T \frac{dS(s)}{S(s)}\right] = E_{t_0}^{\mathbb{Q}}\left[\int_{t_0}^T (r - q)ds + \int_{t_0}^T \sigma(s, S, \dots)dz(s)\right].$$

The expectation of the second term in square brackets is zero. Hence, the expectation evaluates to $(r - q)(T - t_0)$.

- What we would like to know is the initial (i.e. time t_0) value of the trading strategy. The terminal value (i.e. at time T) is $(r - q)(T - t_0)$. Hence, the initial (i.e. time t_0) value of the trading strategy is $\exp(-r(T - t_0))(r - q)(T - t_0)$.
- If we look at the second term in the equation

$$\frac{1}{2} \int_{s=t_0}^T \sigma^2(s, S, \dots)ds = \int_{t_0}^T \frac{dS(s)}{S(s)} - \log(S(T)/S(t_0)),$$

We see it is a static position in a contract which pays the log of the stock price at time T (normalised by its time t_0 price). In other words, it is a static position in an exotic derivative which we call a log contract. What is the value of the log contract?

”Log-contract” replication approach 6

- The price of the log contract, at time t_0 , is:

$$\exp(-r(T - t_0))E_{t_0}^{\mathbb{Q}}[\log(S(T)/S(t_0))].$$

- In principle, we can calculate this expectation.
- For example, if the stock actually follows geometric Brownian motion with constant volatility σ , then:

$$E_{t_0}^{\mathbb{Q}}[\log(S(T)/S(t_0))] = E_{t_0}^{\mathbb{Q}}[(r - q - \frac{1}{2}\sigma^2)(T - t_0) + \sigma \int_{t_0}^T dz(s)].$$

The expectation of $\sigma \int_{t_0}^T dz(s)$ is clearly zero. Hence, the price of the log contract, at time t_0 , is:

$$\exp(-r(T - t_0))(r - q - \frac{1}{2}\sigma^2)(T - t_0).$$

- On the other hand, this is not very useful. We essentially needed to compute $E_{t_0}^{\mathbb{Q}}[\sigma^2]$ which is essentially what we needed to compute to value the variance swap in the first place. Furthermore, the value of the variance swap is trivial to compute under geometric Brownian motion - the (undiscounted) value of the floating leg is simply σ^2 .
- We can also value the log contract under the Heston (1993) stochastic volatility model in which the instantaneous stochastic variance $\Sigma \equiv \Sigma(t)$ follows the SDE:

"Log-contract" replication approach 7

-

$$d\Sigma = \kappa(\theta - \Sigma)dt + c\sqrt{\Sigma}dz_{\Sigma}, \quad \text{with } \Sigma(t_0) \equiv \Sigma_0,$$

- As an exercise (during the lunch break or at the computing lab) I would like you to prove that:

$$E_{t_0}^{\mathbb{Q}}\left[\int_{s=t_0}^T \Sigma(s)ds\right] = \frac{(\Sigma_0 - \theta)}{\kappa}[1 - \exp(-\kappa(T - t_0))] + \theta(T - t_0).$$

- This immediately gives the value of the variance swap. Why is this an intuitive result? What happens when $T \rightarrow t_0$?
- The last result is dependent on the model (Heston (1993)).
- What would be more interesting to know is, what is the price of the log contract (and hence the variance swap) under our stated assumptions (which apart from assuming no jumps allows for quite a rich specification of dynamics eg. local volatility, stochastic volatility, a combination of the two). Motivation for finding results which are only weakly dependent on the model comes from the fact while the result above is dependent on the model (Heston (1993)), it is not strongly so to the extent that the result above does NOT depend on the volatility of volatility nor on the correlation between the instantaneous variance and the stock price.

”Log-contract” replication approach 8

- A key first-step is the following argument. If a trader has a short position in $2/\delta K^2$ vanilla call options with strike K and a long position in $1/\delta K^2$ vanilla call options with strike $K - \delta K$ and a long position in $1/\delta K^2$ vanilla call options with strike $K + \delta K$ (all the options have the same maturity) where $\delta K > 0$, then, if we let δK tend to zero, the payout at maturity of the trader’s portfolio is the same as that of the Dirac delta function. In words, the payout is zero if the stock price is not equal to K and the payout is $+\infty$ if the stock price equals K at maturity.
- In maths, the Dirac delta function is a building block function - we can make other functions by integrating (summing) Dirac delta functions.
- In mathematical finance, we can replicate any European style (path-independent) payoff by recognising that, since it can be represented as a sum (in practice, infinite sum) of Dirac delta functions, it can be represented as a sum (with possibly negative weights) of vanilla options (not necessarily calls) with different strikes.
- Strictly speaking, the step from the first to the second requires the absence of arbitrage (which we assume throughout) and the existence of a market for vanilla options of all strikes (which, in practice, is only an approximation to reality - we discuss this later).

”Log-contract” replication approach 9

- The following result is key. For any generalized function $f(S)$ and any scalar $\kappa \geq 0$:

$$\begin{aligned} f(S) &= f(\kappa) + f'(\kappa)(S - \kappa) \leftarrow \text{tangent approximation} \\ &+ \int_{\kappa}^{\infty} f''(K)(S - K)^+ dK \leftarrow \text{tangent correction} \\ &+ \int_0^{\kappa} f''(K)(K - S)^+ dK \leftarrow \text{tangent correction.} \end{aligned}$$

- This decomposition may be interpreted as a Taylor series expansion with remainder of the final payoff $f(\cdot)$ about the expansion point κ .
- The first two terms give the tangent to the payoff at κ ; the last two terms continuously bend this tangent so it conforms to the nonlinear payoff.
- The payoff of an arbitrary claim has been decomposed into the payoff from $f(\kappa)$ bonds, $f'(\kappa)$ forward contracts with delivery price κ , $f''(\kappa)dK$ calls struck above κ , and $f''(\kappa)dK$ puts struck below.

”Log-contract” replication approach 10

- The proof is as follows:
- Note that S is non-negative. For any fixed κ , the fundamental theorem of calculus implies:

$$\begin{aligned} f(S) &= f(\kappa) + 1_{S>\kappa} \int_{\kappa}^S f'(u) du + 1_{S<\kappa} \int_{\kappa}^S f'(u) du \\ &= f(\kappa) + 1_{S>\kappa} \int_{\kappa}^S f'(u) du - 1_{S<\kappa} \int_S^{\kappa} f'(u) du \\ &= f(\kappa) + 1_{S>\kappa} \int_{\kappa}^S \left[f'(\kappa) + \int_{\kappa}^u f''(v) dv \right] du \\ &\quad - 1_{S<\kappa} \int_S^{\kappa} \left[f'(\kappa) - \int_u^{\kappa} f''(v) dv \right] du. \end{aligned}$$

- Noting that $f'(\kappa)$ is independent of u , Fubini’s theorem implies:

$$\begin{aligned} f(S) &= f(\kappa) + f'(\kappa)(S - \kappa) + 1_{S>\kappa} \int_{\kappa}^S \int_v^S f''(v) dudv \\ &\quad + 1_{S<\kappa} \int_S^{\kappa} \int_S^v f''(v) dudv. \end{aligned}$$

”Log-contract” replication approach 11

- Integrating over u yields:

$$\begin{aligned} f(S) &= f(\kappa) + f'(\kappa)(S - \kappa) + 1_{S > \kappa} \int_{\kappa}^S f''(v)(S - v)dv \\ &\quad + 1_{S < \kappa} \int_S^{\kappa} f''(v)(v - S)dv \\ &= f(\kappa) + f'(\kappa)(S - \kappa) + \int_{\kappa}^{\infty} f''(v)(S - v)^+ dv \\ &\quad + \int_0^{\kappa} f''(v)(v - S)^+ dv. \end{aligned}$$

- Q.E.D.
- Note the result is completely model independent.

"Log-contract" replication approach 12

- Recall the decomposition of the payoff function $f(S)$:

$$f(S) = f(\kappa) + f'(\kappa)(S - \kappa) + \int_0^\kappa f''(K)(K - S)^+ dK + \int_\kappa^\infty f''(K)(S - K)^+ dK.$$

- No arbitrage implies that the initial (i.e. time t_0) price $V_{t_0}[f(S)]$ of $f(S(T))$, payable at time T , can be expressed in terms of the initial (i.e. time t_0) price $\exp(-r(T - t_0))$ of a bond maturing at time T and the initial prices $C(t_0, K)$ and $P(t_0, K)$ of vanilla calls and puts respectively maturing at time T :

$$V_{t_0}[f(S)] = f(\kappa) \exp(-r(T - t_0)) + f'(\kappa)[C(t_0, \kappa) - P(t_0, \kappa)] + \int_0^\kappa f''(K)P(t_0, K)dK + \int_\kappa^\infty f''(K)C(t_0, K)dK.$$

- When $\kappa = S(t_0) \exp((r - q)(T - t_0)) \equiv F_0$, the forward stock price, the second term vanishes by put-call parity (because $C(t_0, K) - P(t_0, K) = 0$ in this special case), and the initial price decomposes as:

$$V_{t_0}[f(S)] = \underbrace{f(F_0) \exp(-r(T - t_0))}_{\text{intrinsic value}} + \underbrace{\int_0^{F_0} f''(K)P(t_0, K)dK + \int_{F_0}^\infty f''(K)C(t_0, K)dK}_{\text{time value}}.$$

”Log-contract” replication approach 13

- Lets apply our general formula for the special case when $f(S) = \log S$. Then:

$$V_{t_0}[\log S] = \exp(-r(T - t_0)) \log \kappa + \frac{[C(t_0, \kappa) - P(t_0, \kappa)]}{\kappa} - \int_0^\kappa \frac{P(t_0, K)}{K^2} dK - \int_\kappa^\infty \frac{C(t_0, K)}{K^2} dK.$$

- The price of the log contract, at time t_0 , is the last expression minus $\exp(-r(T - t_0)) \log S(t_0)$:
- Note that the term $\frac{[C(t_0, \kappa) - P(t_0, \kappa)]}{\kappa}$ is simply $1/\kappa$ forward contracts struck at κ . It is a static position.
- The term $\exp(-r(T - t_0)) \log \kappa$ (likewise $\exp(-r(T - t_0)) \log S(t_0)$) is simply $\log \kappa$ (likewise $\log S(t_0)$) in cash (or, equivalently, in bonds).
- We have replicated the payoff of a log contract and hence, by no-arbitrage, priced a log contract.
- Note that the log contract is replicated by static positions in bonds and vanilla options (and possibly forward contracts).
- In practice, we have to replace the integrals by discrete summations since vanilla options will not be traded with literally all strikes.

”Log-contract” replication approach 14

- The position in $1/S$ units of stock is a dynamic position and is continuously rebalanced. We gave the value of this position a few slides ago.
- Taking into account both the log contract and the dynamic position in $1/S$ units of stock, we have priced a variance swap.
- The price, at time t_0 , of the floating leg of the variance swap is:

$$\begin{aligned} & \frac{2}{(T - t_0)} (\exp(-r(T - t_0))(r - q)(T - t_0) \\ & + \exp(-r(T - t_0)) \log S(t_0) - \exp(-r(T - t_0)) \log \kappa \\ & - \frac{[C(t_0, \kappa) - P(t_0, \kappa)]}{\kappa} \\ & + \int_0^\kappa \frac{P(t_0, K)}{K^2} dK + \int_\kappa^\infty \frac{C(t_0, K)}{K^2} dK). \end{aligned}$$

- We have focussed on replication but hedging is the same as replication with a minus sign.
- In practice, κ is often set equal to the forward stock price as this generally delineates between whether puts or calls have the greatest liquidity in the market.

”Log-contract” replication approach 15

- In practice, we only have options traded in the market for a discrete set of strikes (rather than a continuum of strikes).
- If we were to ignore all options with strikes outside a particular range (equivalently set the weights to zero), then it is clear from the pricing formula above that we will always price the variance swap at below fair value.
- In practice, there will be a benefit to a trader trading variance swaps in the context of a very large vanilla options book. Options with strikes so high or so low that they have no liquidity today may have been traded when the spot price was much higher or lower in the past and as such may be on the trader’s book. These can be aggregated with the variance swap trades which produces an economy of scale.

"Log-contract" replication approach 16

- With only a discrete set of strikes available, the hedge for the log contract will not be perfect.
- However, there is an easy and intuitive way to account for this.
- The log contract is always a concave function of the stock price. Hence, we can construct chords or tangents which always lie below or above the log contract payoff. We can then solve analytically for the weights for vanilla options which exactly match the chords or tangents. This will perfectly sub-replicate or super-replicate the log contract and at the same time give something akin to a bid-offer spread in the price of the variance swaps. The paper by Demeterfi, Derman, Kamal and Zou (1999, "More than you ever wanted to know about volatility swaps") illustrates this very well.

"Log-contract" replication approach 17

- There is a second equally intuitive way of accounting for a discrete set of strikes:
- Evaluate the log contract at some pre-specified stock prices. Take as given the positions in bonds and forward contracts from the portfolio constructed on the slides above. Then solve for the weights of the call and put options with the available strikes which minimise the sum of the squares of differences between the log contract and the "almost-replicating" portfolio at the pre-specified stock prices.
- Because the "almost-replicating" portfolio is linear in these weights, this problem is easily solvable by Tikhonov regularisation (which means that one only has to invert a matrix - it does NOT involve non-linear least squares fits ("calibration")).
- Furthermore, one can use the weights (dK/K^2) from the slides above as initial guesses in the Tikhonov regularisation.

"Log-contract" replication approach 18

- This second way has the disadvantage of not sub-replicating or super-replicating the log contract. On the other hand, it may give the trader an "almost-replicating" portfolio at lower cost than perfect sub-replication or super-replication. Furthermore, while the "almost-replicating" portfolio will have residual risks, the trader may be content to have these risks in the context of having a view on which parts of the implied volatility surface are cheap or expensive - a view which can also be easily incorporated into the Tikhonov regularisation.

Interview questions

- Let me ask you two questions which you might be asked at job interviews.
- Do the prices of vanilla options depend only on the (risk neutral) distribution of the (log of the) stock price at maturity (as opposed to the (risk neutral) distribution of the (log of the) stock price at any other times)?
- Do vanilla option prices contain information about the prices of any path-dependent derivatives?

Interview questions 2

- The answer to the first question is yes. The prices of vanilla options depend only on the (risk neutral) distribution of the (log of the) stock price at maturity? This is a well-known result.
- Perhaps, initially surprisingly, the answer to the second question is also yes. In fact, we have seen this today: We have priced variance swaps whose payoff is clearly path-dependent. We can price variance swaps in terms of vanilla options. Hence, vanilla option prices do contain information about the prices of path-dependent derivatives, namely, variance swaps.
- To score an additional bonus point, you should mention that this conclusion only holds to the extent that the assumptions that we made hold. The assumptions include that there are no jumps in the stock price which is somewhat restrictive. However, apart from that, the assumptions we have made are actually quite weak.

Discretely monitored variance swaps

- Recall that a (discretely monitored) variance swap has a payoff defined as follows: It is written at time t_0 and matures at time T . The time interval $[t_0, T]$ is partitioned into N time periods t_i , $i = 1, 2, \dots, N$ where $t_N = T$. The time periods do not have to be equal although they are often approximately equal. The payoff of a (discretely monitored) variance swap at time T is:

$$\frac{1}{(T - t_0)} \sum_{i=1}^N ((\log(S(t_i)/S(t_{i-1})))^2 - K^2),$$

where K is a constant (the fixed leg) (usually chosen so that the initial (i.e. time t_0) price of the variance swap is zero).

- We will focus on the floating leg.
- Consider a process for the stock price as follows:

Discretely monitored variance swaps 2

- Under the risk-neutral equivalent martingale measure \mathbb{Q} (which may, in fact, not be unique)

$$S(t) = S(t_0) \exp((r - q)(t - t_0) + X_t),$$

where $X_{t_0} \equiv 0$ and X_t is such that $\mathbb{E}_{t_0}^{\mathbb{Q}}[\exp(X_t)] = 1$ for all $t \geq t_0$. Clearly, $\exp(X_t)$ is a martingale.

- Here r and q are the risk-free rate and the dividend yield which we will assume are constant for notational convenience. However, one nice feature of the methodology we will now discuss is that it is easy to relax this assumption and have either deterministic term-structures or have stochastic interest-rates and dividend yields.
- Actually, the only assumption we need to make is that we have a market with no-arbitrage.

Discretely monitored variance swaps 3

- Introduce z (which may be real or complex). We define:
$$\mathbb{E}_{t_0}^{\mathbb{Q}}[\exp(iz \log(S(t)/S(t_0)))] =$$
$$\mathbb{E}_{t_0}^{\mathbb{Q}}[\exp(iz((r - q)(t - t_0) + X_t))],$$
to be the characteristic function of $\log(S(t)/S(t_0))$.
- Mathematically, the characteristic function is the Fourier Transform of the probability density function of $\log(S(t)/S(t_0))$.
- The characteristic function is known in essentially closed form for many stochastic process including when the stock price follows:
- The Black and Scholes (1973) geometric Brownian motion model, the Heston (1993) stochastic volatility model, the Merton (1976) jump-diffusion model, models of the affine jump-diffusion type (which covers many models with stochastic interest-rates, stochastic interest-rates AND jumps), all Levy process models (see the book by Schoutens (2003), "Levy processes in finance: Pricing financial derivatives" for reading), Levy process models with stochastic time-changing (stochastic time-changing generalises the idea of stochastic volatility) or processes more suitable for other asset classes such as the CEE2 process of Carr and Crosby (2008) or the commodities model of Crosby (2008).

Discretely monitored variance swaps 4

- In fact, the characteristic function is known in essentially closed form for nearly every model used in finance except for local volatility models.
- This is true even though the probability density function is typically not known in closed form.
- Fourier inversion methods can then be used to price vanilla options.

Discretely monitored variance swaps 5

- Recall the sequence of dates at which the variance swap payoff is determined:

$$t_0 < t_1 < \dots < t_{i-1} < t_i < \dots < t_N = T.$$

- Define the extended characteristic function $\Phi(z; t_i, t_{i-1})$ as follows:

$$\begin{aligned}\Phi(z; t_i, t_{i-1}) &\equiv \mathbb{E}_{t_0}^{\mathbb{Q}}[\exp(iz[(r - q)(t_i - t_{i-1}) + X_{t_i} - X_{t_{i-1}}])] \\ &= \mathbb{E}_{t_0}^{\mathbb{Q}}[\exp(iz \log(S(t_i)/S(t_{i-1})))].\end{aligned}$$

- Essentially any model which has an analytic characteristic function also has an analytic extended characteristic function.

- Then note:

$$\frac{\partial^2 \Phi(z; t_i, t_{i-1})}{\partial z^2} = -\mathbb{E}_{t_0}^{\mathbb{Q}}[(\log(S(t_i)/S(t_{i-1})))^2 \exp(iz[(r - q)(t_i - t_{i-1}) + X_{t_i} - X_{t_{i-1}}])].$$

- Hence, evaluating the last equation at $z = 0$:

$$\frac{\partial^2 \Phi(0; t_i, t_{i-1})}{\partial z^2} = -\mathbb{E}_{t_0}^{\mathbb{Q}}[(\log(S(t_i)/S(t_{i-1})))^2]$$

Discretely monitored variance swaps 6

- The price of any derivative is the expected discounted payoff.
- Hence, the price, at time t_0 , of the floating leg of the variance swap is:

$$\begin{aligned} & \mathbb{E}_{t_0}^{\mathbb{Q}}\left[\exp(-r(T - t_0)) \frac{1}{(T - t_0)} \sum_{i=1}^N [\log(S(t_i)/S(t_{i-1}))]^2\right] \\ &= -\frac{1}{(T - t_0)} \exp(-r(T - t_0)) \sum_{i=1}^N \frac{\partial^2 \Phi(0; t_i, t_{i-1})}{\partial z^2}. \end{aligned}$$

- But we know $\Phi(z; t_i, t_{i-1})$ and hence $\frac{\partial^2 \Phi(0; t_i, t_{i-1})}{\partial z^2}$ in essentially closed form for many models. Hence, we can price the variance swap.

Discretely monitored variance swaps 7

- This methodology is very generic and can be used for almost all stochastic processes that have been used in mathematical finance (with the exception of local volatility models because neither the characteristic function nor the extended characteristic function are known).
- The disadvantage of this methodology is that it says nothing about hedging.
- There's no doubting this is a big practical disadvantage. However, one would typically be interested to use this methodology when there are jumps in the stock price process. In this case, the market is incomplete and hence perfect hedging or replication is not possible anyway.
- In practice, one would choose a stochastic process. Then one calibrates the parameters of the stochastic process by finding those parameter values which minimise the sum of squares of differences between the market prices and model prices of vanilla options. Using these parameters, one can then price variance swaps using the formula on the last slide.

Discretely monitored variance swaps 8

- The methodology allows us to highlight some features of variance swaps.
- Question: Is a continuously monitored variance swap worth more or less than a discretely sampled variance swap? To answer this question, we will answer a slightly more generic question first.
- Consider two variance swaps based on the realised variance observed between t_0 and T . The times at which the stock price is observed to compute the payoff are equally spaced (ie $t_i - t_{i-1}$ is the same for all i). The difference is that for the first variance swap, the number of monitoring times is N_1 , and for the second variance swap, the number of monitoring times is N_2 , with $N_2 = 2N_1$. Which variance swap is worth more?
- We assume that the “extra” monitoring times of the second variance swap lie exactly in the middle of the intervals between the monitoring times of the first variance swap and that the “other” monitoring times of the second variance swap coincide with those of the first variance swap.

Discretely monitored variance swaps 9

- The payoffs of the (floating legs of the) variance swaps are:

$$\frac{1}{(T - t_0)} \sum_{i=1}^{N_1} (\log(S(t_i)/S(t_{i-1})))^2, \quad \frac{1}{(T - t_0)} \sum_{j=1}^{N_2} (\log(S(t_j)/S(t_{j-1})))^2,$$

respectively.

- The answer to our question is clearly going to be somewhat dependent on the stochastic process X_t .
- Suppose X_t is a Levy process (a process with stationary and independent increments eg. Brownian motion).
- It is not difficult to see that the extended characteristic function for a Levy process is of the form:

$$\begin{aligned} & \mathbb{E}_{t_0}^{\mathbb{Q}}[\exp(iz[(r - q)(t_i - t_{i-1}) + X_{t_i} - X_{t_{i-1}}])] \\ &= \exp((t_i - t_{i-1})[iz(r - q) + \psi(z) - iz\psi(-i)]), \end{aligned}$$

for some function $\psi(z)$, independent of t_{i-1} and t_i .

- For example, if the Levy process is Brownian motion with volatility σ , then $\psi(z) = -\sigma^2 z^2/2$.

Discretely monitored variance swaps 10

- Applying our formula, we have that the price, at time t_0 , of the floating leg of the first variance swap is:

$$\begin{aligned} & \exp(-r(T - t_0))[-\psi''(0)] \\ & + \frac{\exp(-r(T - t_0))}{(T - t_0)} \sum_{i=1}^{N_1} [-(\psi'(0) - i\psi(-i) + i(r - q))^2 (t_i - t_{i-1})^2], \end{aligned}$$

with a similar expression for the second variance swap.

Note that in the last expression we used the result:

$$\frac{1}{(T - t_0)} \sum_{i=1}^{N_1} (t_i - t_{i-1}) = 1.$$

- The first term is independent of the monitoring frequency N_1 but the second term is not.
- Note that $[-\psi''(0)]$ is non-negative (actually strictly positive except in degenerate cases) (because it is the expectation of a non-negative quantity) and $[-(\psi'(0) - i\psi(-i) + i(r - q))^2 (t_i - t_{i-1})^2]$ is non-negative (actually strictly positive except in a special case) because it can be shown that the quantity $(\psi'(0) - i\psi(-i) + i(r - q))$ is imaginary with zero real part.

Discretely monitored variance swaps 11

- We can see that the second term gets larger when N_1 gets smaller (because $t_i - t_{i-1}$ gets larger).
- This means that, because $N_2 = 2N_1$, the price of the (floating leg of the) first variance swap is always greater than or equal to the price of the (floating leg of the) second variance swap. Furthermore, equality only occurs in the special case that the imaginary part of $(\psi'(0) - i\psi(-i) + i(r - q))$ is zero (the real part is always zero).
- However, this term is essentially the drift term of $\log(S(t))$ (multiplied by i). So we conclude that equality occurs only when the drift of $\log(S(t))$ is exactly equal to zero.
- Note also that if the payoff of the variance swap were to subtract the square of the mean, then this term would always be cancelled out and we could conclude that the price of a variance swap on a Levy process would be completely independent of the monitoring frequency.

Discretely monitored variance swaps 12

- In any event, the second term

$$\frac{\exp(-r(T - t_0))}{(T - t_0)} \sum_{i=1}^{N_1} [-(\psi'(0) - i\psi(-i) + i(r - q))^2 (t_i - t_{i-1})^2]$$

will typically be tiny compared to the first term

$$\exp(-r(T - t_0))[-\psi''(0)].$$

- For example, if the Levy process is Brownian motion with volatility $\sigma = 0.2$ and $r = 0.03$, $q = 0$, $T - t_0 = 1$, $N_1 = 252$ (which corresponds to daily monitoring of a one year swap), then the second term is less than one ten thousandth of the first term. (Note: As an exercise (during the lunch break or at the computing lab) I would like you to prove this mathematically). This means that the second term is completely negligible (especially relative to the likely bid-offer spread - typically around 0.5 to 1.0 percentage points).
- This suggests that, although it is true that for $N_2 = 2N_1$, the price of the (floating leg of the) first variance swap is always greater than or equal to the price of the (floating leg of the) second variance swap, in practice (for daily monitoring, say), any difference between the two will typically be very small - at least for processes with stationary and independent increments (i.e. Levy processes).

Discretely monitored variance swaps 13

- Further intuition can be gleaned by considering, firstly, a two year variance swap with only one monitoring date and, secondly, a two year variance swap with two monitoring dates at year one and at year two. The price of the first involves the expectation of $[\log S(2) - \log S(0)]^2$ and the price of the second involves the expectation of $[\log S(1) - \log S(0)]^2 + [\log S(2) - \log S(1)]^2$.
- Straightforward algebra shows the first quantity is greater than the second if, and only if, the expectation of $2[\log S(2) - \log S(1)][\log S(1) - \log S(0)]$ is positive.
- If $\log S(t)$ has zero drift, this expectation is identically equal to zero for a process with independent increments (by definition). Hence, we see again that the two variance swaps have the same price in this special case.
- For processes that have neither stationary nor independent increments such as in the model of Heston (1993), this expectation will (typically) be positive (this is true even if, somehow, $\log S(t)$ has zero drift). Hence, for such processes, the prices of variance swaps may be much more sensitive to the monitoring frequency.
- In the Heston (1993) model, the prices of variance swaps will be most sensitive to the monitoring frequency when the mean reversion rate is large and when the correlation is far from zero.

Discretely monitored variance swaps 14

- The answer to the question "Is a continuously monitored variance swap worth more or less than a discretely sampled variance swap?" is obtained by letting the number of monitoring times tend to infinity in our arguments above:
- The price of a discretely sampled variance swap is greater than or equal to the price of a continuously monitored variance swap.
- Strict equality will only hold under special circumstances.
- However, generally speaking, in practice, any differences will be small.
- As a final comment, we note that, in the limit that $t_i - t_{i-1} \rightarrow 0$, for all i , i.e. for a continuously monitored variance swap, under a Levy process, we have that the price, at time t_0 , of the floating leg of the variance swap tends to:

$$\exp(-r(T - t_0))[-\psi''(0)].$$

This is because the second term (see two slides ago) tends to zero.

- As a sanity check on the last formula, for the case of Brownian motion with volatility σ , $\psi(z) = -\sigma^2 z^2/2$. Hence, $-\psi''(0) = \sigma^2$, which agrees with our intuition.

From variance swaps to volatility swaps

- Volatility swaps also trade - although less frequently. The payoff of a (discretely monitored) volatility swap is:

$$\frac{1}{(T - t_0)} \sum_{i=1}^N \left(\sqrt{(\log(S(t_i)/S(t_{i-1})))^2} - K_v \right),$$

where K_v is a constant (the fixed leg) (again usually chosen so that the initial price of the volatility swap is zero).

- Is there a simple, if approximate, way to relate volatility swap rates to variance swap rates?
- Suppose that future realised variance $V(T)$ has (under the risk neutral measure \mathbb{Q}) mean μ_V and variance Σ_V^2 . In other words, μ_V is the fixed rate on a variance swap with zero initial price.

$$\mu_V = \mathbb{E}_{t_0}^{\mathbb{Q}}[V(T)] \quad \text{and} \quad \Sigma_V^2 = \mathbb{E}_{t_0}^{\mathbb{Q}}[(V(T) - \mu_V)^2].$$

- Doing a Taylor series expansion of $\sqrt{V(T)}$ around its mean implies (correct to second order):

$$\sqrt{V(T)} = \sqrt{\mu_V} + \frac{(V(T) - \mu_V)}{2\sqrt{\mu_V}} - \frac{(V(T) - \mu_V)^2}{8\mu_V^{3/2}}.$$

From variance swaps to volatility swaps 2

- Taking expectations under \mathbb{Q} and, observing that $\mathbb{E}_{t_0}^{\mathbb{Q}}[(V(T) - \mu_V)] = 0$ and that $\sqrt{\mu_V} = \sqrt{\mathbb{E}_{t_0}^{\mathbb{Q}}[V(T)]}$, implies that (correct to second order):

$$\mathbb{E}_{t_0}^{\mathbb{Q}}[\sqrt{V(T)}] = \sqrt{\mathbb{E}_{t_0}^{\mathbb{Q}}[V(T)]} - \frac{\Sigma_V^2}{8\mu_V^{3/2}}.$$

- The term on the left is the fixed rate on a volatility swap such that it has zero initial price. The first term on the right is the square root of the fixed rate on a variance swap such that it has zero initial price.
- Note that the former ($\mathbb{E}_{t_0}^{\mathbb{Q}}[\sqrt{V(T)}]$) is certainly less than or equal to the latter ($\sqrt{\mu_V} = \sqrt{\mathbb{E}_{t_0}^{\mathbb{Q}}[V(T)]}$) (with equality only in the degenerate case that $\Sigma_V^2 = 0$). This is to be expected from Jensen's inequality.
- We stress the last result is only an approximation.

Are implied volatilities predictions of future realised volatilities

- One occasionally hears it said that implied volatilities are the market's best guesses of future realised volatilities.
- Is this true? What does it mean (if anything)?
- Consider a stock price process of the form:

$$\frac{dS}{S} = (r - q)dt + \sigma(t, S, \dots)dz,$$

where $\sigma(t, S, \dots)$ might be stochastic but, if it is stochastic, it is independent of dz .

- We consider the price, at time t_0 , of a vanilla (standard European) option, maturing at time T , which is struck at the forward price $F(t_0) \equiv S(t_0) \exp((r - q)(T - t_0))$. We denote realised volatility, over the time period $[t_0, T]$, by $RV(t_0, T)$.

$$RV(t_0, T) = \sqrt{\frac{1}{(T - t_0)} \int_{t_0}^T \sigma(s, S, \dots)^2 ds}.$$

Are implied volatilities predictions of future realised volatilities 2

- Since the volatility is independent of the stock price (this is a key part of the argument), we can compute the price of the vanilla option by, firstly, conditioning on the realised volatility and using the Black and Scholes (1973) formula and then, secondly, taking expectations over the realised volatility (in other words, by using the tower law i.e. the law of iterated expectations).

- Hence, the price, at time t_0 , of the vanilla option is:

$$\mathbb{E}_{t_0}^{\mathbb{Q}}[\exp(-r(T - t_0))[F(t_0)N(RV(t_0, T)\sqrt{T - t_0}/2) - F(t_0)N(-RV(t_0, T)\sqrt{T - t_0}/2)]].$$

- Suppose the option maturity $T - t_0$ is very small. We can do a Taylor series expansion of the term in the inner square brackets to deduce that the price, at time t_0 , of the vanilla option is approximately (correct to terms in $(T - t_0)$):

$$\begin{aligned} & \mathbb{E}_{t_0}^{\mathbb{Q}}[\exp(-r(T - t_0))[F(t_0)RV(t_0, T)\sqrt{T - t_0}/\sqrt{2\pi}]] \\ &= \frac{\exp(-r(T - t_0))F(t_0)\sqrt{T - t_0}}{\sqrt{2\pi}}\mathbb{E}_{t_0}^{\mathbb{Q}}[RV(t_0, T)]. \end{aligned}$$

Are implied volatilities predictions of future realised volatilities 3

- On the other hand, we can compute the price of the vanilla option using the implied volatility appropriate for an at-the-money-forward strike and a maturity of $T - t_0$. This is simply the Black and Scholes price with implied volatility $IV(t_0, T)$. We can do the same Taylor series expansion for small $T - t_0$ to conclude the price of the option is approximately (correct to terms in $(T - t_0)$):

$$\frac{\exp(-r(T - t_0))F(t_0)\sqrt{T - t_0}}{\sqrt{2\pi}}IV(t_0, T).$$

- If we equate these two vanilla option prices and cancel terms, we obtain:

$$\mathbb{E}_{t_0}^{\mathbb{Q}}[RV(t_0, T)] = IV(t_0, T).$$

- Hence, we see that the risk-neutral expectation of future realised volatility, over a short time period, is approximately (correct to terms in $(T - t_0)$) equal to the at-the-money-forward implied volatility of options maturing at the end of the short time period.

Are implied volatilities predictions of future realised volatilities 4

- So the claim that implied volatilities are the market's best guesses of future realised volatilities is true - at least for very short time periods.
- Or is it?
- Carr and Wu (2006) show that the sample average difference between the 30-day realized variance on the S+P 500 and the VIX squared is more than -150 bp and highly significant. The variance risk premium in excess returns form is -40 per cent, for being long a 30-day variance swap and holding it to maturity. In other words, shorting variance swaps and hence receiving the fixed leg generates positive excess returns on average.
- Does this contradict our result on the last slide?
- No. As Carr and Wu (2006) point out, the highly negative variance risk premium indicates that investors are averse to variance risk and the compensation for bearing variance risk can come in the form of a lower mean variance level under the real world empirical measure than under the risk neutral measure \mathbb{Q} .

Summary and General Conclusions

- Variance swaps can be priced and hedged or replicated by synthetically creating log contracts.
- They are very actively traded in the markets as are futures and options on the CBOE VIX index. The VIX index is the market price of a portfolio of vanilla options which has weights derived from those required to replicate log contracts.
- The extended characteristic function approach prices discretely monitored variance swaps. It is very simple and generic but it has the disadvantage that it says nothing about hedging or replicating variance swaps.

References

- Trading variance and log contracts was introduced by Anthony Neuberger (Neuberger A. (1990) "Volatility trading" Working paper, London Business School; Neuberger, A. (1994) "The Log Contract: A new instrument to hedge volatility", Journal of Portfolio Management, Winter, p74-80; Neuberger, A. (1996) "The Log Contract and Other Power Contracts", in The Handbook of Exotic Options, edited by I. Nelken, p200-212).
- The paper by Kresimir Demeterfi, Emanuel Derman, Michael Kamal and Joseph Zou (Demeterfi K., Derman E., Kamal M. and Zou J. (1999) "More than you ever wanted to know about volatility swaps" Journal of Derivatives 6(4), p 9-32; also a Goldman Sachs Quantitative Strategies Note available on Emanuel Derman's website <http://www.ederman.com>) is an excellent and very readable article.
- A paper by Peter Carr and Liuren Wu (Carr P. and L. Wu (2006) "A Tale of Two Indices", Journal of Derivatives, 13(3), p13-29) examines VIX futures and options in depth.
- The extended characteristic function approach can be found in a seminar presentation given by George Hong of UBS at Cambridge University in 2004. (Hong G. (2004) "Forward Smile and Derivative Pricing" Summer 2004, available on the website of the Centre for Financial Research, Judge Business School, Cambridge University).