Variance Swaps

- **Variance swaps offer straightforward and direct exposure to the volatility of an underlying**

- **Returns from variance swaps can act as a diversifying asset**

- **Variance swaps can be used for hedging volatility exposures or generating alpha**

Overview

In this note we discuss the variance swap market, mechanics, pricing and uses.

Variance swaps offer straightforward and direct exposure to the volatility of an underlying asset. They are liquid across major equity indices and large cap stocks, and increasingly across emerging market indices and other asset classes.

Major uses include taking a volatility view, diversifying returns, hedging and relative value trading. Variance swaps can also be used to trade forward volatility and correlation.

Variance swaps can be replicated by a delta-hedged portfolio of vanilla options, so that pricing reflects volatilities across the entire skew surface. In practice this means that variance swaps trade at a small premium to ATM implied volatilities.

Variance swap cash-flows
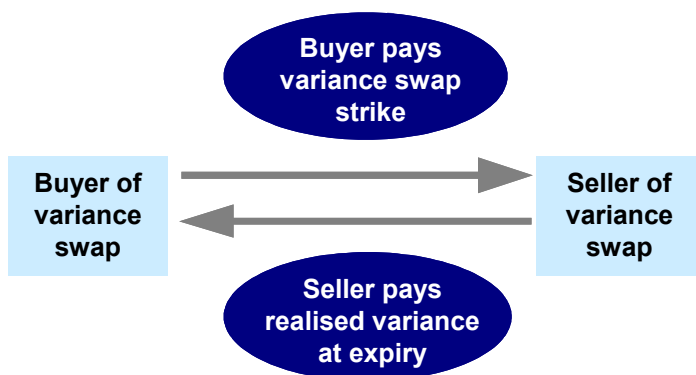
Source: JPMorgan

Peter Allen
(44-20) 7325-4114
peter.allen@jpmorgan.com

Stephen Einchcomb
(44-20) 7325-9064
stephen.einchcomb@jpmorgan.com

Nicolas Granger[AC]
(44-20) 7325-7033
nicolas.m.granger@jpmorgan.com

The certifying analyst is indicated by an [AC]. See page 102 for analyst certification and important legal and regulatory disclosures.

www.morganmarkets.com

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

JPMorgan

Table of Contents

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Introduction

Variance swaps **are instruments which offer investors straightforward and direct exposure to the volatility of an underlying asset** such as a stock or index. They are swap contracts where the parties agree to exchange a pre-agreed variance level for the actual amount of variance realised over a period. Variance swaps offer investors a means of achieving direct exposure to realised variance without the path-dependency issues associated with delta-hedged options.

Buying a variance swap is like being **long volatility** at the strike level; if the market delivers more than implied by the strike of the option, you are in profit, and if the market delivers less, you are in loss. Similarly, selling a variance swap is like being short volatility. However variance swaps are *convex* **in volatility**: a long position profits more from an increase in volatility than it loses from a corresponding decrease. For this reason variance swaps normally trade above ATM volatility.

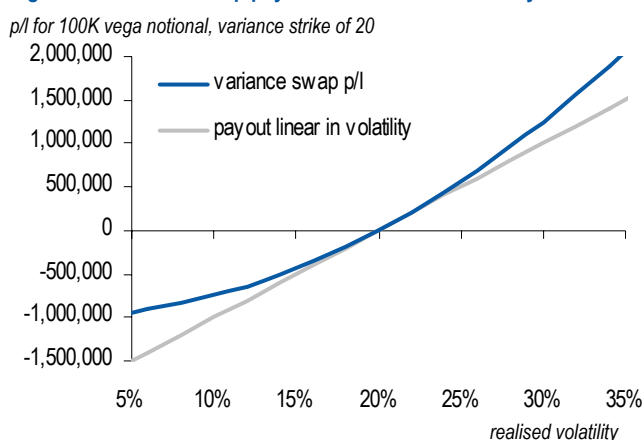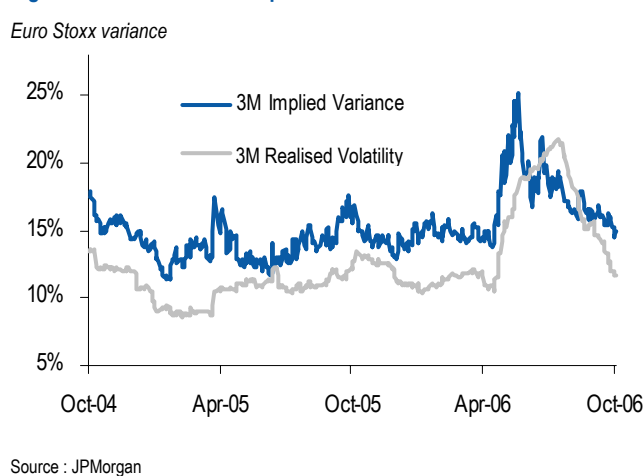**Figure 1 : Variance swap payoffs are convex in volatility**

p/l for 100K vega notional, variance strike of 20

Source : JPMorgan

Figure 2 : Euro Stoxx 50 implied and realised variance levels

Euro Stoxx variance

Source : JPMorgan

Market development

Variance swap contracts were first mentioned in the 1990's, but like vanilla options only really took off following the development of robust pricing models through replication arguments. The directness of the exposure to volatility and the relative ease of replication through a static portfolio of options make variance swaps attractive instruments for investors and market-makers alike.

The variance swap market has grown steadily in recent years, driven by investor demand to take direct volatility exposure without the cost and complexity of managing and delta-hedging a vanilla options position. Although it is possible to achieve variance swap payoffs using a portfolio of options, the variance swap contract offers a convenient package bundled with the necessary delta-hedging, offering **investors a simple and direct exposure to volatility**, without any of the path dependency issues associated with delta-hedging an option.

Variance swaps were initially developed on index underlyings. In Europe, variance swaps on the Euro Stoxx 50 index are by far the most liquid, but DAX and FTSE are also frequently traded. Variance swaps are also tradable on the more liquid stock underlyings – especially Euro Stoxx 50 constituents, allowing for the construction of variance dispersion trades. Variance swaps are tradable on a range of indices across developed markets and increasingly also on developing markets.

Bid/offer spreads have come in significantly over recent years and in Europe they are now typically in the region of 0.5 vegas for indices and 1–2 vegas for single-stocks – although the latter vary according to liquidity factors. We expect spreads to come in further in future as liquidity continues to improve. Spreads for indices in the US and Japan are similar with single stock bid/offers in the 2 – 2.5 vegas range. Spreads are naturally higher in emerging markets although these too are becoming more liquid. With volatilities often much higher and less stable than in developed markets, emerging market variance can offer interesting opportunities.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

The most liquid variance swap maturities are generally from 3 months to around 2 years, although indices and more liquid stocks have variance swaps trading out to 3 or even 5 years and beyond. The most liquid maturities generally coincide with the quarterly options expiry dates, meaning that they can be efficiently hedged with exchange-traded options of the same maturity. Good liquidity is also seen in the front 3 months for short-dated index variance. The **VIX, VSTOXX and VDAX indices** represent the theoretical prices of 1-month variance swaps on the S&P500, Euro Stoxx and DAX indices respectively, and are calculated by the exchanges from listed option prices, interpolating to get 1-month maturity. These volatility indices are widely used as benchmark measures of equity market risk, even though they are only short-dated measures and are not directly tradable.

Uses of variance swaps

Increasingly, investors have come to **view volatility itself as an asset class**, one that can **diversify** investment returns or **hedge** unwelcome investment scenarios. For example, volatility typically rises in a bear market, so holding a long volatility position (e.g. being long variance swaps) can help to hedge an equity portfolio. However, as with any asset, the path to riches is to buy when cheap and sell when expensive. And here investors can look to exploit the mean-reverting nature of volatility to add alpha to portfolios. For example, systematically selling volatility has historically been profitable, with returns akin to selling credit protection.

Variance swaps can be traded outright or spread against other products to create a wide variety of exposures. Below we list some common uses.

- **Exploiting a volatility view:** Variance swaps are ideal for taking a direct view on the volatility of an underlying without the path-dependency issues of a delta-hedged option.

- **Specific hedging purposes:** Variance swaps can be used for macro-hedging and also for hedging specific volatility exposures, such as that resulting from structured products or life assurance policies.

- **Rolling short variance:** Short variance swaps can be used to capture the observed equity index volatility risk premium. Rolling short index variance is an attractive systematic volatility strategy from a risk-return perspective.

- **Diversification:** Volatility can be thought of as an asset class in its own right, and as such can act to diversify returns within a portfolio.

- **Index variance spreads:** Variance swaps can be used to trade the spread of volatilities between two indices. Such trades can be thought of as either relative value volatility trades or as 'volatility-beta' trades aiming to profit from a spread of volatilities widening as volatility increases.

- **Relative value single-stock volatility:** Use volatility pairs, or cross-sectional regression volatility models to find rich/cheap single-stock volatilities

- **Variance dispersion and correlation trading**: Trading variance swaps on an index against variance swaps on its constituents provides exposure to equity correlation.

- **Forward variance and volatility spikes:** Long forward volatility can avoid potentially negative carry at the cost of slide down the term structure, and can be a useful way of positioning for volatility spikes.

- **Trading the variance term structure:** Variance swaps can be used to trade the shape of the variance term structure, analogous to the use of spreads and butterflies in fixed income to trade the shape of the yield curve.

- **Skew and convexity trades:** Variance swaps are long skew and convexity. Trading variance against (delta-hedged) vanilla options provides interesting exposures to skew and/or convexity.

- **Cross asset class trades:** Equity Volatility and credit spreads are correlated, both being measures of corporate risk. Variance swaps are useful instruments in debt/equity trades, either at the index or single name level.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Outline of the document

Part 1: Variance Swap Mechanics

We consider the **mechanics of variance swaps**. We begin by defining **realised volatility** and consider some of its commonly observed properties. We go on to define a variance swap, including a sample term sheet and specifications, and look at the different measures of notional: *vega* **and** *variance* **notionals**. Variance swap **mark-to-markets** are easily calculated, since variance is additive. This leads to the simple valuation of **forward-starting variance swaps**.

Part 2: The Variance Swap Market

We look at the variance swap **market development**, the **pricing of variance swaps** (including rules of thumb) and variance swap drivers. Variance swap strikes are well correlated with standard Black-Scholes implied volatility derived from options prices. This should not be surprising, as both can be interpreted as market estimates of future volatility. Variance swaps trade slightly above ATM volatility levels, essentially due to the fact that variance swaps are long convexity (or volatility of volatility).

Realised volatility and risk-aversion are major **drivers of variance swap levels**, with structured product flows impacting the long end of the maturity curve. Short-dated variance swaps tend to do a relatively good job **of predicting future realised volatility**, at least better than using previous realised.

Part 3: Uses of Variance Swaps

We discuss the **uses of variance swap**s, ranging from simple expressions of a **volatility view**, to **relative value volatility** trades, to more complex trades **using variance to trade correlation, skew, convexity** etc. One important use of variance swaps, at least at the index level, is to exploit the **volatility risk premium** by systematically rolling a short variance swap position. The volatility risk premium reflects the observed tendency of implied volatility to trade above realised volatility. Even though returns from a short variance or volatility position have more downside than upside risk, historically investors have been more than compensated for this risk. Such strategies are now being incorporated into structured products such as **JPMorgan's Yield Alpha** product.

Part 4: Replication and Hedging

Part 4 is more technical in nature. We outline the mechanism whereby **a variance swap can be replicated by a portfolio of options**. The main drawback of a delta-hedged option as a vehicle for trading volatility is that the dollar-gamma, which characterises the exposure to realised volatility, is **path dependent.** It is possible to remove this path dependency by trading a carefully chosen portfolio of options of different strikes. A variance swap consists of just this portfolio, bundled up with the necessary delta-hedging to give direct exposure to realised volatility. We give a concrete example of such a portfolio and consider some practical issues arising from variance swap hedging.

In particular we discuss **the potential impact of variance swap hedging on the underlying market**, driven by the need of replicating option portfolios to be delta-hedged on the close (unlike other options). Depending on market positioning, such hedging activity may exacerbate large moves into the close, potentially adding to market volatility. Hedging flows, and anticipation thereof, have been much talked about, although in an efficient market such flows should be propagated back throughout the trading day.

Part 5: Future Developments

We discuss **3rd generation variance swap products and future developments**. Commonly watched volatility indices such as VIX and VSTOXX are in fact calculated as variance swap strikes, and futures exchanges have listed futures on these indices. As yet there is no official variance swap fixing mirroring deposit or swap rate fixings. **Conditional variance swaps** have also increased in liquidity, and **variance swaps on other asset classes** are starting to be traded. We expect variance swap markets to continue to gain liquidity, fuelled by investor demand to trade pure volatility and by the desire of investment banks to recycle their exotic and structured product risks.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

JPMorgan

1. Variance Swap Mechanics

- 1.1   Realised volatility
- 1.2   The variance swap contract
- 1.3   Vega notional and variance notional
- 1.4   Variance swap convexity
- 1.5   Variance swap example
- 1.6   Variance swap mark-to-market
- 1.7   Forward variance
- 1.8   Variance swap contract specifications
- 1.9   Example variance swap term sheet

Part 1: Variance Swap Mechanics

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Part 1: Variance Swap Mechanics

Variance swaps are instruments which offer investors straightforward and direct exposure to the volatility of an underlying asset such as a stock or index. They are swap contracts where the parties agree to exchange a pre-agreed variance level for the actual amount of variance realised over a period (variance is the square of volatility). Variance swaps offer investors a means of achieving direct exposure to realised variance without the path-dependency issues associated with delta-hedged options.

Before discussing the mechanics of the variance swap contract we begin by defining the notion of **realised volatility**, and look at some of its characteristics. We then define a **variance swap** and describe some common conventions such as the means of **expressing trade notionals**. A variance swap pays out linearly in terms of the realised variance, but when considered in terms of volatility (as it more commonly is) the payout is **convex**. We give a concrete example of a variance swap trade to illustrate how the p/l from realised variance accrues over time. This naturally leads into a discussion of how variance swaps are marked-to-market, through changes both to implied and realised variance. The linearity of variance which makes **marking-to-market** so straightforward also allows us to construct exposure to forward variance, through **forward starting variance swaps** which offer an interesting alternative to the usual spot variance exposure. Finally we consider some technicalities associated with variance swap contracts, such as the treatment of dividends and M&A events, and conclude Part 1 by giving an example of a variance swap term sheet.

1.1: Realised volatility

Volatility measures the variability of returns of an underlying asset and in some sense provides a measure of the **risk of holding that underlying**. In this note we are concerned with the volatility of equities and equity indices, although much of the discussion could apply to the volatility of other underlying assets such as credit, fixed-income, FX and commodities.

Figure 3 shows the history of realised volatility on the Dow Jones Industrial Average over the last 100 years. Periods of higher volatility can be observed, e.g. in the early 1930's as a result of the Great Depression, and to a lesser extent around 2000 with the build-up and unwind of the dot-com bubble. Also noticeable is the effect of the 1987 crash, mostly due to an exceptionally large single day move, as well as numerous smaller volatility spikes.

Figure 3 : Long term history of realised volatility
3-month realised volatility (Dow Jones Industrial Average)
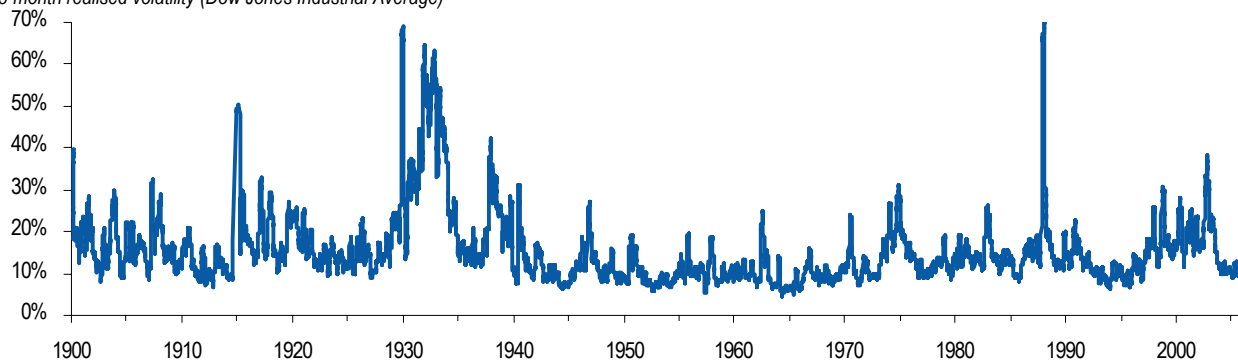
Source: JPMorgan

The following are some of the commonly observed properties of (equity market) volatility:

- Volatility tends to be **anti-correlated with the underlying** over short time periods (Figure 4)
- Volatility can **increase suddenly in 'spikes'**
- Volatility can be observed to **experience different regimes** (Figure 5)
- Volatility tends to be **mean reverting** (within regimes)

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 1: Variance Swap Mechanics

This list suggests some of the reasons why investors may wish to trade volatility: as a partial hedge against the underlying – especially for a volatility spike caused by a sudden market sell-off; as a diversifying asset class; to take a macro view e.g. for a potential change in volatility regime; or to trade a spread of volatility between related instruments.

Figure 4 : in the short term volatility is anti-correlated with index level
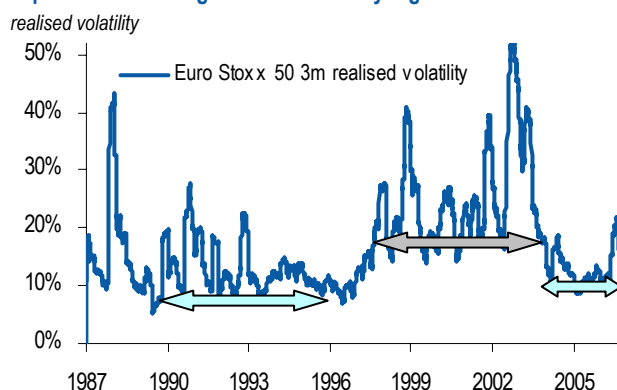

Source : JPMorgan

Figure 5 : 20 years of Euro Stoxx realised volatility: the index has experienced both high and low volatility regimes

Source : JPMorgan

Measuring realised volatility and variance

Realised (or historical) volatility can be measured in a variety of ways. One of the simplest measures is the annualised standard deviation of daily (log) returns of an underlying, calculated from a data set spanning some fixed period of time – for example 1-month historical volatility is calculated as the annualised standard deviation of daily returns occurring in one month. This type of measure has the advantage of being very easy to understand and calculate and it is this which we used throughout this note.

One issue with using this fixed time-period volatility is that the realised volatility calculated in this way can suddenly drop as a large magnitude return falls out of the window of observation. Other, more sophisticated, measures of realised volatility exist, which seek to alleviate this problem and provide a potentially better means of forecasting future volatility – for example by giving more weight to recent returns (exponentially weighted measures) and/or attempting to model the mean-reverting and regime changing nature of volatility (GARCH models). See *"Options, Futures and Other Derivatives", J. Hull,* for further details.

Volatility can also be calculated using returns of various frequencies (weekly, daily, intra-day returns etc). If the returns are truly independent, then, over a long enough time-frame, the realised volatility should be independent of the frequency of the returns used to calculate it. However, in practice, volatility is usually calculated using daily close-to-close returns. This is almost always the type of volatility which dictates the payoff of a variance swap contract.

Volatility is usually quoted as an *annualised* volatility and expressed as a percentage. The annualised volatility represents the expected annual standard deviation in distribution of the underlying, assuming returns are independent. To move from the standard deviation of daily log returns to annualised volatility, simply multiply by the square-root of the number of trading days in the year, usually taken to be 252. Note that since the square-root of 252 is about 16, a volatility of 16% equates to a standard deviation of daily moves of around 1%.

Volatility can also be defined by an RMS (root-mean-squared) measure, which is like a standard deviation but assuming a zero mean (Box 1). This simplifies calculations, and in practice, the difference from the traditional standard deviation is very small, for all but the shortest-dated volatility (Figure 6, Figure 7). The average return is the drift, which should be near zero on a daily basis. Furthermore, **it is this RMS volatility measure which is used to define the payout of the variance swap contract.**

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Box 1: Definition of realised (RMS) volatility

Volatility, σ, is defined by:

$$\sigma^2 = \frac{252}{T}\sum_{i=1}^{T}\left[\ln\left(\frac{S_i}{S_{i-1}}\right)\right]^2$$

where $\quad S_i$ is the stock price on day i, and

$\quad\quad\quad\quad T$ is the number of days

Table 1: Calculating realised volatility (RMS)

	Index level	Log return	Log return squared
13-Oct-05	3331.4		
14-Oct-05	3349.6	0.5%	0.000030
17-Oct-05	3356.8	0.2%	0.000005
18-Oct-05	3334.8	-0.7%	0.000043
19-Oct-05	3279.6	-1.7%	0.000278
20-Oct-05	3284.8	0.2%	0.000002
21-Oct-05	3271.0	-0.4%	0.000018
24-Oct-05	3310.1	1.2%	0.000141
25-Oct-05	3292.5	-0.5%	0.000028
26-Oct-05	3304.3	0.4%	0.000013
27-Oct-05	3241.1	-1.9%	0.000372
Variance of log returns			0.000093
Annualised variance			0.0234394
Annualised volatility (RMS)			15.3%

Average the squared log returns

Annualise (x 252)

Take square root

Source : JPMorgan

Volatility conveys information about the usual size of moves in the underlying in relation to its absolute price. **Typical stocks have (annualised) volatilities of around 15-35%, with equity indices generally somewhat lower.** Maturities considered usually range from a few days to a few years. Clearly short-dated realised volatility will be much 'noisier' than longer dated, but will convey more up-to-date information. Longer dated volatility represents a smoother 'average' volatility but will include older data possibly outside of the current regime – e.g. 5-year Euro Stoxx volatility currently remains some 9% above 1-year volatility due to the inclusion of data from the high volatility regime in 2002-2003.

Figure 6 : The difference between RMS volatility and standard realised volatility can be noticeable for short-dated volatility…

5-day realised volatility (Euro Stoxx 50)
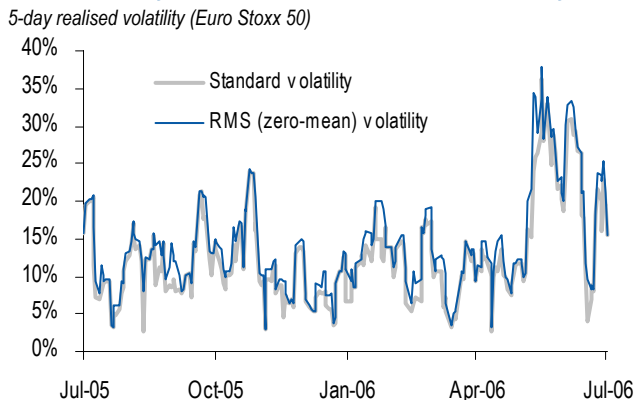
Source : JPMorgan

Figure 7 : …but the difference is negligible for longer-dated volatility

6-month realised volatility (Euro Stoxx 50)
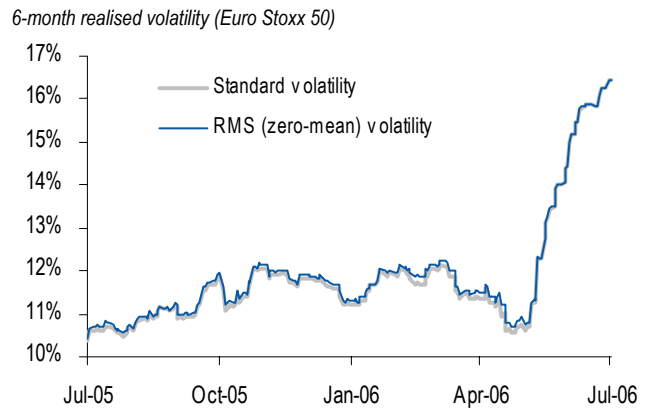
Source : JPMorgan

Although it is volatility which is the more intuitive concept, especially since it is measured in the same units as the underlying, variance (volatility squared) is in some sense more fundamental. Most importantly variance is additive, whereas volatility is not. It is realised variance which the delta-hedging of options seeks to capture, although the exposure is complicated by path dependency issues. **Variance swaps pay-out purely based on realised variance.**

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan ⬧

Part 1: Variance Swap Mechanics

1.2: The variance swap contract

Variance swaps **are instruments which offer investors straightforward and direct exposure to the volatility of an underlying asset** such as a stock or index. They are swap contracts where the parties agree to exchange a pre-agreed variance level for the actual amount of variance realised over a specified period (Figure 8). Variance swaps offer investors a means of achieving direct exposure to realised variance without the path-dependency issues associated with delta-hedging options.

The *strike* of a variance swap, not to be confused with the strike of an option, represents the level of volatility bought of sold and is set at trade inception. The strike is set according to prevailing market levels so that the swap initially has zero value. If the subsequent realised volatility is above the level set by the strike, the buyer of a variance swap will be in profit; and if realised volatility is below, the buyer will be in loss. A buyer of a variance swap is therefore long volatility. Similarly, a seller of a variance swap is short volatility and profits if the level of variance sold (the variance swap strike) exceeds that realised.

Figure 8: P/L of a variance swap
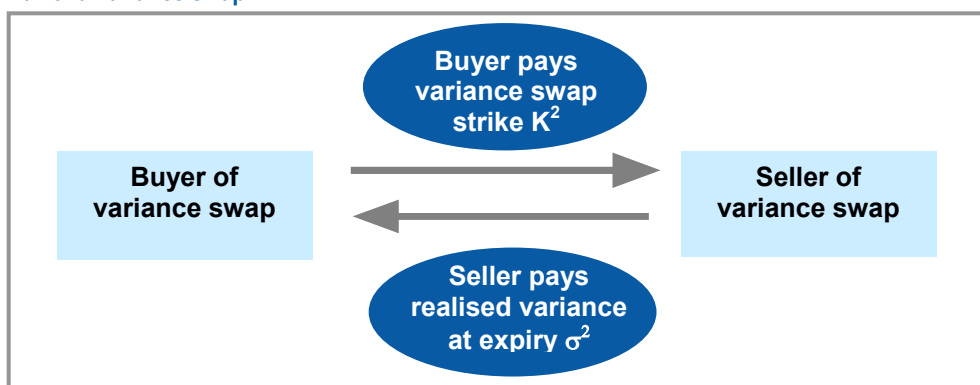
By convention, volatility is scaled by a factor of 100, for example a strike of 20 represents a volatility of 20%. Note that variance swap strikes are quoted in terms of volatility, not variance; but pay out based on the difference between the level of variance implied by the strike (in fact the strike squared) and the subsequent realised variance. This means the p/l of a variance swap is non-linear (convex) with volatility, although of course it is linear in terms of variance.

Box 2: Variance swap p/l

The p/l for a (long) variance swap is given by:

$$p/l = N_{Variance} \times \left(\sigma^2 - K^2 \right)$$

where K is the variance swap strike, σ^2 is realised variance, and $N_{Variance}$ is the ***variance notional.***

Example 1: Variance swap p/l

An investor wishes to gain exposure to the volatility of an underlying asset (e.g. Euro Stoxx 50) over the next year.

The investor buys a 1-year variance swap, and will be delivered the difference between the realised variance over the next year and the current level of implied variance, multiplied by the variance notional.

Suppose the trade size is €2,500 variance notional, representing a p/l of €2,500 per point difference between realised and implied variance.

If the variance swap strike is 20 (implied variance is 400) and the subsequent variance realised over the course of the year is $15\%^2 = 0.0225$ (quoted as 225), the investor will make a loss because realised variance is below the level bought.

Overall loss to the long $= €437,500 = €2,500 \times (400 - 225)$.   The *short* will profit by the same amount.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan ⬦

Part 1: Variance Swap Mechanics

1.3: Vega notional and variance notional

The notional for a variance swap can be expressed either as a *variance notional* or a *vega notional*. The variance notional is in some sense the 'true' notional – representing the p/l per point difference between the strike squared (implied variance) and the subsequent realised variance.

However, since variance is a somewhat un-intuitive quantity, and most market participants are used to thinking in terms of volatility, trade size is typically expressed in *vega notional*. **Vega notional represents the average profit or loss for a 1% (1 vega) change in volatility**, and thus gives a more economically meaningful idea of the variance swap's exposure to volatility. The vega notional is equal to the variance notional multiplied by twice the variance strike.

The p/l of a (long) variance swap can be calculated by multiplying the vega notional by the difference between the realised variance and the strike squared, and then dividing by twice the strike. This ensures that when realised volatility is close to the variance strike, the p/l is close to the difference between implied and realised volatility multiplied by the vega notional. That is, the variance swap payout, expressed in vega notional, is *locally linear* around the strike.

Box 3: Vega and variance notionals

$$p/l = N_{Vega} \times \left(\frac{\sigma^2 - K^2}{2K} \right) = N_{Variance} \times \left(\sigma^2 - K^2 \right)$$

where N_{Vega} is the vega-notional, and

$N_{Variance}$ is the variance notional.

$N_{Vega} = N_{Variance} \times 2K$

Example 2: Vega notional

Suppose a 1-year variance swap is stuck at 20 with a vega notional of €100,000.

If the index realises 25% volatility over the next year, the long will *receive* €562,500 = €100,000 x $(25^2 - 20^2)$ / (2 x 20).

However if the index only realises 15%, the long will *pay* €437,500 = €100,000 x $(15^2 - 20^2)$ / (2 x 20).

Therefore the average exposure for a realised volatility being 5% away from the strike is €500,000 or 5 times the vega notional, as expected.

Note that the variance notional is €100,000 / (2 x 20) = €2,500, giving the same calculation as that used in Example 1.

The p/l of a variance swap is often expressed in terms of vega notional. In the example above, a gain of €562,500 is expressed as a profit of 5.625 vegas (i.e. 5.625 times the vega notional). Similarly a loss of €437,500 represents a loss of 4.365 vegas. The average exposure to the 5% move in realised volatility is therefore 5 vegas, or 5 times the vega notional.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
**Nicolas Granger**
(44-20) 7325-7033

1.4: Variance swap convexity

Although variance swap payoffs are linear with variance they are **convex** with volatility. The vega notional represents only the *average* p/l for a 1% change in volatility. In fact **a long variance swap position will always profit more from an increase in volatility than it will lose for a corresponding decrease in volatility** (as in Example 2). This difference between the magnitude of the gain and the loss increases with the change in volatility. This is the convexity of the variance swap.

Ultimately this convexity is the reason that variance swaps strikes trade above at-the-money volatility. Because of the convexity, a variance swap will always outperform a contract linear in volatility of the same strike (e.g. Figure 9). The price paid for this convexity should depend on the expected variability of the (realised) volatility. The higher the variability of volatility, the larger the likely discrepancy between the variance strike and the realised volatility, and hence the more beneficial the convexity becomes (Figure 9, Figure 10). In the limiting case, with volatility of volatility zero, the p/l of the variance swap will always be zero (volatility is constant) and hence the variance swap strike and hypothetical volatility swap strike must coincide. See Section 2.8 for a detailed discussion of the value of the variance swap's convex payoff.

Figure 9 : When volatility remains close to the variance swap strike, the variance swap payout is similar to a linear "volatility swap"…
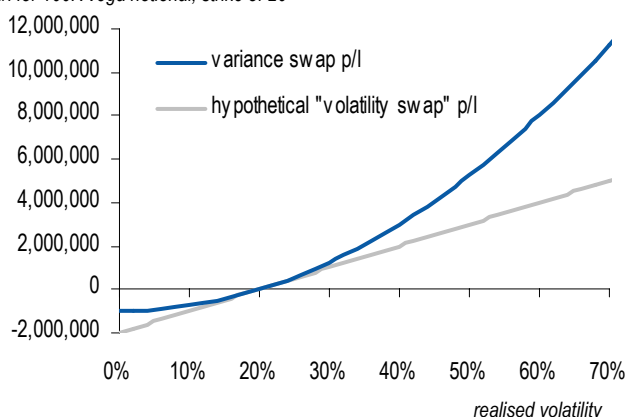
p/l for 100K vega notional, strike of 20
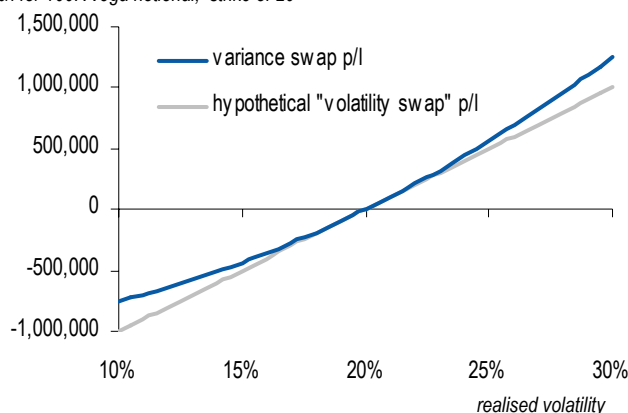
Source : JPMorgan

Figure 10 : … but for large swings in volatility, the convexity of the variance swap payout becomes important

p/l for 100K vega notional, strike of 20

Source : JPMorgan

In common with long options positions, the loss from a long variance swap is limited, whereas the loss from a short variance swap is not (unless capped). In the case of options, the loss from the long is limited to the premium paid. In the case of variance swaps the maximum loss occurs when realised volatility is zero, but in this case the convexity of the payout works in favour of the long with the result that **the maximum vega loss on a long variance swap is half the strike** (Example 3). For the short, losses are potentially unlimited, unless the variance swap payout is capped. The standard cap, if applied, is 2.5x the variance strike (see Section 1.8 for details of caps). In this case, it turns out that the maximum vega loss on a short 2.5x capped variance swap is 2.625 times the strike (see Example 3).

Example 3: Maximum loss of a variance swap

Suppose a 1-year variance swap is stuck at 20 with a vega notional of €100,000.

For a long variance swap the maximum loss will occur when the realised volatility is zero.

In this case p/l = €100,000 x $(0^2 - 20^2)$ / (2 x 20) = €100,000 x -10 = - €1M

This is equal to a loss of ½ times the 20 strike or 10 vegas, as expected.

For a 2.5x capped short variance swap the volatility on which the variance swap pays out will be capped at 50 = 2.5 x strike. The maximum loss will occur when realised volatility is at least 50.

This gives a maximum loss of €5.2M = €100,000 x $(50^2 - 20^2)$ / (2 x 20) = 2.625 x strike x vega notional.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

1.5: Variance swap example – accruing realised volatility

In this example we show the evolution of the p/l of a short Euro Stoxx 50 variance swap through the accrual of realised variance, using prevailing market prices. The variance swap was sold on 14 October 2005, with a strike of 16.5 for €100,000 of vega notional (equivalent to €3,030 variance notional). The variance swap maturity was 20 (business) days, expiring on the 10th November 2005.

Over the 20-day period realised volatility was 14.3%, 2.2% lower than the level of implied sold. Since the vega notional gives the expected p/l for a 1 point (1%) move in volatility, we would expect a p/l from the short position of €220,000 (=2.2 times the vega notional). However, due to the convexity of the payout working in favour of the long, and against the short, the actual p/l is slightly less: €206,714.

Table 2: Daily p/l and cumulative realised volatility from a 20-day short variance swap struck at 16.5

Day	Date	Index Level	return	log return squared	daily volatility (annualised)	daily p/l (€)	accrued realised volatility (annualised)	accrued p/l (€)
	13-Oct-05	3331.4						
1	14-Oct-05	3349.6	0.5%	0.000030	8.6%	29,962	8.6%	29,962
2	17-Oct-05	3356.8	0.2%	0.000005	3.4%	39,494	6.6%	69,457
3	18-Oct-05	3334.8	-0.7%	0.000043	10.4%	24,777	8.1%	94,234
4	19-Oct-05	**3279.6**	**-1.7%**	**0.000278**	**26.5%**	-65,070	**15.0%**	29,165
5	20-Oct-05	3284.8	0.2%	0.000002	2.5%	40,304	13.4%	69,469
6	21-Oct-05	3271.0	-0.4%	0.000018	6.6%	34,554	12.6%	104,023
7	24-Oct-05	3310.1	1.2%	0.000141	18.9%	-12,622	13.6%	91,401
8	25-Oct-05	3292.5	-0.5%	0.000028	8.5%	30,377	13.1%	121,778
9	26-Oct-05	3304.3	0.4%	0.000013	5.7%	36,403	12.5%	158,181
10	27-Oct-05	**3241.1**	**-1.9%**	**0.000372**	**30.6%**	-100,824	**15.3%**	57,357
11	28-Oct-05	3246.0	0.1%	0.000002	2.4%	40,400	14.6%	97,757
12	31-Oct-05	3320.1	2.3%	0.000510	35.9%	-153,615	17.4%	-55,858
13	01-Nov-05	3312.5	-0.2%	0.000005	3.7%	39,194	16.8%	-16,664
14	02-Nov-05	3320.6	0.2%	0.000006	3.9%	38,935	16.2%	22,271
15	03-Nov-05	3361.6	1.2%	0.000151	19.5%	-16,312	16.4%	5,959
16	04-Nov-05	3355.3	-0.2%	0.000004	3.0%	39,880	15.9%	45,839
17	07-Nov-05	3362.8	0.2%	0.000005	3.6%	39,318	15.5%	85,157
18	08-Nov-05	3361.7	0.0%	0.000000	0.5%	41,210	15.0%	126,367
19	09-Nov-05	3355.8	-0.2%	0.000003	2.8%	40,040	14.6%	166,407
20	10-Nov-05	3361.0	0.2%	0.000002	2.5%	40,307	14.3%	**206,714**
						206,714		

Note how most days were low volatility and produced modest positive returns for the short variance position. However, a few days with relatively large moves caused significant daily losses (Table 2, Figure 12). This pattern is common: short variance will profit modestly most of the time, but will lose heavily on large moves.

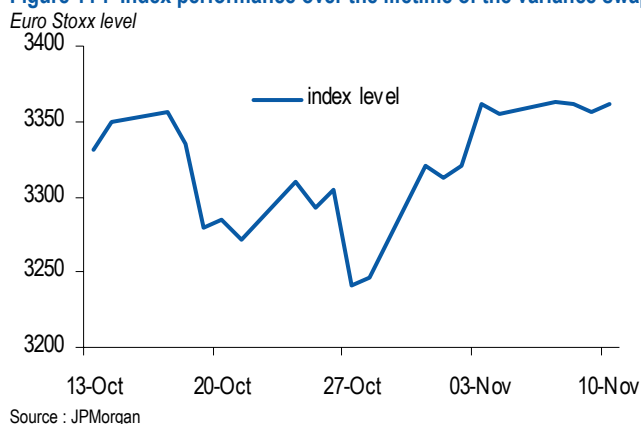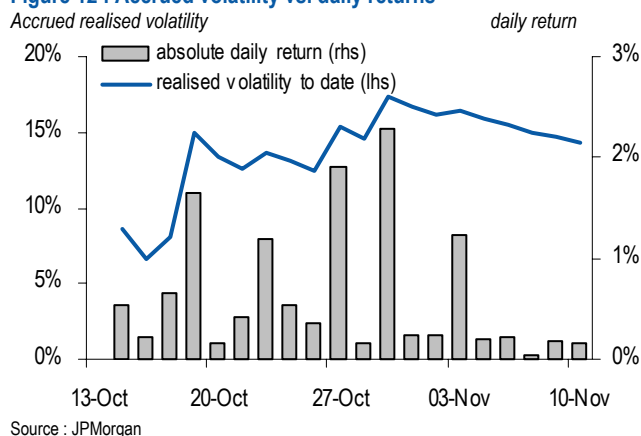**Figure 11 : Index performance over the lifetime of the variance swap**
Euro Stoxx level

Source : JPMorgan

Figure 12 : Accrued volatility vs. daily returns
Accrued realised volatility   *daily return*

Source : JPMorgan

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 1: Variance Swap Mechanics

1.6: Variance swap mark-to-market

Although the example on page 14 illustrates how the variance swap accrues realised volatility over time, **if an investor wishes to close out a variance swap position before maturity they must also consider any change in value of the exposure to volatility over the remainder of the variance swap term.** Indeed for variance swaps valued relatively soon after inception, this change in value of future expected variance can be the main driver of p/l, an especially important consideration for long dated variance swaps.

Marking to market of variance swaps is easy: variance is additive. At an intermediate point in the lifetime of a variance swap, the expected variance at maturity is simply the time-weighted sum of the variance realised over the time elapsed, and the implied variance over the remaining time to maturity.

All that is needed to compute the mark-to-market of a variance swap is:

- The realised variance since the start of the swap; and
- The implied variance (variance strike) from the present time until expiry

Since the variance swap is usually settled at maturity, a discount factor between the present time and expiry is also required.

Example 4: Variance swap mark-to-market

Suppose a 1-year variance swap is stuck at 20 with a vega notional of €100,000 (variance notional of €2,500).

If the volatility realised over the first 3 months is 15%, but the volatility realised over the following 9 months is 25%, then, since variance is additive, the variance realised over the year is:

Variance = [¼ x 15^2] + [¾ x 25^2] = 525   (≈22.9 volatility).

At expiry the p/l would be €2,500 x (22.9^2 – 20^2) = €312,500

Now, suppose again that realised volatility was 15% over the first 3 months. In order to value the variance swap after 3 months we need to know both the (accrued) realised volatility to date (15%) and the fair value of the expected variance between now and maturity. This is simply the prevailing strike of a 9-month variance swap. If this is currently trading at 25, then the same calculation as above gives a fair value *at maturity* for the 1-year variance swap of €312,500.

Although the fair value at maturity (now 9 months in the future) is €312,500, we wish to realise this p/l now (after 3-months). It is therefore necessary to apply an appropriate interest rate discount factor.

If, after 3-months, 9-month LIBOR is 4%, the discount factor applied would be 1/(1+¾ x 0.04) = 0.97. Hence the mark-to-market p/l after 3 months would be equal to about €303,400.

Box 4: Variance swap mark-to-market

$$p/l = N_{var} \times ([\sigma_t^2 - K_{old}^2] \quad x \quad \%time\ elapsed \quad + \quad [K_{new}^2 - K_{old}^2] \quad x \quad \%time\ remaining)$$

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 1: Variance Swap Mechanics

Box 5: Calculating the variance swap mark-to-market

Let $K_{s,t}$ represent the strike for a variance swap at time s maturity at time t.

Let $\sigma_{s,t}$ be the realised volatility between time s and time t.

Then the p/l (in terms of variance notional) at time t, of a variance swap maturing at time T, is given by:

Variance (forward) p/l at time t (per unit variance notional): $\quad = \dfrac{t}{T}\left(\sigma_{0,t}^2 - K_{0,T}^2\right) + \dfrac{T-t}{T}\left(K_{t,T}^2 - K_{0,T}^2\right)$

$$= \frac{1}{T}\left[t\sigma_{0,t}^2 + (T-t)K_{t,T}^2\right] - K_{0,T}^2$$

Example 5: Variance swap mark-to-market continued

Returning to Example 4 we can calculate the mark-to-market p/l (at maturity) directly from the formula above as follows:

p/l = €2,500 $\quad$ x ($[15^2 - 20^2]$ x ¼ $\quad+\quad$ $[25^2 - 20^2]$ x ¾)

$\quad$ = €2,500 $\quad$ x (-175 x ¼ + 225 x ¾)

$\quad$ = €2,500 $\quad$ x $\quad$ 125

$\quad$ = €312,500

Applying the interest rate discount factor of 0.97 as explained in Example 4 gives a mark-to-market p/l of €303,400

Box 6: Variance swap mark-to-market through daily accruals

An alternative approach to calculating mark-to-market is to consider a variance swap as a sequence of daily variance swaps. On each day the p/l accrued by a long variance swap is given by:

Daily accrual on day i (per unit of variance notional) $= \dfrac{1}{T}(\sigma_{i,i-1}^2 - K_{0,T}^2)$,

where $\sigma_{i,i-1}$ is the daily annualised realised volatility equal to $252\ln^2\left(\dfrac{P_i}{P_{i-1}}\right)$ and $K_{0,T}^2$ is the initial strike.

Each daily p/l is given by a daily variance exposure equal to the total variance notional divided by the total number of days between the start and end dates of the swap. It is important to notice that **the total remaining exposure to variance decreases linearly with time**. Calculating a mark-to-market or closing out an existing position requires entering into a new but opposing variance swap position in the remaining variance notional, *not* the initial notional. This is done in order to cancel out the remaining exposure to the daily realised volatility component. The remaining variance notional size will be equal to *(T-t)/T* times the original notional, proportional to the time remaining to maturity on the swap. Thus p/l of the trade at time t will be equal to:

$$\frac{t}{T}(\sigma_{0,t}^2 - K_{0,T}^2) + \sum_{t+1}^{T}\frac{1}{T}(K_{t,T}^2 - K_{0,T}^2)$$

The first term represents the realised variance accrued by the variance swap up to time t, whilst the second term represents a fixed stream of daily accruals up to expiry at time T. This is exactly equivalent to the mark-to-market equation given above and presents a useful framework for thinking about forward variance, the subject of the next section.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan ⬤

Part 1: Variance Swap Mechanics

1.7: Forward variance

The reasoning used to compute the variance swap mark-to-market in the preceding section can also be used to calculate the expected variance over a forward starting window. Suppose we know the strikes for a short-maturity variance swap expiring at time t, and a longer maturity variance swap expiring at time T. We want to find the expected realised variance, F, between time t and time T. Since variance is additive, the longer maturity variance swap is simply the time-weighted sum of the short maturity variance and the expected variance over the forward period, thus enabling us to compute this expected forward variance. This equates to the fair strike of the **forward-starting variance swap**.

Example 6: Calculating forward variance

Suppose that one year variance is struck at 20, but 3-month variance is struck at 15. What is the fair strike for a variance swap starting in 3 months time, and expiring 9 months later?

Since variance is additive, the variance expected over the year is the time weighted sum of the variance expected in the first three months, and that expected in the remaining nine months. This second value is the forward strike F. Then, using the additivity of variance as for the mark-to-market calculation:

$$20^2 = [¼ \times 15^2] + [¾ \times F^2]$$

Solving for F, gives a strike for the forward variance of 21.4.

Box 7: Calculating forward variance through daily accruals

Using the approach that a variance swap can be thought of as a strip of daily variance swaps (see p16), a combination of long and short variance swaps can be used to cancel out exposure to realised volatility over a particular time period. This shows how a forward variance swap can be constructed, giving exposure to realised volatility starting after a future date.

To illustrate this we being by considering the p/l of a *long* variance swap, which because of the additivity of variance can be spilt into two periods (1 → t and t+1 → T) as follows

$$\sum_1^t \frac{N_T}{T}(\sigma_{i,i-1}^2 - K_T^2) + \sum_{t+1}^T \frac{N_T}{T}(\sigma_{i,i-1}^2 - K_T^2) \quad \text{where } N_T \text{ is the variance notional and } \sigma_{i,i-1} \text{ is realised daily volatility}$$

Consider the p/l for a *short* variance swap, maturity t, of variance notional $\frac{t}{T}N_T$ (with delayed payment). This is equal to:

$$-\sum_1^t \frac{t}{T}\frac{N_T}{t}(\sigma_{i,i-1}^2 - K_t^2) \quad \text{which will offset the realised variance from the second term in the p/l from the long above.}$$

Combining the long and short variance positions gives an overall p/l of: $\sum_1^t \frac{N_T}{T}(K_t^2 - K_T^2) + \sum_{t+1}^T \frac{N_T}{T}(\sigma_{i,i-1}^2 - K_T^2)$

The first p/l term represents a fixed quantity which can be subtracted from the fixed strike in the second p/l term. We then define a forward variance to be equal to this combined strike, such that

$$F_{t,T}^2 = K_T^2 - \frac{1}{T-t}\sum_1^t (K_t^2 - K_T^2) = K_T^2 - \frac{t}{T-t}(K_t^2 - K_T^2) = \frac{T}{T-t}K_T^2 - \frac{t}{T-t}K_t^2$$

The overall p/l becomes $\sum_{t+1}^T \frac{T-t}{T}\frac{N_T}{T-t}(\sigma_{i,i-1}^2 - F_{t,T}^2)$

This just represents the p/l of a variance swap starting a time t, ending at time T of notional size: $\frac{T-t}{T}N_T$.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 1: Variance Swap Mechanics

The calculations in Box 7 show that a long position in forward starting variance (of unit variance) can be constructed from:

- Long *T/(T-t) variance* notional of spot variance maturity *T*; and

- Short *t/(T-t) variance* notional of spot variance maturity *t*, but with payment delayed until maturity *T*.

Box 8: Calculating forward starting variance

Suppose K_t is the strike of the short-maturity variance swap (of maturity *t*),

K_T is the strike of the longer-maturity variance swap (maturity *T*), and

$F_{t,T}$ is the strike of the forward starting variance swap (with exposure to variance between *t* and *T*), then:

$$K_T^2 = \frac{t}{T}K_t^2 + \frac{T-t}{T}F_{t,T}^2 \quad \text{(as in the mark-to-market calculation in Box 4)}$$

Solving for *F* gives a value for the forward variance:

$$F_{t,T}^2 = \frac{T}{T-t}K_T^2 - \frac{t}{T-t}K_t^2 \quad \text{or} \quad F_{t,T} = \sqrt{\frac{T}{T-t}K_T^2 - \frac{t}{T-t}K_t^2}$$

Example 7: Forward variance construction

As in Example 5, suppose a 3-month variance swap is struck at 15, and a 1-year variance swap is struck at 20.

The 9-month variance swap, 3-month forward variance will be struck at 21.4.

Suppose we wish to trade €100,000 of forward variance swap notional.

As in the case of a spot variance swap, this equates to €100,000 / (2 x 21.4) = €2,336 variance notional.

Hence we require -3/9 x €2,336 = -€778 variance notional of the 3-month variance swap (with payment delayed); and

12/9 x €2,336 = 3,115 variance notional of the 12-month variance swap.

This in turn equates to short 2 x 15 x €778 = €24,000 vega notional of the 3-month variance swap; and

long 2 x 20 x €3,115 = €125,000 of the 12-month variance.

One of the things worth noting on forward starting variance, is that more variance is need on the longer leg (which tends to be less liquid) than the shorter leg and that **the total notional of the two legs will be greater than the notional of the forward.** This can have a practical impact of increasing the bid/offer of forward starting variance swaps. However, as the variance swap market evolves and matures, forward starting variance is increasingly quoted as a single level rather than a weighted spread with the consequence that bid/offers are converging to similar levels as seen for spot variance.

The ability to re-express the spot variance term structure into a strip of forward starting variances raises natural analogies with fixed income – in particular the swap curve and a strip of forwards/futures. Forward variance swaps are useful for taking a direct view on the future value of implied variance and/or the future shape of the variance term structure curve in the same way that Eurodollar and Euribor futures are used to take view on the future level of interest rates. Furthermore the fixed income concepts of slide and carry are useful concepts to apply to both spot and forward starting variance swaps. See Sections 3.8 and 3.9 for examples of the uses of forward variance swaps.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 1: Variance Swap Mechanics

1.8: Variance swap contract specifications

Variance swap conditions are set out on term sheets such as the one shown in Section 1.9. Strikes are typically expressed as 100 times the annualised volatility represented – for example 25 will represent a breakeven realised volatility of 25%. For this reason the calculated value of the realised variance (using an RMS calculation, see p10) must also be multiplied by 100.

Margins and collateral

Variance swaps are usually margined in a similar manner to options, with an initial amount to be posted as collateral (3 times the vega amount in the example term-sheet below). Further margin calls will be made during the course of the trade as necessary. Settlement is calculated at maturity and cash-flows exchanged shortly afterwards (T+2 in the example below).

Disrupted days

The realised variance used in a variance swap payout is calculated from closing (or official observed) prices on observation dates over a specified period. Observation dates are defined as all those scheduled trading days which are not disrupted. The occurrence of a disrupted day could potentially work for or against a long variance position. Suppose a 5% loss is followed by a 6% gain over 2 consecutive trading days. If the first day is declared as disrupted, its closing price will not be used in the calculation of the realised volatility and only the single combined return of 0.7% will be counted (see Example 8). Conversely, if a 5% loss on a day declared disrupted is followed by a 6% loss the next day which is non-disrupted, this will count as a single, compounded 1-day loss of 10.7%. Such a move will act to increase realised volatility more than by the separate 5% and 6% moves, and will work in the favour of the long.

Example 8: Disrupted days

On Wednesday, 18 January 2006 the Tokyo Stock Exchange (TSE) closed early following revelations of the investigation into Livedoor. At the point when trading was suspended on the TSE the Nikkei was at 15341, down 2.9% on the previous day's close of 15806. However, the day was declared *disrupted*, and the provisional closing price was not accepted for variance swap calculations.

The next day, Thursday, January 19th, was considered a normal trading day, and the Nikkei recovered most of the previous day's losses, closing at 15696, up 2.3% on the last level on the disrupted Wednesday.

For the purposes of the variance swap calculation, only the change from the close on the Tuesday to the close on the Thursday, a return of -0.7% counted towards the calculation of realised volatility. A long variance swap position would therefore have lost out due to the Wednesday being declared a disrupted day.

Index reconstitution risk

Variance swaps on indices are defined to pay out on the returns of the index and *not* on the weighted returns of the basket of current constituents. This means that index variance swaps (especially long-dated ones) will be exposed to the risk of index reconstitution, and the variance swap may end with an exposure to a very different set of stocks, potentially with different volatility characteristics than when it was originally traded. For example, in the Hang Seng, the cap on the weight of HSBC, a low volatility stock, and the forthcoming addition of a number of potentially higher volatility Chinese H-Shares may act to increase the overall volatility of the index.

Dividend adjustments

Variance swaps on single names are typically adjusted for dividends, both special and regular. This means that the return on the ex-dividend date is calculated after adjusting for the dividend. For example, if a stock is worth €100 on the day before the dividend, pays a dividend of €5 and closes at €94 on its ex-dividend day, the return used in calculating the realised volatility for the variance swap payout will be 94/95-1 = -1.05% (*not* 94/100-1 = -6%). This **means that changes to the underlying due to dividend payments do not count towards the realised volatility calculation for the variance swap**. This can be important for stocks with high dividend yields which pay entire annual dividend in a single payment – e.g. Italian stocks such as Telecom Italia and ENEL (Figure 13).

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 1: Variance Swap Mechanics

By convention index variance swaps are not usually adjusted for dividends – since the payments are usually more spread out and the impact of a few index points worth of dividend is generally small in comparison to the average daily move. However, there are cases of indices containing high dividend paying stocks some of which pay on the same day – e.g. the SPMIB and the SXKE (the Eurozone Telecoms index) – where index moves as a result of dividend payments can be significant. In these cases the extra volatility resulting from dividend payments will count in the calculation for the variance swap payout, potentially giving an artificial boost to the realised volatility.

Variance swap caps

Variance swaps, especially on single-stocks (and sector indices), are usually sold with caps. These are often set at 2.5 times the strike of the swap capping realised volatility above this level. Variance swap caps are useful for short variance positions, where investors are then able to quantify their maximum possible loss (to 2.625 times the strike, see Box 9).

Box 9: Variance swap caps

Suppose a variance swap is traded with a cap set at 2.5 times the strike. Then the p/l (for the long) is given by:

$$p/l = N_{Vega} \times \frac{(Min\{\sigma, 2.5K\})^2 - K^2}{2K}$$

Note that this sets the maximum loss for the short to $N_{Vega} \times \frac{(2.5K)^2 - K^2}{2K} = N_{Vega} \times 2.625K$

In practice caps are rarely hit – especially on index underlyings and on longer-dated variance swaps. When caps are hit, it is often due to a single large move – e.g. due to an M&A event or major earning surprise on an individual name, or possibly from a dramatic sell-off on an index. Single-day moves needed to cause a variance swap cap to be hit are large and increase with maturity. For example a 3-month variance swap struck at 20 and realising 20% on all days except for one day which has a one-off 24% move, will hit its cap. A similar 1-month maturity swap would need only a 14% move to hit the cap, whereas the required 1-day move on a 1-year swap would be 46%. For lower strikes the required moves are also lower.

Closing out variance swaps

Investors wanting to realise mark-to-market p/l by closing out a variance swap before expiry generally have two choices. They can either enter into an opposing offsetting contract, potentially with a different counterparty, or they can agree with the original counterparty on a level at which to close out the contract. In theory this second alternative is simpler as it removes all collateral obligations, frees up capital and avoids any technical issues with changes in cap level (see Example 9). However, in practice it is possible to negotiate bespoke contracts specifically to close out existing positions in a process known as *novation* which allows for the different counterparties to close out the contract fully.

Example 9: Offsetting capped variance swaps

Suppose that in an investor buys a 6-month variance swap with a strike of 20%.

This has the standard 2.5x cap meaning the exposure to realised volatility will be capped at 50%.

1-month later, the (5-month) variance swap is now trading at a strike of 30% leading to a significant mark-to-market p/l – the exact amount depending on the level of volatility realised over the month. The investor wants to lock in this profit.

With the strike now at 30%, the cap on an new variance swap contract will by default be set at 2.5 x 30% = 75%

Then if the investors sells this 30-strike variance swap in an attempt to close out his position the difference in caps will mean he takes on a short volatility exposure if the subsequent realised volatility is above 50% (although capped at 75%)

In effect, in the course of trying to close out his position, he will have sold a '50%/75% call spread' on volatility. Whilst the price he gets for selling the variance swap will reflect this higher cap, the residual volatility exposure is presumably unwanted, and the investor would be best either trading directly with the original counterparty or negotiating a bespoke contract with another counterparty in order to fully close out his outstanding contract.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

M&A events

M&A events present particular issues for variance swaps, as an underlying being acquired can effectively cease to exist. The various options exchanges have rules for how such an event is dealt with, and variance swap contracts generally follow these rules. Investors should be aware in which exchange the variance swap is based since M&A rules are exchange specific. For share offers, or composed offers consisting of less than 66% cash, the options, and hence the variance swap, will usually transfer into a variance swap on the new underlying after the transaction date.

In the case of cash bids, options are unwound immediately after the offer has been declared unconditional. In Europe, some bourses unwind vanillas at intrinsic and others at the theoretical fair value (generally taking the implied volatility to be the average over the 5 days before the bid). The Euronext exchange unwinds at fair value; Eurex unwinds at intrinsic but is switching to fair value from the beginning of 2007; IDM in Italy unwinds at intrinsic and MEFF in Spain works on a case by case basis. In any case documentation should either specify the exact rules in the case of an M&A event or should refer to the rules of the appropriate exchange.

One further issue with cash bids, is that the underlying will usually trade at close to the bid level after the offer has been accepted, often with close to zero volatility.   However, the options are not unwound until the offer is declared unconditional. For example when Telefonica bid for O2 on 31 October 2005, it took until 3 January 2006 for the offer to go unconditional (Figure 14). Although outstanding options were settled at the implied volatility prevailing before the October bid, they were not unwound until the beginning of January, leading to two months of virtually zero accrued volatility on any outstanding variance swaps. In this case, given the size of the daily move on the day of the bid (+ 25%) the 3-month realised volatility from the beginning of October up until the unconditional date on January 3rd, was around 50%, well in excess of levels of implied variance trading before the bid. However, in cases where the move on the bid is smaller, and the time up until unconditional acceptance is longer, the low realised volatility over this period could hurt long variance swaps.

Figure 13 : The apparent drop in spot and consequent spike in volatility was entirely due to the payment of a single 6.3% dividend.

Share price (€)   *1m realised volatility*
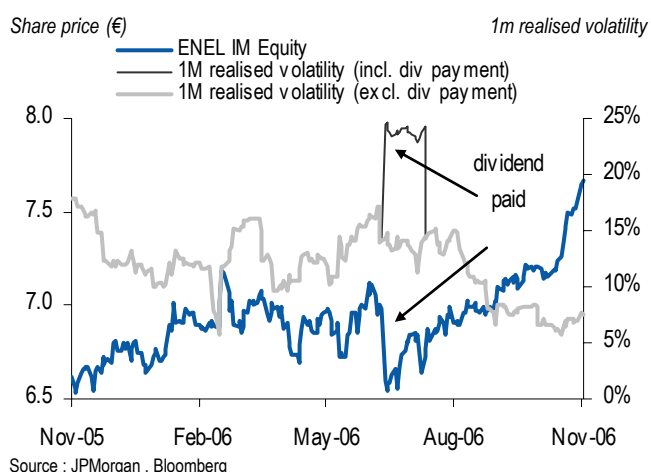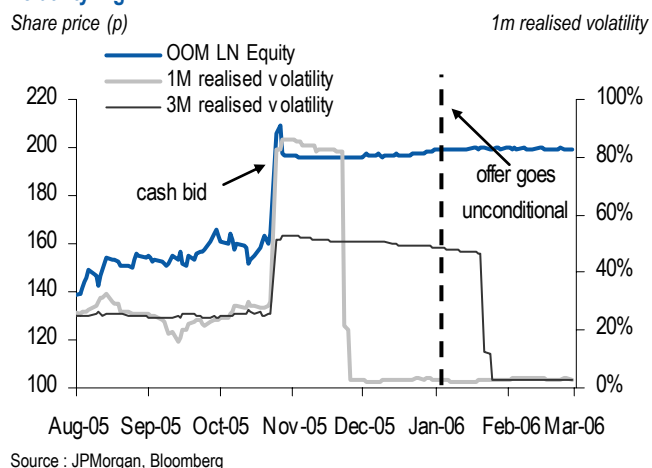
Source : JPMorgan , Bloomberg

Figure 14 : O2 realised volatility collapsed following the bid. But the large 1-day move from the bid was large enough to keep realised volatility high

Share price (p)   *1m realised volatility*

Source : JPMorgan, Bloomberg

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 1: Variance Swap Mechanics

1.9: Example variance swap term sheet

Sep06 OTC Variance Swap: FTSE Index Variance

Instrument:	Variance Swap
Trade Date:	10th August 2006
Maturity Date:	15th September 2006
Variance Buyer:	JPMORGAN
Variance Seller:	*******
Underlying:	FTSE100 INDEX (*Bloomberg: UKX Index*)
Denominated Currency:	Great British Pounds ("GBP")
Vega Amount:	GBP 100,000
Strike Price:	**16.95**
Cap Level:	**36.95**
Equity Payment:	**T+2** after the Observation End Date, the Equity Amount will be calculated and paid in accordance with the following formula:
	$$\left(\frac{\text{Vega Amount}}{2\times \text{Strike}}\right)\times\left[\text{Min}\left(\text{FRV}^2, CapLevel^2\right) - \text{Strike Price}^2\right]$$
	If Equity Amount is positive the Variance Seller will pay the Variance Buyer the Equity Amount. If Equity Amount is negative the Variance Buyer will pay the Variance Seller an amount equal to the absolute value of the Equity Amount.
Where FRV =	$$\sqrt{\frac{252\times\sum_{t=1}^{N}\left(\ln\frac{P_t}{P_{t-1}}\right)^2}{Expected\_N}}\times 100$$

t means each Observation Day.
N means number of Observation Days.
Expected_N means 25, the number of days which, as of the Trade Date, are expected to be Scheduled Trading Days in the Observation Period.
ln means the natural logarithm.
P_t means in respect of: (i) the final Observation Day, the Official Settlement Price of the Exchange-traded Contract on the Observation End Date; and (ii) each other Observation Day, the official level of the Index at the Valuation Time on such Observation Day.
P_{t-1} means in respect of: (i) the first Observation Day, the Settlement Price per Underlying on the Observation Start Date and; (ii) each subsequent Observation Day, the Settlement Price per Underlying on the Observation Day immediately preceding such Observation Day.
Observation Day Each Scheduled Trading Day that is not a Disrupted Day during the Observation Period.
Observation Period The period from, but excluding, the Observation Start Date to, and including, the Observation End Date.
Observation Start Date Trade Date.
Observation End Date Maturity Date.
Settlement Price means the official price or, if there is no official price, the mid-market price per Underlying on the Exchange at the scheduled weekday closing time (without regard to after hours or any other trading outside of the regular trading session hours) on an Observation Day.

Calculation Agent:	JPMorgan Securities Ltd.
Documentation:	ISDA

2. The Variance Swap Market

Part 2: The Variance Swap Market

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

JPMorgan

Part 2: The Variance Swap Market

Part 2 begins with an overview of the development and current state of the variance swap market. We then look at where **variance swap levels have traded historically**, concentrating particularly on the Euro Stoxx 50 and S&P500 indices. We leave the detailed theoretical pricing argument until later (see Part 4), but discuss briefly the **relationship between variance levels and traded option volatilities** and consider some **pricing rules of thumb**.

Variance swaps can be thought of as the **market's best estimate of future realised variance** and we consider how good a job they do of predicting this. We discuss the **convexity premium** for variance swaps in terms of the spread of variance swap prices over ATM options implied volatility.

We also consider the **drivers of variance swaps prices**: principally **realised volatility and risk aversion** at shorter maturities and **structured product flow** at the longer end of the curve. This leads in to a discussion of the **variance swap term structure and curve dynamics**.

2.1: Market development

Variance swap contracts were first mentioned in the 1990's, but like vanilla options only really took off following the development of robust pricing models through replication arguments. (See e.g. *Derman, 1999* or *"Towards a Theory of Volatility Trading, Carr and Madan', 2001*). The directness of the exposure to volatility and the relative ease of replication through a static portfolio of options make variance swaps attractive instruments for investors and market-makers alike.

The variance swap market has grown steadily in recent years, driven by investor demand to take direct volatility exposure without the cost and complexity of managing and delta-hedging a vanilla options position. Although it is possible to achieve the variance swap payoff through a portfolio of options, the variance swap contract offers a convenient package bundled with the necessary delta-hedging, offering **investors a simple and direct exposure to volatility**, without any of the path dependency issues associated with delta-hedged options.

As liquidity in the variance swap market has increased, to a certain extent levels now deviate from the theoretical price obtained from a replicating portfolio of options. This is partly because the construction of the replicating portfolio includes a relatively large contribution from the deep out-of-the-money options and the reduced liquidity in these options gives rise to some discrepancy in the overall pricing of variance swaps. We discuss practical replication issues in more detail in Section 4.9.

Variance swaps were initially developed on index underlyings. In Europe, variance swaps on the Euro Stoxx 50 index are by far the most liquid, but DAX and FTSE are also frequently traded. Variance swaps are also tradable on the more liquid equity underlyings – especially Euro Stoxx 50 constituents, allowing for the construction of variance dispersion trades. Variance swaps are tradable on a range of indices across developed markets and increasingly also on emerging markets.

Table 3: Global variance swap liquidity

	Major indices on which variance is traded	Typical index bid offers ,vegas (most liquid maturities)	Typical single-stock bid offers, vegas (most liquid stocks and maturities)
Developed Europe	Euro Stoxx 50, FTSE, DAX, CAC	0.4 – 0.6	1 – 2
U.S.	S&P500, Nasdaq, Russell 2000	0.4 – 1	1 – 2
Japan	Nikkei 225, Topix	0.4 – 0.6	2 – 2.5
Asia Ex.	HSI, HSCEI, Kospi	1 — 2	2 - 4
Emerging Markets	CECE, Wig20, RDX, Top40, Bovespa	3 – 5	5 – 6

Source: JPMorgan estimates, November 2006

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 2: The Variance Swap Market

Bid/offer spreads have come in significantly over recent years and in Europe they are now typically in the region of 0.5 vegas for indices and 1 – 2 vegas for single-stocks – although the latter especially vary according to liquidity factors. We expect spreads to come in further in future as liquidity continues to improve. Spreads for indices in the US and Japan are similar with single stock bid/offers in the 2 – 2.5 vegas range. Spreads are naturally higher in emerging markets although these too are becoming much more liquid. With volatilities often much higher and less stable than in developed markets, emerging market variance can offer interesting opportunities.

The most liquid variance swap maturities are generally from 3 months to around 2 years, although indices and more liquid stocks have variance swaps trading out to 3 or even 5 years and beyond. Maturities generally coincide with the quarterly options expiry dates, meaning that they can be efficiently hedged with exchange-traded options of the same maturity. Good liquidity is also seen in the front 3 months for short-dated index variance. The **VIX, VSTOXX and VDAX indices** represent the theoretical prices of 1-month variance swaps on the S&P500, Euro Stoxx and DAX indices respectively, and are calculated by the exchanges from listed option prices, interpolating to get 1-month maturity. These volatility indices are widely used as benchmark measures of equity market risk, even though they are only short-dated measures and are not directly tradable. See p94 for further details.

Typical sizes of variance swap trades are around € 100,000 – 200,000 of vega for index underlyings and € 50,000 vega notional for single-stock underlyings. As with any OTC market, estimating total traded volumes is problematic, with recent estimates for daily traded volumes on indices in the region of $3-5 million vega (*source: Risk Magazine*). This represents an estimated 5- or 10-fold increase on volumes three years ago. Traders in Europe and Japan estimate that approximately 10-30% of the vega traded in the market is done through variance swaps, with a slightly lower proportion in the U.S.

Forward variance and index variance spreads (e.g. Euro Stoxx 50 vs. S&P500) are increasingly quoted with a single bid/offer spread, often comparable to, or only slightly higher than, the spread on a single leg, meaning the second component of the spread can effectively be traded at mid-market levels. These developments make forward variance and index spreads much more attractive for investors who may previously have found the transaction costs prohibitive. JPMorgan publishes its variance swap runs, including prices for forwards and spreads, and conditional variance prices, on Bloomberg JPVS.

The advent of the variance swap market has allowed the investor community to trade variance directly, without the onerous human costs (and transaction costs) of delta-hedging a portfolio of options. Of course, the market-making community can offset variance swap exposure with the advantage of economy of scale. As dedicated 'volatility-arbitrage' funds have grown, variance swaps have allowed them to directly express their views, encouraged in part by the attractive returns available from shorting variance in 2003-2004 as equity volatility fell dramatically. In addition, selling stock correlation using variance swaps (so-called 'variance dispersion' trades) has also been a popular trade in the hedge fund community.

More recently, the investor base has widened further with the advent of structured products which embed variance swap returns. Such products may be leveraged and/or offer capital guarantees, making them attractive to a broader range of institutions including private banks and sophisticated private investors. For example, **JPMorgan's recently launched Yield Alpha and Variance Alpha indices incorporate variance swap exposure, with a variety of structured products offering access to these indices**. See Section 3.3 and JPMorgan's *Yield Alpha Strategy Guide* for further details.

Variance swap liquidity has so far been concentrated on the equity markets, in part fuelled by the need for investment banks to recycle volatility and correlation exposures that are embedded in structured products. However variance swaps do trade on non-equity underlyings – for example FX, interest rates, commodities and credit – albeit with less liquidity. In theory there is no reason why variance swaps should not gain liquidity in these asset classes, depending on investor demand and providing that the underlying option markets are liquid enough.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 2: The Variance Swap Market

2.2: Historical prices

Figure 15 shows Euro Stoxx 50 implied variance levels between 2000 and 2006 for maturities of 1-month, 6-months and 3-years. Short-dated variance swap prices tend to be more variable than those for longer maturities, with 1-month variance experiencing highs above 60 and lows of almost 10 over the six year period. **Variance swap prices can be observed to follow high and low regimes** in a similar manner to realised variance (Figure 5), and within this period different volatility regimes are clearly discernible:

- A 'medium' volatility regime up until mid 2002, with implied variance priced in the 20-30 range, interrupted by a large spike in September 2001;

- A 'high' volatility regime from mid 2002 – mid 2003 with short-dated variance priced consistently above 40;

- A 'low' volatility regime since mid 2004 with short-dated variance below 20 (excepting the May 2006 spike).

Longer maturity variance swap levels vary less, with 3-year implied variance remaining in the 15-35 range, and in particular reacting less to spikes. A sudden unexpected event likely to dramatically increase short-term volatility dramatically, is perhaps less likely to cause the same level of elevated volatility over the next few years.

Figure 15 : Euro Stoxx implied variance levels since 2000
Implied variance
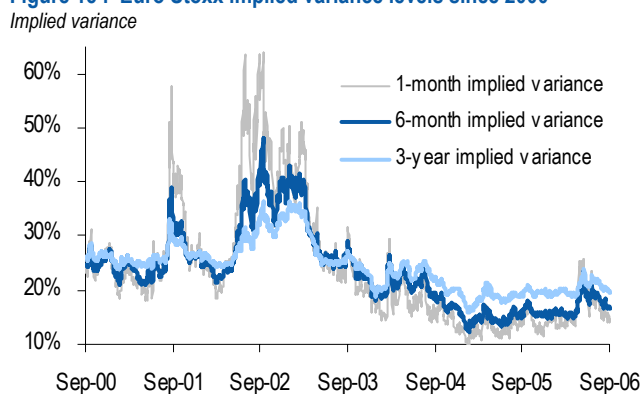
Source : JPMorgan

Figure 16 : VIX (short-dated S&P500 variance) since 1990
VIX level *realised volatility*
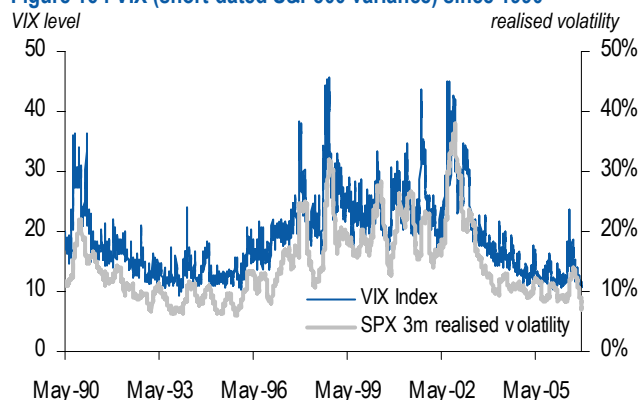
Source : JPMorgan

Single-stock variance swaps generally trade at higher levels than index variance, due to the fact that index volatility will be less that the average of its constituents because of diversification effects. Single stock variance can increase substantially in times of high company uncertainty, for example ahead of an M&A event or following a profit warning or unexpected earnings announcements. See Table 4 below for some examples of index and single-stock variance prices.

Figure 16 shows VIX levels since 1990. The VIX gives a proxy for rolling 1-month maturity variance swap levels on the S&P 500 (see p94). Again **variance swap levels can be seen to be well correlated with realised volatility** both in the short term (including spikes) and in the long run, following the underlying pattern of volatility regimes.

Note also that **variance swaps tend to trade somewhat above levels of realised volatility** – in part reflecting the **volatility risk premium**, with investor risk aversion and hedging programmes in the market supporting demand for index options leading to a persistent overpricing of (index) variance. Below we discuss the two principal reasons why variance swaps tend to trade above realised volatility.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 2: The Variance Swap Market

In general, index variance swap strikes tend to trade above the level of prevailing realised volatility, for two distinct reasons.

1. **Variance swaps are convex in volatility** and the variance swap buyer should fairly pay for this convexity, meaning that equity variance swaps trade above implied ATM volatility.

2. **The volatility risk premium** mentioned above, which is particularly noticeable at the index level, and is manifest in the premium of implied ATM volatility to realised volatility.

Of these two factors, the second applies both to variance swaps and vanilla options, whereas the first applies only to variance swaps. Thus a simple framework has (ATM) vanilla option implied volatilities trading above prevailing levels of realised volatility due to the volatility risk premium; and then variance swap levels trading above ATM options implied volatility due to the additional premium required to pay for the convexity. Typically, at least for short-dated index variance, **the convexity premium is worth perhaps 1-2 vegas** with the **volatility risk premium around 2-3 vegas,** representing the spread of ATM implied volatility over realised volatility.

The following sections consider the drivers of variance swap prices, with the discussion centring around attempts to quantify, and to some extent value, the contributions of these two factors – convexity and volatility risk premium – presented above. We begin by considering the relationship between ATM options prices and variance swap levels for various underlyings, demonstrating that they are well correlated (both across underlyings and over time). The difference comes down to value of the convex payout which is represented by the convexity of the volatility skew. This leads into a discussion of rules of thumb for pricing variance swaps based on vanilla options prices – using assumptions about the skew to produce an estimate for the variance swap price which in theory is priced using the entire skew curve.

We go on to explicitly consider the drivers of variance swap prices: principally realised volatility/variance and the convexity and risk-premium discussed above as the main reasons for variance swaps to trade above these levels. Risk aversion (driving the risk premium) can perhaps be measured in the market as the excess demand for put options or its natural effect – increasing put skews. Note that this kind of increase in risk aversion also increases skew convexity, leading to a higher implied volatility of volatility and more value being assigned to the convex variance swap payout.

In Section 2.6 we explicitly discuss and attempt to measure the volatility risk premium. This risk premium can partly be explained by the asymmetry of returns of volatility (and variance) products: longs occasionally make large gains but mostly suffer small losses, whereas a short position can be likened to selling insurance (similar to selling credit protection) with a steady income punctuated with occasional large drawdowns. Indeed, the **volatility risk premium can be viewed as a potential source of alpha,** exploited by strategies which systematically short volatility, for example selling index variance swaps or traditional index covered call strategies.

Although convexity and the risk premium cause variance swaps to trade above prevailing levels of realised volatility, a simple regression demonstrates that variance swaps tend to give a better estimate of future volatility (after adjusting for the bias) than that given by prevailing levels of realised. This is especially true at short maturities where market information priced into the variance swap market, but not represented in realised volatility can give variance swaps some real forward looking attributes.

In Section 2.8 we turn our attention to convexity, the other cause of variance swaps trading above realised volatility, and the reason why they trade above ATM implied volatility. The value of this convex payout depends on how much future realised volatility is likely to deviate from the variance swap strike, with large potential moves in volatility increasing the value of the variance swap. The convexity premium should thus be linked to the 'volatility of volatility' and by comparing the spread of variance over ATM implied volatility with realised volatility of volatility we find that this is the case. Nevertheless, at least for S&P500 variance, we find like volatility, that 'volatility of volatility' also tends to be overpriced: perhaps of 'volatility of volatility risk premium'?

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 2: The Variance Swap Market

2.3: Variance swaps and option volatilities

Variance swap strikes are well correlated with standard Black-Scholes implied volatility derived from options prices.
This should not be surprising, since both can be interpreted as market estimates of future volatility. However, ATM option volatility reflects the market estimate of future volatility realised around the current level, whereas variance swaps represent the market estimate of variance, independent of future market level. As explained below this means that variance swaps tend to trade somewhat above levels of ATM volatility.

Table 4 shows ATM volatility and variance swap levels for a number of European stocks and indices. We have used indicative mid-market levels as of 24th October 2006. Note that these **variance swap levels are typically about 1-2 vegas above ATM volatility**. Figure 17 and Figure 18 show the relationship between ATM volatility and variance swap strikes for the members of the Euro Stoxx 50. Whilst all variance swaps price above their corresponding ATM volatility, there is variation in the spread, in part due to the differing skews. This spread of variance above ATM volatility becomes larger and more unpredictable at longer maturities, as the effect of the particular shape of the skew becomes more important.

Table 4: ATM volatility and variance levels (as of COB 24 Oct 2006)

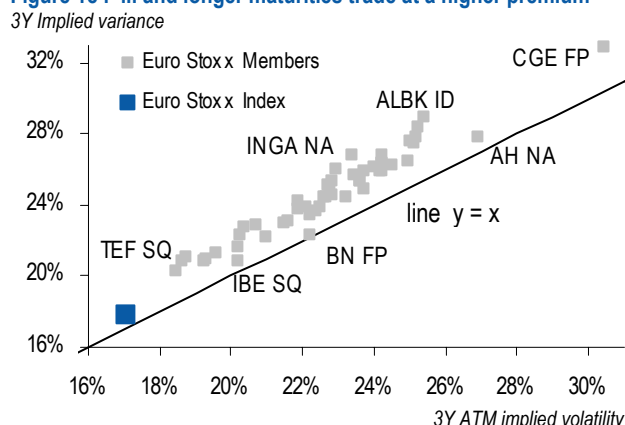|	Ref	Dec-06 ATM implied volatility	Variance swap strike	Mar-07 ATM implied volatility	Variance swap strike	Dec-07 ATM implied volatility	Variance swap strike	Dec-09 ATM implied volatility	Variance swap strike	Dec-11 ATM implied volatility	Variance swap strike
Euro Stoxx 50	4,014.0	12.9%	14.2%	14.7%	15.9%	16.2%	17.3%	18.3%	19.1%	19.8%	20.4%
FTSE	6,182.5	11.3%	12.4%	12.6%	13.8%	14.3%	15.1%	16.5%	16.9%	18.4%	18.4%
SMI	8,662.6	11.9%	12.8%	12.8%	14.0%	13.8%	15.0%	15.3%	16.1%	16.5%	17.1%
DAX	6,247.5	13.8%	14.5%	15.8%	16.6%	17.7%	18.1%	20.5%	20.1%	22.6%	21.6%
Nestle	424.0	13.9%	14.3%	14.3%	14.8%	15.4%	16.3%	16.7%	17.9%	17.0%	18.4%
BP	6.08	17.2%	17.9%	18.3%	19.2%	19.9%	21.2%	21.3%	23.2%	22.0%	24.5%
Axa	30.44	21.9%	23.0%	23.2%	24.5%	24.3%	26.3%	25.7%	29.0%	27.7%	32.4%
Ahold	8.26	27.7%	28.5%	26.4%	27.2%	27.1%	28.1%	27.3%	28.2%	28.2%	29.1%
Rio Tinto	28.50	32.9%	34.6%	32.9%	34.8%	33.2%	35.5%	34.1%	37.0%	34.7%	38.2%

Source : JPMorgan

For example note that at 3-year maturity (Figure 18) the variance strikes of Ahold and ING are similar (around 27) even though their ATM implied volatilities differ by more (24% vs. 27%). Whilst this is partly due to the different forwards (ING has a dividend yield of 3.6%; Ahold doesn't pay a dividend) the skew surface of Ahold is relatively flat, whereas that of ING is more convex, due to much steeper put skews. Note also that at 6-month maturity the index variance appears somewhat rich in comparison to its ATM volatility. This is caused by the index having higher put skews than the single-stocks. However, at 3-year maturities, the greater upside convexity of the single-stocks becomes more important and the index variance appears to be priced in line with its ATM volatility in comparison to the single stocks.

Figure 17 : 6m variance swaps trade above 6m ATM implied volatility
6M implied variance
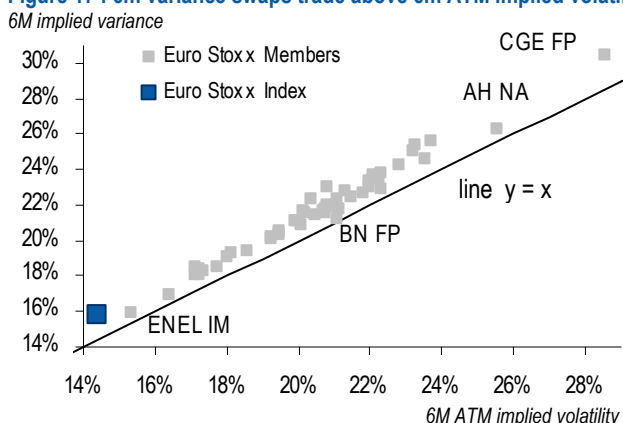
Figure 18 : ... and longer maturities trade at a higher premium
3Y Implied variance

-Source : JPMorgan

Source : JPMorgan

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 2: The Variance Swap Market

Although they move in close step with changes in ATM volatility, as noted above, **variance swaps invariably trade *above* the level of ATM implied volatility** of the corresponding maturity. Why is this? There are two ways of looking at this which both turn out to be the same explanation.

- Variance swaps trade above ATM volatility because you pay extra for the convexity of the variance swap: **the gain from an increase in volatility is more than the corresponding loss from a decrease in volatility** (Figure 9, Figure 10). This does not come free.

- The theoretical price of the variance swap is calculated from the prices of a replicating portfolio of options, so the variance swap strike can be thought of as a weighted average of vanilla option volatilities (see e.g. p82). In the presence of **skew and skew convexity**, average volatilities will usually be above ATM volatility, making the variance swap more expensive.

Ultimately both of these explanations come down to the fact that **a long variance swap is also long volatility of volatility**, due to the convex nature of the payoff discussed previously. This convex payoff is paid for through a replicating portfolio of options, across the whole range of strikes, which in effect price in this volatility of volatility (see Part 4 for details).

Variance swap prices arise from a kind of weighted average of volatilities across the skew surface. For short dates where the skew is relatively linear around ATM volatility, and put skews are more significant than call skews, the price can be thought of as a function of ATM volatility level and the slope of skew. This assumption lies behind *Derman's approximation* to the variance swap level, which is particularly useful for estimating variance swaps where a limited number of option volatilities are available. See the next section for more details.

In practice, the contribution of the skew component means that variance swap strikes tend to trade at similar levels to OTM puts – usually around 95-100 strike for maturities of around 1-3 months (Figure 19), and 90-95 strikes for maturities of about a year (Figure 20). The skew, and convexity (in terms of the overall shape of the skew surface) become more important factors for pricing at longer dates, where the probability of reaching more OTM strikes increases (Figure 18).

Figure 19 : The VSTOXX (theoretical interpolated mid variance price) generally trades just below the Euro-Stoxx 1M 95% put volatility

Implied volatility/variance
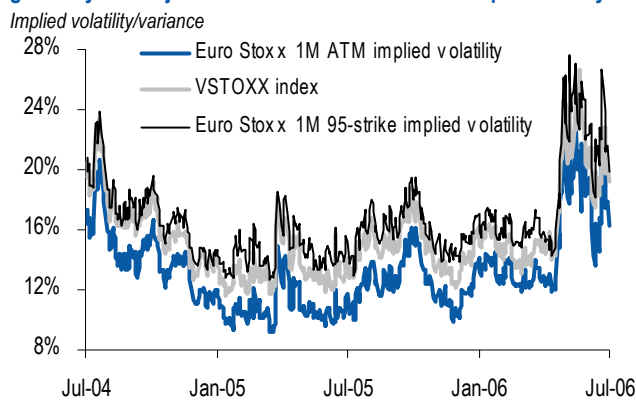
Source : JPMorgan

Figure 20 : 1Y Euro-Stoxx variance swaps trade closer to 90-strike put volatility

Implied volatility/variance
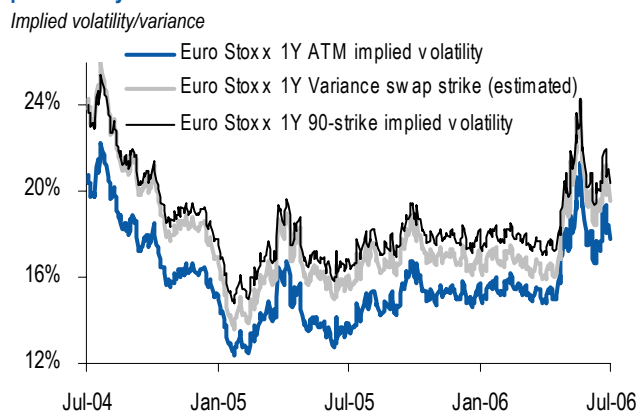
Source : JPMorgan

Variance swap pricing is explored in much greater detail in the Part 4, where we outline how the theoretical variance swap price is arrived at by considering a replication argument using a static portfolio of vanilla options. This enables the variance swap price to be expressed as a weighted average of option volatilities across the skew surface, with the weighting relating to the probability of the corresponding option expiring in the money. **This makes explicit the precise impact of skew and convexity on variance levels**. See p82 for details.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

2.4: Pricing rules of thumb

Whilst it is necessary to have prices available for the entire strip of (OTM) options in order to calculate the true theoretical price of a variance swap, **reasonable approximations for variance swap prices can be made under certain assumptions about the skew**. In this section we introduce some commonly used approximations obtained by assuming **linear, log-linear** and **quadratic** skews curves. We discuss their effectiveness and applicability and give examples of their use.

Flat skew

In the hypothetical case where the skew surface is flat meaning that all strikes trade at identical implied volatilities – the variance swap level will be the (constant) implied volatility level. However, this is not a particularly useful observation because a flat skew implies zero volatility of volatility (i.e. constant volatility) in which case volatility cannot change and the p/l of a variance swap could only be zero.

Linear skew

If the skew is assumed to be linear, at least for strikes relatively close to the money, then **Derman's approximation** can be used (*See "More than you ever wanted to know about volatility swaps", Derman, 1999*). This presupposes a linear (put) skew, effectively assuming the call skew is flat, and calculates the variance swap strike as a function of three variables: ATM (forward) volatility, the slope of the skew and the maturity of the swap.

Box 10: Derman's approximation

$$K_{VAR} \approx \sigma_{ATMF} \sqrt{1 + 3T \times skew^2}$$

where K_{VAR} is the strike of the variance swap

σ_{ATMF} is the implied volatility of the forward strike

T is the maturity in years

skew is generally taken to be the slope of 90/100 skew

In practice, **this approximation tends to work best for short-dated index** variance (up to about 1-year), where put skews are often relatively linear and call skews relatively flat, at least close to the money. As maturity increases and the OTM strikes have a greater effect on the variance swap price, the contribution of the skew becomes more important (Figure 21), but the inability of the approximation to account for the skew convexity can make it less accurate. Similarly, for single-stocks, where the convexity can be much more significant, even at shorter dates, the approximation can be less successful. Where the estimate is accurate it tends to be closer to the traded bid than the offer (Figure 22).

Figure 21: Variance swap prices are more sensitive to skew at longer maturities

Variance swap price, Derman's approximation (ATM volatility = 20%)

Source : JPMorgan

Figure 22: Derman's approximation has been relatively accurate for index variance, but can under-estimate the variance swap mid.

Indicative Euro Stoxx 6-month variance swap prices

Source : JPMorgan

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Example 10: Derman's approximation

An index is trading at 100, with the 6-month forward price at 102.5. The 90-strike volatility is 26%, 100-strike volatility is 22% and the ATM forward volatility (102.5-strike volatility) is 21%.

$$6\text{-month volatility surface:} \quad 90\text{-strike:} \quad 26\%$$

$$100\text{-strike:} \quad 22\%$$

$$102.5\text{-strike:} \quad 21\%$$

Slope of the skew = (26% - 22%)/10% = 0.4.

Maturity, T=0.5,

$1 + 3T$ x $skew^2$ = 1.24

Therefore K_{VAR} = 21 x $sqrt(1.24)$

= 23.38, which is 2.4 vegas above the ATM forward volatility

Table 5-Table 8 show volatilities, variance swap levels and their approximations for a number of European stocks and indices as of 23 October 2006. Note that consistent with the discussion above, Derman's approximation is relatively accurate for indices in the shorter maturities. For single-stocks and longer-dated index estimates, the approximation is less successful.

Table 5: Dec-06 (2-month) volatility, skew and variance

	ATM volatility	ATMF volatility	Skew (90/100)	Variance estimate (Derman)	Actual Variance Swap
SX5E	12.9%	12.8%	5.8%	13.7%	14.2%
FTSE	11.3%	11.2%	5.9%	12.0%	12.4%
SMI	11.9%	11.8%	5.1%	12.5%	12.8%
DAX	13.8%	13.5%	5.3%	14.3%	14.5%
NESN VX	13.9%	13.8%	3.1%	14.1%	14.3%
BP/ LN	17.2%	17.2%	2.4%	17.4%	17.9%
CS FP	21.9%	21.9%	3.1%	22.3%	23.0%
AH NA	27.7%	27.7%	1.8%	27.8%	28.5%
RIO LN	32.9%	32.8%	1.9%	33.1%	34.6%

Source : JPMorgan; Data from 23 October 2006.

Table 6: Mar-07 volatility (5-month), skew and variance

	ATM volatility	ATMF volatility	Skew (90/100)	Variance estimate (Derman)	Actual Variance Swap
SX5E	14.7%	14.4%	3.7%	15.5%	15.9%
FTSE	12.6%	12.4%	3.9%	13.5%	13.8%
SMI	12.8%	12.7%	3.6%	13.6%	14.0%
DAX	15.8%	15.3%	3.7%	16.4%	16.6%
NESN VX	14.3%	14.2%	1.9%	14.5%	14.8%
BP/ LN	18.3%	18.2%	1.6%	18.5%	19.2%
CS FP	23.2%	23.0%	2.1%	23.6%	24.5%
AH NA	26.4%	26.3%	1.1%	26.5%	27.2%
RIO LN	32.9%	32.8%	1.3%	33.1%	34.8%

Source : JPMorgan; Data from 23 October 2006.

Table 7: Dec-07 (c. 1-year) volatility, skew and variance

	ATM volatility	ATMF volatility	Skew (90/100)	Variance estimate (Derman)	Actual Variance Swap
SX5E	16.2%	16.0%	2.2%	17.3%	17.3%
FTSE	14.3%	13.8%	2.5%	15.2%	15.1%
SMI	13.8%	13.7%	2.1%	14.6%	15.0%
DAX	17.7%	16.7%	2.3%	18.1%	18.1%
NESN VX	15.4%	15.4%	1.1%	15.7%	16.3%
BP/ LN	19.9%	19.7%	1.2%	20.2%	21.2%
CS FP	24.3%	24.2%	1.4%	25.0%	26.3%
AH NA	27.1%	26.9%	0.8%	27.2%	28.1%
RIO LN	33.2%	32.9%	0.9%	33.4%	35.5%

Source : JPMorgan; Data from 23 October 2006.

Table 8: Dec-09 (c. 3-year) volatility, skew and variance

	ATM volatility	ATMF volatility	Skew (90/100)	Variance estimate (Derman)	Actual Variance Swap
SX5E	18.3%	18.1%	1.4%	19.6%	19.1%
FTSE	16.5%	15.7%	1.6%	17.5%	16.9%
SMI	15.3%	15.1%	1.3%	16.3%	16.1%
DAX	20.5%	18.8%	1.4%	20.4%	20.1%
NESN VX	16.7%	16.6%	0.7%	17.0%	17.9%
BP/ LN	21.3%	21.0%	0.9%	21.9%	23.2%
CS FP	25.7%	25.7%	1.0%	26.8%	29.0%
AH NA	27.3%	26.9%	0.5%	27.2%	28.2%
RIO LN	34.1%	33.6%	0.6%	34.2%	37.0%

Source : JPMorgan; Data from 23 October 2006.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 2: The Variance Swap Market

Log-linear skew

In reality the volatility skew is not linear across all option strikes – and more accurate variance swap approximations can be used. The **log-linear model** described below is one possibility.

Box 11: Log-linear skew

Suppose that the skew curve is log-linear of the form:

$$\sigma(K) = \sigma_{ATMF} - \beta \ln(K/F),$$

       where F is the forward price and K is the strike

       and β represents the slope of the skew

Then: $K_{VAR} \approx \sqrt{\sigma_{ATMF}^2 + \beta\sigma_{ATMF}^3 + \dfrac{\beta^2}{4}\left(12\sigma_{ATMF}^2 T + 5\sigma_{ATMF}^4 T^2\right)}$

See "Just what you need to know about variance swaps", JPMorgan, February 2005 for further details

Example 11: Log-linear skew

An underlying is trading at 100, with 6-month forward at 102.5. 90-strike volatility is 26%, 100-strike volatility is 22% and the ATM forward volatility (102.5-strike volatility) is 21%.

      6-month volatility surface: 90-strike: 26%

                                          100-strike: 22%

                                          102.5-strike: 21%

$$\beta = \frac{-4\%}{\ln(0.9)} \approx 0.38, \, T = 0.5$$

Then the log-linear approximation gives

K_{VAR} = 23.55%, in this case slightly higher than the value given by Derman's approximation in Example 10

Gatheral's formula

A further approximation to the variance swap price can be obtained from *Gatheral's formula* which expresses the variance swap strike as an integral (or summation) of the implied volatilities across the entire range of strikes. This formula characterises the skew curve in terms of the Black-Scholes *d2* parameter, which measures the 'moneyness' of the associated OTM option. Implied volatilities of traded options can be used to parameterise the skew curve in this form, by converting the strikes into their *d2* equivalents. This leads to potentially powerful methods of variance swap approximation: by fitting a quadratic, or higher order polynomial to the skew surface parameterised in terms of *d2*, it is then possible to directly calculate a theoretical variance swap price from this parameterisation.

For example, if the skew curve is quadratic in the variable $z = d2 : \sigma^2(z) = \sigma_0^2 + \alpha z + \beta z^2$ then it can be shown that the theoretical variance swap strike is $K_{VAR} = \sigma_0^2 T + \beta T$, i.e. in 'd2-space' the linear component of the skew, α, has no effect on the variance swap price: only on the base level of volatility σ_0, and the convexity β affect the swap strike See p82 and *"The Volatility Surface", Gatheral, 2006 for further details*

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

2.5: What drives variance swap levels?

Since **variance swap levels move in close step with options implied volatility**, they are in large part driven by the same factors which drive these volatilities, as well as by the shape of the skew curve. Thus the drivers of variance swap levels (at least theoretical levels) are essentially the same drivers as for options volatilities and associated skews. However, supply and demand issues particular to the variance swap market can cause levels to differ slightly from their theoretical values, usually within the arbitrage corridor defined by bid/offer spreads.

In common with options volatilities, **one of the most important drivers of variance swap levels, at least at shorter maturities, is historical realised volatility** (e.g. see Figure 16). The correlation between strikes of 1-month variance swaps and the volatility realised over the previous month is 0.91 ($R^2 = 0.83$, Figure 23). However, as maturity increases, the correlation (and beta) between variance swap prices and realised volatility decreases (Figure 24). Even at short maturities, realised volatility is not the only driver; evidenced by the fact that short-dated variance is superior to realised volatility as a predictor of future volatility (see p37).

Figure 23: Short-dated implied variance is driven by short-dated realised volatility

1M Implied/variance

$y = 0.848x + 0.057$

$R^2 = 0.829$

trailing 1M realised volatility

Source : JPMorgan (Euro Stoxx 50 data since August 2000)
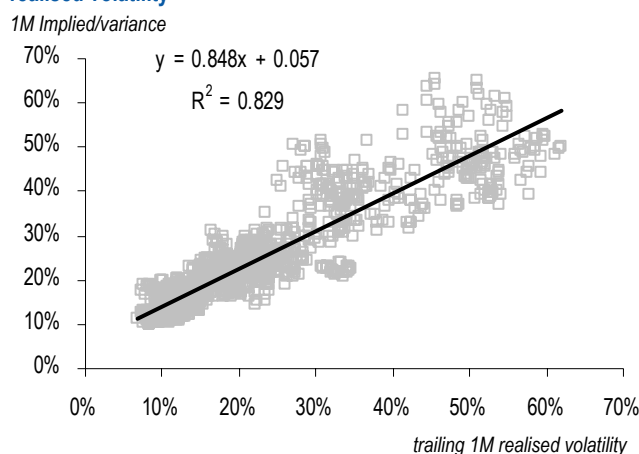
Figure 24 : Short-dated variance is most sensitive to realised volatility

Correlation / beta

R-squared of implied v ariance v s 1M realised v olatility

beta of implied v ariance v s 1M realised v olatility

variance swap maturity

Source : JPMorgan (Euro Stoxx 50 data since August 2000)
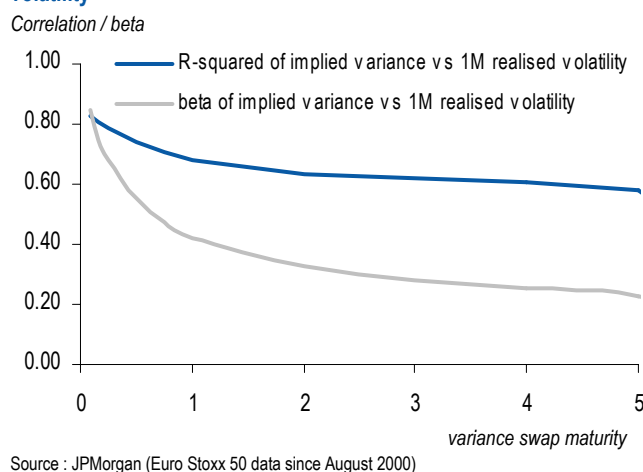
Clearly **risk-aversion is also an important driver**, dictating supply and demand factors for volatility products, variance swaps included. Since variance swap prices in effect carry a significant skew component, demand for protection in the form of put options will clearly drive up put skews, as well as volatilities, and thus have a feed-through into variance swap prices. The skew (in fact more correctly the skew convexity) directly influences variance swap levels, with its effect becoming more pronounced at longer maturities (Figure 21). In addition, at longer dates **structured product flow** and the longer-term **macro-economic outlook** also become important considerations.

Ultimately volatility itself is directional, at least over short time frames, with volatility tending to increase if the market sells off. In this sense, variance swaps can be driven in the short term by the (inverse of) market level (Figure 25, Figure 26). This is most clearly seen during market shocks when volatility tends to spike as the underlying sells off.

Finally, it should be added that variance swap flows can have a feedback effect into the rest of the market. With variance swaps currently accounting for perhaps 10-30% of the traded vega demand, variance swap hedging by market makers can influence implied volatilities and particularly skews. Indeed under certain conditions the effects of variance swap hedging can potentially influence the underlying market. See p90 for a more detailed explanation.

Whatever the drivers of variance swap levels, we do observe that they tend to trade above realised volatility, most notably at the index level. Whilst this variance premium is partly due to the convexity of the variance swap payout, the remainder is the so-called **volatility risk premium**, which is the subject of the next section.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Figure 25 : Over short time periods, variance swap levels are inversely correlated with market level ...
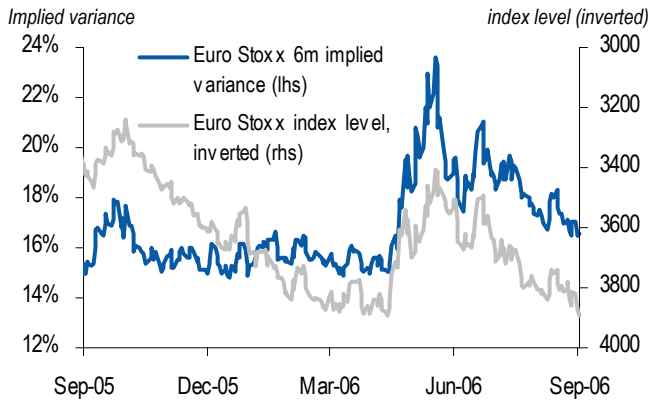
Figure 26 : ... with the relationship changing as volatility switches regimes
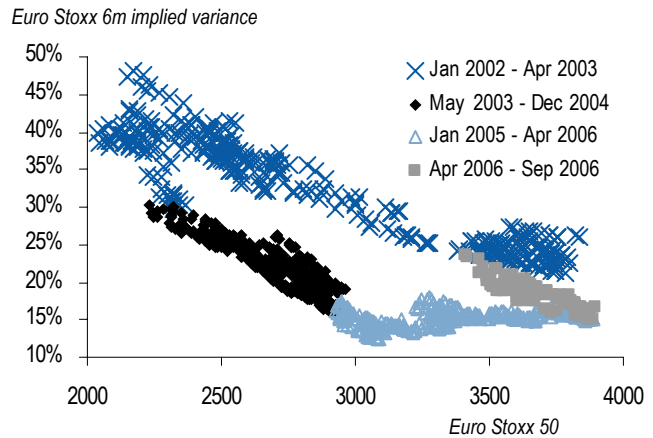
2.6: The volatility risk premium

The variance swap strike represents the market's fair price for exposure to realised variance over a specified period of time. However **variance swap levels tend to trade consistently above comparable realised volatility** (Figure 23), with only part of the spread attributable to the convexity of the variance swap payout. Back-testing shows that **long volatility positions, at least on index underlyings, are biased to make a loss** (Figure 27). Conversely, short index volatility positions are, on average, profitable. **This bias is referred to as the volatility risk premium.** In this section we investigate why this bias exists, what it represents and how it changes over time.

Figure 27 : Average returns from long variance are negative…
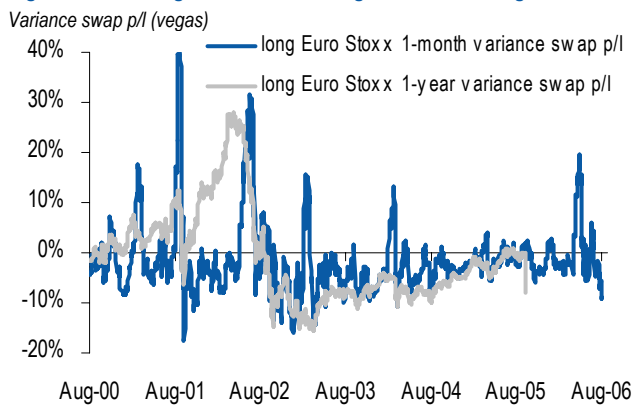
Figure 28 : …but the tail is heavily skewed to the upside
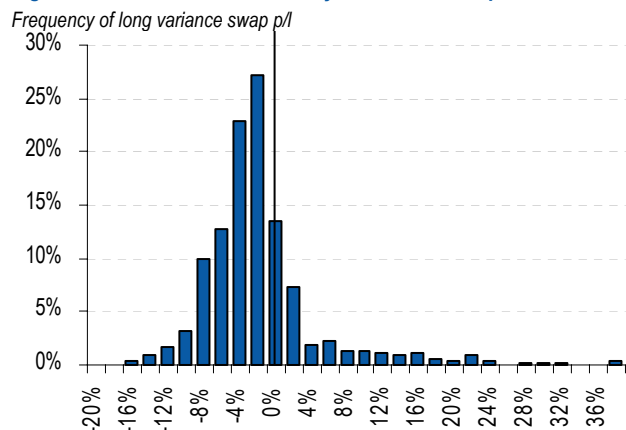
Investors in the real world are not risk-neutral and, in general, are prepared to pay slightly higher than the theoretical fair value for exposure to long volatility, often in the form of protection. This is in some sense understandable, since returns from short volatility positions are limited to the premium received in the case of options, or half the strike multiplied by the vega-notional in the case of short variance swaps. On the other hand, **returns from long volatility positions can be potentially unlimited**. Thus although expected returns from long volatility positions are, on average, negative, the potential rewards can be high if timed correctly. In this sense **long volatility is like buying insurance**: paying a relatively small premium for a potentially large payout if things go wrong, but expecting to forfeit some, or all, of the premium on most occasions. Conversely selling volatility can be likened to selling insurance (or credit protection): potentially very

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 2: The Variance Swap Market

risky, but in the long run likely to make lucrative returns, providing enough capital is available to absorb the occasional inevitable losses.

Example 12: The insurance-like nature of long volatility positions

A variance swap is struck at 10% for €100K of vega notional.

Realised volatility → 0%:   loss for the long is €0.5M (this is the maximum loss, not €1M, due to the favourable effect of convexity for the long, see Example 3)

Realised volatility → 20%:   profit from the long is €1.5M

Realised volatility → 40%:   profit from the long is €7.5M

Example 12 above shows why sellers often fail to drive variance swap prices below certain thresholds (e.g. around 12% on the Euro Stoxx), even when short-dated realised volatility drops to very low levels. For example, the y-intercept on Figure 23 suggests an absolute floor of 6% for implied variance in the hypothetical case that realised volatility drops to zero.

The above discussion illustrates that the returns from long variance positions, although on average negative, have long upside tails (Figure 28). That is, **returns are *not* normally distributed: long positions tend to lose a little and often, but can occasionally make large profits.** Short positions usually make modest gains, but can be exposed to very large losses if volatility spikes up significantly (see Section 3.3).

From the point of view of a short volatility position, the observed premium of implied over realised volatility is often referred to as the *carry* of the short position and can be a good source for generating alpha, especially in the form of strategies continuously selling variance through variance swaps (see p47). This strategy can work best as a rolling short-maturity strategy, with the occasional (large) losses from being short volatility diversified away over time. This is somewhat analogous to the vanilla options equivalents: straddle overlays or covered call overwrite strategies (See e.g. *"Calling All Overwriters" JPMorgan, August 2002*), or to selling credit protection, which usually generates positive carry but has a small chance probability of large losses.

We can attempt to quantify the volatility risk premium by computing the spread of implied variance over contemporaneous realised. This is complicated by a number of factors.

Firstly, real variance swap data is limited, though indicative variance swap levels back in the 1990's can be inferred from levels of the corresponding volatility indices (VIX for the S&P, VSTOXX for the Euro Stoxx – see page 94 for more details). These volatility indices represent rolling 1-month mid implied variance levels on their respective indices but tend to be slightly higher (around 0.5 – 2 vegas) than real variance swap prices.

Secondly, part of the risk premium is attributable to the convexity of the variance swap payout. Whilst we make no attempt here to remove this from this characterisation of the risk premium, any backtest of variance swap payoffs will implicitly take this into account (e.g. Figure 28). We have estimated the convexity premium for 1-month variance swaps at somewhere around 1% (1 vega) under normal market conditions (see p 38). Finally, since volatility can spike from time to time, an average spread of implied to realised volatility needs to be taken over a suitable time period to smooth out some of the noise.

Given these caveats, we characterise the risk premium as the 1-year rolling average of the spread between (estimated) 1-month variance swap strikes and 1-month realised volatility. Values seem to be relatively stable over time with the S&P risk premium varying between 2% and 9% since 1990, with an average value of 5%. Euro Stoxx since 1999 is similar with an average risk-premium of about 4%. The recent peak in risk premiums for both indices occurred in 2004, when volatility had subsided from its highs but risk-aversion remained elevated. Since then the risk premium has fallen from about 7% to 2.4% over 2 years on the Euro Stoxx (Figure 29), and from 5.8% to 2.8% on the S&P (Figure 30). Note that these measures include a convexity premium compensating for the negative convexity of a short variance swap as mentioned above.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 2: The Variance Swap Market

Figure 29 : Variance risk-premium on the Euro Stoxx 50
Volatility/ points
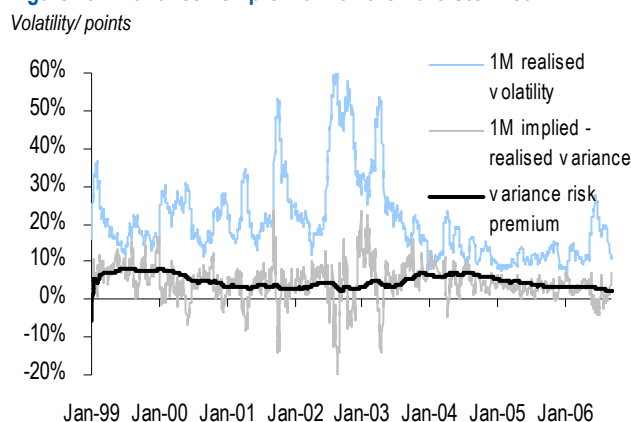
Source : JPMorgan

Figure 30 : Variance risk-premium on the S&P500
Volatility points
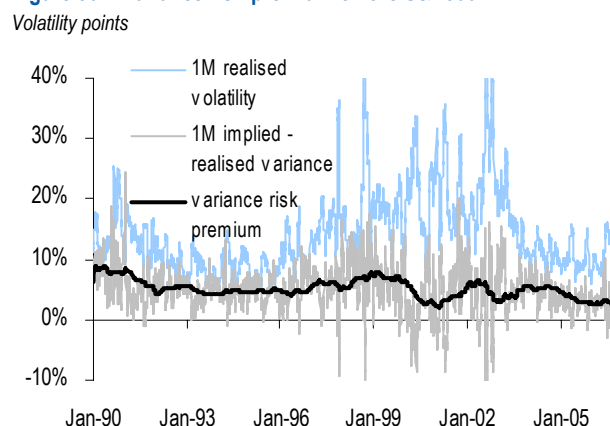
Source : JPMorgan

2.7: Variance swaps as predictor of future volatility

In theory the variance swap strike represents the market prediction of realised variance over the term of the swap. This is complicated somewhat by supply and demand factors, especially at longer maturities where prices can be driven by structured product flow. At shorter maturities risk-aversion and protection buying can play a role in biasing the variance swap strike to be above a 'fair' expectation of future realised variance as explained in the section above.

Nevertheless, at least **at short-dates, implied variance tends to do a relatively good job of forecasting future realised variance.** Using data since August 2000, the R^2 between 1-month Euro Stoxx variance swap strikes and subsequent 1-month realised volatility is 0.71 (Figure 32). This compares with an R^2 of 0.61 between 1-month *realised* volatility and that realised volatility 1-month later (Figure 31). This suggests that implied variance better predicts changes in future realised volatility in comparison with previous realised volatility. Note that here we are correlating the square roots of implied and realised variance, not least because units are more tractable, although results using variance itself are similar.

As discussed in the previous section, variance is itself a biased estimator, as it generally overestimates realised volatility. Given historical data, **a simple model of future realised volatility as 0.9x variance strike provides the best fit with the data** (Figure 32). We find similar results (implied variance a good predictor, with a slight edge over trailing realised) using VIX, VSTOXX and VDAX levels as predictors for future realised volatility on the S&P, Euro Stoxx and DAX respectively, using data back to the inception of the relevant volatility indices.

Given that implied variance appears to have better predictive power than prevailing realised volatility, what extra information could the variance swap and option markets be pricing in?

- **Mean reversion of volatility**: the implied variance can take account of the current medium term "average" of volatility and the fact that volatility tends to be mean-reverting within a regime. If volatility over the last year has been say 30%, but over the last month 10%, 1-month implied variance will likely be priced well above 10%.

- **Knowledge of forthcoming events**: volatilities (especially short-dated) will tend to be marked up ahead of earnings seasons or other events likely to increase volatility. In this sense implied variance can be truly forward looking, whereas past realised volatility is not.

- **Risk aversion in the market**: If investors are nervous and willing to pay extra for protection, they are perhaps also prepared to liquidate their positions quickly in a downturn and increase market volatility. Thus an increase in implied variance could precede an increase in realised volatility if the market sells off.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Figure 31 : 1-month realised volatility is relatively well correlated with realised volatility over the next month…

subsequent realised 1-month volatility

$$y = 0.7843x + 0.0452$$
$$R^2 = 0.6157$$

1-month realised volatility
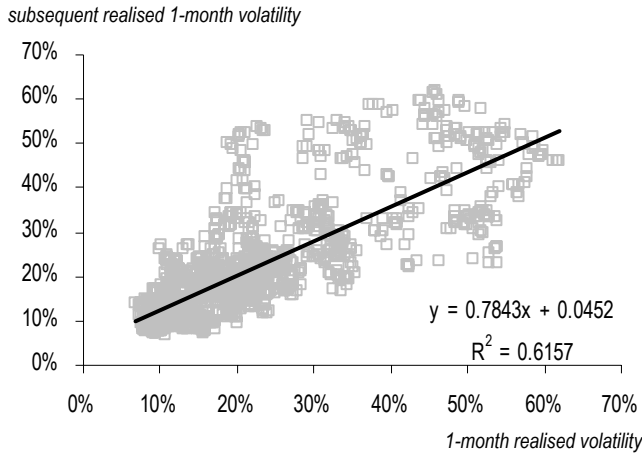
Source : JPMorgan (Euro Stoxx 50 data since Aug 2000)

Figure 32 : … but the correlation of 1-month implied variance with subsequent realised variance is somewhat stronger

subsequent realised 1-month volatility

$$y = 0.9037x - 0.0022$$
$$R^2 = 0.7087$$

1-month implied variance
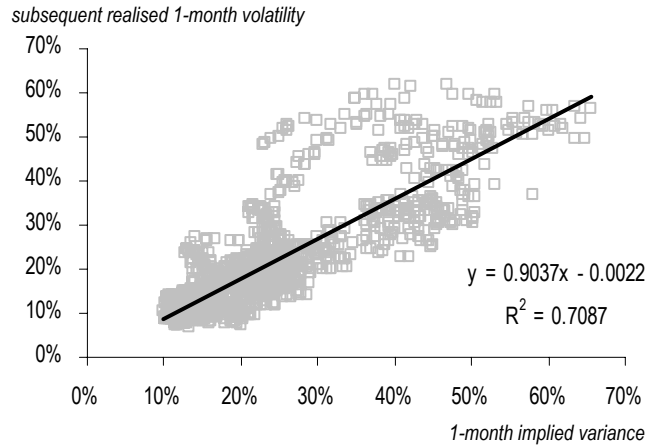
Source : JPMorgan (Euro Stoxx 50 data since Aug 2000)

Instead of trying to predict future 1-month realised volatility using volatility realised over the past month, we could use a combination of realised volatilities of different maturities (e.g. 5-day, 1-month, 3-month and 1-year). Even choosing the weightings of these by performing an in-sample regression still does not give as good a correlation with future realised as using the implied variance. **That is, variance swap prices really do seem to be adding additional information to that encapsulated in backward looking volatility measures**.

Although short-dated implied variance is a relatively good predictor of subsequent realised volatility, the same is not true for longer maturities. Further in the future is harder to predict. Moreover, any edge of implied over realised volatility as a predictor for future realised effectively disappears for maturities longer than 3-months (Figure 34).

Figure 33 : Correlation between implied variance and past and future realised variance
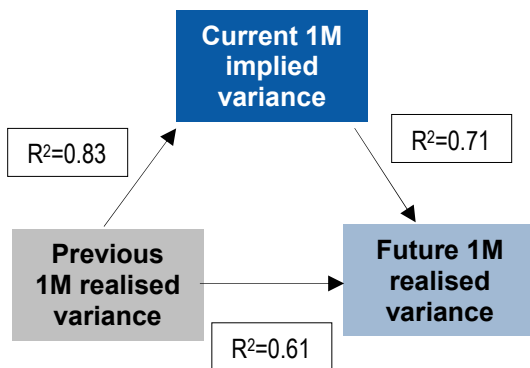
Source : JPMorgan (Euro Stoxx 50 data since Aug 2000)

Figure 34 : Correlation of implied variance with subsequent realised of similar tenor decreases with maturity.

r-squared vs. future realised volatility

maturity (months)

Source : JPMorgan (Euro Stoxx 50 data since Aug 2000)

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033
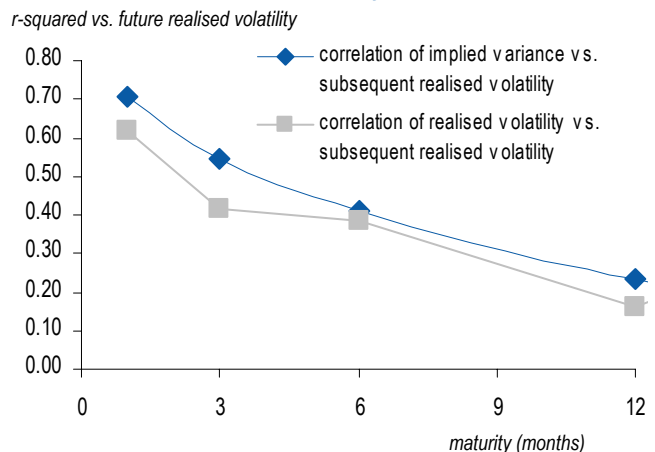
2.8: Is variance swap convexity fairly priced?

As previously discussed (see p29) **variance swap strikes usually trade** *above* **corresponding at-the-money implied volatilities**. This is essentially due to convexity, which can be thought of in the form of the extra value of the convex variance swap payout. Equivalently, convexity can be thought of as the extra volatility paid for the OTM options needed to replicate the variance swap, as represented by a convex skew. In this section we attempt to quantify the value of the convex variance swap payout in order to account for the observed premium of variance over ATM volatility.

The convex payout means that a variance swap must be worth more than a hypothetical 'volatility swap' struck at the same level. How much more the variance swap is worth depends on how much volatility is likely to change (volatility of volatility) and hence how much effect variance swap convexity will have. If volatility is unlikely to change much, say a vega or two either way (e.g. Figure 9), then the convexity of the variance swap will be worth little and the variance swap should trade only just above ATM volatility. In contrast **if the likely range of volatility is high, then the convexity is likely to be valuable** (e.g. Figure 10), and the variance swap should price this in by trading significantly above ATM volatility.

We can **attempt to quantify the price paid for the convexity by computing the difference between the variance swap strike and the ATM volatility**, implicitly assuming that we could trade a "volatility swap" linear in volatility, struck at ATM implied volatility. We can then compare the (linear) payoff from the volatility swap with the (convex) payoff of the variance swap to estimate the value of convexity.

Figure 35 : Measuring the value for variance swap convexity
Extra vega p/l due to convexity (var swap – vol swap p/l)

change in volatility from strike level
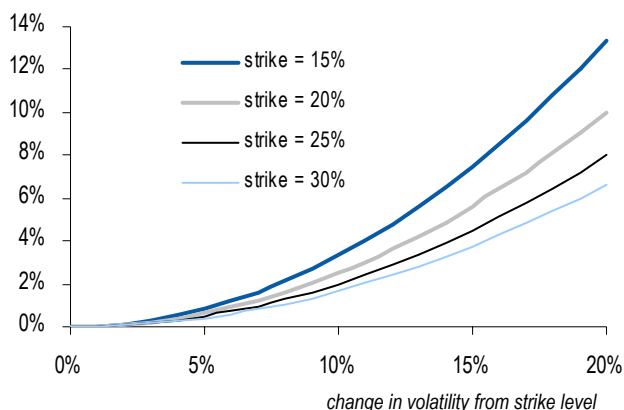
Source : JPMorgan

Figure 36 : 1-month implied variance has some predictive power for future realised variance
subsequent realised 1-month volatility

$$y = 0.9037x - 0.0022$$
$$R^2 = 0.7087$$

1-month implied variance
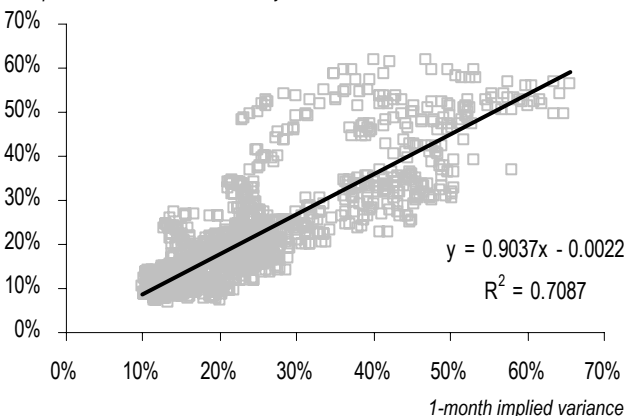
Source : JPMorgan (Euro Stoxx 50 data since Aug 2000)

As demonstrated above, **short-dated variance swaps have some value as a predictor of future volatility**. As well as looking at the regression R^2, we can calculate the standard error of the future volatility as estimated by the variance swap strike (Figure 36). This tells us the average RMS error in the 1-month variance strike as a predictor of realised volatility over the coming month. The idea is that this 'average error' is giving information about the historical variability of the future realised volatility around the estimate provided by (0.9x) the variance swap strike. Taking data from the last 6-years, this standard error is around 6% (6 vegas), that is, the standard deviation of the difference between implied variance and future realised variance is 6%.

This 6 vegas represents an 'expected' range of values for future volatility, thus determining the magnitude of the convexity of the variance swap payoff. How much is the 6% volatility of volatility worth? It depends on the strike of the variance swap, **but for strikes in the region of 15-20% the convexity is worth an extra 1-1.5%** (Figure 35), which corresponds fairly closely to what is observed in practice for short-dated Euro Stoxx variance.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

We can also calculate an **'implied' variability of volatility** by considering the spread between short dated ATM volatility and variance strikes. By assuming that a hypothetical 'volatility swap' would be struck at ATM volatility, and the only difference between the price of this instrument and a standard variance swap of the same maturity is driven by the expected volatility of volatility, we can back out a value for this 'volatility of volatility'. Over the last six years**, the difference between Euro Stoxx ATM volatility and variance swap strikes suggests an average implied 'volatility of volatility' of about 7% - fairly close to the 6% average error of the variance swap in predicting volatility.**

Although, this figure of 6% volatility of volatility for the Euro Stoxx is a long term average (Aug 2000 – Sep 2006), like volatility, it changes over time. By estimating this standard error, on a rolling 1-year basis, we see that it has a minimum of just below 2% and a maximum around 10%, the changes roughly reflecting ambient levels of realised volatility (Figure 39). We observe a similar pattern, but slightly lower range for the S&P (Figure 40). Notice also that over the last few years, Euro Stoxx implied volatility of volatility has been more or less in line with the actual variability of volatility (variance appears fairly priced with respect to ATM volatility), but in the US, S&P implied volatility of volatility has consistently over-estimated the actual variability of volatility (variance is overpriced with respect to ATM volatility; skew/convexity is too high)

Figure 37 : Variance swaps trade at a premium to ATM volatility

Implied volatility / variance
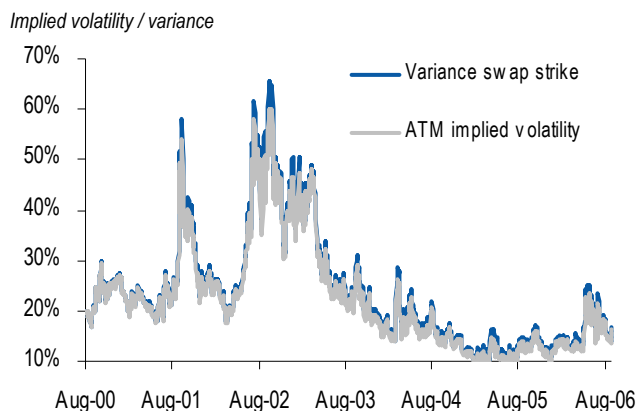
Source : JPMorgan (Euro Stoxx 50 data since Aug 2000)

Figure 38 : The spread between variance and ATM volatility reflects the implied volatility of volatility

Spread / implied volatility of volatility
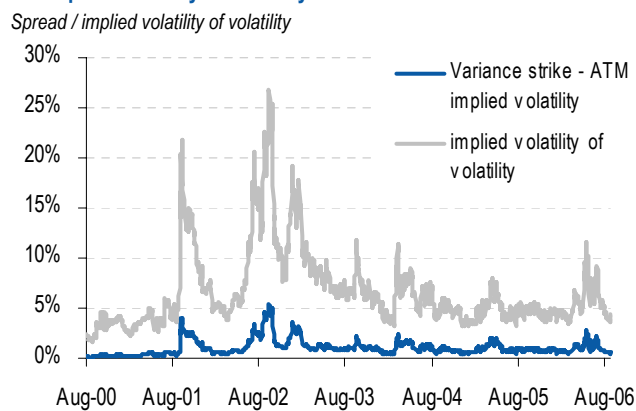
Source : JPMorgan (Euro Stoxx 50 data since Aug 2000)

Figure 39 : Euro Stoxx 50: implied and realised volatility of volatility

Implied/realised volatility of volatility
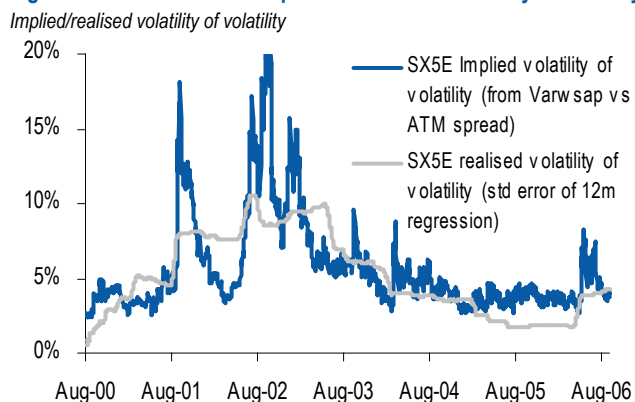
Source : JPMorgan

Figure 40 : S&P 500: implied and realised volatility of volatility

Implied/realised volatility of volatility

Source : JPMorgan

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 2: The Variance Swap Market
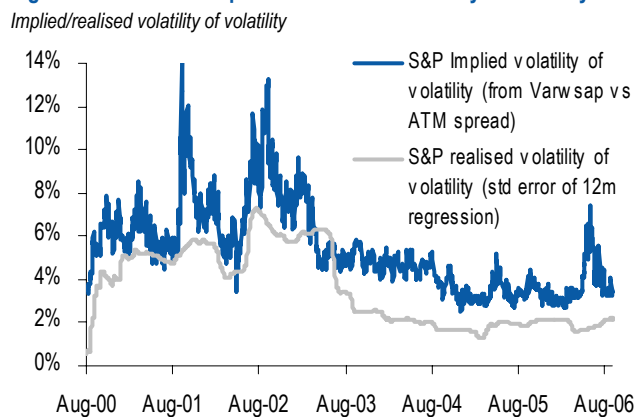
Note that there is a **strong correlation between volatility and the spread between variance swap levels and ATM volatility**. This makes sense, if we believe that the volatility of volatility is correlated with volatility (Figure 41, Figure 42). That is, large (absolute) changes in volatility are more likely when volatility is high that when volatility is low. This is not surprising: when volatility is low, it is generally stable (give or take the odd spike); when volatility is high it tends to swing about more wildly.

Figure 41 : Volatility is correlated with volatility of volatility

Volatility / volatility of volatility
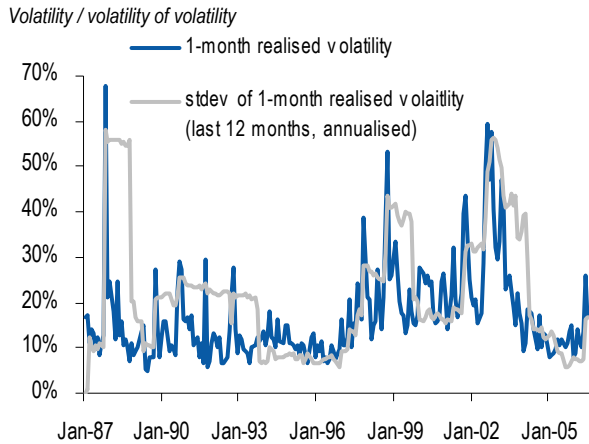
Source : JPMorgan (Euro Stoxx 50 data since 1987)

Figure 42 : Correlation of 1-year realised volatility, with the variability of the volatility realised over each of the 12 months

Stddev of monthly realised volatilities over the year
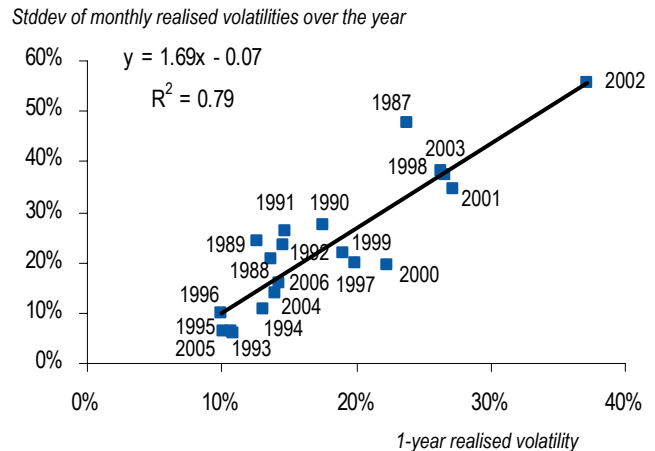
Source : JPMorgan (Euro Stoxx 50 data since 1987)

Estimates for the expected volatility of volatility are used for pricing options on variance. However in the case of variance swaps they are also of use for pricing variance swap caps. **A long capped variance swap can be thought of as a standard variance swap plus a short call on variance, stuck at the cap level** (Figure 43). A standard cap of 2.5x current implied variance strike is relatively far *out-of-the-money*, assuming that the volatility of volatility is not too large. This means that the value of the cap should be relatively small compared to the variance swap strike and should not have a major effect on pricing. However, for a long position, a variance swap with a cap will always be worth less than an uncapped variance swap of the same strike. Therefore capped variance swaps must trade with strikes slightly below their uncapped equivalents – the difference, in theory, representing the current value of the call on variance.

Figure 43 : A capped variance swap is equivalent to a covered call on variance

p/l (vegas)

realised volatility
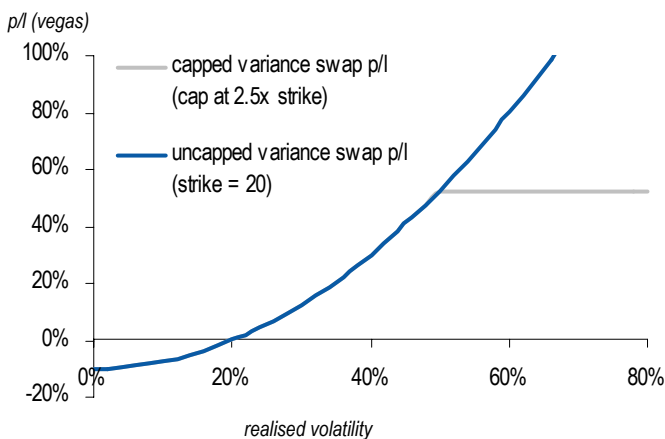
Source : JPMorgan

Table 9: P/L from a long variance swap @ 20 strike

Realised volatility	Realised minus implied volatility	Uncapped variance swap p/l (vegas)	Capped variance swap p/l (vegas)
0	-20%	-10.0%	-10.0%
10%	-10%	-7.5%	-7.5%
20%	0%	0.0%	0.0%
30%	10%	12.5%	12.5%
40%	20%	30.0%	30.0%
50%	30%	52.5%	52.5%
60%	40%	80.0%	52.5%
80%	60%	150.0%	52.5%
100%	80%	240.0%	52.5%

Source : JPMorgan

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

2.9: Variance term structure

The variety of drivers of short and long-dated variance can lead to interesting dynamics of the variance swap term structure curve. **The shape and potential movement of the implied variance curve is important in determining the most promising parts of the curve to buy or sell variance**, including positioning for volatility exposure through forward variance.

Variance swap term structures (in common with ATM volatility term structures) are usually **upwards sloping** (Figure 44) and have a **tendency to flatten, or even invert at the short end, following increases in volatility**. Term structures for variance are generally steeper than the equivalent ATM volatility curves due to the increasing effect of the skew at longer maturities (Figure 44).

The variance swap curve can also be thought of representing the **mean-reverting nature of volatility**. The short end is most sensitive to prevailing levels of realised volatility (Figure 45), whereas the long end (5-years+) tends to be more anchored to some long-term estimate of average volatility. However, longer-dated volatility is also driven by structured product flows and tends to be more susceptible to supply/demand dynamics.

Figure 44 : Euro Stoxx term structure COB 23 October 2006

Implied volatility/variance
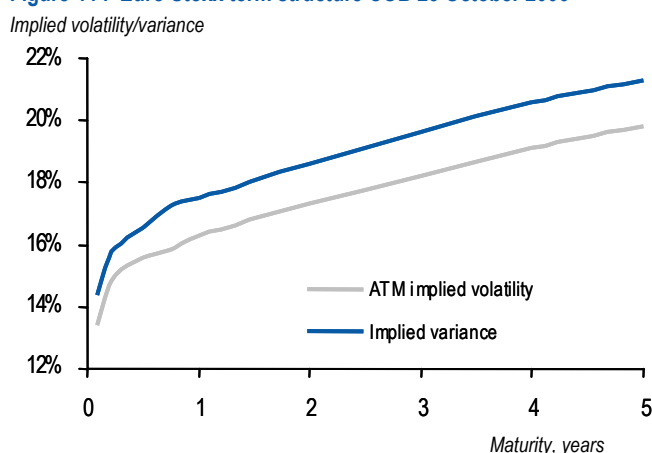
Source : JPMorgan

Figure 45 : Short-dated volatility is driven by short-dated realised

Implied volatility/variance

Source : JPMorgan

An upward sloping curve in a low-to-medium volatility environment can also be thought of pricing in the chance of a change of volatility regime sometime in the future. Volatility regimes represent a shift in the medium/long-term average of realised volatility and are generally considered to last a number of years, sometimes correlated to economic cycles. Over a short time period the chance of a switch to a higher volatility regime may be negligible, whereas even if volatility has been low for 2-3 years, the chance of a change in volatility regime of the course of a 5-year trade could be significant.

Another important point to consider, especially for longer maturity trades, is the effect of the mark-to-market p/l. Whilst a short-dated variance swap is principally exposed to realised variance, **longer dated variance swaps take on significant exposure to changes in implied variance before expiry**. Therefore, if an investor has just bought or sold a 5-year variance swap, their principal exposure will initially be to 5-year implied variance. As observed above, this is driven by factors not necessarily correlated with current realised variance (structured product flows etc.) and the p/l can be unpredictable.

These observations help to explain why investors can be reluctant to sell long-dated variance at seemingly high levels (e.g. 25% for Euro Stoxx 5-years) when the relevant indices have almost never realised these levels over any 5 year period (Figure 46). For example, in 1998 LTCM experienced crippling mark-to-market losses on short positions on long-dated volatility. As it happened, 1998-2003 was probably the worst 5-year period in the last half-century to be short realised volatility, so this particular trade may well have lost even if it had been held to maturity.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Figure 46 : Realised index volatility has rarely exceed 25% over 5 years

5-year realised volatility
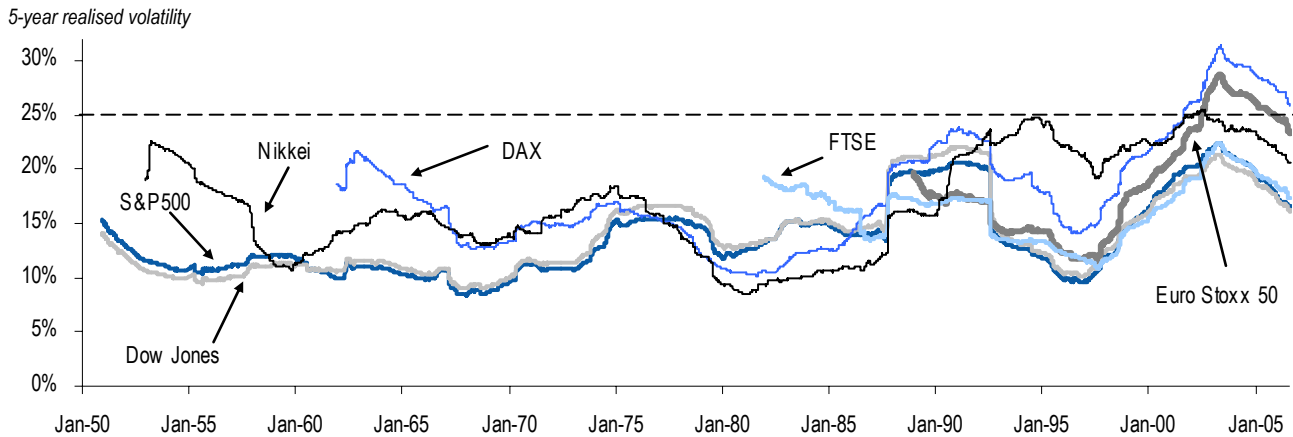
Source : JPMorgan

Investors looking to buy (or sell) volatility through variance swaps will clearly need to consider the best placed part of the term structure curve to target. In general **short-dated variance will have the greatest exposure to realised volatility** (gamma), whereas **longer data variance will initially be more exposed to changes in implied variance** (vega), and movements in the curve. See Section 3.8 for details of using forward variance to position for an increase in volatility.

Typical movements of term structures can be explained in part by the "root-time" rule sometimes used in the market for characterising curve changes. That is for a 'normal' move in volatility the change in implied variance at a given point on the curve will be proportional to the inverse square root of the maturity. For example, if 1-year variance increases by 1%, 3-month variance moves up by 2% and 4-year variance moves up by 0.5%. These kind of movements will have the effect of flattening a normal upwards-sloping term structure as volatility increases, and steepening it if volatility decreases, and changes will be most visible at the shorter end of the curve. Since the relative price for an ATM option is approximately $0.4\sigma\,sqrt(T)$ a root-time move in volatility ensures that prices for (ATM) vanilla options will change in the same proportion across the maturity curve.

Analysis of term structure movements is very important especially for long forward volatility trades. Here, the loss due the slide from moving down a normal term-structure is traded-off against the savings in carry due to the elimination of exposure to realised volatility. Direct trades on the predicted evolution of the variance maturity curve are also possible (Section 3.9).

Figure 47 : Variance term structure moves more at the short end…

Implied variance
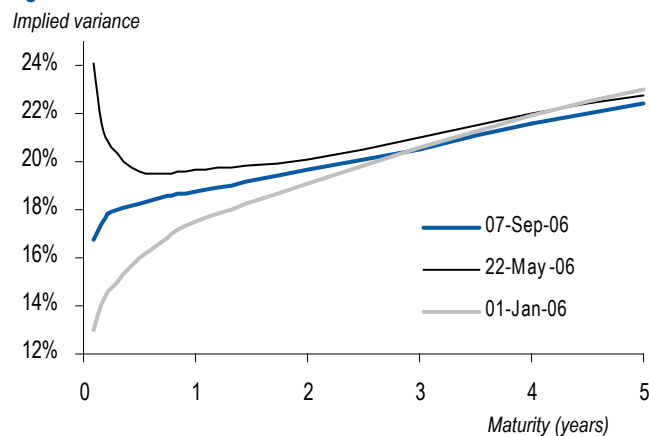
Source : JPMorgan (Euro Stoxx 50 data)

Figure 48 : …as demonstrated by the May 2006 volatility spike when the short end inverted and the longer end flattened

Variance spread

Variance swap maturity
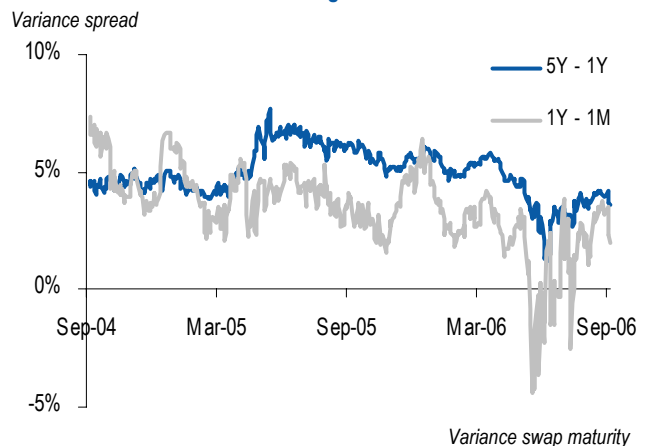
Source : JPMorgan (Euro Stoxx 50 data)

3. Uses of Variance Swaps

Part 3: Uses of Variance Swaps

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 3: Uses of Variance Swaps

Part 3: Uses of Variance Swaps

Overview

In Part 3 we consider uses of variance swaps, outlined briefly in the list below. Taken at face value, a variance swap is just a contract providing direct exposure to (realised) volatility. But why is it useful to have the ability to trade such an instrument? Variance swaps can be used for trading macro volatility views, diversifying returns, hedging specific exposures, and exploiting market inefficiencies. By trading combinations of variance swaps either of different maturities (e.g. forward variance) or on different underlyings (variance spread trades and dispersion trades) it is possible to use variance swaps to create exposure to factors other than realised volatility – e.g. future implied volatility, moves in the volatility term structure or correlation. By combining variance swaps with other (volatility) instruments trading on the same underlying, such as straddles, delta-hedged options, or conditional variance swaps, it is possible to take exposures to other factors such as skew and convexity. Finally, variance swaps can be used together with credit default swaps (CDS) to trade equity volatility against credit.

- **Exploiting a volatility view:** Variance swaps are ideal for taking a direct view on the volatility of an underlying without the path-dependency issues of a delta-hedged option.

- **Specific hedging purposes:** Variance swaps can be used for macro-hedging and also for hedging specific volatility exposures, such as that resulting from structured products or life assurance policies.

- **Rolling short variance:** Short variance swaps can be used to capture the observed equity index volatility risk premium. Rolling short index variance is an attractive systematic volatility strategy from a risk-return perspective.

- **Diversification:** Volatility can be thought of as an asset class in its own right, and as such can act to diversify returns within a portfolio.

- **Index variance spreads:** Variance swaps can be used to trade the spread of volatilities between two indices. Such trades can be thought of as either relative value volatility trades or as 'volatility-beta' trades aiming to profit from a spread of volatilities widening as volatility increases.

- **Relative value single-stock volatility:** Use volatility pairs, or cross-sectional regression volatility models to find rich/cheap single-stock volatilities

- **Variance dispersion and correlation trading**: Trading variance swaps on an index against variance swaps on its constituents provides exposure to equity correlation.

- **Forward variance and volatility spikes:** Long forward volatility can avoid potentially negative carry at the cost of slide down the term structure, and can be a useful way of positioning for volatility spikes.

- **Trading the variance term structure:** Variance swaps can be used to trade the shape of the variance term structure, analogous to the use of spreads and butterflies in fixed income to trade the shape of the yield curve.

- **Skew and convexity trades:** Variance swaps are long skew and convexity. Trading variance against (delta-hedged) vanilla options provides interesting exposures to skew and/or convexity.

- **Cross asset class trades:** Equity Volatility and credit spreads are correlated, both being measures of corporate risk. Variance swaps are useful instruments in debt/equity trades, either at the index or single name level.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan ⬤

Part 3: Uses of Variance Swaps

3.1: Exploiting a volatility view

The most obvious use for variance swaps is to express a direct view on the volatility of the underlying. Delta-hedged options can be used for this purpose but have two principal drawbacks:

- The delta-hedging requires active management (and therefore incurs human and transaction costs).

- The volatility exposure is path-dependent: the p/l of a delta-hedged option depends not only on the amount of volatility realised by the underlying, but where and when this volatility is realised.

Variance swaps have none of these issues and provide a pure exposure to volatility.

At the index level investors may use variance swaps to express a macro view – for example to buy volatility ahead of possible recession, or sell volatility into a temporary spike. Note that the lack of path dependency can make this a more attractive alternative than buying/selling the underlying (or options thereon) since the p/l depends on the outcome over the entire lifetime of the trade and not just where the underlying ends up at maturity.

For example, in a bull market, suppose an investor feels there will be a sharp downwards correction at some stage in the future, say within the next year. The investor does not want to sell his position in the underlying as the market may rally further. Put options are a possibility to use for hedging, but will need to be re-struck in order to roll up the protection as the market rallies. Buying variance can prove profitable if the correction is relatively sudden (volatile) and occurs any time within the maturity of the variance swap. However the investor must be careful to maintain the required volatility exposure and must also consider whether an exposure to realised volatility is desirable, where the usually negative carry from holding a long variance position can be compared to the put premium. To deal with these issues, it is sometimes preferable to use forward starting variance swaps (Section 3.8) to position for a sudden increase in volatility.

Similarly, investors who expect a quiet market, which may gradually trend up, may want to find ways of boosting their alpha based on this view. Straddles are not perfect since a low volatility market often displays trending behaviour; even delta-hedging the straddles will prove sub-optimal as the market moves away from the strikes and the (favourable) exposure to volatility is reduced through the decreased gamma. However, selling variance is an efficient, and non-path-dependent, way of capitalising on this low volatility view. **Indeed back-testing shows that short-variance swap strategies have outperformed short straddle strategies (both delta-hedged and not) in the low volatility, rising market environment since 2004.** See *European Equity Derivatives Weekly Outlook, 6 November 2006* for more details.

On single-name underlyings**, variance swaps could be used to take a view on the uncertainty of a company.** This can be particularly valuable if an event timeline for a company is known, but the likely outcome of the event is unknown, especially if an investor feels that the options market has not correctly priced in the risk associated with such an event. For example, suppose a pharmaceutical company is due to announce the results of a drug trial, but the result of the trial is in the balance. This could be a catalyst for large moves in the stock price in either direction depending on the outcome, which is difficult to capitalise on using the underlying alone. Whilst straddles could be attractive, they may fail to capitalise e.g. if the underlying trends up before the event and then afterwards sells-off suddenly back towards the straddle strike. In this instance variance swaps can be used to implement a view on the uncertain outcome, but otherwise known timing, of an event.

Another use of variance swaps is to take advantage of an expected structural change in a company. For example, suppose an investor expects an industrial conglomerate to sell a more volatile subsidiary. If this part of the business was sold at a fair level, the spot price of the company may remain relatively unchanged, but the volatility characteristics of the company may alter significantly. Hence, whilst it may be far from obvious how to play such an event using the underlying alone, the likely long term volatility expected should fall and selling long dated variance may be profitable.

Variance swaps can also be used by investors to express their view on the likelihood of an M&A event. The nature of variance swap payouts means that large single-day stock moves will generate large p/l for long variance positions, even in comparison to vanilla options. Depending on what is priced in, and the investor's view, long or short variance swap positions may be attractive, either outright or traded against vanilla options such as straddles.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

3.2: Specific hedging purposes

Variance swaps can be useful for hedging purposes. Since volatility is directional (see Section 2.5, Figure 25) at least over short time horizons, variance swaps will likely profit if markets sell-off significantly and especially if this happens suddenly. Being long a variance swap can be compared to owning a put option: a variance swap will likely profit in a market downturn, as will a long put position. But if the market is benign, a variance swap will likely suffer negative carry as realised volatility is below the level bought; likewise a long put will lose its premium.

There are problems however with this approach of attempting to hedge an underlying using variance swaps. Firstly it is not clear how much variance should be used to hedge a fixed quantity of the underlying. If the sell-off is rapid, the variance protection will be effective and only a small notional amount will be required. However, if the underlying slowly drifts down, without any large moves, the long variance swap may fail to provide protection altogether. Sometimes variance swaps do offer more attractive risk-return characteristics for hedging, depending on both hedging need and relative value. In Section 3.10 we analyse the valuation and payoffs of long variance swaps combined with short put positions, effectively analysing the relative attractiveness of long variance and long put positions. In practice, other types of protection, which include a delta exposure, can be more efficient, particularly using combinations of vanilla options to provide protection within a targeted range – e.g. put-spreads, put-ladders, collars etc.

Forward variance swaps can offer a useful alternative to holding spot variance. Forward variance swaps are not directly exposed to realised volatility so do not suffer negative carry in a benign market environment. They do however suffer slide as the position ages (assuming term structure is positively sloped). The p/l on forward variance is determined by changes in the implied volatility term structure, and to the extent that implied volatility reacts to market downturns, they can provide a hedge against an equity portfolio. Forward variance can also be useful for hedging against volatility spikes. See Section 3.8 for a discussion of spot and forward variance swaps including carry and slide analysis.

Variance swaps can also be useful for hedging out specific volatility exposures. For example, **life assurance companies** offer many products which offer some form of guaranteed benefits (e.g. variable annuities, with-profits funds), returns for which they attempt to generate by holding equities. These companies have essentially sold equity put options to their policy holders and hence have taken on a short volatility exposure. Changes to accounting standards are encouraging mark-to-market valuation of these embedded liabilities, whilst capital adequacy requirements (Solvency II) require portfolios to be reserved against changes in market level, volatility, interest rates etc.

Although assurers may hedge embedded optionality by delta-hedging, this will not hedge mark-to-market volatility risk or shifts in implied volatility. Increasingly, insurers are looking to use variance swaps (either spot or forward) to hedge volatility exposure, increasing the demand for long-dated variance – and potentially steepening the volatility term structure. Note also that long-dated variance swaps offer particular exposure to changes in implied volatility, which is often what insurance companies need to hedge. For any long variance swap position, realised volatility losses will accrue if the market is relatively benign, although longer-dated maturities accrue realised volatility losses at a relatively slow rate. Alternatives such as forward starting variance and conditional variance can also be explored as ways of targeting hedge exposure. See Section 2.10 for a discussion of forward starting variance swaps and *Conditional Variance Swaps, April 2006* for a discussion of conditional variance swaps.

Similarly, **structured products** desks provide investors with a wide variety of equity-based investments which incorporate volatility (and correlation) exposures, ranging from simple capital guaranteed products to more complex multi-asset payoffs. The aggregate risks of their positions needs to be managed, hedged or sold-on into the market, and variance swaps provide a very useful way for desks to manage their volatility exposures, which will itself change as the market moves. Such hedging flows, or anticipation thereof, can have a large effect on the volatility surface, particularly at longer dates.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 3: Uses of Variance Swaps

3.3: Rolling short variance

Historically one of the most successful volatility strategies has been the systematic selling of short-dated index variance. This strategy takes advantage of the volatility risk premium described in Section 2.6. Returns from this strategy are on average positive, but with a large negative tail from volatility shocks. By implementing a long-term rolling strategy we attempt to diversify away these occasional large losses over time.

A diversification argument also argues for using short-dated variance swaps, even though term structures are generally upward sloping, meaning the shorter-term trades earn less carry. From a risk-return perspective, it may be better to regularly take in a relatively small premium (e.g. 2.3 vegas) every month than to take in a larger premium (e.g. 5 vegas) once a year. In practice 1-month maturity variance swaps appear to be a good maturity to target – shorter dated options are available but are liable to incur relatively higher transaction costs. In addition 1-month trades correspond to VIX/VDAX/VSTOXX levels allowing for greater visibility and ease of back-testing.

The back-test presented here comprises the short variance component of JPMorgan's "**Yield Alpha**" product (see "*Yield Alpha Strategy Guide*", JPMorgan). The various forms of the Yield Alpha product aim to take advantage of different types of of carry trade across a spectrum of asset classes, by choosing a basket of strategies with attractive 'carry-to-risk' characteristics. Short variance presents a good example of such a carry trade, by taking advantage of the observed equity volatility risk premium (Section 2.6).

We have back-tested the strategy of selling 1-month S&P variance on the first day of every month since March 1990. As variance swap prices were not available for the whole period, we have estimated the variance swap strike from VIX levels. Recall that the VIX represents the theoretical 1-month variance swap price interpolated from exchange-traded S&P options (see p94). Since the VIX tends to trade somewhat above real variance swap prices (typically 0.5–1 vega) we have estimated variance swap mid levels at 75 bps below the VIX. Then to get a realistic variance swap bid level we have deducted a further 50bps. Finally, to limit the risk of the strategy, we have only sold capped variance swaps (capped at 2.5x the strike) and assumed that these trade 50 bps below uncapped variance swaps. To summarise, we have modelled the selling price for capped variance swaps at 1.75% below prevailing VIX levels, a level we believe is somewhat conservative.

Results from back-testing are very encouraging (Figure 49). Average **returns are 2.3 vegas per month, with 161 of the 196 months yielding positive returns** (Table 10). Nevertheless, the maximum gain (14.8%) is much less than the maximum loss (22.6%) and there were four months with losses greater than 10%, compared with only two months with gains greater than 10% – indicating the negative tails of the return distribution (see e.g. Figure 28). This is also seen when comparing the standard deviation of negative returns (5.2%) with that of the positive returns (only 2.5%). Although positive returns dominate, **losses are relatively infrequent but can be large and unpredictable.**

Figure 49 : Performance of rolling short variance on the S&P

Vega points

Source : JPMorgan

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 3: Uses of Variance Swaps
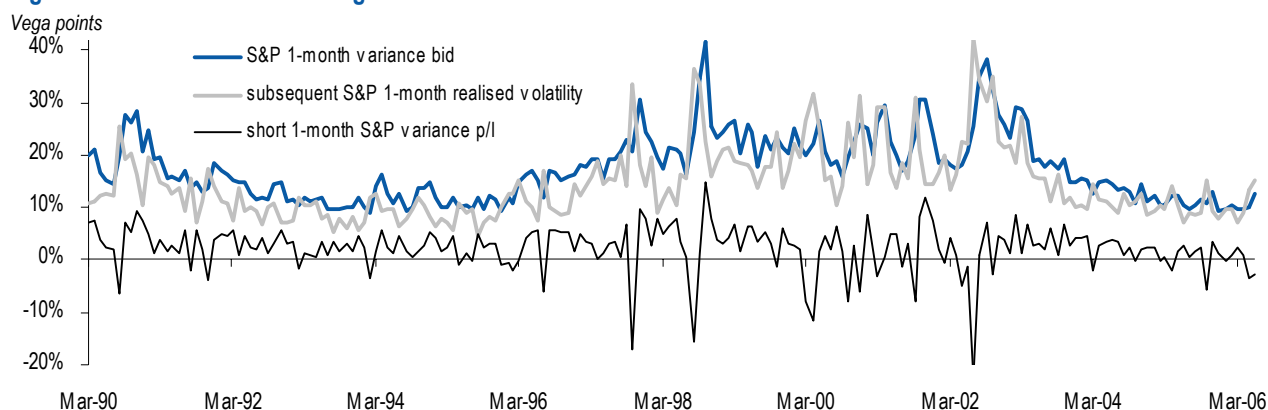
The annualised return is 27.3% (un-compounded), which when compared to the annualised standard deviation of returns of 15.3% gives **an information ratio of 1.8**. Incidentally, over the course of the back-test, the 2.5x cap was never even close to being hit. In fact the greatest level of realised compared to the level of implied bought at was in July 2002 when the realised volatility of 42% was about 1.7x the 25% level of variance sold. Note that if we assume we can actually trade at prevailing VIX levels, the average monthly return rises to 4% (48% annualised) and the information ratio to about 3.1.

Given this pattern of returns we can consider creating an index based on the performance of shorting variance swaps. We assume that a fixed amount of capital, say $100M, can be risked (with the 2.5x cap allowing the maximum losses to be known). Each month we sell a 1-month variance swap with vega notional equal to 0.5% of the available capital (e.g. $500K). This means that as long as the variance swap strike does not exceed 76, the maximum loss cannot be greater than the capital held due to the cap. At the end of each month when the variance swap expires, the capital will increase or decrease depending on the p/l of the short variance swap, and a new short variance swap position for the next month is initiated using 0.5% of the new capital amount.

We consider such an index, starting in 1990 with a value of 100 and each month entering into a short 1-month variance swap with vega notional equal to 0.5% of the current index level (Example 13). By July 2006, 16 years since inception, the index level has reached 741 (Figure 50).

Example 13: Calculation of the 'Variance Alpha' index

Suppose the index is at 100, and the 1-month variance swap strike is 20.

The vega-notional will be 0.5 index points (ip),

Given the 2.5x cap, the maximum loss will occur if realised volatility is 50% or more

In this case the monthly loss will be capped at $0.5 \times (20^2 - 50^2)/(2 \times 20) = 26.25$ip, and if this loss occurred, the subsequent vega-notional will be 0.36875 index points.

In contrast if realised volatility over the next month was only 10%, the p/l would be: $0.5 \times (20^2 - 10^2)/(2 \times 20) = 0.5 \times 7.5 = 3.75$ip, the new index level would be 103.75, and the new vega notional equal to 0.51875 index points.

Figure 50 : Index performance from rolling short variance strategy
Index performance
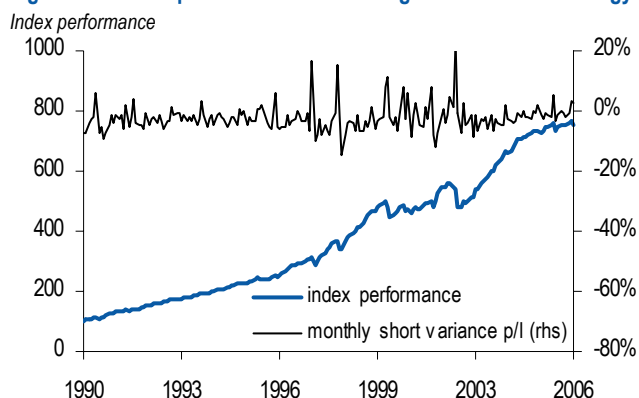
Source : JPMorgan, results from Mar 1990 – Jul 2006

Table 10: Performance statistics for rolling short variance strategy

Average Return (monthly)	2.3%
Average Return (annualised)	27.3%
Annualised Stdev	15.3%
Risk-return	1.78
Index performance (1990 = 100)	741

	All	Positive	Negative
Number of months	195	161	34
Percentage	100%	82.6%	17.4%
Standard deviations	4.4%	2.5%	5.2%
Returns In excess of +/- 10%	6	2	4
Largest return		14.8%	-22.6%

Source : JPMorgan, results from Mar 1990 – Jul 2006

Similar results are obtained from this strategy using the Euro Stoxx 50 or DAX indices, with the VSTOXX and VDAX as the basis for variance swap prices. Using slightly longer maturity variance swaps, e.g. 3-months, works well although the timing of the index roll-dates then becomes more important.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 3: Uses of Variance Swaps

It is an interesting exercise to compare the performance of rolling short volatility strategies using variance swaps with similar strategies using short straddles, both delta-hedged or not (see *European Equity Derivatives Weekly Outlook, 6 November 2006*). Whilst all three strategies perform well, short variance swaps appear to deliver the most consistent alpha across different volatility environments, and most effectively capture the premium of implied to realised volatility. Rolling short 1-month naked straddle strategies perform well, profiting from the somewhat range-bound behaviour of equity indices, and avoiding the negatively convex payoff which can hurt short variance positions in times of high volatility. However, their p/l is extremely path dependent and they can suffer in times of low-volatility trending markets (e.g. 2004-2006). Delta-hedged straddles do a better job than un-hedged straddles of capturing the premium of implied to realised volatility, but also suffer in trending markets, as otherwise positive returns from selling volatility above the level subsequently realised are eroded as the market trends away from the straddle strike, and the gamma exposure of the straddle is reduced.

With returns from short volatility trades somewhat un-correlated with the underlying, these types of strategy work well as overlay strategies aimed at boosting alpha and diversifying returns.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

3.4: Diversification

One of the notable features of the returns of a rolling short volatility index is its similarity to returns of a bond index (Figure 51, Figure 52). In rolling short variance, regular periods of positive p/l resulting from the spread between implied and realised are punctuated by large losses arising from spikes in volatility. In bond indices, the regular p/l resulting from coupon payments and accrued interest is punctuated by capital losses caused by rising yields, or, in the case of credit portfolios, defaults. In the same way as interest is the payment required for providing financial capital, the implied - realised volatility premium can be thought of as the payment required for providing "equity-insurance" capital. **It is this "equity-insurance" capital that can be considered an alternative asset class** that can be invested in and the payment received is captured through a rolling short variance swap strategy.

Figure 51 : Index performance from rolling short variance strategy
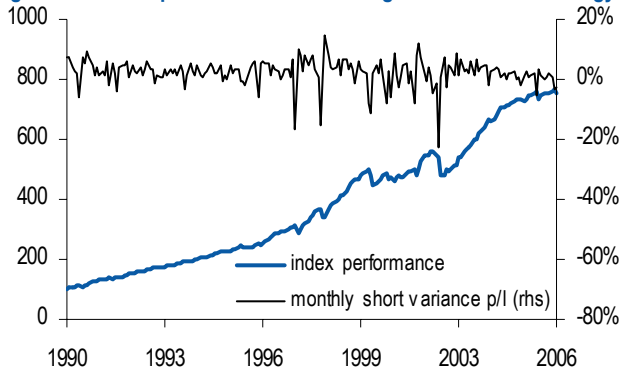
Source : JPMorgan

Figure 52 : Bond index performance – JPMorgan USD traded bonds
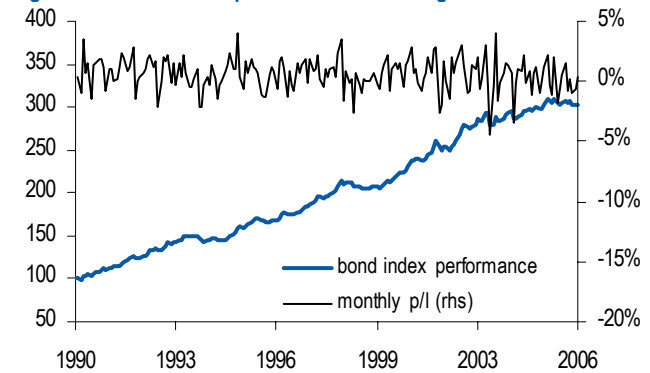
Source : JPMorgan

Hence, volatility (or more correctly the implied – realised spread) can be thought of as an asset class in its own right. Moreover since there is relatively low correlation between the p/l from short variance and the underlying equity market, it can be used within a portfolio to diversify returns (Figure 53, Figure 54).

Figure 53 : Efficient frontier of stock plus bond moves further out by adding short variance…

Portfolio return
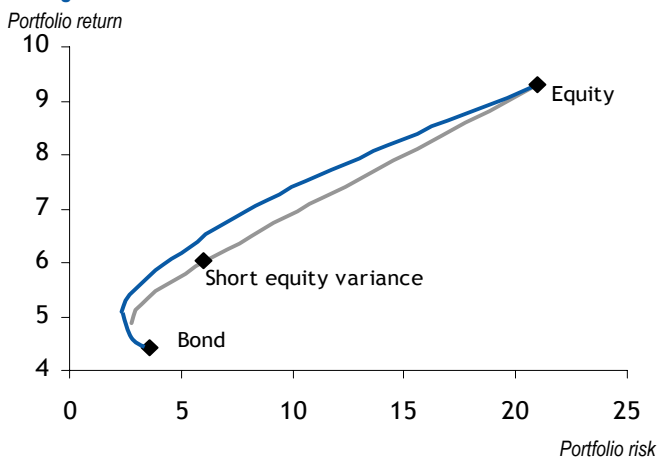
Source : JPMorgan

Figure 54 : … and variance generally replaces bonds in asset allocation along the efficient frontier.

Asset allocation, white dashed line is optimal allocation split for bond + equity only
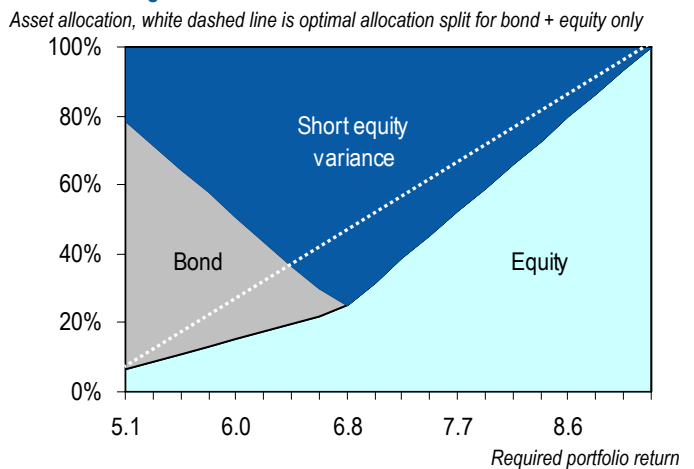
Source : JPMorgan

In Figure 53 we show the risk/return for each individual strategy – equity (long Euro Stoxx), bonds (long Euro Governments), short Euro Stoxx 50 variance – overlaid with an efficient frontier for a two-asset bond/equity mix (grey line) and an efficient frontier for all three asset classes (blue line). Because variance returns are not fully correlated to any combination of bonds and equities, including returns from short variance swaps helps to push out the efficient frontier. Moreover when used in conjunction with bonds and equities, short variance swaps generally replace bonds within an efficiently allocated portfolio (Figure 54), due to somewhat bond-like nature of short variance swap returns.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan ●

Part 3: Uses of Variance Swaps

3.5: Index variance spreads

One use of variance swaps is to trade the spread of volatility between two indices. These type of trades fall into two main categories:

1. **Relative value volatility trades**, where the variance of one index appears to be trading at an unwarranted premium or discount to another.

2. '**Volatility-beta' trades** which are long variance on a higher beta index and short variance on a well-correlated but lower beta index. These trades aim to be net long volatility, particularly as volatility increases, with the short variance leg hedging out some or all of the otherwise negative carry if volatility remains at prevailing levels.

Opportunities of the first type can be evaluated by looking at current spreads and comparing these to historic spreads of realised (and implied) variance, at least in similar volatility environments, to see if an opportunity exists. The volatility pairs regression analysis discussed in the following section can also be used to screen for these kind of opportunities. See also "*Volatility Pairs*, JPMorgan, August 2004". Some index variances have a tendency to appear either structurally rich or cheap as a result of local activity, such as protection buying, or other supply and demand dynamics.

Note though that some indices have local 'regime-changes' as their composition or behaviour of the members changes. For example the Nasdaq was a very high volatility index in 2000-2001 realising well over 50% volatility. However, following the unwind of the dot-com bubble, many companies of the companies in the index have now 'matured' and the Nasdaq behaves much more like other developed market indices – realising around 10-20% volatility in 2005-2006.

Volatility-beta trades

The second type of spread trading involves taking opposing volatility positions in two well-correlated indices, but where one index tends to be more sensitive to increases in volatility than the other. We have seen previously that there exists a premium of implied to realised volatility, which will probably exist for both indices. Thus a spread trade may to some degree cancel this extra cost, whilst maintaining some overall exposure to increased volatility through the higher sensitivity of the buy side index. Such trades, as previously stated, are often referred to as **volatility-beta** trades.

One of the characteristics of a low market volatility regime is the compression of volatility spreads across stocks and indices. This is highlighted by the contraction of cross-sectional deviation between different index realised volatilities whenever overall market volatility becomes more subdued (Figure 55).

Figure 55: Spreads between different indices have compressed as implied volatility has fallen…

6M realised volatility

Source: JPMorgan
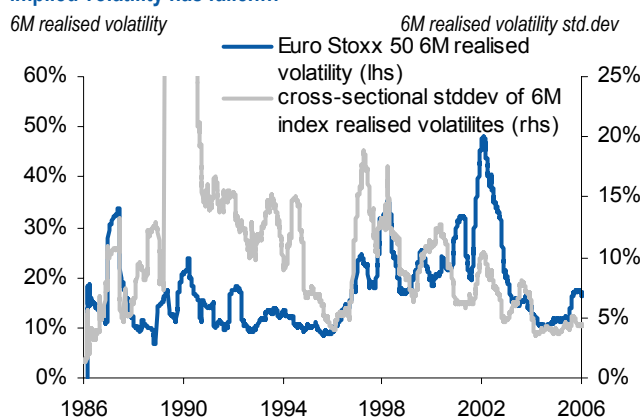
Figure 56: …while there has been a general trend of increasing correlation between indices

Index-Index, 6M Average pairwise correlation

Source: JPMorgan
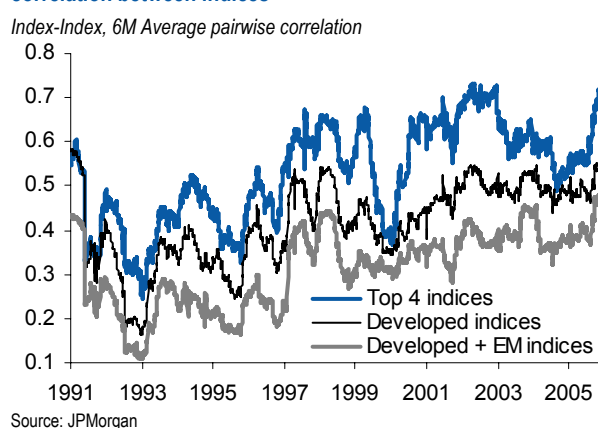
Moreover, in the current period of low volatility (2004-2006), this deviation of index volatility is now close to its all-time lows, though this may also be a consequence of the trend of increasing correlation between indices globally over the last few years (Figure 56).

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 3: Uses of Variance Swaps

The high correlation between indices, coupled with the likelihood that index volatilities will probably diverge in a rising volatility market, gives opportunity to investors to position themselves more efficiently for an increase in volatility using volatility spread trades rather than outright long volatility positions.

As an example of this type of trade, we look at the behaviour of DAX volatility compared to that of Euro Stoxx 50. Given that the Euro Stoxx 50 is a pan-region index with a greater number of members and a more balanced sector composition than the DAX, it should be expected to realise less volatility than the DAX on a purely correlation effect alone. This is indeed the case, with the average spread of realised volatility at around 2% over the last fifteen years (Figure 57).

Another important feature of the DAX / Euro Stoxx volatility spread is that it tends to widen as volatility increases. This is a consequence of the *volatility beta* effect, which sees DAX realised volatility increase (and then decrease) by more than Euro Stoxx realised volatility. We can quantify this effect by looking at the beta of a regression of DAX on Euro Stoxx 6-month realised volatility, which has been about 1.25 over the last three years and 1.05 over the last fifteen years (Figure 58).

Figure 57: DAX tends to realise more volatility than the Euro Stoxx …

6M realised volatility
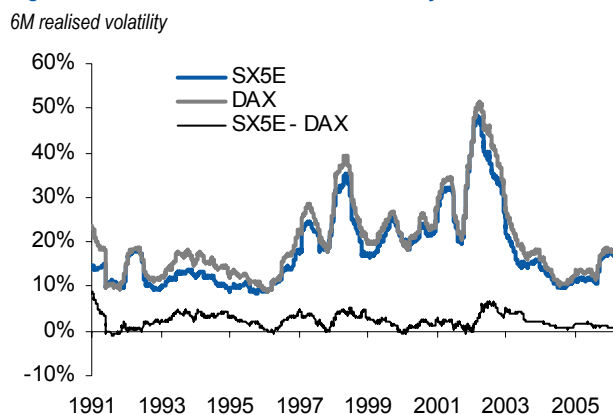
Source: JPMorgan

Figure 58: … and spread tends to expand as volatility increases

6M DAX realised volatility

$$y = 1.0493x + 0.0142$$
$$R^2 = 0.9721$$

6M Euro Stoxx 50 realised volatility
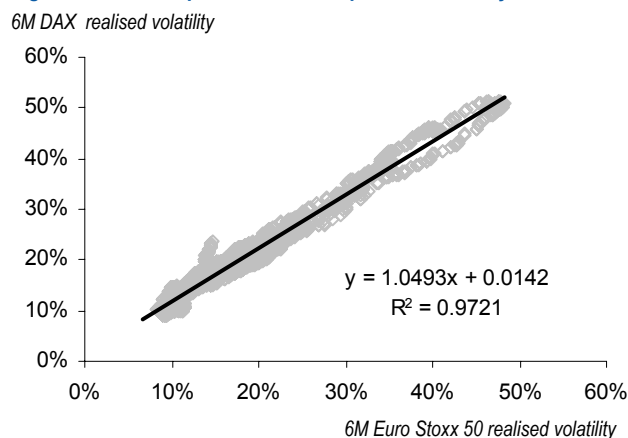
Source: JPMorgan, data since 1991

Given the directionality shown in Figure 58, it makes sense to compare a long variance strategy on the DAX with a spread trade which is long DAX / short Euro Stoxx 50 variance. We can backtest this directly by comparing the prevailing levels of implied variance of the two indices, and the spread between them, with the historical realised volatility over recent years.

Example 14: DAX vs. Euro Stoxx volatility beta trade

On 25 September 2006, 6-month DAX and Euro Stoxx implied variance were trading at 18.6% and 17.2% respectively, leading to a (mid) spread of 1.4% of DAX over Euro Stoxx.

Suppose it had been possible to own DAX variance at the 18.9% level (mid + 30bps) throughout the last 15 years, then the average profit would have been 3.6% (3.6 vegas) with a standard deviation of 13%. However, the majority of profits would have been made in the high volatility period between 1997 and 2003, and outside of this period the average p/l would have been -3.8%, reflecting the usually negative carry of a long volatility position.

Suppose instead we consider the p/l of being long the DAX/Euro Stoxx variance spread. This would have lead to an average p/l of about 0.25%, but with a standard deviation of returns of only 2%.

To compare these two alternatives fairly, we scale the p/l of the spread trade by its standard deviation of returns. In this case, this leads to a p/l of 1.7%, less than the 3.6% of the long DAX variance trade. However, much of this difference is due to the higher assumed bid offer costs for executing the spread trade. Also note that the returns for the spread trade are more normally distributed than those of the long DAX trade (Figure 60).

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Figure 59: Long DAX variance p/l is similar to the spread trade against Euro Stoxx …

6M DAX variance p/l   *6M DAX-Euro Stoxx 50 variance p/l (6.5X)*
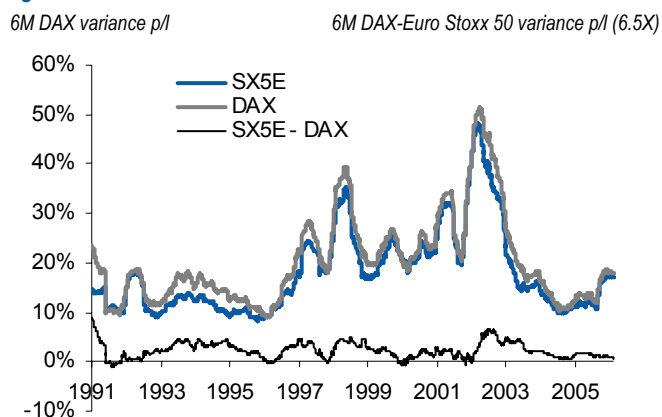
Source: JPMorgan; assuming prevailing variance levels as of 25 Sep 2006

Figure 60: … but the spread trade does have a more normally distributed p/l

Population of p/l scenarios

6M variance p/l
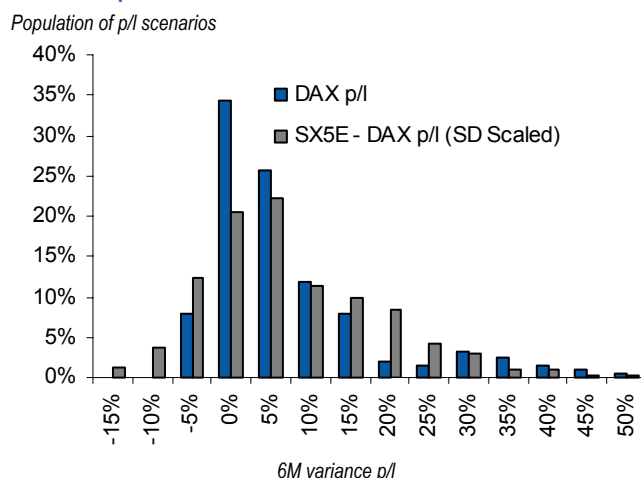
Source: JPMorgan; assuming prevailing variance levels as of 25 Sep 2006

The above example shows that in this case the spread trade had no real benefit compared to just holding long DAX variance. However, in some sense the indices were *too* correlated in this case. Scaling up the spread trade by 6.5 times also necessitated effectively scaling the 60bps transaction costs up to 3.9 vegas in comparison with the long DAX trade (transaction cost 30bps). Without these transaction costs the average long DAX p/l would have been 3.9 vegas and the spread traded p/l, 5.6% for the same level of risk.

In general we can look for these kinds of spread trades across a wide range of indices, to find index pairs where the volatilities are well correlated and which display the volatility beta effect described above. We can asses the attractiveness of these pairs by backtesting the variance spread trade against going long variance on the higher beta index, and scaling the spread p/l to give a similar level of risk.

One point worth making in this analysis of volatility-beta trades is that it does not attempt to directly measure the relative value between index volatilities. What we are attempting to achieve is to get a long volatility exposure, but additionally using a correlated index with lower beta to mitigate the carry of such positions. Hence, indices with a low correlation between their realised volatilities are not considered as a direct volatility-beta trade.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

3.6: Relative value single-stock volatility

Variance swaps can also be used to trade spreads between single-stock volatilities – either as a way of partially hedging the variance of a single-name against wider market/sector moves, or in order to directly trade relative value single-stock volatility. Two possible methods for assessing relative value single stock variances are *temporal regression* and *cross-sectional regression*.

Volatility pairs analysis

In this case, temporal regressions work by looking at the evolution of the ratio (or spread) of the implied variances between two stocks. A *z-score* can be computed to express how far the ratio is from its long term average. Assuming these ratios of implied variance have mean-reverting tendencies, then pairs with z-scores of large magnitude, expressing a large deviation from their long term average, could potentially represent opportunities for volatility pairs trades. This applies particularly to well correlated stocks, typically from the same sector with similar volatility characteristics.

Of course sometimes a divergence in implied variances of two stocks whose variances were previously highly correlated may simply reflect real changes in the underlying market. For example, consider the case of two similar companies with similar volatility characteristics where one company issued an unexpected profit warning raising concerns about its future viability. This could lead to a large change in ratio of implied variances as the variance of one company spikes but the other remains relatively stable. Clearly in this case, the ratio of implied volatilities would not necessarily be expected to mean-revert and selling the ratio in the hope that it would close up would not necessarily be a sensible trading strategy.

To some extent we can screen out these type of scenarios by also considering the z-score arising from the ratios of realised volatilities. If the realised volatility ratio has recently diverged from its long-term average in the same way as the ratio of implied variances, then the change in the implied variance ratio may be justified, and no trading opportunity exists. Thus the trick is to look for pairs with a high z-score in their implied volatility ratio, which is not reflected in a corresponding high z-score for their realised volatility ratio (Figure 61, Figure 62).

Figure 61: Alcatel implied volatility increased in relation to Cap Gem...

1Y Implied volatility   *implied volatility ratio*
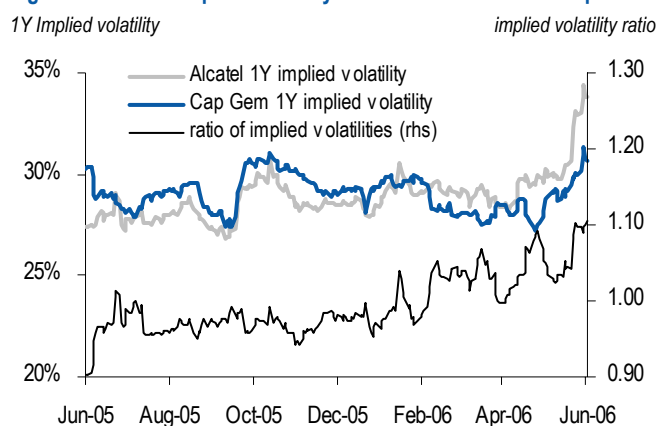
Source: JPMorgan

Figure 62:...but this is not reflected in the ratio of realised volatilities

3M realised volatility   *realised volatility ratio*
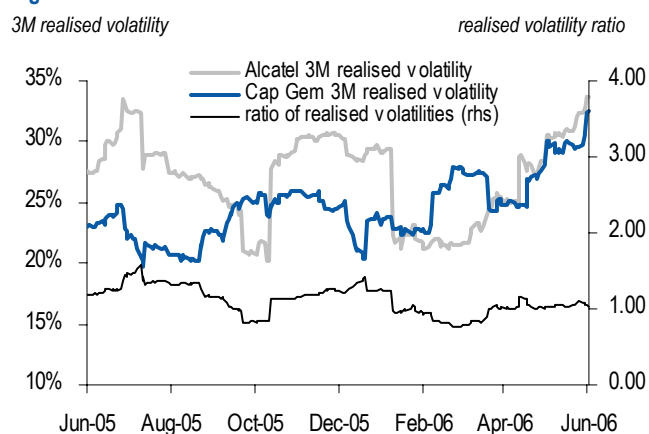
Source: JPMorgan

JPMorgan produces a daily Volatility Pairs Report considering all possible pairs within 11 broad-based sectors and ranks them according to the difference between the implied and realised z-scores. We have back-tested the results of trading volatility on the top 3 ranked pairs in sector, with encouraging results. See *Volatility Pairs, August 2004* and "*Identifying and exploiting dislocations in single-stock volatility*", *June 9 2006* for further details. Note that variance swaps are particularly suited to this type of trade due to the direct nature of the exposure to volatility they offer.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Cross-sectional regression

An alternative means of screening for rich/cheap single-stock variance is to use a **cross-sectional regression methodology**. The idea here is to compare the current implied volatilities (or variances) of a universe of stocks against a number of other factors which may have some bearing on their volatilities – for example their historic volatilities and credit spreads. It is important to note that unlike the temporal regression method explained above, this regression is not trying to model the historical relationship of stock volatilities, but is simply trying to model all current stock implied volatilities as a multi-linear function of a number of various other stock properties. The following variables are used in our regression model:

- Historic volatility (should usually be the main factor determining a stock's implied volatility)
- Stock beta vs. the market (high beta should translate into high volatility)
- Dividend yield (high yielding should perhaps be less volatile)
- 3-month return (poor performers may tend to have higher volatility)
- CDS spread (reflects a company's credit risk and leverage, and should correlate with implied volatility)

Figure 63 and Table 11 show the results of the cross-sectional regression model for 112 major European names as of 9 November 2006. The implied volatilities of Centrica, EADS, BG Group, Scottish Power and Rio Tinto appear rich, all with actual implied volatilities well above those suggested by the regression model. Nestle, UBS, Ciba, Iberdrola and MAN appear below the regression line, having apparently cheap volatilities.

We have backtested this method by trading selected stock variances from a universe of major European names - buying variance on the 10 stocks with the cheapest looking variance according to this metric and selling variance on the 10 richest looking names. Results are good, with around 6 vegas average p/l (pre-bid/offer) since 2003 from trading the spread between the long and short variance baskets. See *European Equity Derivatives Weekly, 21 November 2005* and *"Introducing JPMorgan's Cross-Sectional Volatility Report", November 2006* for further details.

Figure 63: Cross-sectional model: stocks with rich volatility appear above line; those with cheap volatility below it

3m implied volatility

cross-sectional model implied volatility
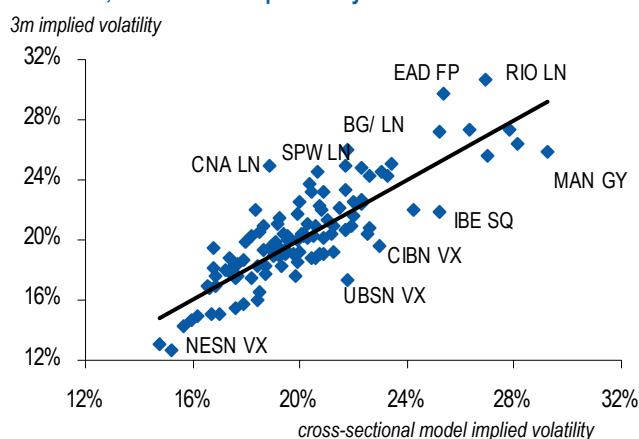
Source: JPMorgan

Table 11: The top 5 richest and cheapest European volatilities according to the cross-sectional model (9 November 2006)

	Bloomberg Ticker	3m realised volatility	3m ATM implied volatility	model 3m implied volatility	z-score
1	CNA LN	16.0%	24.9%	18.9%	2.2
2	EAD FP	30.8%	29.7%	25.4%	1.6
3	BG/ LN	22.8%	25.9%	21.8%	1.6
4	SPW LN	23.5%	24.5%	20.7%	1.4
5	RIO LN	33.4%	30.7%	26.9%	1.4
...					
108	NESN VX	11.2%	12.6%	15.2%	-1.0
109	IBE SQ	33.9%	21.9%	25.2%	-1.2
110	CIBN VX	25.5%	19.6%	23.0%	-1.2
111	MAN GY	34.3%	25.8%	29.3%	-1.3
112	UBSN VX	20.4%	17.3%	21.8%	-1.7

Source: JPMorgan

This method can also be used to select stocks with cheap implied volatility/variance to trade against index variance in a bespoke dispersion trade. See e.g. *European Equity Derivatives Weekly, 12 June 1006.*

On both the temporal and cross-sectional regressions investors should consider whether they wish to be exposed to the relative *realised* or *implied* variance of the stocks. Investors believing that the current spread between the implied volatilities of a pair of stocks represents an unreasonable view of the difference between their future *realised* volatilities, should trade this spread through relatively short dated variance swaps with the intention of holding them to expiry. Investors who believe that the spread between the *implied* volatilities may mean-revert in the short term, may be better trading long-dated variance with the intention of closing out their positions early, or even using forward starting variance.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan 🟦

Part 3: Uses of Variance Swaps

3.7: Variance dispersion and correlation trading

An extension to trading volatility pairs is to **trade the volatility of an index against a the volatility of a basket of stocks**. This naturally leads to the concept of trading correlation. Index variance should always be less than the average variance of the constituents, due to the diversifying effects of the index, the size of this discount effectively being related to the correlation between the constituents of the index. However, back-testing shows that the premium of index implied over realised variance tends to be greater than the corresponding premium for single-stock variance. That is, **selling variance on an index and buying variance on its members has generally been a profitable strategy**. This can be partly explained by demand for protection at the index level, whilst fund managers seek to sell options on single-stocks on which they have a view in order to raise premium.

Realised index volatility can be thought of as arising from two factors: (a) the volatilities of the constituents, together with (b) the correlation between them. Greater correlation means less diversification and hence higher index volatility.

Realised index variance is described by the equation:

$$\sigma_I^2 = \sum_i \omega_i^2 \sigma_i^2 + 2 \sum_{i<j} \omega_i \omega_j \sigma_i \sigma_j \rho_{ij}$$

where ω_i and σ_i are respectively the weight and volatility of the i^{th} stock in the index and ρ_{ij} is the pairwise correlation of the i^{th} and j^{th} stocks.

Since these are realised quantities, they can be observed in the market and the pairwise correlations directly calculated, and averaged to give an overall correlation figure for the index. Another way of expressing this 'average correlation' is to assume that the correlation is the same for all distinct pairs of stocks and then rewrite the above equation to solve for this average correlation.

$$\rho = \frac{\sigma_I^2 - \sum_i w_i^2 \sigma_i^2}{\left(\sum_i w_i \sigma_i\right)^2 - \sum_i w_i^2 \sigma_i^2} \quad \text{or approximately (for a well diversified index): } \rho \approx \frac{\sigma_I^2}{\left(\sum_i w_i \sigma_i\right)^2}$$

The advantage of this approach is that the average *implied correlation,* which we cannot directly measure, can be backed out in exactly the same way from the implied variance of the index and its constituents. As in the case of realised correlation, for a well diversified index, this turns out to be approximately equal to the square of the ratio of index volatility to average single-stock volatility.

This implied correlation can be traded by selling index variance and buying single-stock variance, the resulting position being short correlation. To hedge out the exposure to volatility, at least initially, the index and single-stock legs of the trade must be weighted with more index vega-notional than single-stock vega notional – the exact amounts depending on the level of implied correlation. Such a trade is known as a **variance dispersion trade**.

Like volatility, correlation usually trades at a premium to realised levels (Figure 64), making it another attractive candidate for carry-style trades: rolling short correlation through variance dispersion performs well (Figure 65). Nevertheless, volatility and correlation are themselves correlated, and tend to spike together in a sharp sell-off. Losses from short correlation through variance dispersion can occasionally be very large, especially since the trade becomes short volatility following adverse moves in correlation. Similarly, the p/l from favourable moves in correlation is diminished as the trade becomes long volatility under these conditions. This 'volatility-scaling effect' helps to account for the premium of variance dispersion levels compared to correlation swaps and the relatively modest performance of variance dispersion trades in the low volatility environment since 2003.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Figure 64 : Implied correlation tends to trade at a premium to realised
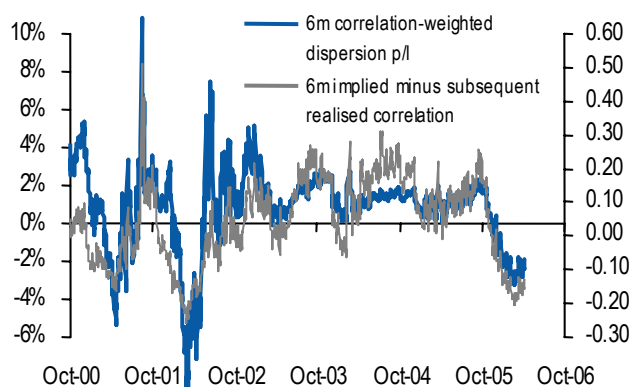
Euro Stoxx 50 6m correlation
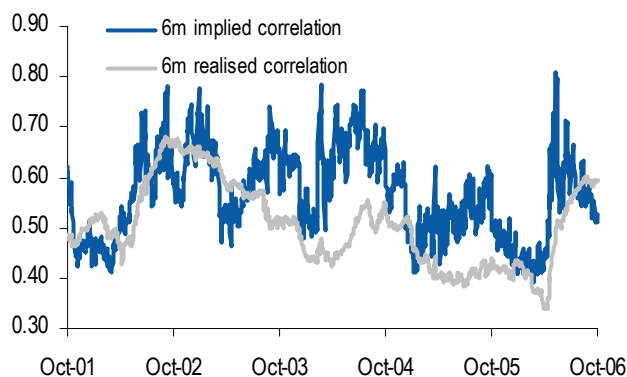
Source : JPMorgan

Figure 65 : rolling p/l of variance dispersion reflects the spread between implied and realised correlation, but is diminished by low volatility

vega payoff   *correlation points*

Source: JPMorgan

Note that variance swaps make possible trading equity correlation in this form. Managing a correlation trade through delta-hedged vanilla options on the index and all its constituents would be onerous, whereas un-hedged vanillas (short index options, long options on the members) would not give a true correlation exposure. However such pseudo-dispersion trades using naked calls or straddles can give interesting and often attractive p/l profiles.

Indeed the demand for variance swaps as a means of trading correlation has been suggested as a reason for the relatively good liquidity of variance swaps on equity underlyings in comparison to other asset classes such as bonds and FX where vanilla options can be even more liquid than for equities.

Equity correlation can also be traded more directly through correlation swaps. For technical reasons, mostly to do with the observed strong positive correlation between volatility and correlation, correlation swaps levels tend to trade somewhat below the levels of correlation associated with variance dispersion trades. This tends to mean that for sellers of correlation, variance dispersion almost always gives more attractive entry levels, at least optically.

See *Correlation Vehicles, May 2005* for further details.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

3.8: Forward variance and volatility spikes

One of the strengths of variance swaps is the ease of pricing and constructing forwards (see p17). This ability to re-express the spot variance term structure into a strip of forward starting variances raises natural analogies with fixed income where the swap curve can be decomposed into strips of forwards or futures.

The first analogy is to **use forward variance to trade the future volatility of an underlying**, in a similar way that investors use Euribor, EuroDollar or Fed Funds interest rate futures to trade the future direction of interest rates. For example, investors expecting a more prolonged hiking cycle from ECB will sell Euribor futures at far out maturities expecting the yield curve to steepen (Figure 66). Likewise investors expecting a pick up in volatility that is greater than that currently priced in, can buy forward starting variance swaps to express their view (Figure 67).

Figure 66 : Forward interest curve compared to history.

3-month Euribor
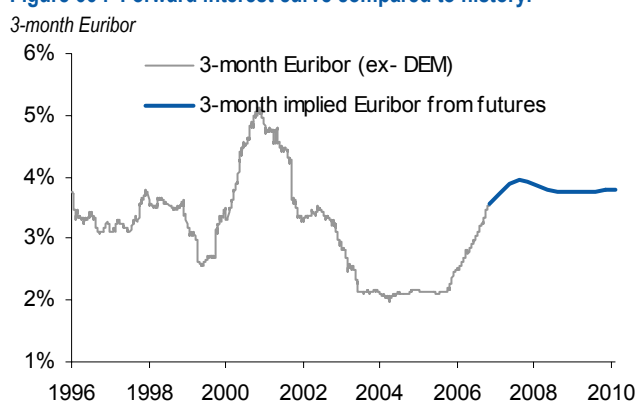
Source : JPMorgan

Figure 67 : Forward variance compared to history

3-month volatility
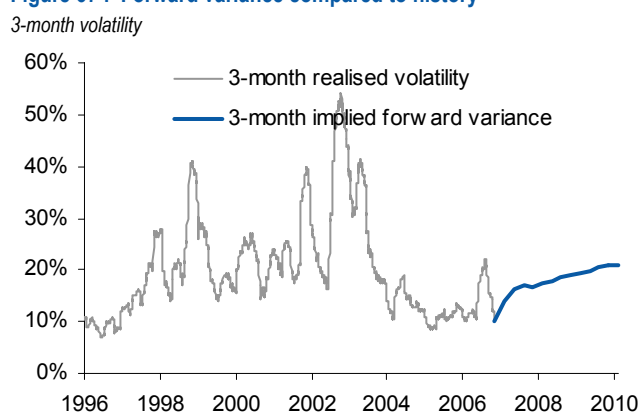
Source : JPMorgan

For a forward volatility position, **p/l before the forward date will be entirely driven by changes in *expectations* of volatility, as captured by movements in the implied volatility term structure**. It is only once the forward date is reached (and the shorter-dated leg of the forward has expired) that there is direct exposure to realised volatility. In the same way, p/l on a Euribor contract is driven by changes in *expectations* of the level of future 3-month rates, and it is only on expiry of the futures contract that p/l is finally settled against prevailing 3-month spot rates.

Forward variance can be an attractive vehicle to be long volatility without accruing realised volatility losses that can occur with a long spot variance position (see volatility risk premium as discussed in Sections 2.6 and 3.3). Given that investors have a choice to trade forward volatility across a different maturities and forward dates, as well as being able to use spot variance swaps, we use a variety of metrics to help us decide where to concentrate exposure. The following discussion concentrates on index variance swaps.

One important point to note is the difference in behaviour between interest rates and realised volatility. Short term interest rates are dictated by central bank action and typically move in a relatively smooth trend, hence forward interest can reasonably accurately capture the future direction of such a trend. In contrast, realised equity volatility can be driven more by short term spikes, with only a fairly slowly underlying "base" level change. This makes using forward variance to trade the future path of volatility more unpredictable, as a correct call on the underlying trend can be ruined through an unfortunate choice of timing. This problem can be mitigated by using longer term variance swaps, which will help to smooth over the effect of any short term spikes.

Below we introduce the fixed income concepts of **slide** and **carry** as applied to volatility. In addition, the performance of a forward volatility position crucially depends on the **dynamics of the volatility term structure**, which we characterise using **principal component analysis** (PCA). Having done this, we will focus on **the ability of long forward variance swaps to hedge against volatility spikes.**

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Volatility carry

Carry in fixed income is taken to be the cost of funding a spot position and reflects the difference between short and long term interest rates. For example, an investor wanting to take a view that long term interest rates will rise will need to either short sell a bond or short a long dated swap. In order to fund this position, they will have to pay the difference between the long and short rate. Since yield curves tend to be upward sloping this is usually a positive spread, so the investor will have to pay the "carry" to be short. Likewise a long investor will receive the carry from holding long maturity positions.

In variance swaps the holder of a long variance position will have to pay the spread of implied over realised every day they hold the swap. Since this spread is typically positive (see Section 2.6), being long variance often means a negative carry position, at least at the index level. Likewise, being short index variance swaps is usually viewed as a positive carry trade. Note that for both fixed income and variance swaps, forward starting instruments have no exposure to carry until the forward starting date is reached.

Volatility slide

Slide in both fixed income and variance swaps occurs because once a trade is implemented the maturity is fixed, so that relative maturity diminishes over time. **Slide** measures the impact of this diminished relative maturity assuming that the shape of the term structure remains unchanged over the period. For an upward sloping variance term structure – typical of low volatility environments – a holder of a long variance swap will lose out through time as the variance swap 'slides' down the term structure curve. Both spot and forward starting variance swaps have exposure to slide.

Principal component analysis

The behaviour of changes in the implied variance curve can be analysed using **principal component analysis** (PCA). PCA decomposes moves in the volatility curve into its major components, or factors. Figure 68 shows the two most important components of weekly changes in the Euro Stoxx volatility curve, using 6 years of data. The first, and most important, factor illustrates how **the dominant curve move is to increase (or decrease) much more at the short-end of the curve**. This illustrates how short dated implied variance tends to move about twice as much as longer dated maturities during a shift up (or down) in the term structure, reflecting the tendency for increases in volatility to flatten the curve and for decreases in volatility to steepen the curve (consistent with the root-time adjustment described above, see p42). **The second component of curve moves is more of a twist component**, with volatility decreasing (or increasing) at the very short-end of the curve and increasing (or decreasing) further out.

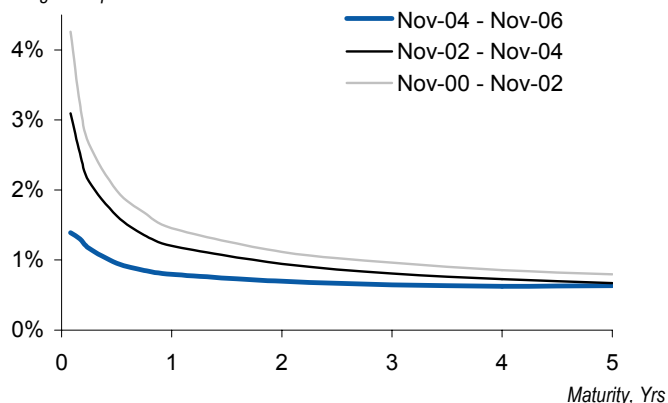**Figure 68 : Principal component decomposition of Euro Stoxx volatility curve since 2000**

Change in implied variance
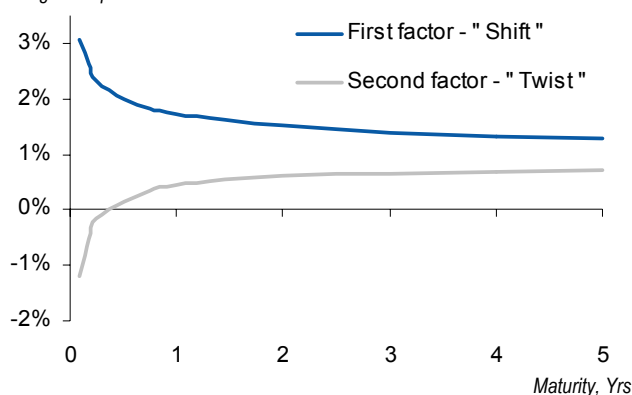
Source : JPMorgan

Figure 69 : Principal component decomposition of Euro Stoxx volatility curve since 2000 … the first component has flattened through time

Change in implied variance

Source : JPMorgan

If we also look at how the first component has changed through time, Figure 69 shows that **primary curve moves have become more parallel over the last 2 years**. This is illustrated by considering the movements in the variance curve following the volatility spikes in April 2005, October 2005 and May/June 2006 (Figure 73). These moves have often been relatively parallel, at least up to the 2-year maturity, consistent with the observed flattening of the first principal component

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006
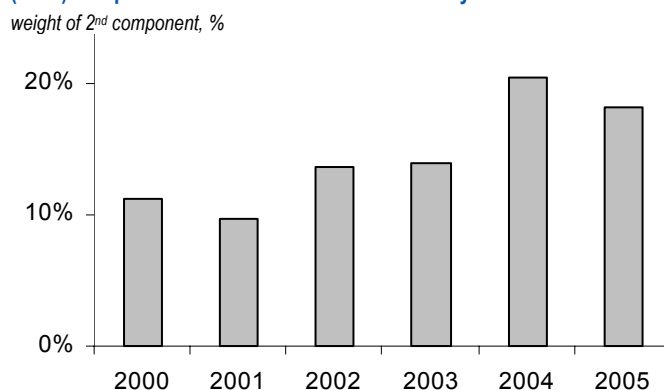
JPMorgan

Part 3: Uses of Variance Swaps

changes. At the same time however, secondary twist-like moves have become relatively more important, as shown in Figure 71. This means that curve twists have become more independent and less well correlated to the direction of volatility.

Figure 70: Movements in variance term structure predicted by the principal components analysis model
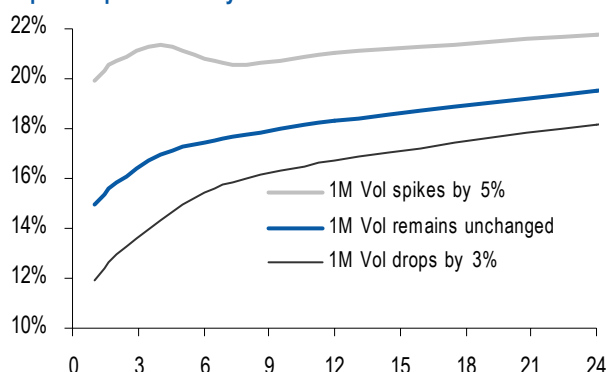
Source : JPMorgan

Figure 71 : Euro Stoxx principal component analysis: weight of second (twist) component has increased over the last 5 years

weight of 2nd component, %

Source : JPMorgan

PCA analysis is useful in calculating expected gains or losses for different parts of the implied variance curve following a general increase (or decrease) in volatility. In conjunction with carry and slide analysis for forward variance they can help identify the best placed part of the curve to be long (or short) volatility.

Volatility Spikes

Below we use carry, slide and PCA analysis to help decide which long variance position (spot or forward, long or short-maturity) is likely perform well as a hedge against a volatility spike. Factors to consider include:

- The likely magnitude of a volatility spike and its effect across the variance term structure;
- The likely timing of a volatility spike;
- The cost of holding the position until the spike occurs;
- The possibility of term structure flattening, which could adversely affect the mark-to-market of forward starting variance.

Forecasting when a spike in volatility might occur is well nigh impossible. However, by decomposing a long history of realised volatility into periodic components using Fourier transforms, **one of higher weighted components of the transform suggests that there is a bias for a spike to occur about every six months**. If the shift up in volatility happens immediately, the strategy that would benefit most is likely to be the spot short-dated variance swap, as the increase in implied variance will likely be greatest at the short end of the curve.

As a final point it is worth considering, if a spike does occur, whether or not to unwind a long volatility position. If an investor holding forward variance into a spike believes that fundamentals do not point to a raised base level of realised volatility they would most likely look to unwind, whereas if they believed that the spike heralded the beginning of a higher volatility regime they may look to profit by continuing to hold their long variance position. Also if a spike has not occurred before the forward start, investors may consider rolling their position to a new forward starting variance position, so as to avoid a realised volatility exposure.

Whilst it is important to consider the likely time of arrival of any spike and the effect such a spike is likely to have on the variance term structure, the behaviour of volatility, both implied and realised, after a spike is also an important consideration. Our analysis of spikes (in realised volatility) on major worldwide indices since the 1950's (See *European Equity Derivatives Weekly, 30 May 2006*), suggests the following results:

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
**Nicolas Granger**
(44-20) 7325-7033

- There is *not* a strong pattern of reduction in volatility in the first month following a spike.

- In the second month after the spike, there is a clear bias for volatility to decrease, on average by about 4.5 vegas.

- For those spikes which subside in this manner (a reduction in volatility within 2 months) the reduction in volatility is almost the same size as the increase seen in its spike. That is, short-lived spikes tend to subside to pre-spike levels within 2 months (Figure 72)

- For those spikes which have not subsided within 2 months, it is likely that there has been a significant further increase in volatility, and although these kind of spikes tend to eventually subside, they take longer and typically subside to a significantly higher base level than the previous spike (on average 5.5 vegas higher after 6 months).

- The short lived spikes tend to outnumber the longer-lived spikes by about 3 to 1.

Figure 72: Volatility tends to reset higher following longer-lasting spikes
1M realised volatility compared to pre-spike level
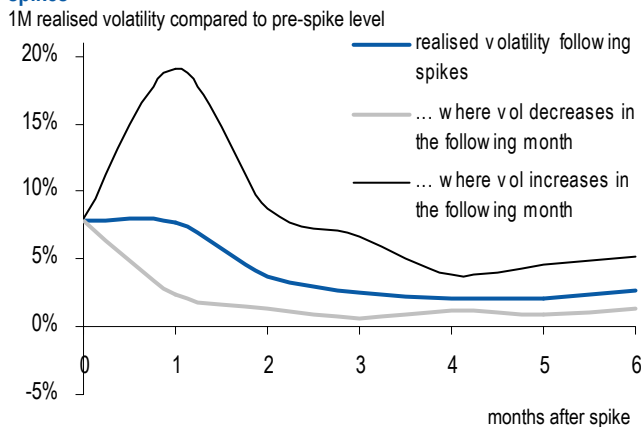
Source: JPMorgan: Data for DJIA since 1900

Figure 73: Changes to the Euro Stoxx 50 variance term structure during the volatility spike in May/June 2006
Implied variance
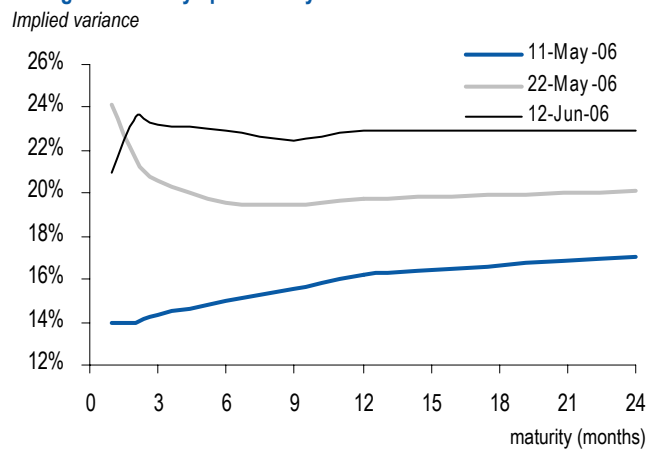
Source: JPMorgan

Given that we can make an assumption on the likely timing of a spike and have a reasonable idea, from PCA, on how it will affect the curve, we can estimate the likely p/l from holding different maturity (spot or forward starting) variance swaps. We can then gauge which variance swap offers the best cost/benefit profile given the prevailing variance term structure.

Figure 74: Variance swap term structure as of 20 October 2006
Implied variance
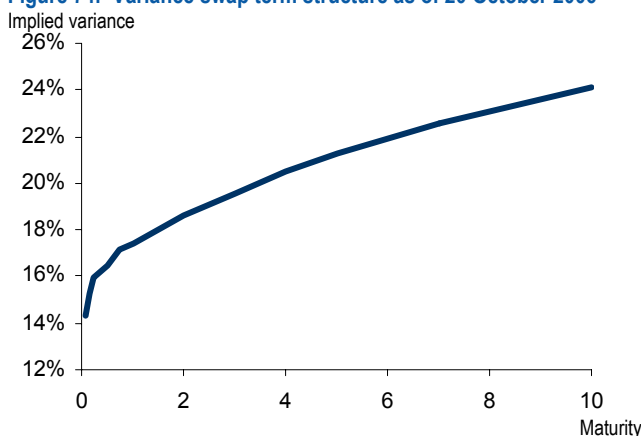
Source: JPMorgan

Figure 75: Expected cost of carry and slide over the next 3-months assuming realised volatility of 12%

Source: JPMorgan

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 3: Uses of Variance Swaps
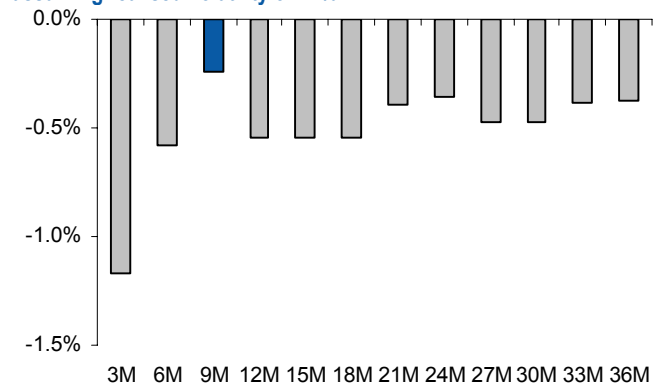
Example 15: positioning for a volatility spike through forward starting variance

We consider the best placed forward-starting 6-month variance to position for a move up in volatility, given the prevailing variance term structure on 20th October 2006 (see Figure 74).

For 6-month variance swaps starting at various forward dates, we calculate the slide expected over the next 3-months, which gives our best estimate of the cost of holding that forward-starting instrument (Figure 75). We find that the 9-month forward starting maturity is the one with the least slide. In fact, it is fairly typical that the 9- or 12- month forward starting maturity has the least slide as it normally represents the place where the term structure begins to flatten out.

If a shift up in volatility happens immediately, the forward starting term that would benefit most is actually the spot 6-month variance swap, as the increase in implied variance will likely be greatest at the short end of the curve. However, if the spike were to happen in three months time, a spot 6-month variance swap would have likely already have accrued a negative p/l from the usual negative spread of realised to implied volatility ("carry") and 3-month variance would then have to spike by about 4 vegas just to break even (consistent with 1-year variance increasing by about 2 vegas, according to our PCA analysis of the volatility term structure). In contrast, the 9- and 12-month forward starting variance works relatively well under several different timings of spikes.

Figure 76 and Figure 77 show the expected p/l of 6-months variance swaps of various forward starts as a function of the magnitude of the move in 1-year variance. Figure 76 shows the projected p/l for a spike after 3-months and Figure 77 for a spike in 6-months. Given any variance move at 1-year maturity, the first principal component tells us how to shift the whole curve and we can calculate p/l of all forward starting swaps. It is worth noting that a 4% move in the 1-year maturity equates to a move of about 7% in the 1-month maturity, which is a fairly typical size for a volatility spike.

The shorter maturity forward starting variance does best if the magnitude of the shift is large and happens within the next few months. If the move up happens later, the low realised volatility before the spike increases the breakeven of this maturity**. The 9-month forward starting maturity does well in most scenarios as the low slide helps mitigate the cost of holding the long volatility position.**

Figure 76: p/l if spike in implied volatility were to happen in 3-months, assuming realised volatility of 12% until then
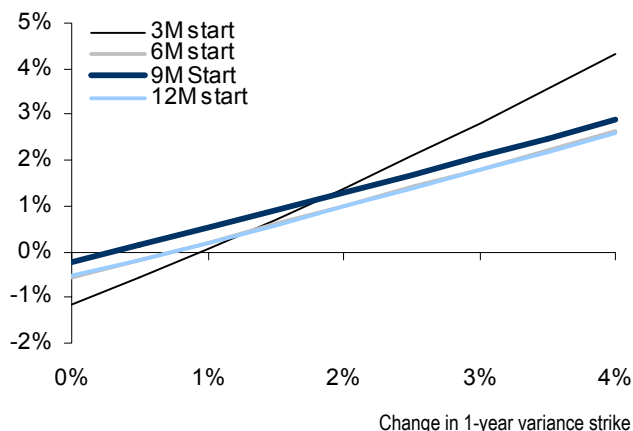
Source: JPMorgan

Figure 77: p/l if spike in implied volatility were to happen in 4-months, assuming realised volatility of 12% until then
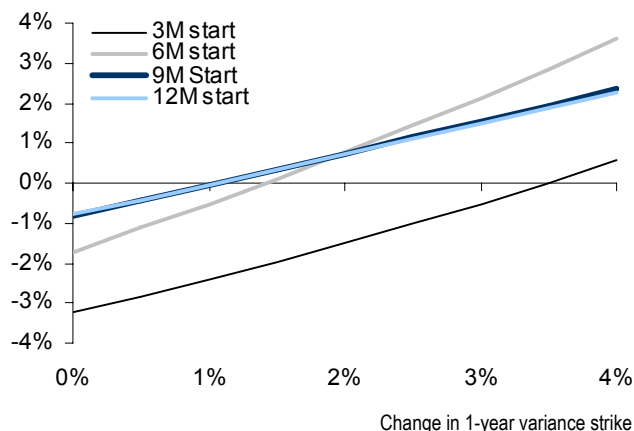
Source: JPMorgan

Thus for investors expecting a rise in volatility even within the next few months, the **forward starting maturities can be a better place to position for this than the spot maturities.** Moreover, the forward starting maturities can provide some cushion against the precise timing of any spike, which by their very nature are difficult to predict. Note though that entering and exiting a forward-starting variance swap can be more costly than running a spot variance swap to maturity, because the forward has to be traded twice and because it has a (slightly) wider bid-offer spreads.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

## Example 16: Performance of forward starting variance swap over volatility spikes

To illustrate the potential advantages of using forward starting variance swaps, we consider two specific examples.

As a way of minimising negative carry, we suggested using forward starting variance swaps for a long volatility view in both January 2006 (*European Equity Derivatives Weekly, 6th January 2006*) and in April 2006 (*European Equity Derivatives Weekly, 18th April 2006*), both of which would have benefited from the May/June volatility spike.

In the later case (April), the spike occurred soon after implementation and hence spot variance performed well; it had not accrued much negative realised carry prior to the spike (Figure 78).

However in the January example, the spot variance performed relatively poorly, since by the time that the volatility spike occurred, previous low realised volatility had increased the required breakeven level, and also the residual exposure to implied volatility had fallen. In contrast the forward starting variance swaps did much better, participating more fully in the spike in implied volatility. Moreover the long maturity 12-month forward did almost as well as the 6-month forward, which had the added advantage of being an almost spot instrument at the time of spike (Figure 79).

Figure 78: The p/l of various long 6-month term variance strategies starting in May
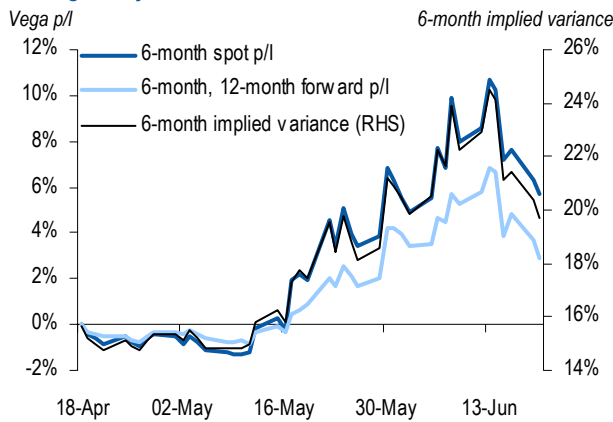
Source: JPMorgan

**Figure 79 The p/l of various long 6-month term variance strategies starting in January**

Source: JPMorgan

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033
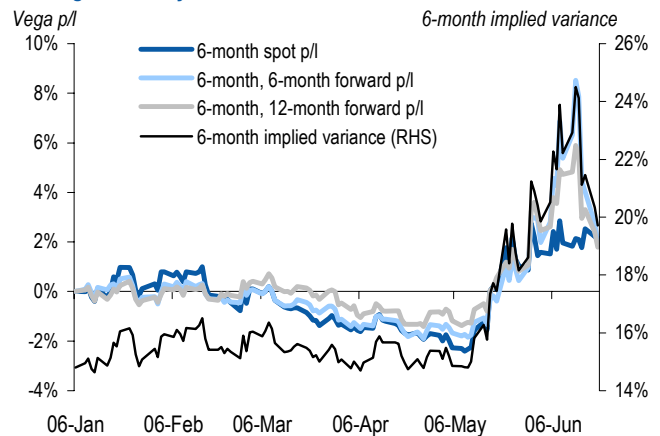
3.9: Trading the variance swap term structure

Whilst we have shown that variance swaps are useful for taking a direct view on the future value of implied and realised variance, they can also be used to trade the shape of the variance term structure curve in analogy with the use of spreads and butterflies in fixed income to trade the shape of the yield curve. One of the main uses of these types of trades is to express a directional view, but at a more advantageous level or with reduced carry. Again this is best described using examples.

The first example attempts to take advantage of a flattening of volatility term structures (Figure 80). In this case, in mid-May 2006, the Euro Stoxx 50 1-year-2-year implied variance spread had come close to its two-year lows (Figure 81), which in turn left 1-year variance 1-year forward close to its lows as seen in early 2005 (Figure 82). This meant that the market was implying a similar outlook for forward realised volatility in May 2006 as it had at the beginning of 2005. This appeared to be somewhat contradictory given the macro economic concerns prevailing at the time.

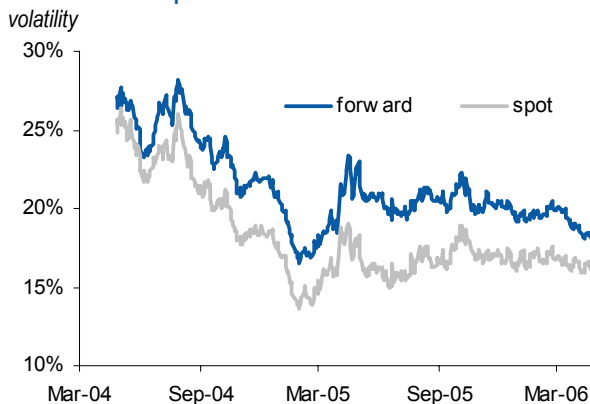**Figure 80 : Euro Stoxx 50 curve had flattened up to 15th May**
Variance swap level
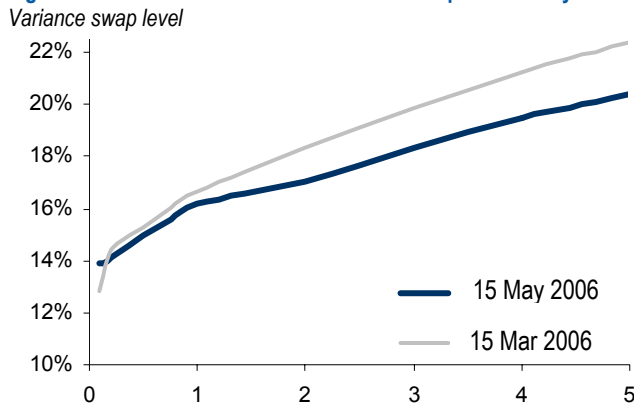
Source : JPMorgan

Figure 81 : ... and particularly in the 1-2 year part of the curve
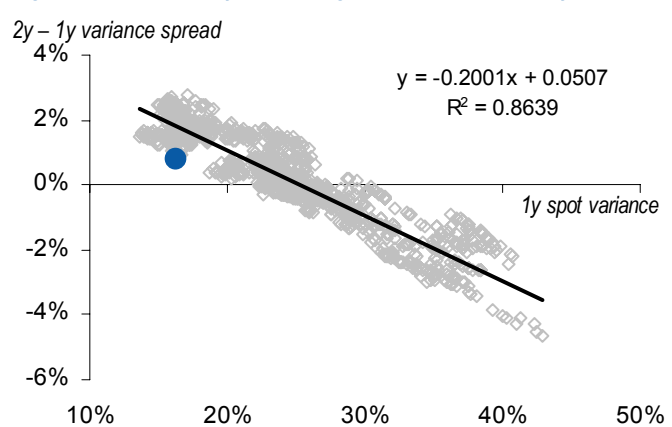*Volatility*   *Volatility*
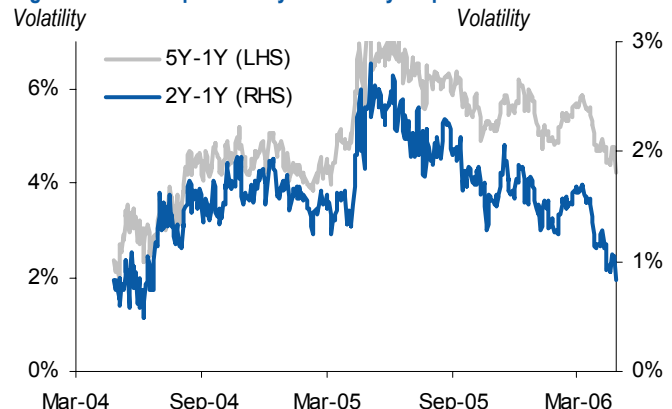
Source : JPMorgan

Figure 82 : Euro Stoxx 50: 1 year 1 year forward variance had come down relative to spot variance
volatility

Source : JPMorgan

Figure 83 : ... especially accounting for the level of volatility
2y – 1y variance spread

$$y = -0.2001x + 0.0507$$
$$R^2 = 0.8639$$

Source : JPMorgan, data since Aug 2000

Investors, believing that these concerns may act to increase volatility, but wishing to avoid a direct long exposure to realised variance, could have used the recent flattening to enter into 1-year, 1-year forward variance swaps. For example, the June-07/June-08 forward variance was trading around 18.6%, compared to spot 1-year variance of around 17%.

Whilst this forward volatility was low by historical standards, it was still well above the recent realised volatility of the Euro Stoxx 50 and would suffer a negative slide of 1.6% over the next year, assuming an unchanged term structure. Hence a long 1-year variance 1-year forward position was only likely to be profitable if implied volatility spikes or realised volatility rises substantially over the next year.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 3: Uses of Variance Swaps

Investors unwilling to take on an outright long volatility exposure, could still take advantage of the relatively low forward volatility strike by recognising that the spread between 1-year and 2-year variance was low compared to the level of 1-year variance (Figure 83). By combining a variance calendar (long 2-year, short 1-year) with some additional long 1-year variance (weighted as per the regression beta in Figure 83), investors could be long the curve whilst minimising outright implied volatility exposure.

For example, the beta of -0.2 in Figure 83 suggested a trade ratio of long 1 unit of 2-year variance and short 0.8 units of the 1-year variance. This would be equivalent to trading the spread between forward and spot variance, hedged for volatility directionality. Note that this position would have positive carry providing the market realises less than the 1-year strike (around 16.4%). Also note the trade ratio will need adjustment through time as the position ages.

The second example aims to take advantage of a Euro Stoxx 50 term structure that appears to be overly convex, (late October 2006) due to the steepness of the volatility term structure between the 1- and 9- month maturity variance swaps. This steepness can be observed by comparing the Euro Stoxx 50 term structure with that of the FTSE on the same date (Figure 84), where a bulge is clearly visible in the Euro Stoxx 50 term structure that is not seen in the equivalent area of the FTSE.

Whilst this could have reflected anticipation on the part of investors of an increase in volatility over the following months, this slope appeared steep compared to the overall level of implied volatility (Figure 85). As the level of implied volatility falls, the spread between the near term forwards and those further out usually widens. However as the May-2006 volatility spike receded this spread of forwards stayed narrow, reflecting an expectation on the part of investors for a near-term move up in implied volatility.

Figure 84: Euro Stoxx term structure bulged compared to FTSE …
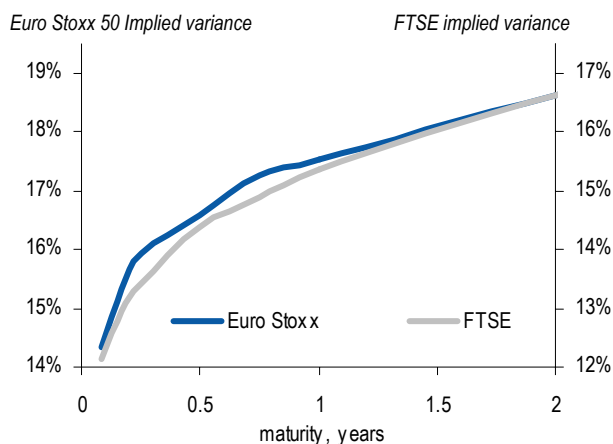
Source: JPMorgan

Figure 85: … meaning that near maturity Euro Stoxx forwards look high compared to further out maturity forwards.
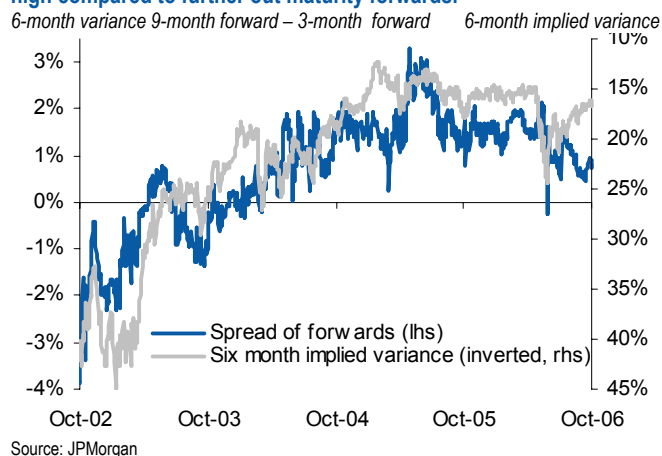
Source: JPMorgan

Nonetheless, by assuming that the relationship between this spread and spot implied variance would hold, it can be seen that a spread which is long 6-month variance 9 months forward and short 6-month variance 3 months forward, is equivalent (at least according to a regression) to selling a spot 6-month variance but at higher than prevailing levels.

Another way to think about a spread of forward variances is as a butterfly of spot variance swaps. In this case the long 1 unit of variance notional of 6-month variance 9 months forward is constructed from long 2½ units of 15-month spot variance and short 1½ units of 9-month variance. Similarly the short 1 unit of variance notional of 6-month variance 3 months forward is constructed from short 1½ units of 9-month spot variance and long ½ units of 3-month variance (Section 1.7). In total this equates to a **variance butterfly trade** of long 2½ units of 15-month spot variance, short 3 units of 9-month spot variance and long ½ a unit of the 3-month spot variance. This example shows how butterflies can sometimes be used to trade the underlying asset (i.e. variance) at more attractive levels than spot, an often used technique in fixed income.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 3: Uses of Variance Swaps

3.10: Skew and convexity trades

Since variance swaps are priced off the entire skew curve, they take on a significant exposure to skew and convexity. This is particularly apparent when calculating the mark-to-market: if skew and/or convexity instantaneously increase and nothing else changes a long variance swap will profit. Variance swaps are long skew and convexity, at least if overall levels of volatility (ATM volatility) do not decrease. Note that in theory exposure to skew is neutral, but typical shapes of skew surfaces mean that increasing the skew increases the convexity, and really it is this which affects the variance swap price, see p82 for further details.

Straddles vs. variance swaps

One possibility for owning convexity is to trade variance swaps against delta-hedged straddles. Whilst variance swaps are in effect long volatility *and* convexity, the exposure to volatility (both implied and realised) can be hedged out by trading a short straddle against a long variance swap position, leaving a residual long convexity exposure. However, this hedge is only effective providing that the underlying remains close to the straddle strike, since the straddle's exposure to volatility, both in the form of gamma and vega, will be reduced as the underlying trends away from the strike. In order to maintain as pure as possible an exposure to convexity, the straddle will need to be re-struck as the underlying moves.

Variance swaps plus naked short puts

One way of trading realised skew is to combine a put with a variance swap. For example consider an investor wanting to express a long volatility position via variance swaps. The p/l at maturity is simply given by:

$$p/l = N_{Vega} \times \frac{(\sigma^2 - K^2)}{2K}$$

However, we know from looking at past history that volatility tends to be directional, in particular increasing – often substantially – in a sell off, and decreasing or maintaining its current level as the underlying rallies. This helps to explain the commonly observed shape of the index skew curve, with relatively steep put skews representing the expectation of increased volatility on the downside, and flatter call skews signalling the likelihood of more benign volatility on the upside. To a first approximation this type of skew curve can look somewhat like the payoff from a long put, suggesting that it should be possible to **partially offset the observed negative correlation of volatility and market direction by trading short naked puts against long variance swap positions.**

Consider an investor who wishes to be long volatility, believing that realised volatility will be above that encapsulated by the skew. To this end, he would be willing to be long volatility at a higher than prevailing variance strike if the market sells off, in return for being long volatility at a lower strike if the market rallies. This exposure can be approximated by combining a long variance swap with a short put.

Suppose we sell a put of strike S_P for absolute premium P for N_{Put} notional of underlying. The combined p/l of the short put and long variance swap at time T is:

$$N_{Vega} \frac{(\sigma^2 - K^2)}{2K} + \frac{N_{Put}}{S_0}(P - MAX(S_P - S_T, 0)) = N_{Vega} \frac{(\sigma^2 - \widetilde{K}(S_T)^2)}{2K}$$

where the adjusted strike is now dependent on the underlying's price at maturity S_T and is written as:

$$\widetilde{K}(S_T) = \sqrt{K^2 - \frac{2KN_{Put}}{N_{Vega}S_0}(P - MAX(S_P - S_T, 0))}$$

We can set the notional size of the put such that on the upside our breakeven is equal to the lowest variance realised if the market has rallied.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Example 17: Long variance swap plus short put

A Euro Stoxx 50 6-month variance swap is trading at 16, the current market level is 4,077 and the 4,050 (99%) put is 175. Since 1996, the lowest 6-month realised volatility in a rally has been about 9%.

For €100K of vega exposure we risk losing €100K x ($9^2 - 16^2$)/ (2 x 16) = –€547K. Thus to offset this loss we need to raise the same amount through put premium. We do so by selling €547K x 4077 / 175 = €12.741M notional of the 4050 puts.

If the market were to fall, the losses on the put would act to increase the breakeven volatility for the long variance swap. For example, if the market falls to 3,600 (-11.7%) the equivalent breakeven realised would be 23, calculated from the equation above:

Breakeven = sqrt{ 16^2 – 2 x 16 x [547 – 12,741 x (4,050 – 3,600) / 4,077] / 100 }
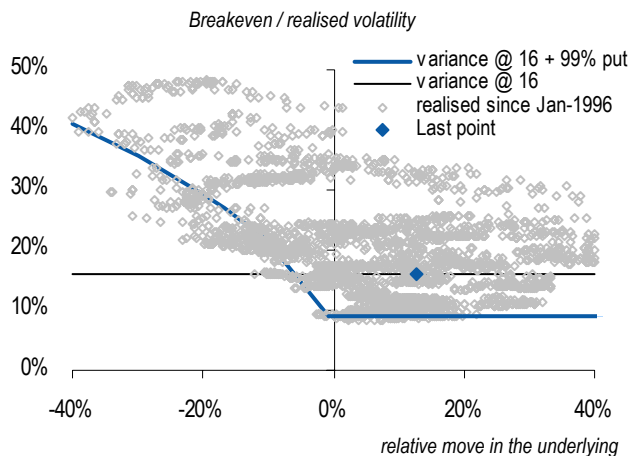
Conversely, if the market rallies, the premium retained from the put would act to lower the variance swap breakeven

For example, if the market rallies to 4400 (+7.9%) the breakeven volatility would be 9:

Breakeven = sqrt(16^2 – 2 x 16 x 547/100), as per construction

Table 12 shows breakeven realised volatility levels for different market levels, for two different put overlays. The first overlay targets an upside breakeven volatility of 9%, as per the calculations above. The second involves selling fewer puts, which lowers the downside volatility breakeven at the expense of an increased upside breakeven – in this case 12%.

Figure 86 : Long variance and short put can better reflect the tendency of volatility to rise is a sell-off.

Breakeven / realised volatility

— variance @ 16 + 99% put
— variance @ 16
◇ realised since Jan-1996
◆ Last point

relative move in the underlying

Source: JPMorgan

Table 12: Volatility breakevens for long variance, short put

| Market Level	Variance	+ 4,050 Put	+ 4,050 Put
Target upside breakeven		**9.0%**	**12.0%**
3,000	16%	33.6%	28.6%
3,200	16%	30.5%	26.2%
3,400	16%	27.0%	23.7%
3,600	**16%**	**23.0%**	**20.8%**
3,800	16%	18.2%	17.4%
4,000	16%	11.4%	13.3%
4,200	16%	9%	12%
4,400	16%	9%	12%
4,600	16%	9%	12%
4,800	16%	9%	12%
5,000	16%	9%	12%

Source: JPMorgan

This concept of shorting puts against a long variance swap position, effectively neutralises some of the long downside puts which are embedded within a variance swap (see Part 4). Taken to its extreme, combinations of short options can be used to remove exposure to realised volatility completely if the underlying moves above or below a certain level. This observation naturally leads to the concept of third generation volatility products such a conditional variance swaps (see *"Conditional Variance Swaps", April 2006*). Since conditional variance swaps are effectively priced from only part of the volatility skew, they can be used to gain more direct skew exposures. For example the popular 'long up-95% conditional variance, short 105% down-conditional variance' trade locks in a premium dependent on the level of skew, providing the underlying remains within 5% either side of current spot.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

## 3.11: Cross asset class trades: trading equity volatility against credit

Investors may wish to consider trading variance swaps against credit default swaps (CDS) to exploit any discrepancies between the equity and credit markets. Both variance swaps and CDS give some measure of the risk associated with a particular company, and should be to some extent correlated. This can be seen when the relationship is aggregated to the index level. Here credit and CDS indices such as Maggie and iTraxx show a strong relationship to Euro Stoxx 50 implied variance in Europe (Figure 87) and in US the CDX HG and HY indices is well correlated to S&P variance (Figure 88).

Figure 87: Implied variance is related to credit at the index level in Europe…

Maggie industrials spread   *Euro Stoxx 3m implied variance*
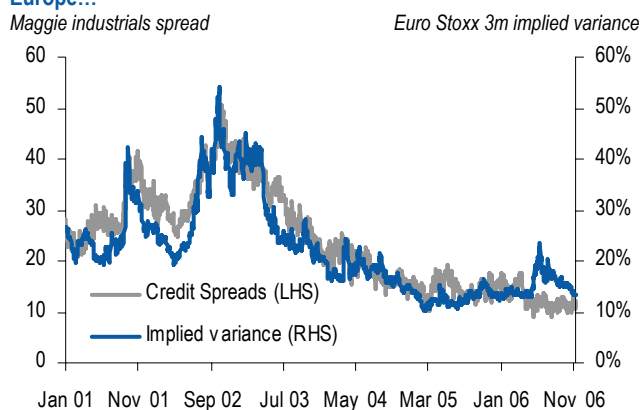
Source: JPMorgan

Figure 88: … and the same is true in the US.

CDX 5Y HG spread   *SPX 3m implied variance*
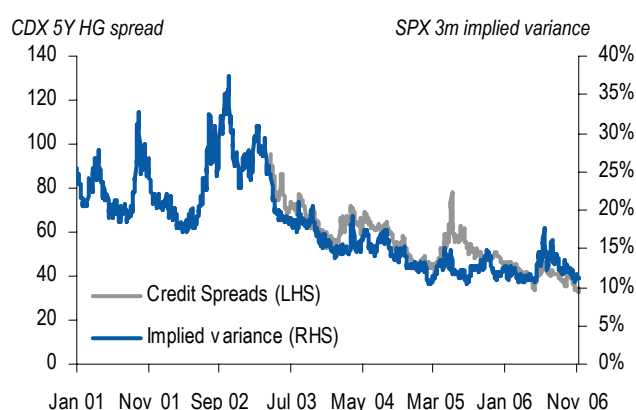
Source: JPMorgan

While credit and equity have a very strong relationship over the long run, there have been several instances where short term opportunities have arisen in one of the asset classes, and where the other asset could have been used as a hedge against a wider macro economic change in regime.

In May 2005, the credit markets sold-off following the downgrade of GM bonds to non-investment grade and correlation worries in structured credit. This credit sell-off was particularly noticeable in the US, although was not reflected in a commensurate increase in S&P equity implied volatility (Figure 88). Credit investors could have sold CDS protection, taking advantage of the dislocation in prices. However, supposing that credit markets were actually heralding some macro change in the economic environment, then it would be likely that equity volatility would in that case also rise. Buying equity volatility against a long credit position could have provided a hedge to possible macro deterioration.

In May 2006, almost the opposite occurred, with equity volatility spiking, mainly in Europe, whilst credit markets remained largely unmoved (Figure 87). Thus a short equity volatility/long CDS position could have taken advantage of the price dislocation, whilst providing a certain degree of protection for a more macro change in regime.

Somewhat counter-intuitively (5Y) CDS levels are more correlated with short-dated variance than with longer-dated variance, likely due to benchmark effects, with the most liquid variance swaps trading at shorter maturities, and longer-dated variance more susceptible to structured product flows.

See *Synthetic Short CDS with Positive Carry, October 2006* for more details on debt/equity trades.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

## Single stock variance swaps and debt / equity

Whilst indices can provide interesting short term opportunities to trade credit against equity variance, they do suffer from basis risk arising from different index members. This can be particularly important in credit indices which may be driven by idiosyncratic default events by single companies. Probably the majority of debt/equity trading is focussed on single name capital structure analysis.

This analysis normally takes the form of attempting to model the relationship of a company's share price to the price of its credit and equity options, for example using reduced form or capital structure models. A typical trade to take advantage of debt versus equity mispricing is a long/short combination of equity puts traded against CDS. One of the main deficiencies of this approach, particularly for reduced-form models based on regressions of share price versus CDS is that a company's share price is non-stationary, compared to CDS which is more mean reverting.

For example, consider the behaviour of Fiat CDS as a function of its share price (Figure 89). In the pre-2004 period, the CDS tended to widen when the share price fell through €15. However more recently it has taken a much lower share price (below €10) to increase CDS. This is because leverage reduction has reduced the price at which the credit market assumes that the company will become distressed. A simple regression model would assume that CDS behaved like a €15 put option and a combined trade using this strike put against CDS would not correctly model the capital structure.

In contrast, because implied variance factors in some of this reduced risk, the relation between Fiat implied variance and credit is more stationary (Figure 90). Hence, variance swaps can be useful instruments as vehicles for expressing a view on the capital structure of a company. In practice, some combination of variance and the underlying equity may be used.

Figure 89: Fiat CDS has changed relationship with stock price …

5Y CDS spread

Stock Price
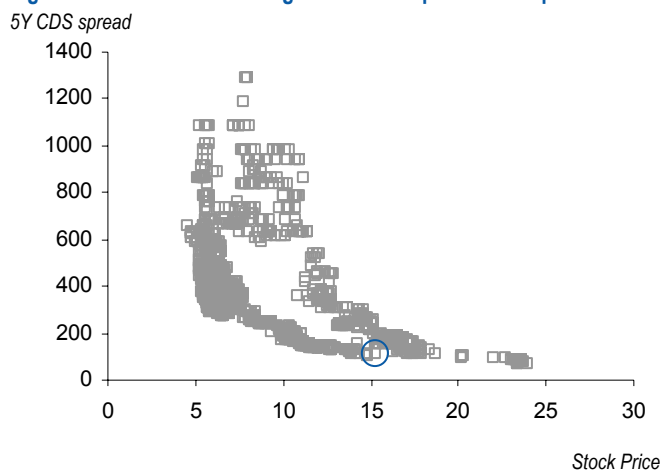
Source: JPMorgan

Figure 90: … whereas Fiat CDS is more linearly related to variance

5Y CDS spread

$$y = 3567.9x - 770.06$$
$$R^2 = 0.652$$

1Y implied variance
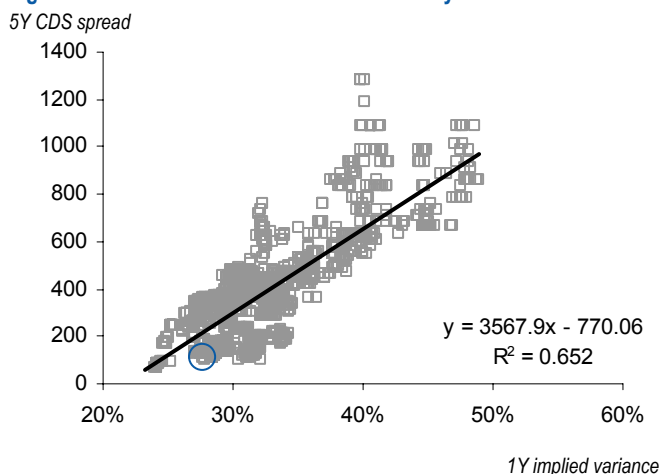
Source: JPMorgan

Although debt/equity trading is probably the major form of cross-asset class trading, the advent of commodity variance swaps introduces the possibility of further opportunities in cross-asset arena. For example one could envisage trading mining company variance against a basket of metals variance or oil company variance against crude oil variance.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

JPMorgan

4. Replication and Hedging

- 4.1 Delta-hedging and dollar-gamma
- 4.2 Theta – the cost of gamma
- 4.3 Options path-dependency
- 4.4 From options to variance swaps
- 4.5 Variance swap replication in one page
- 4.6 Sensitivity to skew and convexity
- 4.7 Variance swap Greeks
- 4.8 Setting up the replicating portfolio
- 4.9 Replicating and hedging in practice
- 4.10 Effects of variance swap hedging
- 4.11 Why not volatility swaps?

Part4: Replication and Hedging

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Part 4: Replication and Hedging

Overview

In Part 4 we consider the more technical issues surrounding variance swaps and particularly how they are priced and replicated. From the perspective of market-makers or others attempting to hedge variance swaps in the options market, the replication argument is crucial, showing how (delta-hedging) a portfolio of vanilla options can be used to reliably capture realised volatility.

From the point of view of an investor trading a variance swap the hedging mechanism may seem unimportant. The investor simply buys or sells variance by entering into the contract and receives or pays out an amount at expiry depending only on the difference between the agreed strike and the level of variance realised. Even if the investor chooses to close out the variance swap before expiry, the mark-to-market is easily calculated (based on realised variance to date and prevailing variance swap prices in the market) without reference to any other market variables.

Nevertheless, a knowledge of the mechanics of variance swap replication is instructive, helping to clarify how realised volatility can be accurately captured, and where the prices of these instruments come from. In terms of pricing, variance swaps can be observed to trade at levels of ATM volatility "plus a bit". The "bit" depends on the steepness of the volatility skew (and convexity) and the time to maturity, with steeper skews (higher convexity) and longer maturities both leading to higher variance strikes. **The replication argument shows unambiguously how the volatilities across the skew contribute to the final variance swap price.**

Perhaps the most important point to come out of the replication argument (for a non-replicating investor) is that **if a market-maker hedges a variance swap, this can, in theory, act against the investor's interests:** investor sells variance, market-maker sells options to hedge their long variance position and the delta-hedging of these short options *on the close* can act to increase the close-close volatility of the underlying. This is in contrast to an investor who sells options to a market-maker who then delta hedges: the market-maker is long options, and to replicate they will buy low / sell high at any time during the trading day and hence act to reduce volatility. In this case the market-maker will act in favour of the investor's interests. See detailed discussion in Section 4.10.

We begin the section with a review of delta-hedged options, and particularly how p/l is generated by 'trading the gamma'. This exposure to gamma is paid for by the time-decay as measured by the theta. We define **dollar gamma**, which allows delta-hedging p/l to be calculated directly from the relative moves in the underlying. The main drawback of delta-hedged options as vehicles for trading volatility is that the dollar-gamma, which characterises the exposure to volatility, changes over time and with the level of the underlying. This means that the exposure to volatility is **path dependent** and unless the volatility is constant *and* the delta-hedging instantaneous, then it is possible that an option could be purchased and delta-hedged at a volatility below that subsequently realised, but the p/l from this delta-hedging program ends up negative.

To remove this path-dependency it is natural to look for **a contract whose dollar-gamma is constant, both with the passage of time and with changes in the underlying.** We demonstrate that such a contract can be created with a static portfolio of options across the continuum of strikes (weighted as 1/strike-squared). This provides an exposure to realised volatility (squared) independent of the path of the underlying. A variance swap consists of this portfolio, bundled up with the necessary delta hedging.

In fact, the variance swap replication argument is valid under much wider conditions than the relatively restrictive Black-Scholes assumptions, in particular it assumes very little about the underlying volatility process. After giving an informal demonstration of the constant dollar-gamma of the replicating portfolio, we show that we can use this portfolio to capture realised variance in a much more general setting. The variance swap replication argument can also be extended to apply to the so-called 3rd generation instruments such as gamma-swaps and conditional variance swaps, helping to understanding how these products can be hedged and to explain their relative attractiveness from a market-maker's perspective. In turn, ease of replication promotes liquidity and the attractiveness of such products to investors.

We go on to give a concrete example of how to construct a replicating portfolio and discuss various practical issues associated with the hedging and replication of variance swaps.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

4.1: Delta hedging and dollar gamma

In this section we outline how vanilla options can be used to trade volatility. Options are exposed to a wide range of factors (performance of underlying, time-to-maturity, volatility, rates, dividends etc). The first-order exposure to moves in the underlying can be hedged out by the familiar delta-hedging process described below. This leaves the exposure to volatility, paid for in time-decay as the most important sensitivity. However, as we shall see, this is not a pure volatility exposure, but is **path dependent**, varying over time and as the underlying moves towards and away from the option strike.

Suppose we hold a (call) option. Then to make the position instantaneously delta-neutral (insensitive to the direction of the underlying) we can sell an amount of the underlying equivalent to the delta of the option. Delta measures the sensitivity of the option to moves in the underlying. For an ATM option delta is about 0.5. For other options the delta is more dependent on the level of volatility, with deltas of in-the-money calls tending towards 1 as the option becomes more in the money; and doing so more quickly at lower levels of volatility. Similarly OTM call deltas tend to zero, more quickly for low volatility (Figure 91). Note by convention deltas are often multiplied by 100, e.g. a delta of 0.30 is quoted as 30.

By frequently re-adjusting this delta-hedge (in theory continuously, in practice usually daily) the sensitivity to direction of the underlying can be hedged out over the lifetime of the option. Note that this is a dynamic hedge which must continuously be monitored and re-balanced. P/l will come from the accumulated action of continuously re-balancing the holding of the underlying, in order to keep the entire portfolio delta-neutral over time. This is referred to as the gamma p/l, and is paid for in the option premium which is marked to market as lost theta.

Figure 91 : Deltas of 1-month maturity ATM call options
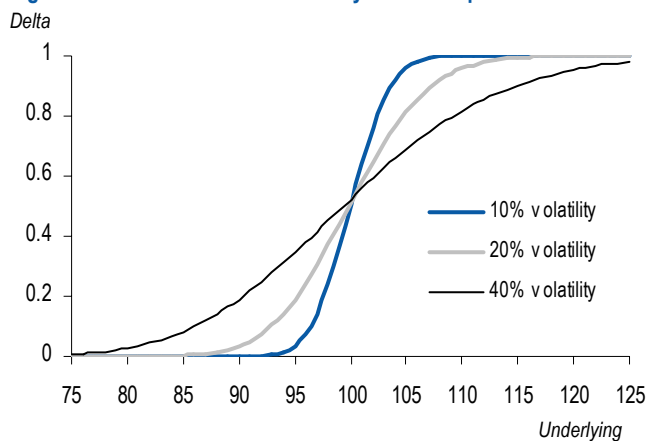
Delta

Source : JPMorgan

Figure 92 : The convexity of the option

p/l
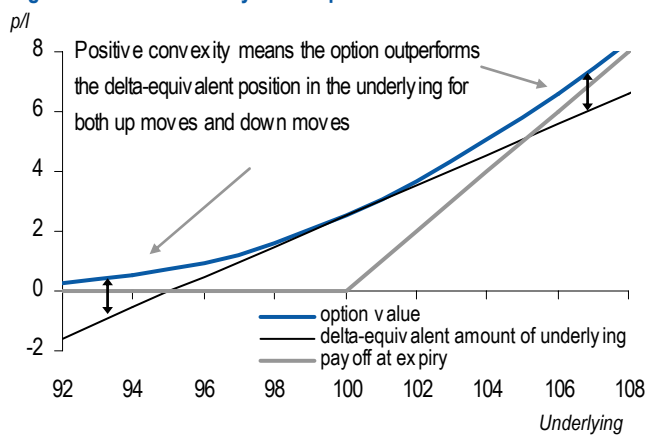
Source : JPMorgan

How is this gamma p/l actually made? In order to understand this, we must consider the p/l made over a single delta-hedging step, typically a day. Essentially the gamma (rate of change of the delta) measures the convexity of the option. This convexity (curvature) always works in favour of long options positions, meaning that although a replicating hedge is accurate for small moves in the underlying, for larger moves the long option will outperform the replicating hedge in both directions (Figure 92). For a delta-hedged option, the gamma p/l will be the outperformance of the option over the replicating hedge. Gamma p/l is largest for short-maturity, ATM options, where the chance of the option ending in-the-money can change rapidly, even for relatively small moves.

As spot moves the delta changes. The gamma measures the *change in delta per unit of the underlying,* and a long option is long gamma. The more the delta changes, the more the replicating delta-hedge will underperform the long options position (Figure 92 again). How much will it underperform? The actual amount depends on the difference between the initial and final delta of the option, which is *gamma x dS*, where *dS* is the change in the underlying. Assuming the gamma remains constant, the option gains or loses value at the *average* delta (*average delta = initial delta + ½ gamma x dS*) whilst the replicating hedge gains or loses value at the initial delta. Since the difference between the two is *½ gamma x dS*, the difference in performance between the option and the replicating hedge is *½ gamma x dS times the move in the underlying*. Overall **option outperformance is therefore *½ gamma x dS²*** (Figure 93).

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 4: Replication and Hedging

That is, **the gamma p/l from a move in the underlying is proportional to the gamma of the option and the square of the move.** Thus, the p/l from gamma hedging will be large when gamma is high, and moves in the underlying are large. For instance the gamma p/l from a 2% move will be four times that of a 1% move.

Figure 93 : Gamma p/l from delta-hedged options
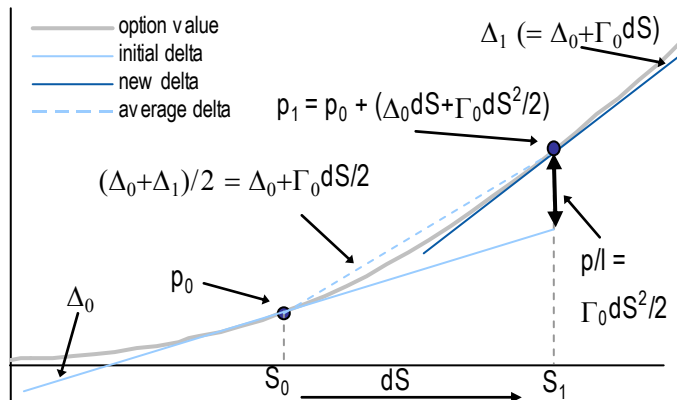
Source : JPMorgan

Box 12: Deriving the gamma p/l

1. Spot is initially S_0, option is worth p_0, delta and gamma are initially Δ_0 and Γ_0 respectively.

2. Spot increases, by dS, to S_1

3. Delta increases, to $\Delta_1 = \Delta_0 + \Gamma_0 dS$

4. Average delta is therefore $\Delta_{Avg} = \Delta_0 + \Gamma_0 dS/2$

5. New premium $p_1 = p_0 + \Delta_{Avg} dS$

6. $p_1 = p_0 + \Delta_0 dS + \Gamma_0 dS^2/2$ *(by 4 and 5)*

7. p/l on long option is $\Delta_0 dS + \Gamma_0 dS^2/2$ *(by 6)*

8. p/l on short futures is $-\Delta_0 dS$

9. total p/l is $\Gamma_0 dS^2/2$ *(by 7 and 8)*

Box 13: Calculating the gamma

The delta of a call option is given by: $\Delta = N(d_1)$

where $d_1 = \dfrac{\ln(S/K) + (r + \sigma^2/2)T}{\sigma\sqrt{T}}$ and $N(x) = \dfrac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-\frac{t^2}{2}} dt$

$\left\{ \begin{array}{l} S: \text{underlying} \\ K: \text{strike} \\ T: \text{maturity} \\ \sigma: \text{implied volatility (at the strike)} \\ r: \text{interest rates} \end{array} \right.$

Since gamma is the rate of change of delta with respect to the underlying: $\Gamma = \dfrac{\partial \Delta}{\partial S}$, differentiating $N(d_1)$ w.r.t. S gives

$\Gamma = \dfrac{N'(d_1)}{S\sigma\sqrt{T}}$ where $N'(x) = \dfrac{1}{\sqrt{2\pi}} e^{\frac{-x^2}{2}}$

Since the delta of a put is given by $\Delta = N(d_1) - 1$, differentiating has the same effect. Therefore the gamma of a put is the same as the gamma of call.

Note that as spot S, moves away from the strike K, $N'(d_1)$ will become small and gamma will decrease very rapidly. Then, since $N'(d_1)$ is independent of the monetary units of the underlying, gamma will be inversely proportional to the value of the underlying.

Gamma is a useful concept, but since it measures change in delta *per unit of underlying,* it is dependent on the absolute level on the underlying. For example the gamma of an option on a stock worth €10 will be double the gamma of the equivalent option on a stock worth €20 (with the same relative strike, volatility etc). Moving to the concept of the *dollar-gamma* is more useful, since it allows us to directly calculate the gamma p/l for a given percentage underlying move, making it much easier to compare gamma exposures across different underlyings.

To compute the dollar gamma and hence the actual cash p/l made from the delta-hedging process we first define the *dollar delta,* which is simply the cash equivalent exposure of the underlying. For example a position which is long $100,000 notional of call options, each with a delta of 30, will be long $30,000 of dollar delta. The *dollar gamma* is then defined to be the change in the dollar delta for a 1% move in the underlying. This dollar gamma is therefore equal to the normal gamma multiplied by the square of the value of the underlying and divided by 100 (Box 14).

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 4: Replication and Hedging

Box 14: Dollar delta and dollar gamma

Dollar delta: $\$\Delta$ = cash value of the underlying delta hedge (for 1 share of the underlying $\$\Delta = \Delta \times S$)

Dollar gamma: **$\$\Gamma$ = change in $\$\Delta$ for a 1% move in the underlying S.**

It can easily be shown that **$\$\Gamma = \Gamma S^2/100$**   ($S => S + S/100$; $\Delta => \Delta + \Gamma S/100$; $\$\Delta => \$\Delta + \Gamma S^2/100$)

Once the dollar gamma has been computed, the gamma p/l from a single delta-hedging step is easy to calculate: **for a return R, the gamma p/l is equal to 50 $\$\Gamma$ x R^2** (Box 15). For example if we hold a position which is long $100,000 of dollar-gamma and the underlying moves by 3%, the p/l will be 50 x $100,000 x 0.03^2 = $4,500.

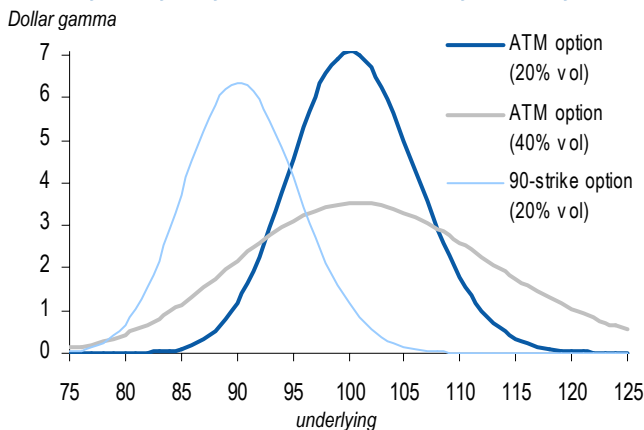### Box 15: From gamma to dollar gamma

For a change in spot $S => S + dS$, the gamma p/l will be $\Gamma \, dS^2/2$

Since $\$\Gamma = \Gamma S^2/100$, the p/l can be written as $(100 \, \$\Gamma)/ S^2 \times dS^2/2 = 50 \, \$\Gamma \, (dS/S)^2$
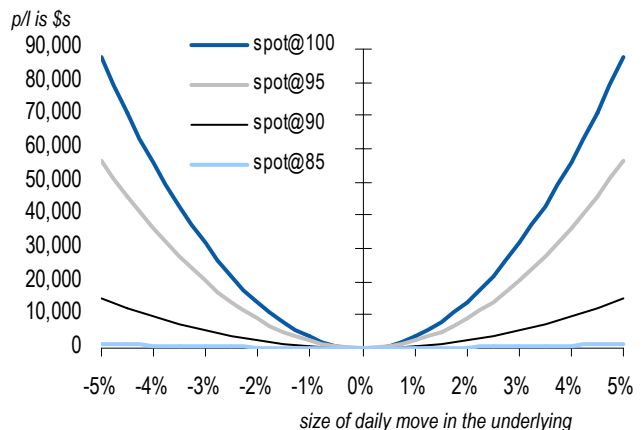
Gamma p/l = 50 $\$\Gamma$ R^2   where R is the return dS/S in percentage terms

Figure 94 : Dollar gamma is highest close to the option strike, and falls away rapidly away from the strike, especially if volatility is low

Dollar gamma

Source : JPMorgan

Figure 95 : Gamma p/l for a 100-strike option with spot at different levels. The p/l is quadratic in the size of the move.

p/l is $s

Source : JPMorgan

Example 18: Dollar Gamma

Suppose an underlying is trading at $100. We buy 1,000 contracts (100 shares each) of ATM call options with an implied volatility 20%. Delta is 0.51. Therefore the equivalent cash holding in the underlying is 10M x 0.51= $5.1M

The gamma is 0.07 which means that the delta will change by approximately 0.07 for a $1 move in the underlying. E.g. if spot moves to $99, delta will be approximately 0.44; if spot moves to $101, delta will be approximately 0.58.

The dollar gamma is given by multiplying gamma by the spot squared and dividing by 100. In this case it will be 100^2 x 0.07/100 = $7 per share. Since we hold options on 100,000 shares this equates to a dollar gamma of $700,000. This leads to a p/l of $3,500 for a 1% move in the underlying, $14,000 for a 2% move and $87,500 for a 5% move (Figure 95).

If the stock drops by 5%, then with spot now at 95, the new dollar gamma (for the next time step) will be only $450,00 (Figure 94) and a subsequent 1% move will now lead to a p/l of only $2,250 (Figure 95).

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan ◆

Part 4: Replication and Hedging

4.2: Theta - the cost of gamma

Of course, this positive convexity from the gamma must be paid for, and it is paid through the time decay or theta. Theta can be thought of as the amortized (non-linear) cost of the option, spread over its lifetime. In order to make money on the delta-hedging strategy over a single delta-hedging step (typically a day), it is necessary to make more on the gamma (from a large daily move) than is lost on the time decay.

Box 16: Gamma vs. theta – calculating the delta-hedging p/l

If delta hedging an option, the p/l in one time step is given by: $\frac{1}{2}\Gamma S^2\left[R^2 - \sigma^2\delta t\right]$ or $50\$\Gamma\left[R^2 - \sigma^2\delta t\right]$

where R is the return on the underlying, σ represents the implied volatility at which the option was purchased, and δt is the length of the time step.

This comes from the relationship between theta and gamma: $\theta = -\frac{1}{2}\Gamma S^2\sigma^2$ at the heart of the Black-Scholes analysis. (See e.g. *Hull* for further details.)

If delta-hedging daily, the strategy will profit if the daily realised volatility, in the form of the square of the return, R, is greater than the 'daily implied volatility' given by the option implied volatility scaled by the squared root of the time-step.

The above equation can be easily extended to give the p/l from delta-hedging over the lifetime of the options, where the summation is taken over all the delta-hedging steps.

$$p/l = \frac{1}{2}\sum_i \Gamma_i S_i^2\left[R_i^2 - \sigma^2\delta t\right] = 50\sum_i \$\Gamma_i\left[R_i^2 - \sigma^2\delta t\right]$$

Note that the only reason the exposure to volatility is non-constant is the changing gamma. See *Carr and Madan* for a thorough dissection of the p/l from delta-hedging a vanilla option, with possibly non-constant volatility.

Example 19: Daily breakevens

Suppose an option is bought at 25% implied volatility and delta-hedged daily. Then the daily breakeven return can be calculated by de-annualising the implied volatility (dividing by the square-root of 252) to get a value of 1.57%.

If the underlying moves by more than 1.57% *in either direction* the long option position will profit. If it moves by less than this breakeven it will lose. The magnitude of the gain or loss is determined by the dollar gamma, with large dollar gammas acting to magnify both gains and losses.

If delta-hedging weekly, the weekly breakeven will be 25%/sqrt(52) = 3.5%. Whilst monthly delta-hedging would give a breakeven of 7.2%.

Vega and other sensitivities

In reality the implied volatility of the option will change throughout the lifetime of the option. This will affect the option p/l directly: if volatility goes up, the option becomes more valuable; but also indirectly since the hedging quantities dictated by the delta will also change. Of course, if holding an option to maturity (and delta-hedging at the original implied volatility) the vega p/l over the lifetime of the option will be zero and the total p/l will the be the gamma vs. theta p/l described above. However, in practice since the changes in implied volatility alter the hedging amounts, the actual effects can be more difficult to quantify.

Option prices are also sensitive to changes in rates (measured by the *rho*) and to changes in dividends (sometimes call *mu*). However, for ease of exposition we assume zero interest rates and an underlying with zero dividend yield, unless stated otherwise.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

4.3: Options path-dependency: can volatility be captured by delta-hedging?

Suppose a market-maker buys and delta-hedges a vanilla option. Assume that whatever volatility is realised is constant and the option is delta-hedged over infinitesimally small time intervals. Then the market-maker will profit if and only if realised volatility exceeds the level of volatility at which the option was purchased. However:

1. The magnitude of the p/l will depend not only on the difference between implied and realised volatility, but where that volatility is realised, in relation to the option strike. If the underlying trades near the strike, especially close to expiry (high gamma) the absolute value (either positive or negative) of the p/l will be larger.

2. If volatility is not constant, where and when the volatility is realised is crucial. The differences between implied and realised volatility will count more when the underlying is close to the strike, especially close to expiry.

Indeed, for non-constant volatility, **it is perfectly possible to buy (and delta-hedge) an option at an implied volatility below that subsequently realised, and still lose from the delta-hedging**. For example consider the performance of delta-hedging 1-year Euro Stoxx 50 options from 1st November 2001 to expiry 1-year later on 1st November 2002. The index was initially at 3500 (with ATM implied volatility at 28.5%) and up until May 2002 remained in the range 3500-3800, realising around 20% volatility. After May, the index fell rapidly to around the 2500 level, realising high (around 50%) volatility on the way. Over the whole year, realised volatility was 36%.

Consider first the performance of the 2500-strike option (Figure 96). This was purchased at 33.4%, (above the 28.5% ATM volatility due to the put skew). Initially the option made a loss due to the low realised volatility, but this was kept small by the low gamma exposure so far away from the strike. However, as the index sold off toward the 2500 strike, increasing the gamma, volatility picked up dramatically and the option ended up making a profit of 120 index points (ip), or about 18 vegas.

In contrast the 4000-strike option (Figure 97) initially had much higher gamma than the 2500-strike option and consequently lost more due to the lower realised volatility. When the volatility did occur, this happened well away from the option strike, and the low gamma meant that the delta-hedging failed to capitalise on this high volatility with the delta-hedged options ending up making a *loss* of 60 ip (about 5 vegas). This is despite the fact that the option was purchased at an implied volatility of 26%, 10% below that subsequently delivered.

Figure 96 : The 2500-strike option profited due to the high volatility and gamma close to expiry…
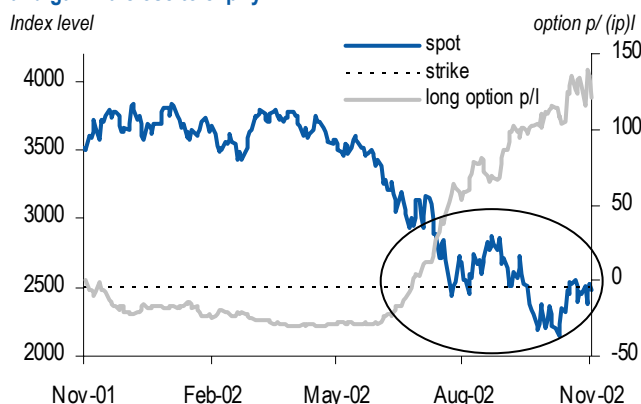
Source : JPMorgan

Figure 97 :…whereas the 4000-strike option made a loss, despite being purchased at implied volatility 10% below that later delivered
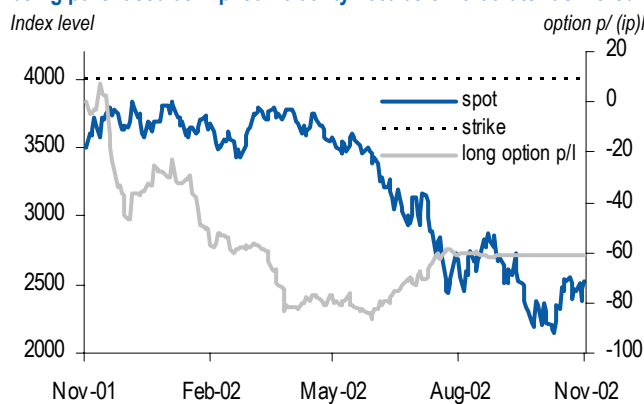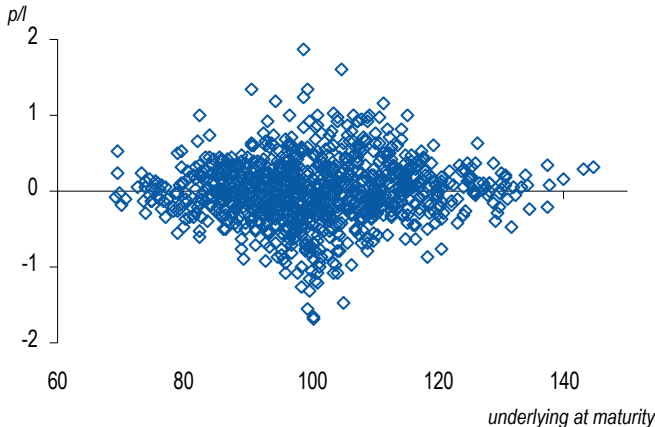

Source : JPMorgan

When delta-hedging an option, the p/l is a (daily) accrual of the difference between the realised volatility (in the form of the return squared) and the implied volatility scaled by the time step. However, the magnitude of the contribution of this (daily) accrual is weighted by the current dollar gamma of the option, which as it depends on the path of the underlying, is inherently unpredictable (see Box 16).

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
**Nicolas Granger**
(44-20) 7325-7033

The non-constancy of exposure to realised – implied volatility (squared) comes from the dollar gamma changing over the life-time of the option. Thus the difference between realised and implied volatility will count more when (dollar) gamma is high. An at-the-money option will have high exposure to the spread between implied and realised volatility whilst the underlying remains close to its initial value (Figure 98), whilst an OTM option will take on no significant exposure to volatility until the underlying trades at around the option strike (Figure 100).

Figure 98 : Simulation of p/l for hedging an ATM option: magnitude of returns is highest if underlying remains close to the money.
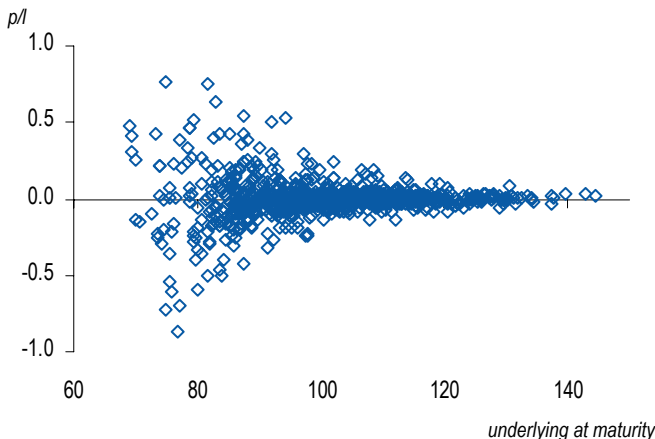

Simulation: option struck at 100, implied volatility=20%, maturity = 100 days, flat skew
Source : JPMorgan

Figure 99 : … returns are correlated with volatility but not perfectly, exposure is locally linear for small changes in volatility.
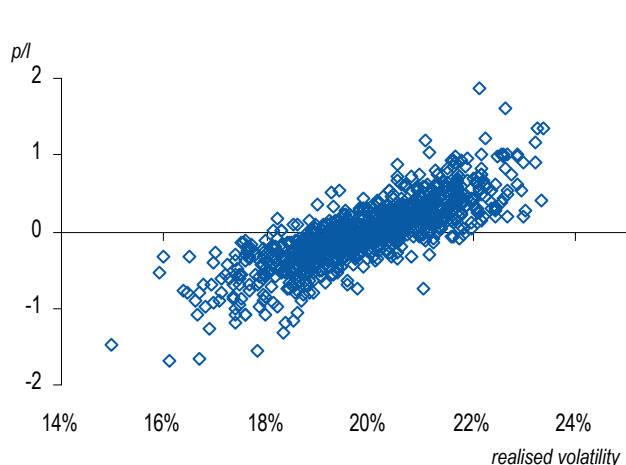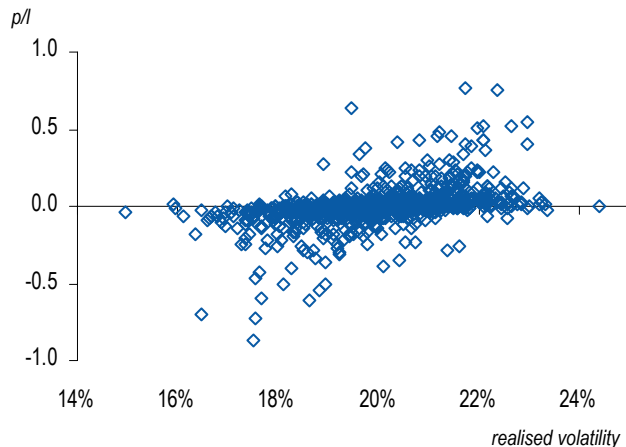

Simulation: option struck at 100, implied volatility=20%, maturity = 100 days, flat skew
Source : JPMorgan

Figure 100 : Simulation of p/l for hedging an OTM (80-strike) option: magnitude of p/l is highest when underlying is close to the strike

Simulation: option struck at 80, implied volatility=20%, maturity = 100 days, flat skew
Source : JPMorgan

Figure 101 : … however, due to the mostly low gamma, the exposure to volatility is poor

Simulation: option struck at 80, implied volatility=20%, maturity = 100 days, flat skew
Source : JPMorgan

To summarise:

- The p/l of a delta-hedged option does (generally) depend on the spread between the (square of the) implied volatility at which the option was purchased and the subsequent (squared) realised volatility

- This dependence is relatively strong if the underlying trades close to the strike of the option

- But the sensitivity to realised volatility tails off rapidly as the underlying moves away from the option strike

- This sensitivity is accounted for by the (dollar) gamma. **If there was such an instrument where the dollar gamma remained constant both over time and for changes in the underlying, then the exposure to volatility would be constant.**

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 4: Replication and Hedging

4.4: From options to variance swaps

In order to create a portfolio of options which has constant exposure to volatility (actually variance) we need to ensure it has a dollar gamma which is constant for moves both in the underlying and over time. The p/l from delta-hedging this portfolio would then be just a constant multiplied by the squared returns summed over the hedging steps, otherwise known as realised variance. The cost of this portfolio represents the price of exposure to realised variance. When considering how to construct such at exposure, there are three possible approaches to take:

1. **Use a single vanilla option**, but buy or sell additional amounts of this option to keep the dollar gamma constant over time. This has the advantage of using only a single option strike, but the option itself needs to be dynamically traded, and the position could end up with enormous amounts of the option as the underlying moves away from the strike and the gamma of a single option decreases; suggesting…

2. **Re-strike the option to maintain a constant gamma**. Start with an ATM option and on each re-hedging step, either sell or hold the previous option, and buy an amount of the new ATM to achieve the constant gamma exposure. This is better than (1) but still requires (potentially costly) dynamic trading of options; finally …

3. **Construct a portfolio of options in such a way that the dollar gamma remains constant over both moves in the underlying and with the passage of time.** This has the advantage of not requiring the dynamic trading of options, but the initial portfolio turns out to require a strip of options across a continuum of strikes.

It is this third approach that is used to construct variance swaps, at least in theory – in practice it is more like a combination of (2) and (3) above (see Section 4.9). This method has the advantage of using only a static options portfolio – although the options themselves must still be dynamically delta-hedged with the underlying. Perhaps more importantly this approach is to some extent **independent of the volatility process driving the underlying**. The first two approaches would require continually calculating the gamma over the course of the trade, and this gamma will be highly (and dangerously) dependent on the assumed implied volatility (and volatility process) at that time.
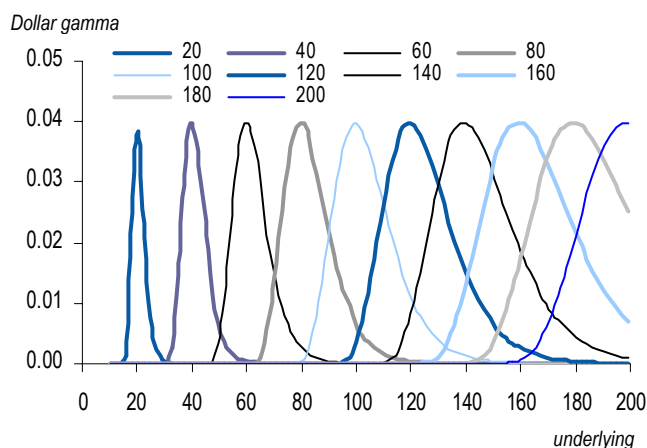
To achieve a constant dollar gamma across strikes what kind of portfolio is needed? One important observation is that (peak) dollar gamma increases linearly with strike (Figure 102). Naively, it may be thought that weighting the options in the portfolio (across all strikes) by the inverse of the strike will achieve a constant dollar gamma. It does have the property that each option in the portfolio has an equal peak dollar gamma (Figure 103). However, the dollar-gammas of the higher strike options 'spread out' more, and the effect of summing these *1/K*-weighted options across all strikes still leads to a dollar-gamma exposure which still increases with the underlying (Figure 105). In fact, in can be shown that this increase is linear, and therefore weighting each option by the inverse of the strike-squared will achieve a portfolio with constant dollar gamma (Figure 104, Figure 105).

Figure 102 : Peak dollar gamma of an option increases linearly with the underlying
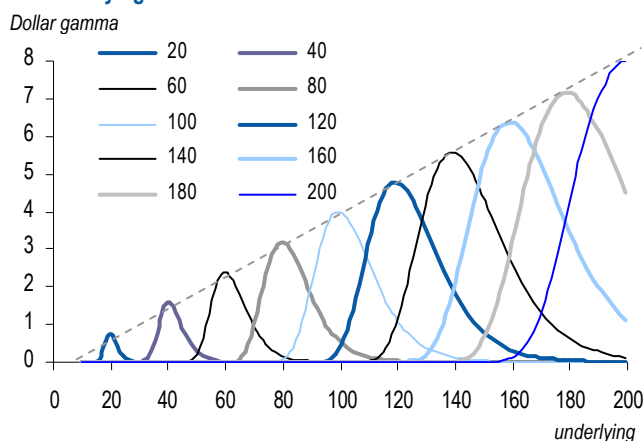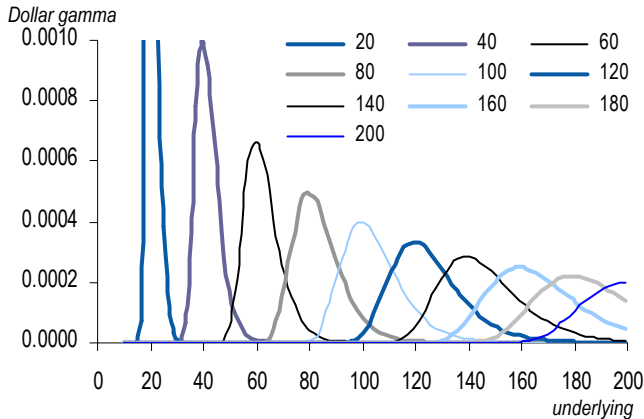

Source : JPMorgan

**Figure 103 : Peak dollar gamma of options divided by strike is constant**

Source : JPMorgan

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Figure 104 : Weighting options by the inverse of the strike-squared gives a portfolio with aggregate constant dollar-gamma

Dollar gamma

Source : JPMorgan

Figure 105 : Weighting options as the inverse of the strike squared gives constant dollar-gamma

Dollar gamma
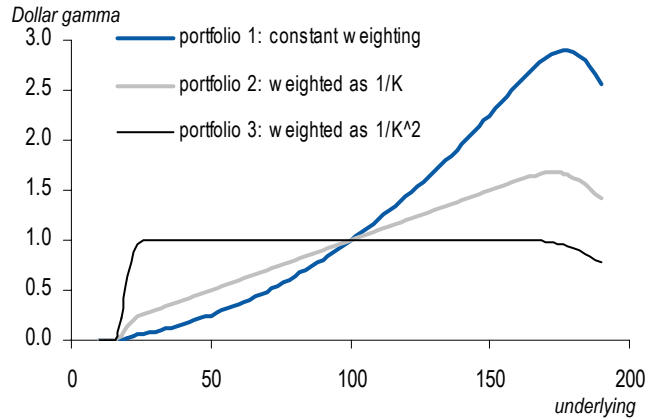
Source : JPMorgan

One way of explaining why a portfolio of (European-style) options weighted as the inverse strike squared gives a constant exposure to volatility is as follows. Recall (Box 16) that the p/l at each hedging interval is given by:

$$p/l = \frac{1}{2}\sum_i \Gamma_i S_i^2 \left[R_i^2 - \sigma^2 \delta t\right].$$ Only the $\Gamma_i S_i^2$ term prevents this being a direct exposure to variance.

Thus if we find a portfolio with gamma proportional to *1/S* then the exposure to implied minus realised variance will be constant. Since the gamma of a portfolio represents the second derivative of the portfolio's value with respect to the underlying, this means that the second derivative of the portfolio's value must be proportional to *1/S²*.

As $\dfrac{d}{dS}\left(-\ln S\right) = \dfrac{-1}{S}$ and $\dfrac{d^2}{dS^2}\left(-\ln S\right) = \dfrac{1}{S^2}$, the negative natural log of the underlying represents such a payoff.

In fact (by integrating 1/S² twice) **any function of the underlying S of the form** $-a\ln S + bS + c$ **will have a constant dollar gamma** (of *a*). The action of delta-hedging this combined portfolio will recover the realised volatility of the underlying. Note that the *bS + c* terms represent a static positions in bonds (*c*) and forwards (*b*) with no volatility exposure.

Unfortunately, log contracts are not traded in the market. But we can replicate such a contract with vanilla options using an infinite sum of calls and puts across the continuum of strikes, each weighted by the inverse square of strike. Integrating the value of this portfolio at expiry, demonstrates that the non-linear part of the payoff is indeed a negative log contract:

$$\int_0^{S_0} \frac{(K-S_T)^+}{K^2}\,dK + \int_{S_0}^{\infty} \frac{(S_T-K)^+}{K^2}\,dK = \int_{S_0}^{S_T} \frac{(S_T-K)}{K^2}\,dK = -\ln(S_T) - \frac{1}{S_0}S_T + \ln(S_0) - 1$$

Note $(S\text{-}K)^+$ is a shorthand for *Max(S-K,0)* – e.g. the payout from a call option on S with strike K. A put is $(K\text{-}S)^+$.

To summarise:

1. A portfolio of calls and puts, weighted as 1/strike-squared, has constant dollar-gamma;

2. Delta-hedging this portfolio provides constant exposure to the difference between implied and realised variance regardless of where the volatility is delivered;

3. Hence the p/l from delta-hedging this portfolio is proportional to difference between realised and implied variance.

A variance swap can therefore be created by replicating a log contract with options which are then delta-hedged. Although this argument is easy and intuitive, it does depend on certain strong (Black-Scholes) assumptions on the underlying. The power of the more general replication argument outlined below lies in the fact that it applies under much less restrictive conditions.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 4: Replication and Hedging

4.5: Variance swap replication in one page

In this section we make few assumptions: namely continuous trading is possible, complete markets (for options of all strikes) exist, no transaction costs, and the underlying (futures contract) follows a diffusion process (with no jumps).

Let F_t for t in the interval $(0,T)$ be the price for a future on the underlying maturing at some time after T.

Let $f(F_t)$ be the function $f(F_t) = \dfrac{2}{T}\left[\ln\left(\dfrac{F_0}{F_t}\right) + \dfrac{F_t}{F_0} - 1\right]$

Therefore $f'(F_t) = \dfrac{2}{T}\left[\dfrac{1}{F_0} - \dfrac{1}{F_t}\right]$ and $f''(F_t) = \dfrac{2}{TF_t^2}$ Note that both $f(F_t)$ and $f'(F_t)$ disappear at $F_t = F_0$

Using Ito's lemma: $\displaystyle\int_0^T \dfrac{F_t^2}{2} f''(F_t)\sigma_t^2 dt = f(F_T) - f(F_0) - \int_0^T f'(F_t)dF_t$ for any smooth function f

Hence substituting f as above gives: $\dfrac{1}{T}\displaystyle\int_0^T \sigma_t^2 dt = \dfrac{2}{T}\left[\ln\left(\dfrac{F_0}{F_t}\right) + \dfrac{F_t}{F_0} - 1\right] - \dfrac{2}{T}\int_0^T\left[\dfrac{1}{F_0} - \dfrac{1}{F_t}\right]dF_t$

The LHS of the above equation represents the average realised variance of the future over the period $(0, T)$. The final term represents a continuously rebalanced futures position, which is easily replicated and costless to initiate. The first term on the RHS represents a static position in a contract paying $f(F_t)$ at expiry. This can be replicated with options:

This replicating portfolio can be found directly by using the fact that every smooth (twice differentiable) payoff of futures f, can be replicated by a portfolio of cash, the underlying and options as follows (see *Carr and Madan* for details):

$f(F_T) = f(\kappa) + f'(\kappa)\left[(F_T - \kappa)^+ - (\kappa - F_T)^+\right] + \displaystyle\int_0^\kappa f''(K)(K - F_T)^+ dK + \int_\kappa^\infty f''(K)(F_T - K)^+ dK$ for any κ.

Using the function f above, and $\kappa = F_0$ gives: $\left[\ln\left(\dfrac{F_0}{F_T}\right) + \dfrac{F_T}{F_0} - 1\right] = \displaystyle\int_0^{F_0} \dfrac{(K - F_T)^+}{K^2}dK + \int_{F_0}^\infty \dfrac{(F_T - K)^+}{K^2}dK$

Thus (at expiry) $\dfrac{1}{T}\displaystyle\int_0^T \sigma_t^2 dt = \dfrac{2}{T}\int_0^{F_0} \dfrac{(K - F_T)^+}{K^2}dK + \dfrac{2}{T}\int_0^{F_0} \dfrac{(F_T - K)^+}{K^2}dK - \dfrac{2}{T}\int_0^T\left[\dfrac{1}{F_0} - \dfrac{1}{F_t}\right]dF_t$

Therefore holding a portfolio of calls and puts, weighted according to the inverse of the strike squared, and simultaneously holding $1/F_t - 1/F_0$ of futures at time t will pay out average realised volatility squared at expiry. (Note, in a Black-Scholes environment, that this continuously rebalanced futures holding encapsulates the delta-hedging of the options.) Therefore the forward cost of this portfolio is simply:

$K_{VAR}^2 = \dfrac{2e^{rT}}{T}\left[\displaystyle\int_0^{F_0} \dfrac{P_0(K)}{K^2}dK + \int_{F_0}^\infty \dfrac{C_0(K)}{K^2}dK\right]$

The strike-squared, K_{VAR}^2, can be thought of as the (future-valued) inverse strike-squared weighted sum of the time values of the options portfolio

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

4.6: Sensitivity to skew and convexity

Skew is commonly thought of as an important component of variance swap prices, with put skews seen as having a much greater impact on prices than call skews. Although, in practice, this is a useful framework for thinking about how variance swap prices behave, it is not theoretically correct.

For example, suppose an underlying has 3-month ATM volatility of 20%, a linear put skew of 5% (per 10 volatility points), with put volatilities capped at 35% and all OTM call volatilities flat at the level of the ATM volatility (Figure 106). The theoretical 3-month variance strike can then be calculated to be 23.05 (pricing OTM options 1 point apart between 1 and 300). Now consider a situation where the skew is a mirror image of that in Figure 106. ATM volatility is 20%, and all OTM put volatilities are flat at 20%, but call volatilities increase linearly by 5% per 10 points as they become more OTM, capped at 35% (Figure 107). In this case the theoretical variance swap price is virtually identical at 23.15.

Figure 106: High put skew increases the variance strike above ATM volatility…
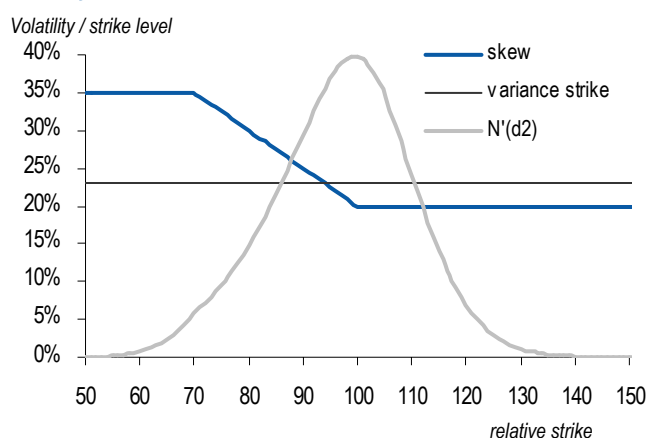
Source : JPMorgan

Figure 107 : …but increasing the call skew has almost exactly the same effect.
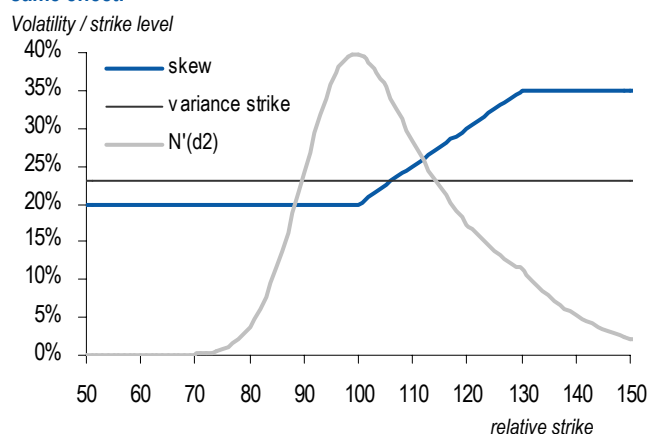
Source : JPMorgan

Why does this symmetry exist? Since the 1/strike-squared replicating portfolio has a much higher weighting of puts, it would be natural to assume that their associated implied volatilities should have a proportionally greater effect on the variance swap strike. Also, in practice, variance swaps *can* effectively be priced and hedged with ATM volatility plus a contribution from the skew.

To answer the first point, the **exposure to the skew curve is symmetrical**. That is, the contribution to the variance swap price of the volatility of an OTM call is exactly the same as from an OTM put with same (risk-neutral) probability, $N(d2)$, of ending in the money. The OTM puts have a greater weighting in the replicating portfolio, simply because their dollar gamma is lower and so they must be scaled up in order to provide constant dollar gamma across the range of strikes.

For the second point, in reality, equity skews look much more like Figure 106 than Figure 107, at least for short dated index volatility. This explains why Derman's approximation assuming a linear skew is most successful in these conditions (See Section 2.4). Directionality of volatility, at least in the case of indices, makes volatility skews highly asymmetric, with OTM calls trading much closer to ATM volatility than the corresponding OTM puts.

So what exactly determines the contribution of volatilities across the skew curve to the variance swap price? Clearly a very out-of-the money option should make a relatively small contribution to the variance swap price. (This must be true in order to replicate variance swaps in practice, otherwise the value of the replicating portfolio will be very sensitive to the exact choice of replicating portfolio.) Alternatively, one can think of the chance of needing the constant dollar gamma in the area represented by the very out-of-the-money option as being very small. It then seems sensible that a variance swap represents a kind of weighted average of volatilities across the skew curve, with the closer-to-the-money volatilities higher weighted.

In fact, this is exactly the case, with the average being taken over the variances rather than the volatilities, and the weighting function simply being the risk-neutral probability density function, $N'(d2)$, that the corresponding OTM option ends up in-the-money.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 4: Replication and Hedging

If we define the variable z to be the standard Black-Scholes parameter $d2$: $z = d2 = \dfrac{\ln(S/K) + (r - \sigma^2/2)T}{\sigma\sqrt{T}}$ then;

It can shown that: $K_{VAR}^2 = \displaystyle\int_{-\infty}^{\infty} N'(z)\sigma_{BS}^2(z)dz$

The $N'(d2)$ term is the probability density function for the underlying at expiry, T. That is, the cumulative distribution $N(d2)$, gives the (risk-neutral) probability that the underlying will be trading above z at time T. Thus the parameter z, simply represents the 'moneyness' of the corresponding OTM option. This means that **the variance swap price is a weighted sum of squared option implied volatilities weighted by the probability that the (OTM) option will end in-the-money.** See *Gatheral* 2006 for a detailed of the proof of this.

Furthermore, if the skew curve is quadratic in the variable z (the moneyness of the option) of the form:

$$\sigma_{BS}^2(z) = \sigma_0^2 + \alpha z + \beta z^2$$

then substituting and integrating gives $K_{VAR} = \sigma_0^2 T + \beta T$, i.e. in '*d2-space*' the variance swap price is not affected by the linear component of the skew α, but only on the base level of volatility σ_0, and the convexity parameter β.

This explains the phenomena seen in Figure 106 and Figure 107: very different (linear) skews, but similar convexities give almost identical variance swap strikes. The reason that there is any difference at all in variance swap strikes, is because in *d2*-space, the convexity in Figure 107 is slightly greater, since the maximum volatility is at a strike of 130, which is slightly less out-of-the money than the 70 strike where the maximum volatility is achieved in Figure 106. Indeed, if we keep the 'linear' skew at a slope of 5% or -5% but do not apply the cap until it reaches 50%, which it will do at strikes of 40 and 160 respectively, the version with the upward sloping call skew (similar to Figure 107) will now price about 80bps above the version with the downward sloping put skews (similar to Figure 106) since a 160 strike is significantly more in-the-money that the 40 strike.

This also helps to explain why the convexity has a greater effect on longer maturity variance. As maturity increases, the probability of far out-of-the-money (in relative strike terms) options ending in the money increases. Therefore the relative weight of e.g. the 160-strike call will increase with maturity. At shorter maturities perhaps only the 80-120 portion of the skew surface significantly affects the variance swap price, and this part of the skew is relatively linear (at least for indices). At 5 years, strikes out to say 50 -180 become relevant to the pricing and the convexity of these strikes can be much higher.

In practice, skew is most often thought about in terms of relative strikes, rather than moneyness – but the point remains: the contribution of a point on the skew curve to the variance swap price depends on the (risk-neutral) probability $N(d2)$ of the associated OTM option ending in-the-money. Therefore ATM volatility will provide the greatest contribution to variance swap prices – particularly for short maturities. Both high put skews, and high call skews (where OTM calls have higher volatilities than ATM) will increase variance swap prices. In fact there is some kind of feedback effect here because the contribution of each volatility is determined by the probability of an option being in-the-money, but, for example high put skews will increase the probability of OTM puts ending in the money (greater volatility on the downside makes OTM puts easier to reach) so both the volatility and the weighting factor will increase.

To conclude, from the equation: $K_{VAR}^2 = \displaystyle\int_{-\infty}^{\infty} N'(z)\sigma_{BS}^2(z)dz$ we can make a number of specific observations:

1. Given a flat skew, variance will price at the same level as ATM volatility.

2. Positive convexity will always act to increase variance swap strikes.

3. With a negatively convex skew – OTM volatilities are (on average) less than ATM volatility – it is theoretically possible that variance could price below ATM volatility.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 4: Replication and Hedging

4.7: Variance swap Greeks

In this section we consider the Greeks for variance swaps, giving information about the sensitivity of variance swaps to various market variables. We work directly by differentiating the mark-to-market value of the variance swap contract. An alternative approach would be to calculate the replicating portfolio and sum the Greeks of the contributing options. Note for simplicity of exposition we assume a world with zero interest rates; alternatively in the presence of non-zero rates, the following parameters represent the *forward* Greeks and would have to be discounted by the appropriate discount factor. In analogy with the Greeks for vanilla options, we return to considering realised volatility and variance swap strikes in percentage terms rather than the standard whole-number units.

The value of a variance swap, per unit vega-notional, at time t is given by:

$$P_t = \frac{1}{2K_0}\left[\sigma^2_{Expected,t} - K_0^2\right]$$

where K_0 is the initial (fixed) variance swap strike and $\sigma_{Expected,\ t}$ is the expected realised volatility, at time t, between trade inception and maturity. This expected volatility is then calculated in the same way as the mark-to-market p/l as a combination accrued realised variance to date and future implied variance. Note also that intra-day there is a term representing the square of the move which will act to give the variance swap delta on an intra-day basis:

$$\sigma^2_{Expected,t} = \frac{t-1}{T}\sigma^2_{0,t-1} + \frac{252}{T}\left[\ln\left(\frac{S_t}{S_{t-1}}\right)\right]^2 + \frac{T-t}{T}K^2_{t,T}$$

where $\sigma_{0,t-1}$ is the realised variance accrued to day *t-1*, and $K_{t,T}$ is the strike of a variance swap on day t expiring at *T*. S_{t-1} is the value of the underlying at the close on day *t-1*, and S_t is the value of the underlying at the valuation time on day *t*.

The Greeks of the variance swap can then be calculated by differentiating P_t.

Gamma:

Gamma comes only from the exposure to realised volatility on each day:

$$\Gamma = \frac{\partial^2 P}{\partial S_t^2} = \frac{252}{K_0 T}\left[\frac{1}{S_t^2}\right]$$

Since the dollar gamma is achieved by scaling the gamma by the spot squared, this gives a constant dollar gamma as expected (ignoring higher order errors).

Theta:

Calculating the theta of the variance swap gives:

$$\theta = \frac{\partial P}{\partial t} = \frac{-K^2_{t,T}}{2TK_0}$$

In particular, if the variance strike doesn't change, theta remains constant. Note, that since T is measured in days this value represents a daily theta. Annualising gives a value for theta of:

$\theta = 252\dfrac{-K^2_{t,T}}{2TK_0}$ which can then be shown to satisfy the formula $\theta = -\dfrac{1}{2}\Gamma S^2\sigma^2$ taking σ^2 as the implied variance from the variance swap strike (see Box 16, p76).

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Vega:

We can calculate exposure to volatility in terms of (i) sensitivity to changes in ATM volatility, (ii) variance strike or (iii) implied variance (strike squared).

(i) To compute the vega in terms of sensitivity to ATM volatility, we must make some assumptions about how the variance swap strike relates to ATM volatility. If we assume the Derman approximation: $K_{t,T}^2 = \sigma_{ATM}^2 \left(1 + 3(T-t)skew^2\right)$ then:

$$\frac{\partial P}{\partial \sigma_{ATM}} = \frac{1}{2K_0}\frac{T-t}{T}\frac{\partial}{\partial \sigma_{ATM}}\left(\sigma_{ATM}^2\left(1+3(T-t)skew^2\right)\right) = \frac{\sigma_{ATM}}{K_0}\frac{T-t}{T}\left(1+3(T-t)skew^2\right)$$

Note also that this approximation allows the calculation of sensitivity to the skew or the skew squared:

$$\frac{\partial P}{\partial skew^2} = \frac{3}{2K_0}\frac{(T-t)^2}{T}\sigma_{ATM}^2 \text{ and } \frac{\partial P}{\partial skew} = \frac{3}{K_0}\frac{(T-t)^2}{T}\sigma_{ATM}^2 skew$$

i.e. increasing skew will increase the value of the variance swap, and do so by more, if there is more time remaining until expiry. (Compare Figure 21)

Computing sensitivities to the strike (or strike squared) is more straightforward and needs no such assumptions about the skew surface:

(ii) $\dfrac{\partial P}{\partial K_{t,T}} = \dfrac{T-t}{T}\dfrac{K_{t,T}}{K_0}$ and (iii) $\dfrac{\partial P}{\partial K_{t,T}^2} = \dfrac{1}{2K_0}\left[\dfrac{T-t}{T}\right]$.

These all tell us that the exposure to implied variance (or volatility) decreases with time as the accrued realised volatility is locked in to the p/l.

Delta:

Firstly, assuming that the variance strike K has no sensitivity to the underlying, the variance swap can be seen to take on delta only intra-day:

$$\Delta = \frac{\partial P}{\partial S_t} = \frac{T-t}{T}\frac{252}{K_0}\left[2\ln\left(\frac{S_t}{S_{t-1}}\right)\frac{1}{S_t}\right] = \frac{252}{K_0 T}\ln\left(\frac{S_t}{S_{t-1}}\right)\frac{1}{S_t}$$

This represents the replication of the log contract which will have to be done at the end of the day to capture that day's realised variance.

If the variance strike itself has a dependency on the underlying (implied variance is directional) then the variance swap acquires other sources of delta, in addition to the intra-day delta discussed above.

For example, using Derman's approximation (see p30) $K_{t,T}^2 = \sigma_{ATM}^2\left(1+3(T-t)skew^2\right)$ where σ_{ATM} is the Black-Scholes ATM implied volatility (or at-the-money *forward* volatility if rates or dividends are non-zero).

Calculating delta (only on the close to avoid the intra-day delta discussed above) gives:

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

$$\Delta = \frac{\partial P}{\partial S_t} = \frac{1}{2K_0}\frac{T-t}{T}\frac{\partial}{\partial S_t}\left(K_{t,T}^2\right)$$

Then since $K_{t,T}^2 = \sigma_{ATM}^2\left(1+3(T-t)skew^2\right)$ we have $\dfrac{\partial K_{t,T}^2}{\partial S} = 2\sigma_{ATM}\dfrac{\partial\sigma_{ATM}}{\partial S}\left(1+3(T-t)skew^2\right)$

Which implies: $\dfrac{\partial K_{t,T}^2}{\partial S} = \dfrac{2K_{t,T}^2}{\sigma_{ATM}}\dfrac{\partial\sigma_{ATM}}{\partial S}$

then $\dfrac{\partial K_{t,T}^2}{\partial S} = \dfrac{2K_{t,T}^2}{\sigma_{ATM}}\times skew$ since skew = $\dfrac{\partial\sigma_{ATM}}{\partial S}$ is skew is linear.

Therefore delta $= \dfrac{\partial P}{\partial S} = \dfrac{1}{2K_0}\dfrac{T-t}{T}\dfrac{2K_{t,T}^2}{\sigma_{ATM}}\times skew$.

Thus with a 'normally shaped' negatively sloping skew, the delta of the variance swap will be negative, at least if the skew curve is sticky with strike. This fits with the intuition that (implied) volatility will tend to go down as the underlying rallies – which is exactly what a (negative) linear skew represents.

Other sensitivities to consider include:

Jumps:

The replication argument assumes a continuous diffusion process without jumps. In reality jumps (discontinuities) can happen, and the volatility associated with them is not guaranteed to be captured by the replication process. However, the error associated with this is of order jump (return) cubed and is therefore small enough in practice to be ignored. Although, in the case of a single-stock where an M&A event is anticipated, the probability of a large jump increases and this may be reflected in an increase in variance swap prices.

Stochastic Interest Rates:

Our replication argument is constructed using forwards. In reality the variance swap pays out on the volatility of the spot underlying. If rates (and dividends) are deterministic, this should make no difference to the variance swap strike. If rates are uncertain, this will act to change the relationship between spot and forward volatility, altering the variance swap strike. Except for very long-dated variance the effect on variance swap prices should be small, and in practice can offset the effect of jumps.

Convexity:

As explained above, it is really convexity and not skew which acts to increase the variance swap strike above ATM volatility. Clearly then, rising convexity will increase variance swap strikes as the 'average' implied volatility used in Gatheral's formula goes up.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 4: Replication and Hedging

4.8: Setting up a replicating portfolio

In the preceding sections we showed that a variance swap could be statically replicated by a portfolio of (out-of-the-money) options, weighted according to 1/strike-squared. How do we weight this portfolio to achieve a specific variance-notional or vega-notional exposure? For example if we wish to replicate a variance swap for €100,000 of vega notional, how many of the calls and puts do we need in the replicating portfolio? We know the *relative* proportions of the puts and calls but what is the absolute amount to hold?

Let Π be forward of the price of the portfolio:

$$\mathbf{P} = e^{rT}\left[\int_0^{F_0} \frac{P_0(K)}{K^2}dK + \int_{F_0}^{\infty} \frac{C_0(K)}{K^2}dK\right]$$

The above portfolio (together with the dynamic futures position) will pay

$$\frac{T}{2}\sigma_R^2 \text{ at expiry, where } \sigma_R^2 = \frac{1}{T}\int_0^T \sigma_t^2 dt .$$

Scaling this portfolio by 2/T (which will cost *2Π/T*) will produce a payout of realised volatility squared, but in percentage terms. To put this into the whole-number units consistent with standard quoted variance swap strikes (e.g. 20% volatility expressed as 20) we must multiply by 100^2. The (square-root of) the cost of this portfolio, the variance strike represents the fair cost of future realised variance.

Therefore the fair price for $2\times100^2/T$ lots of the portfolio **P**, paying out realised volatility (squared) in whole number amounts is given by:

$$K_{VAR}^2 = \frac{2\times100^2\times\Pi}{T}$$

Therefore, to be long N_{VAR} of variance notional we must buy $\dfrac{2\times100^2\times N_{VAR}}{T}$ lots of the portfolio **P,**

equivalent to $\dfrac{2\times100^2\times N_{VAR}}{T\times K^2}$ of each OTM K-strike option.

This cost of this portfolio will therefore be $N_{VAR}\times K_{VAR}^2$. The value of the portfolio and its associated delta hedging at expiry will be $N_{VAR}\times\sigma^2$, giving an overall p/l of $N_{VAR}\times\left(\sigma^2 - K_{VAR}^2\right)$ as required.

At inception if the cut-off between puts and calls is taken to be the current forward price of the underlying, no net delta-hedge is needed. However, during the lifetime of the variance swap, the delta hedge will need to be adjusted such that, against the long portfolio of options, there is a constant monetary (dollar-) value of a long forward position but short a constant number of shares of a forward position. Thus at the end of each day's close the notional value of the delta hedge is adjusted so that it equals

$$\frac{2\times100^2\times N_{VAR}}{T}\left(\frac{F_0 - F_t}{F_0}\right) \text{ where } F_0 \text{ was the original forward price at time 0, and } F_t \text{ is the current forward price.}$$

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 4: Replication and Hedging

To achieve a specific *vega* notional, $N_{VAR}=N_{Vega}/2K_{VAR}$ we must trade $\dfrac{100^2 \times N_{Vega}}{T \times K_{VAR}}$ lots of the portfolio **P**.

Since $K_{VAR} = 100\sqrt{\dfrac{2\Pi}{T}}$ the required amount of the portfolio is $\dfrac{100 \times N_{Vega}}{\sqrt{2\Pi T}}$ or $\dfrac{100 \times N_{Vega}}{K^2 \sqrt{2\Pi T}}$ of each (OTM) option.

In practice it is not possible to trade a continuum of option strikes, so we now examine how to (approximately) construct a variance swap from a finite, tradable strip of OTM options. For simplicity, we assume the strikes are equally spaced.

Suppose we have available liquid OTM options of strikes $K_0...K_r$, $K_{r+1}...K_n$ with all options up to K_r puts, and all subsequent options calls. Suppose the forward F_0 such that $K_r < F_0 \le K_{r+1}$ and suppose that the options are equally spaced so that $K_{r+1}-K_{r+1}= \Delta_K$ for all i. The replicating variance swap portfolio is then given by

$$\sum_{i=0}^{r} \frac{\Delta_K P(K_i)}{K_i^2} + \sum_{i=r+1}^{n} \frac{\Delta_K C(K_i)}{K_i^2} \quad \text{approximating the true replicating portfolio:} \quad \int_0^{F_0} \frac{P_0(K)}{K^2}dK + \int_{F_0}^{\infty} \frac{C_0(K)}{K^2}dK$$

To get a variance notional exposure we will need $\dfrac{2 \times 100^2 \times \Delta_K \times N_{Var}}{T \times K^2}$ of each option (contract).

There is a slight complication here with contract sizes. If one option contract represents e.g. 10 shares (or €10 per index point) then we need only 1/10 as many option contracts as we would otherwise. Therefore we must divide the above amount by the contract size, to compute the actual number of options required.

This kind of approximation slightly over-values the true theoretical variance swap strike due to convexity issues. Derman (*Derman, 1999*) discusses this effect and also presents a robust method for using whatever available strikes are to replicate a variance swap.

Example 20: Constructing a variance swap

Suppose we are replicating a 6-month maturity Euro Stoxx 50 variance swap (variance notional of €2,500). Strikes are 200 ip apart and the forward is 3868. Each option represents €10 per index point. The replicating portfolio will consist of puts at strikes 3800, 3600, 3400…1200 and calls at 4000, 4200, 4400…6000. How much of each option contract is needed?

$$\text{No of contracts} = \frac{2 \times 100^2 \times \Delta_K \times N_{Variance}}{T \times K^2 \times Contract\_size} = \frac{2 \times 100^2 \times 200 \times 2,500}{0.5 \times K^2 \times 10} = \frac{2 \times 10^9}{K^2}$$

For example, the number of 3600 puts required is $\dfrac{2 \times 10^9}{K^2} = \dfrac{2 \times 10^9}{3600^2} = 154.3$

The prices, and numbers of options required for each strike are given in Table 13 below:

Suppose the market rallies by 1% over the following day. In order to maintain a "flat-delta" for the overall portfolio, we will need to sell on the close the following notional size:

$$\frac{2 \times 100^2 \times N_{VAR}}{T}1\% = \frac{2 \times 100^2 \times 2,500}{0.5}1\% = €1M \text{ notional}$$

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 4: Replication and Hedging

Table 13: Replicating portfolio for a 6-month Euro Stoxx variance swap, variance notional €2,500.

Option strike	Option type	Spot premium (€)	Fwd. premium x bucket size/strike^2	# Option contracts	Price of options (€)
1200		5.90E-12	8.36E-16	1389	-
1400		5.90E-09	6.14E-13	1020	-
1600		1.20E-06	9.56E-11	781	0
1800		7.10E-05	4.47E-09	617	0
2000		1.80E-03	9.18E-08	500	9
2200		0.024	1.01E-06	413	100
2400		0.195	6.90E-06	347	675
2600	PUTS	1.054	3.18E-05	296	3,118
2800		4.165	1.08E-04	255	10,624
3000		12.777	2.90E-04	222	28,393
3200		26.645	5.31E-04	195	52,041
3400		47.968	8.46E-04	173	82,989
3600		83.143	1.31E-03	154	128,308
3800		140.932	1.99E-03	139	195,197
4000		103.483	1.32E-03	125	129,354
4200		40.164	4.64E-04	113	45,538
4400		11.899	1.25E-04	103	12,292
4600		2.853	2.75E-05	95	2,696
4800	CALLS	0.69	6.11E-06	87	599
5000		0.145	1.18E-06	80	116
5200		0.027	2.04E-07	74	20
5400		4.40E-03	3.08E-08	69	3
5600		6.60E-04	4.29E-09	64	0
5800		9.00E-05	5.46E-10	59	0
6000		1.10E-05	6.23E-11	56	0
Price of one portfolio Π			0.00706	**Price of replicating portfolio for €2,500 of variance notional**	692,074
x 40,000 (2 x 100^2 / T)			282.31		
Square rooted gives….			16.80		

4.9: Replication and hedging in practice

As demonstrated above, a variance swap can be statically hedged with a portfolio of out-of-the-money (European-style) options, weighted according to the inverse squares of their strikes. This makes it easy, in theory, to calculate the fair value of a variance swap, assuming option prices are available across the entire range of strikes. In practice, traded strikes are not continuous, although for major liquid indices they are closely spaced (0.4% notional apart for the S&P, 1% for the FTSE, 1.4% for the Euro Stoxx). A more serious limitation is the lack of liquidity in OTM strikes, especially for puts, as these provide a relatively large component of the variance swap price in the presence of steep put skews. S&P options are listed down to a strike of 600 (c45% of current spot), FTSE to 3525 (c60% of spot) and Euro Stoxx down to 600 (c15% of spot), although in reality, liquidity does not even reach this far.

The problem with the lack of OTM puts can be seen from following through the practical example of setting up the replicating portfolio above. The long futures position is used to create a pay-out which is equivalent to a long log contract *plus* realised variance. In opposition the long options/short forward position is used to create a short log contract and pay the fixed strike. Supposing the market falls significantly, the delta-hedge will be long the log contract (a will hence lose), while the options should counteract this by being short the log contract. However, if not enough downside puts were used, the options portfolio will not fully reflect the short log-contract exposure needed and hence the overall hedge will lose money. This lack of liquidity at the wings has led to the development of conditional variance swaps which can remove exposure to volatility once the underlying moves into areas where vanilla options are illiquid.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

In practice, market-makers will not attempt to hedge with the entire strip of options but typically will use only two or three – including one close to the money and one or more OTM, but liquid, puts. Alternatively they could approach the replicating by hedging the vega with an OTM put whose implied volatility coincides with the variance swap strike – close to the money for 1-month maturity, 95% for 3-6 months, 90% for 1-year, 85% for 2-3 years etc. In this case they would also look to buy back the wings/convexity separately. Other proxies, such as CDS or EDS could also be used instead of the more deeply out-of-the-money puts.

One problem with this kind of approach is that the partial hedge is no longer static, and must be dynamically managed. For example if the market sells off towards the strike, the market maker will have to trade further OTM puts to ensure that their exposure to volatility, in the form of dollar-gamma, remains constant. This makes the actual variance swap replication more akin to a combination of alternatives 2 and 3 listed on p79. Here the constant dollar gamma would be maintained by a combination of holding a portfolio which has roughly constant dollar gamma if the underlying does not move too much, and re-hedging by trading more options if the underlying does move significantly.

In practice, pricing models will often only price contributing options for the variance swap hedge between 1 and 99 delta. For example, on the Euro Stoxx (6-month maturity) this would only price puts down to 2700 (ref 3822) or around 70% of spot. The difference from this truncated and the full replication pricing is typically in the region of 5-25bps.

Another limitation comes from the discrete nature of adjusting the delta hedge on the close, which introduces possible errors due to large daily moves. However these moves are actually of order the cube of the move (as in jumps for the continuous case, see section 4.7) and hence are negligible for all but very large moves. Also as previously noted if interest rate changes are related to changes in the spot underlying, this can also have an impact on the ability to accurate replicate the log-contract and realised volatility

This and other differences between the theoretical and mid-market price of a variance swap contribute to the so-called *variance swap basis*, see Section 4.7. This basis tends to vary across maturities, being larger for longer maturities. At shorter dates a value of around 0.5 vegas would be typical.

4.10: Effects of variance swap hedging

Market-makers who trade variance swaps may hedge their positions by replicating the opposite variance swap position through the replicating options portfolio. This replicating portfolio then needs to be delta hedged. The effects of delta-hedging this portfolio are different to that from normal delta-hedged options for two principal reasons:

1. The actions of delta-hedging the options could potentially act to the disadvantage of the counterparty's position.

2. Since variance swap contracts typically measure close-close realised volatility, the options must be delta-hedged on the close only to capture this.

We investigate both of these properties below.

Firstly, consider a market-maker who has sold a call option to a counterparty. The market-maker will be short the option (and delta) and will therefore buy back the delta in the underlying. If the underlying rallies, the short-delta exposure from the short option will increase, and the market-maker will therefore have to buy more of the underlying in the market. Similarly, if the underlying sells off, the market-maker will have to sell the underlying. In both cases – if the position being hedged is big enough – the action of delta-hedging will have the effect of increasing volatility in the underlying: magnifying both up-moves and down-moves. The argument for a put option is analogous. Similarly, if the market-maker is long options and delta-hedging, the hedging will act to suppress volatility in the underlying, potentially to the advantage of the counterparty who is short the option.

The situation with variance swaps is different. Suppose that the market is such that hedge-funds, or other market participants, have generally sold index variance swaps to market-makers (as was the case before the May 2006 correction). Note that no exchange of options has taken place here – the parties have just taken opposite sides in a contract for difference. Suppose that the volatility sellers do not hedge their variance swaps (they have sold the variance swaps specifically for the direct volatility exposure they offer). But assume that market-makers hedge their short volatility

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 4: Replication and Hedging

exposure. A market-maker who is long the variance swap can offset the risk by shorting the replicating portfolio of options, and delta-hedging. They will therefore be short gamma in the options market (Figure 108).

As described above, a delta-hedger who is short options will act to increase volatility in the underlying – buying as it rallies, and selling as it sells-off. However, the action of these market makers hedging their short options will not necessarily act to increase volatility in the underlying, as the counterparties they have sold options to may be counteracting this effect by themselves hedging their long volatility positions. However, the important difference between these two groups of hedgers is that the variance swap market-makers who are short options, **must hedge only on the close** to capture the close-close realised variance specified in the variance swap contract. In contrast, the hedgers who are long the options, will generally be free to choose when to delta-hedge, as they attempt to capture the true volatility of the underlying process.

Figure 108: Flows in the market as a result of a market-maker buying a variance swap and replicating it in the market
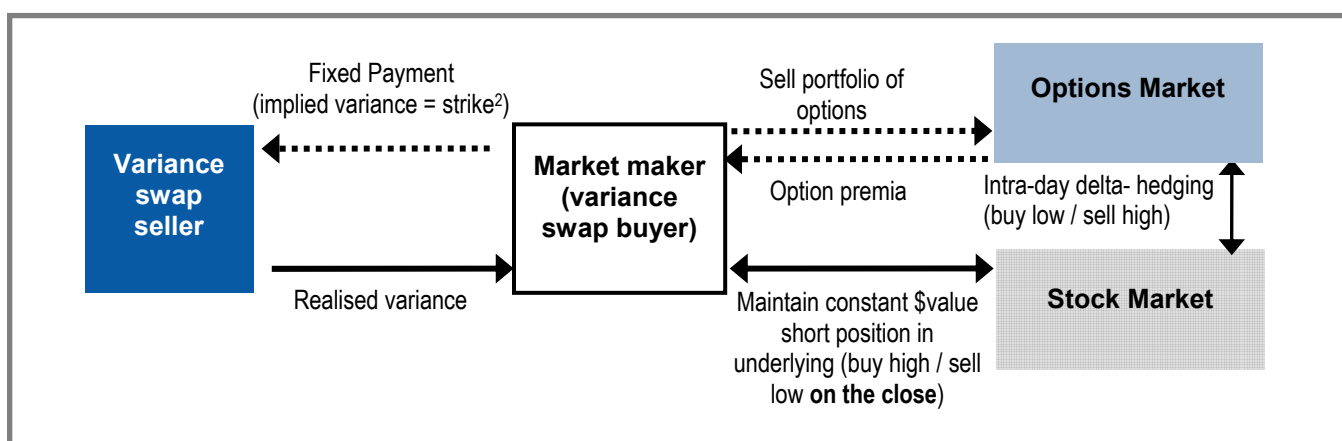
Therefore, the overall effect of hedging these variance swaps need not have the effect of increasing *overall* market volatility, although it may if the long options positions are not being hedged (e.g. they are sold on to end-investors). However, the important point is **that the hedging of long variance swap positions may act to increase close-to-close volatility**, with option hedges on the close having the potential effect of magnifying daily moves. In practice, this only becomes an issue if hedgers (variance swap market-makers) are the same way round and daily moves are large enough.

In fact the *expectation* of these delta-hedging flows could itself act to further amplify moves into the close through possible 'feedback' effects. For instance, on big down days, the market could anticipate a sell-off into the close as a result of long variance swap hedging and take short positions into the close to profit from this anticipated effect. This then drives the underlying market down further and could increases the amount that the variance swap hedgers have to sell on the close.

These feedback effects were much talked about during the volatile market period in May 2006, when the hedging into the close of long variance swap positions, and anticipation thereof, was blamed for the amplification of some already large daily moves. In fact, in an efficient market, such effects should be propagated back through the trading day, and may be offset by intra-day delta-hedging of opposing options.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

4.11: Why not volatility swaps?

A common complaint about variance swaps is that they pay-off based on realised variance (volatility squared) and not simply realised volatility. **Why are there not *volatility swaps*, products which payout linearly on the difference between implied and realised volatility?** If these are more easily understandable why are they not regularly traded?

Whilst volatility can be seen as more of an intuitive measure (being a standard deviation it is measured in the same units as the underlying), variance is in some sense more fundamental – especially because it is additive. For example, the variance of the sum of two independent distributions is just the sum of each of their variances (e.g. see discussion on variance swap mark-to-market, p15). This linearity of variance makes calculating variance swap mark-to-markets and forwards particularly simple.

As pointed out above, the exposure of delta-hedged options to volatility, after accounting for the gamma, is actually an exposure to the difference between implied and realised volatility *squared*. In this sense, **a variance swap mirrors a kind of ideal delta-hedged option whose gamma remains constant**. Furthermore, variance swaps are relatively easy to replicate. Once the replicating portfolio of options has been put in place, only delta-hedging is required. No further buying or selling of options is necessary.

All this explains why variance swaps are attractive instruments to trade, but still does not explain why volatility swaps are not also frequently traded. The main theoretical difficulty with volatility swaps is that they cannot be statically replicated through options. A replicating portfolio must dynamically trade options, and make relatively strong assumptions about the underlying volatility process – in particular about the volatility of volatility. This makes any replication process model dependent, and therefore much more prone to errors than the theoretically robust variance swap replication. See *Derman, 1999* for further details about construction volatility swaps.

In fact **the convention of quoting variance swap notionals in vega, rather than variance, amounts can be seen as an attempt to treat variance swaps like volatility swaps**. The vega notional represents the *average* p/l for a 1 vega change in volatility, but with the convexity meaning that longs will profit by more if volatility increases, and lose by less if volatility decreases. Thus for small changes in volatility, where the effect of the variance swap convexity is relatively limited, variance swaps (measured in vega notional) locally approximate volatility swaps. Volatility swaps can then be thought of as variance swaps without the convexity. The discount of the volatility swaps to variance swaps should therefore reflect the value of this convexity, which in turn is determined by the volatility of volatility (see Section 2.8). Seen in this light, the ability to calculate a fair price for volatility swaps, requires not only a price for the volatility (variance), but also a price for the volatility of volatility – i.e. the means of valuing options on volatility.

Since it is variance (volatility squared) which arises naturally from delta-hedging options, in terms of volatility products, variance should be thought of as the true underlying. In this framework, volatility swaps are naturally thought of as derivatives of variance – paying the square-root of the variance swap contract.

In fact we could dynamically trade a long variance swap (buying more as volatility decreases and selling as volatility increases) to hedge out this convexity bias. Similar to delta-hedging an option the p/l made from the resulting buy low – sell high strategy (for variance) will lead to a p/l based on the volatility of volatility: the larger this volatility-of-volatility, the bigger the discount of the dynamically replicated volatility swap to the variance swap. However, besides issues arising from the transaction costs of dynamically trading variance swaps, estimating this volatility-of-volatility contribution from the dynamic hedging of the variance swap is problematic and model-dependent, making the volatility swap contracts difficult to price.

5. Future Developments

- Variance swap indices
- Variance swap fixings
- Benchmark volatility return index
- Exchange-traded variance swaps
- 3rd generation products
- Derivatives on variance
- Variance swaps on other asset classes

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

Part 5: Future developments

With the increasing liquidity and breadth of the variance swap market, apart from the obvious benefits of reducing transaction costs and the ability to trade variance on more stock and index underlyings, other developments are underway which will in turn help both the liquidity and transparency of the variance swap market.

We highlight a few examples of recent developments, such as:

- Variance swap indices

- Variance swap fixings

- Benchmark volatility return indices

- Exchange traded variance swaps

- 3rd generation products

- Derivatives on variance

- Variance swaps on other asset classes.

Variance swap indices

The volatility indices associated with the major indices (VIX for S&P500, VXN for the Nasdaq, VSTOXX for Euro Stoxx 50, VDAX for DAX and VSMI for the SMI) are now calculated according to a variance swap style calculation. They represent the theoretical level of a rolling 1-month (30 calendar day) maturity variance swap, based on traded option prices. In fact, theoretical variance swap levels are first calculated for listed option maturities, and then the 30-day index level is interpolated. As such this gives the indices a real interpretation: **each index represents the risk-neutral expected variance of the underlying over the next month**. This also makes the indices directly comparable.

Table 14: Main variance swap based volatility index details

	Underlying	Reuters	Bloomberg	Datastream	Calculation details
VIX	S&P 500	.VIX	VIX Index	CBOEVIX	www.cboe.com
VXN	Nasdaq 100	.VXN	VXN Index	CBOEVXN	www.cboe.com
VDAX-NEW	Dax 30	.V1X	V1X Index	VDAXNEW	www.deutsche-boerse.com
VSTOXX	Euro STOXX 50	.V2TX	V2X Index	VSTOXXI	www.stoxx.com
VSMI	SMI	.VSMI	V3X Index	VSMIIDX	www.swx.com

Source: JPMorgan

All these volatility indices are well correlated with traded variance swap levels on their respective indices. However, the level of the volatility index is typically somewhat above the actual traded levels for variance swaps of the same maturity. There are two reasons for this:

- A convex skew should lead to mid-point discrete strike variance swap approximation (as in the VSTOXX index calculation) overestimating the level of implied variance.

- Given the typical bid/offer on equity options, the width of the arbitrage-corridor of trading a variance swap against a portfolio of options can give rise to a fairly large margin of error. The increasing demand to sell index variance and correlation over the last few years, as equity volatility has fallen, helped to cheapen index variance to its theoretical value, but not to the extent at which an arbitrage has been possible.

This 'basis' between the volatility index and actual traded (mid) variance swap levels is usually about 0.5-1 vegas. However this can vary much more widely during volatility spikes when supply and demand issues become important.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 5: Future Developments

These volatility indices, due to their visibility and wide availability, are often referred to by those outside the equity derivatives space as a 'market price of risk' index. While this is to some extent true, the short-dated nature of these variance swaps means the principal driver of the volatility index level is recent realised volatility (see Section 2.5). In reality, longer dated (e.g. 1-year) variance, spreads of implied to realised variance, skew levels or even ratios of put to call open-interest would perhaps be a better proxy for the level of risk-aversion present in the market.

Variance swap fixings

One of the biggest deficiencies of using the portfolio replication method to calculate a variance swap level is its lack of tradability. As noted above there exists a basis between the actual variance swap level and theoretical levels calculated from the portfolio of options. This basis exists probably due either to differences in market-makers' volatility skews or in the supply and demand from end users.

A way to capture a more tradable level for variance swaps is to construct a broker-dealer poll, in a similar way to that used in interest rates markets. For example, each day at 11.00am an interest rate fixing for different maturity deposits is made using a trimmed mean average of several banks' contributed rates. Most notably, the 3- and 6-month fixings are then often used as reference rates which form the basis of many longer dated instruments.

It could be envisaged that the 3-month variance swap could itself become a similar benchmark for volatility. Moreover, longer dated variance swaps could themselves be easily constructed from a strip of forward variance swaps. With structured products increasingly incorporating payoffs derived from variance swaps, having an independent method of valuing these instruments at realistic tradable levels will become more desirable to provide transparent mark-to-market valuations.

Benchmark volatility return index

With the availability of tradable and transparent variance swap levels for different maturities, transparent benchmark total return volatility indices can be constructed using a similar methodology to that described in Section 3.3. Though there will still need to be some development of the strategies involved (i.e. long/short, spot/forward, maturity etc), it will aid the development of volatility funds in countries where an appropriate benchmark is necessary. For example, see the variance component of JPMorgan's Yield Alpha product, as described in *JPMorgan's Yield Alpha Strategy Guide*.

Exchange-traded variance

Following the change in calculation method of the VIX to the variance swap calculation, the CBOE launched futures contracts on the VIX index in March 2004. This was followed in September 2005 by the launch of futures on the European VSTOXX, VDAX and VSMI indices.

Note that these futures contracts are *forward* volatility products, paying out on the value of the relevant volatility index at expiry, a value which is itself an implied variance. Thus the futures contract p/l is based on the change in expected future volatility, not on realised volatility or the spread of realised to implied volatility.

These futures do not expire on the normal index (futures) expiry dates, but 30 calendar days beforehand. For example a Nov-06 VXTOXX future expires on 15th November, 30 days before the 15th December options expiry. Expiry on 15th November is chosen because on that date, the Dec-06 listed options have exactly 30 calendar days remaining maturity and the VSTOXX calculation does not need to interpolate from any other maturities.

Note that there is a convexity bias inherent in marking to market these futures, since they are margined linearly in the level of forward starting volatility (*See VDAX-NEW, VSTOXX and VSMI Futures, JPMorgan September 2005*).

In April 2005, options on the VIX index were launched. These represented the first available exchange traded options on variance. As for the futures, these expire 30 days before an index expiry and are listed to expire 30 days before the

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 5: Future Developments

corresponding quarterly options expiry dates for the underlying. E.g. in November 2006, the listed VIX options are for Nov06, Dec 06, Feb07, May07 and Aug07 expiries.

More recently, beginning September 2006, Euronext announced the ability to list variance swaps on a limited number of indices through its BClear mechanism. This can be useful for end users who wish to take a volatility view with variance, but can not trade OTC products directly.

3rd generation volatility products

Generalised variance swaps or "3rd generation products" include **gamma swaps, corridor swaps** and **conditional variance swaps**. Gamma swaps are very similar to variance swaps but maintain a dollar gamma linear with spot, rather than a constant dollar gamma, leading to lower a skew exposure and a slightly lower strike compared to standard variance swaps. The original corridor variance swaps only accrue variance within a pre-specified range, meaning a long would make a maximal loss when the underlying fails to trade within the range.

Conditional variance swaps have been the most popular (and liquid) of these products, allowing investors to take exposure to realised variance contingent upon the underlying instrument trading within a pre-specified range. Outside this range no p/l accrues. Conditional variance swaps can be useful for hedging complex volatility exposures, taking a view on the volatility levels encapsulated in the skew or buying/selling variance at more attractive levels given a view on the underlying.

As with variance swaps, the sign of the p/l for a conditional variance swap is controlled by the difference between the volatility realised (in the range) over the lifetime of the swap and the pre-agreed fixed strike level. **In the case of the conditional, p/l is accrued** *only* **when the underlying is within a pre-specified range.** In addition, the magnitude of the final p/l is scaled by the proportion of time the underlying has spent in the pre-specified range. If the underlying trades entirely outside the range, the p/l will be zero, so the maximal loss for a long will be when the underlying trades within the range but with very low volatility.
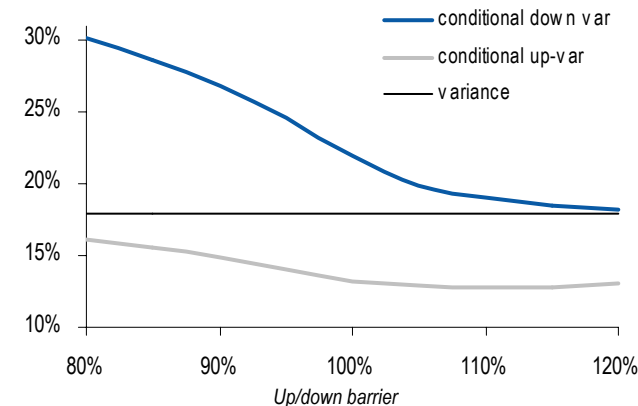
Whilst investors are free to specify the range associated with a conditional variance swap, the two principal types are up- and down- conditional variance swaps (**up-variance** and **down-variance**). Up-variance accrues realised volatility only when the underlying is above a pre-specified level (i.e. no upper barrier), while down-variance is accrued only when the underlying is below the specified barrier (i.e. no lower barrier). In the presence of a positive put-skew, down-variance will normally price above up-variance for close to ATM barrier levels.

Figure 109 : Payout from a conditional variance swap @ 20 strike
p/l, different lines for % days in the range
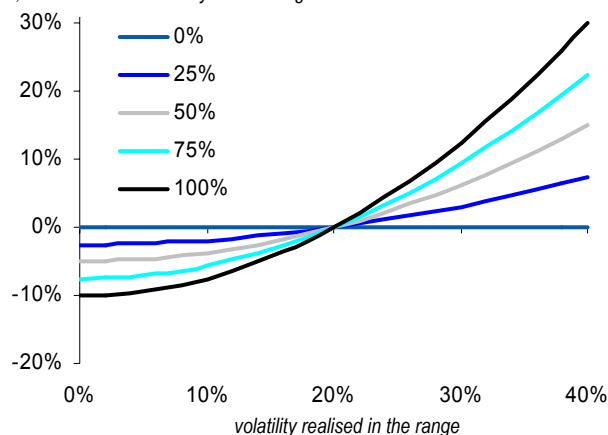
Source : JPMorgan

Figure 110 : Dec-07 conditional variance swaps levels as of Apr 06
Strike

Source : JPMorgan, data as of Apr 2006

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Part 5: Future Developments

Conditional variance swaps can be useful for expressing views on volatility contingent on market level. For example investors seeking crash protection may purchase conditional down-variance, which only becomes activated in the event of a market sell-off. That is, if the market stays above the down-barrier, p/l is zero. Conversely investors who believe volatility will be realised on the upside can buy conditional up-variance swaps, which are usually cheaper than a standard variance swap for which a significant amount of the premium is used to fund the downside skew exposure.

Conditional variance swaps have typically been traded on index underlyings. More recently, **single-stock conditionals** have gained liquidity, allowing investors to trade the volatility of a stock, contingent on stock price.

See *Conditional Variance Swaps, April 2006* for further details.

Derivatives on variance

Obviously with the increased liquidity and breadth of the variance swap market, the development of a liquid market in derivatives based on variance becomes more plausible. This is because the ability to use variance swaps themselves to delta-hedge variance derivatives becomes easier as liquidity increases and transaction costs decrease. For example, **options on variance** now trade in the market as investors look to tailor their volatility exposures.

However, as we have highlighted throughout the note, volatility is a process very different from the underlying instrument itself. For example volatility tends to be a mean-reverting process characterised by fairly long periods of subdued volatility with intermittent spikes of high volatility. Hence, using the simple delta-hedging replication option pricing schema that is basis of Black-Scholes becomes an impractical model to use in this case.

The key principle behind Black-Scholes, is that exposure to the underlying can be removed through delta-hedging, leaving behind the sole residual risk of exposure to the volatility of the underlying. If a non Black-Scholes approach is used another risk factor, such as the market-price-of-risk, is introduced in addition to implied volatility.

Nonetheless, this problem of modelling the underlying process has failed to halt the development of derivatives markets where the Black-Scholes approach is also considered invalid, for example: interest rate options, credit options, out-performance options etc. Derivatives on variance look to be no different.

Variance swaps on other asset classes

While variance swap liquidity has increased massively on equity underlyings over the last few years, variance swaps have yet to really take hold in other asset classes. Nonetheless, variance swaps do trade on assets such as commodities, FX and interest rates.

One issue for variance swaps on other asset classes is the dependence of the payoff on movements in the price of a spot underlying. While this is appropriate for equities, it can be troublesome for commodities which tend to be based upon a rolling forward price. In fact, commodity variance swaps are usually referenced to a futures contract of the same maturity as the swap.

Moreover in fixed income volatility based on yields is often more useful. This would add additional complexity in constructing a replicating portfolio, which would have to take into account the convexity of bond prices to yields. Furthermore, bonds have a defined maturity, unlike equities, which leads to problems defining a meaningful payoff for longer dated derivative products.

Nonetheless, we expect variance swap markets to continue to develop on other assets, particularly commodities and FX, driven by both investor demand to trade pure volatility, and at least in some cases the desire for investment banks to recycle exotic and structured product risk.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

References

General References

More Than You Ever Wanted to Know About Volatility Swaps, Demeterfi, Derman, Kamal & Zou, 1999

Options, Futures and Other Derivatives, John C. Hull, Fifth Edition, 2003

The Volatility Surface, Jim Gatheral, 2006

Towards a theory of Volatility Trading, Peter Carr & Dilip Madan - 2002

JPMorgan Publications

Calling All Overwriters, August 2002

Conditional Variance Swaps, April 2006

Correlation Vehicles, May 2005

European Equity Derivatives Weekly Outlook , various publications

Introducing JPMorgan's Cross-Section Volatility Report, November 2006

Just What You Need to Know About Variance Swaps, February 2005

Synthetic Short CDS with Positive Carry, October 2006

VDAX-NEW, VSTOXX and VSMI Futures, September 2005

Yield Alpha Strategy Guide, July 2006

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

JPMorgan

Risks of common option strategies

Please also refer to option risk disclosure documents

Put Sale. Investors who sell put options will own the underlying stock if the stock price falls below the strike price of the put option. Investors, therefore, will be exposed to any decline in the stock price below the strike potentially to zero, and they will not participate in any stock appreciation if the option expires unexercised.

Call Sale. Investors who sell uncovered call options have exposure on the upside that is theoretically unlimited.

Call Overwrite or Buywrite. Investors who sell call options against a long position in the underlying stock give up any appreciation in the stock price above the strike price of the call option, and they remain exposed to the downside of the underlying stock in the return for the receipt of the option premium.

Booster. In a sell-off, the maximum realised downside potential of a double-up booster is the net premium paid. In a rally, option losses are potentially unlimited as the investor is net short a call. When overlayed onto a long stock position, upside losses are capped (as for a covered call), but downside losses are not.

Collar. Locks in the amount that can be realized at maturity to a range defined by the put and call strike. If the collar is not costless, investors risk losing 100% of the premium paid. Since investors are selling a call option, they give up any stock appreciation above the strike price of the call option.

Call Purchase. Options are a decaying asset, and investors risk losing 100% of the premium paid if the stock is below the strike price of the call option.

Put Purchase. Options are a decaying asset, and investors risk losing 100% of the premium paid if the stock is above the strike price of the put option.

Straddle or Strangle. The seller of a straddle or strangle is exposed to stock increases above the call strike and stock price declines below the put strike. Since exposure on the upside is theoretically unlimited, investors who also own the stock would have limited losses should the stock rally. Covered writers are exposed to declines in the long stock position as well as any additional shares put to them should the stock decline below the strike price of the put option. Having sold a covered call option, the investor gives up all appreciation in the stock above the strike price of the call option.

Put Spread. The buyer of a put spread risks losing 100% of the premium paid. The buyer of higher ratio put spread has unlimited downside below the lower strike (down to zero), dependent on the number of lower struck puts sold. The maximum gain is limited to the spread between the two put strikes, when the underlying is at the lower strike. Investors who own the underlying stock will have downside protection between the higher strike put and the lower strike put. However, should the stock price fall below the strike price of the lower strike put, investors regain exposure to the underlying stock, and this exposure is multiplied by the number of puts sold.

Call Spread. The buyer risks losing 100% of the premium paid. The gain is limited to the spread between the two strike prices. The seller of a call spread risks losing an amount equal to the spread between the two call strikes less the net premium received. By selling a covered call spread, the investor remains exposed to the downside of the stock and gives up the spread between the two call strikes should the stock rally.

Butterfly Spread. A butterfly spread consists of two spreads established simultaneously. One a bull spread and the other a bear spread. The resulting position is neutral, that is, the investor will profit if the underlying is stable. Butterfly spreads are established at a net debit. The maximum profit will occur at the middle strike price, the maximum loss is the net debit.

Pricing Is Illustrative Only:
Prices quoted in the above trade ideas are our estimate of current market levels, and are not indicative trading levels

Risks to Strategies:
Not all option strategies are suitable for investors; certain strategies may expose investors to significant potential losses. We have summarized the risks of selected derivative strategies. For additional risk information, please call your sales representative for a copy of "Characteristics and Risks of Standardized Options". We advise investors to consult their tax advisors and legal counsel about the tax implications of these strategies.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

European Equity Derivatives Strategy
17 November 2006

JPMorgan

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

JPMorgan

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

JPMorgan

Analyst Certification:

The research analyst(s) denoted by an "AC" on the cover of this report certifies (or, where multiple research analysts are primarily responsible for this report, the research analyst denoted by an "AC" on the cover or within the document individually certifies, with respect to each security or issuer that the research analyst covers in this research) that: (1) all of the views expressed in this report accurately reflect his or her personal views about any and all of the subject securities or issuers; and (2) no part of any of the research analyst's compensation was, is, or will be directly or indirectly related to the specific recommendations or views expressed by the research analyst(s) in this report.

Important Disclosures

Explanation of Equity Research Ratings and Analyst(s) Coverage Universe:

JPMorgan uses the following rating system: **Overweight** [Over the next six to twelve months, we expect this stock will outperform the average total return of the stocks in the analyst's (or the analyst's team's) coverage universe.] **Neutral** [Over the next six to twelve months, we expect this stock will perform in line with the average total return of the stocks in the analyst's (or the analyst's team's) coverage universe.] **Underweight** [Over the next six to twelve months, we expect this stock will underperform the average total return of the stocks in the analyst's (or the analyst's team's) coverage universe.] The analyst or analyst's team's coverage universe is the sector and/or country shown on the cover of each publication. See below for the specific stocks in the certifying analyst(s) coverage universe.

JPMorgan Equity Research Ratings Distribution, as of September 29, 2006

	Overweight (buy)	Neutral (hold)	Underweight (sell)
JPM Global Equity Research Coverage	42%	41%	17%
IB clients*	44%	46%	35%
JPMSI Equity Research Coverage	38%	48%	15%
IB clients*	62%	55%	47%

*Percentage of investment banking clients in each rating category.
For purposes only of NASD/NYSE ratings distribution rules, our Overweight rating falls into a buy rating category; our Neutral rating falls into a hold rating category; and our Underweight rating falls into a sell rating category.

Valuation and Risks: Please see the most recent JPMorgan research report for an analysis of valuation methodology and risks on any securities recommended herein. Research is available at http://www.morganmarkets.com , or you can contact the analyst named on the front of this note or your JPMorgan representative.

Analysts' Compensation: The equity research analysts responsible for the preparation of this report receive compensation based upon various factors, including the quality and accuracy of research, client feedback, competitive factors, and overall firm revenues, which include revenues from, among other business units, Institutional Equities and Investment Banking.

Other Disclosures

Options related research: If the information contained herein regards options related research, such information is available only to persons who have received the proper option risk disclosure documents. For a copy of the Option Clearing Corporation's Characteristics and Risks of Standardized Options, please contact your JPMorgan Representative or visit the OCC's website at http://www.optionsclearing.com/publications/risks/riskstoc.pdf.

Legal Entities Disclosures
U.S.: JPMSI is a member of NYSE, NASD and SIPC. J.P. Morgan Futures Inc. is a member of the NFA. J.P. Morgan Chase Bank, N.A. is a member of FDIC and is authorized and regulated in the UK by the Financial Services Authority. **U.K.:** J.P. Morgan Securities Ltd. (JPMSL) is a member of the London Stock Exchange and is authorised and regulated by the Financial Services Authority. **South Africa:** J.P. Morgan Equities Limited is a member of the Johannesburg Securities Exchange and is regulated by the FSB. **Hong Kong:** J.P. Morgan Securities (Asia Pacific) Limited (CE number AAJ321) is regulated by the Hong Kong Monetary Authority and the Securities and Futures Commission in Hong Kong. **Korea:** J.P. Morgan Securities (Far East) Ltd, Seoul branch, is regulated by the Korea Financial Supervisory Service. **Australia:** J.P. Morgan Australia Limited (ABN 52 002 888 011/AFS Licence No: 238188, regulated by ASIC) and J.P. Morgan Securities Australia Limited (ABN 61 003 245 234/AFS Licence No: 238066, a Market Participant with the ASX) (JPMSAL) are licensed securities dealers. **Taiwan:** J.P.Morgan Securities (Taiwan) Limited is a participant of the Taiwan Stock Exchange (company-type) and regulated by the Taiwan Securities and Futures Commission. **India:** J.P. Morgan India Private Limited is a member of the National Stock Exchange of India Limited and The Stock Exchange, Mumbai and is regulated by the Securities and Exchange Board of India. **Thailand:** JPMorgan Securities (Thailand) Limited is a member of the Stock Exchange of Thailand and is regulated by the Ministry of Finance and the Securities and Exchange Commission. **Indonesia:** PT J.P. Morgan Securities Indonesia is a member of the Jakarta Stock Exchange and Surabaya Stock Exchange and is regulated by the BAPEPAM.

Peter Allen
(44-20) 7325-4114
Stephen Einchcomb
(44-20) 7325-9064
Nicolas Granger
(44-20) 7325-7033

JPMorgan

Philippines: This report is distributed in the Philippines by J.P. Morgan Securities Philippines, Inc. **Brazil:** Banco J.P. Morgan S.A. is regulated by the Comissao de Valores Mobiliarios (CVM) and by the Central Bank of Brazil. **Japan:** This material is distributed in Japan by JPMorgan Securities Japan Co., Ltd., which is regulated by the Japan Financial Services Agency (FSA). **Singapore:** This material is issued and distributed in Singapore by J.P. Morgan Securities Singapore Private Limited (JPMSS) [mica (p) 069/09/2006 and Co. Reg. No.: 199405335R] which is a member of the Singapore Exchange Securities Trading Limited and is regulated by the Monetary Authority of Singapore (MAS) and/or JPMorgan Chase Bank, N.A., Singapore branch (JPMCB Singapore) which is regulated by the MAS. **Malaysia:** This material is issued and distributed in Malaysia by JPMorgan Securities (Malaysia) Sdn Bhd (18146-x) (formerly known as J.P. Morgan Malaysia Sdn Bhd) which is a Participating Organization of Bursa Malaysia Securities Bhd and is licensed as a dealer by the Securities Commission in Malaysia

Country and Region Specific Disclosures
U.K. and European Economic Area (EEA): Issued and approved for distribution in the U.K. and the EEA by JPMSL. Investment research issued by JPMSL has been prepared in accordance with JPMSL's Policies for Managing Conflicts of Interest in Connection with Investment Research which can be found at http://www.jpmorgan.com/pdfdoc/research/ConflictManagementPolicy.pdf. This report has been issued in the U.K. only to persons of a kind described in Article 19 (5), 38, 47 and 49 of the Financial Services and Markets Act 2000 (Financial Promotion) Order 2001 (all such persons being referred to as "relevant persons"). This document must not be acted on or relied on by persons who are not relevant persons. Any investment or investment activity to which this document relates is only available to relevant persons and will be engaged in only with relevant persons. In other EEA countries, the report has been issued to persons regarded as professional investors (or equivalent) in their home jurisdiction **Germany:** This material is distributed in Germany by J.P. Morgan Securities Ltd. Frankfurt Branch and JPMorgan Chase Bank, N.A., Frankfurt Branch who are regulated by the Bundesanstalt für Finanzdienstleistungsaufsicht. **Australia:** This material is issued and distributed by JPMSAL in Australia to "wholesale clients" only. JPMSAL does not issue or distribute this material to "retail clients." The recipient of this material must not distribute it to any third party or outside Australia without the prior written consent of JPMSAL. For the purposes of this paragraph the terms "wholesale client" and "retail client" have the meanings given to them in section 761G of the Corporations Act 2001. **Hong Kong:** The 1% ownership disclosure as of the previous month end satisfies the requirements under Paragraph 16.5(a) of the Hong Kong Code of Conduct for persons licensed by or registered with the Securities and Futures Commission. (For research published within the first ten days of the month, the disclosure may be based on the month end data from two months' prior.) J.P. Morgan Broking (Hong Kong) Limited is the liquidity provider for derivative warrants issued by J.P. Morgan International Derivatives Ltd and listed on The Stock Exchange of Hong Kong Limited. An updated list can be found on HKEx website: http://www.hkex.com.hk/prod/dw/Lp.htm. **Korea:** This report may have been edited or contributed to from time to time by affiliates of J.P. Morgan Securities (Far East) Ltd, Seoul branch. **Singapore:** JPMSI and/or its affiliates may have a holding in any of the securities discussed in this report; for securities where the holding is 1% or greater, the specific holding is disclosed in the Legal Disclosures section above. **India:** For private circulation only not for sale. **New Zealand:** This material is issued and distributed by JPMSAL in New Zealand only to persons whose principal business is the investment of money or who, in the course of and for the purposes of their business, habitually invest money. JPMSAL does not issue or distribute this material to members of "the public" as determined in accordance with section 3 of the Securities Act 1978. The recipient of this material must not distribute it to any third party or outside New Zealand without the prior written consent of JPMSAL.

General: Additional information is available upon request. Information has been obtained from sources believed to be reliable but JPMorgan Chase & Co. or its affiliates and/or subsidiaries (collectively JPMorgan) do not warrant its completeness or accuracy except with respect to any disclosures relative to JPMSI and/or its affiliates and the analyst's involvement with the issuer that is the subject of the research. All pricing is as of the close of market for the securities discussed, unless otherwise stated. Opinions and estimates constitute our judgment as of the date of this material and are subject to change without notice. Past performance is not indicative of future results. This material is not intended as an offer or solicitation for the purchase or sale of any financial instrument. The opinions and recommendations herein do not take into account individual client circumstances, objectives, or needs and are not intended as recommendations of particular securities, financial instruments or strategies to particular clients. The recipient of this report must make its own independent decisions regarding any securities or financial instruments mentioned herein. JPMSI distributes in the U.S. research published by non-U.S. affiliates and accepts responsibility for its contents. Periodic updates may be provided on companies/industries based on company specific developments or announcements, market conditions or any other publicly available information. Clients should contact analysts and execute transactions through a JPMorgan subsidiary or affiliate in their home jurisdiction unless governing law permits otherwise.

Revised September 30, 2006.

Investment Strategies Series

This series aims to offer new approaches and methods on investing and trading profitably in financial markets.

1. *Rock-Bottom Spreads*, Peter Rappoport, Oct 2001

2. *Understanding and Trading Swap Spreads*, Laurent Fransolet, Marius Langeland, Pavan Wadhwa, Gagan Singh, Dec 2001

3. *New LCPI trading rules: Introducing FX CACI*, Larry Kantor, Mustafa Caglayan, Dec 2001

4. *FX Positioning with JPMorgan's Exchange Rate Model*, Drausio Giacomelli, Canlin Li, Jan 2002

5. *Profiting from Market Signals*, John Normand, Mar 2002

6. *A Framework for Long-term Currency Valuation,* Larry Kantor and Drausio Giacomelli*, Apr 2002

7. *Using Equities to Trade FX: Introducing LCVI,* Larry Kantor and Mustafa Caglayan, Oct 2002

8. *Alternative LCVI Trading Strategies,* Mustafa Caglayan, Jan 2003

9. *Which Trade,* John Normand, Jan 2004

10. *JPMorgan's FX & Commodity Barometer*, John Normand, Mustafa Caglayan, Daniel Ko, Nikolaos Panigirtzoglou and Lei Shen, Sep 2004.

11. *A Fair Value Model for US Bonds, Credit and Equities*, Nikolaos Panigirtzoglou and Jan Loeys, Jan 2005.

12. *JPMorgan Emerging Market Carry-to-RiskModel*, Osman Wahid, February 2005

13. *Valuing cross-market yield spreads*, Nikolaos Panigirtzoglou, January 2006

14. *Exploiting cross-market momentum*, Ruy Ribeiro and Jan Loeys, February 2006

15. *A cross-market bond carry strategy*, Nikolaos Panigirtzoglou, March 2006

16. *Bonds, Bubbles and Black Holes*, George Cooper, March 2006

17. *JPMorgan FX Hedging Framework*, Rebecca Patterson and Nandita Singh, March 2006

18. *Index Linked Gilts Uncovered*, Jorge Garayo and Francis Diamond, March 2006

19. *Trading Credit Curves I*, Jonny Goulden, March 2006

20. *Trading Credit Curves II*, Jonny Goulden, March 2006

21. *Yield Rotator*, Nikolaos Panigirtzoglou, May 2006

22. *Relative Value on Curve vs Butterfly Trades*, Stefano Di Domizio, June 2006

23. *Hedging Inflation with Real Assets,* John Normand, July 2006

24. *Trading Credit Volatility*, Saul Doctor and Alex Sbityokov, August 2006

25. *Momentum in Commodities*, Ruy Ribeiro, Jan Loeys and John Normand, September 2006

26. *Equity Style Rotation*, Ruy Ribeiro, November 2006

27. *Euro Fixed Income Momentum Strategy*, Gianluca Salford, George Cooper, November 2006