Variance swap volatility dispersion

Izzy Nelken

Super Computer Consulting, Inc., 3943 Bordeaux Drive, Northbrook, IL 60062, USA
Tel: +1 847 562 0600, e-mail: izzy@supercc.com, websites: www.supercc.com and
www.optionsprofessor.com

Received: 24th March, 2005

Izzy Nelken is President of Super Computer Consulting, Inc which specialises in complex derivatives, structured products, risk management and hedge funds. He holds a PhD in computer science from Rutgers University and now lectures in the Department of Mathematics at the University of Chicago. Super Computer Consulting clients include several regulatory bodies, major broker-dealers, large and medium-sized banks as well as hedge funds. Dr Nelken has published a number of books, including 'Volatility in the Capital Markets'.

Practical applications

Volatility dispersion strategies involve selling volatility on the index and buying volatility on the components, traditionally using at the money (ATM) straddles. A problem with this approach, however, is that as the original ATM options move out-of-the-money, they lose their vega exposure. To correct this problem, a new style of trading has emerged the 'variance swap' approach, which eliminates the constant rebalancing requirement associated with ATM options. Moreover, under variance swap methodology, the volatility difference appears more tame.

Abstract

Several trading institutions are actively engaged in 'volatility dispersion' strategies. These involve selling volatility on the index and buying volatility on the components. This trade was traditionally done using at the money (ATM) straddles. An important practical problem with this approach is that market prices move and cause the original ATM options to become out of the money (OTM) and lose their vega exposure. Even if volatility moved as expected by the trader, the profit potential of the trade would be greatly diminished as the options lost their vega. To correct this problem, a new style of trading has emerged in which some practitioners are

trading this strategy using a 'variance swap' approach. This has the advantage that both legs of the trade have relatively constant vega exposure, regardless of stock market movements.

INTRODUCTION

This paper reviews the volatility dispersion trade and compares the two approaches. It gives a heuristic motivation to the variance swap style of the trade. It also provides some empirical evidence that seems to indicate that the variance swap approach is more malleable to trading.

Derivatives Use,
Trading & Regulation,
Vol. 11 No. 4, 2006,
pp. 334-344
© Palgrave Macmillan
Ltd
1747-4426/06 \$30.00

STYLES OF PROPRIETARY TRADING

It is useful to distinguish between two 'extreme' types of trading.

- (1) Simple relationships: these seek to exploit slight price differences between identical products on different exchanges or products whose prices are tied to one another by an exact, pre-defined formula. These trades typically have very low risk and low expected profits. They rely on extremely efficient execution. Owing to the proliferation of fast electronic trading tools in the last several years, the opportunities for these trades have been rapidly diminishing. Examples are:
 - (a) The existence of several different option exchanges in the USA, which, at times, allows one to find zero cost 'butterflies' or other combinations of options that have a small chance of a positive payoff;
 - (b) Trading an index (eg the NDX) against its component stocks.
- (2) Historical relationships — these types of trades rely on historical price correlations that have been observed in the past but have no inherent reason. There is no formula that dictates the relationship between the products. Such trading is quite risky as 'past performance is no guarantee of future results'. Examples of this are:
 - (a) Long short hedge funds that purchase one share and sell another based on historical, observed price relationships.
 - (b) All kinds of technical trading systems. These have worked well in

historical back tests but are not guaranteed to work in the future.

Somewhere between these two extreme types of trading lies the field of statistical trading that is based on a formula. The formula is not exact and depends on several unobservable variables. This type of trading can be characterised by medium risk and medium returns. There is anecdotal evidence from the markets that this type of trading is becoming more and more commonplace. Examples of this type of trading include:

- Convertible bond 'arbitrage' Here, the trader will purchase a convertible bond and short shares against it. Obviously, the price of the bond depends on the price of the underlying share. This is not an exact formula however, since it also depends on many unobservable parameters: implied volatility, credit spreads, recovery values etc.
- Volatility dispersion.

WHAT IS VOLATILITY DISPERSION?

Assume that one sells index options. Obviously, one can hedge the exposure to the index price by trading the underlying index ('dynamic delta hedging'). But one is still exposed to the risk that the index implied volatility will drop and the price of the sold options will also decline. One method of hedging the exposure to implied volatility is to purchase options on the shares that compose the same index.

Obviously, there is a strong relationship

between the implied volatility of the index and the implied volatility of the components. This is not an exact formula however, as it depends on an implied correlation matrix which can only be approximated.

Most dispersion strategies sell index volatility and purchase component volatility, but not the reverse. That is, very rarely do people sell individual stock options and purchase the index options. This is for two main reasons:

- (1) It is well known that index options have high implied volatility numbers by comparison with their historical volatility. Schneeweis and Spurgin¹ have observed that the volatility implied by index option prices is too high relative to the realised volatility of the index. Bollen and Whaley² argue that there is excess buying pressure on S&P500 Index put options (as fund managers seek to purchase insurance against a decline in the stock market).
- (2) One can be surprised by an 'event' in a single stock. A sudden takeover or bankruptcy can cause an individual stock price to react very strongly and very quickly. Selling single stock options could be quite dangerous.

Market makers in index options who sell index options when the premiums are high may choose to engage in volatility dispersion strategies to reduce their volatility risk. Many hedge fund managers also employ these strategies on an opportunistic basis.

THE TWO STYLES

Assume that a trader identified that the implied volatility of a particular index is high relative to its components. Traditionally, one would trade this by selling at the money (ATM) index options (for example an ATM straddle) and purchasing ATM component options. It is common to trade in ATM options since:

- (1) The ATM options have the largest vega exposure.
- (2) If one trades in out of the money (OTM) options, they typically establish a position with a delta exposure. The trader will be obliged to delta hedge that exposure from the initial trade.
- (3) If one were to sell OTM index options, one would have to decide which are the corresponding OTM component options. There are many potential choices,
 - (a) options which are OTM by a similar percentage;
 - (b) options which are OTM by a similar number of standard deviations;
 - (c) options which have similar deltas;
 - (d) and more.
- (4) The ATM options are typically the most liquid.

Assume that the trader might decide to sell an ATM index straddle and purchase ATM component straddles. Of course, for each stock in the index, the trader will have to determine how many component straddles he/she will need to sell. If the index is cap weighted and has many components, the trader might choose not

to trade options on the smaller cap stocks at all.

The trader hopes to profit when implied volatilities come back into line. Perhaps the index implied volatility declines or the implied volatility of the components increases. In any case, the trade will be a winner. Another source of potential profits is that each of the component stocks might experience a price ‘jump’. A jump may occur owing to a takeover, bankruptcy or other material change in the company. In the case of a jump, one of the options in the component straddle suddenly becomes deep in the money and rises in value.

But what if stock prices change slowly? In this case, the straddle that was originally purchased at the money, with close to zero delta, now has a delta exposure. If the traders did not hedge the delta exposure, random changes in stock price would overwhelm the profit potential (which is due to volatility) of the trade. Therefore, traders will typically delta hedge the stock exposure of each component and the index exposure of the index straddle. They do so because they trade attempts to profit from volatility changes.

Nevertheless, a serious problem remains. Once stock prices have moved, even if implied volatility were to return to normal levels, the trade will still not be profitable, as the options have become OTM and have no vega exposure. Some traders resort to exiting OTM straddle positions and entering new ATM straddles — re they keep rebalancing (or ‘chasing the stock price’) to maintain a vega exposure.

What traders require is an instrument that maintains a vega exposure for a wide

range of stock prices. This can be created using the ‘variance swap’ technique.

Nowadays, some practitioners are trading volatility dispersion using variance swap volatility.

In what follows, the volatility dispersion trade is described in more detail and the two styles of trade are explained. The paper concludes with some experimental evidence that seems to suggest that the variance swap volatility is more amenable to trading.

THE PORTFOLIO EQUATION

The well-known Markowitz

Mean–Variance Portfolio Equation expresses the volatility of a portfolio as a function of the volatility of its component stocks.³

$$\text{Var}_p = \sum_{i=1}^n \sigma_i^2 w_i^2 + \sum_{i=1}^n \sum_{j=1}^n \rho_{i,j} \sigma_i \sigma_j w_i w_j$$

$$\sigma_p = \sqrt{\text{Var}_p}$$

Here, a portfolio of n assets with a variance of Var_p and a volatility of σ_p is considered. For asset i , σ_i is the volatility of the i th asset and w_i is its weight in the portfolio. The quantity $\rho_{i,j}$ is the pairwise correlation between assets i and j .

For historical volatility and correlation numbers, this equation is exact. Many traders combine historical correlation estimates with implied volatility numbers. The resulting theoretical volatility for the index is then compared with the actual implied volatility.

A hypothetical example: The two stock index

Assume that one had an index composed of two stocks. One could observe the implied

volatility of options on the index and on the two component stocks. Plugging these numbers into the portfolio equation and solving for the correlation would give the ‘implied correlation number’ between the two stocks. If the implied correlation number is quite high, it may be a good time for a trade. Certainly, if the implied correlation number is close to 1, there is little risk that the correlation will climb even further.

Even for a two stock index though, the trade requires some assumptions.

Suppose one sells 100 ATM straddles on the index. One must determine how many ATM straddles to purchase on each of the components.

The choice of quantities and strikes is typically accomplished by the technique of equating ‘Greeks’.

- Deltas are hedged separately in each instrument.
- The trader then calculates how many gamma or vega units he/she has sold and then purchases an equal amount of units in the components.

The potential profit of this trade comes from several sources:

- If the implied correlation comes back to normal level, the trade will make money. This will happen either because the implied volatility of the index option has decreased, or because the implied volatility of the components has increased, or both.
- A ‘windfall’ gain. This happens when a sudden idiosyncratic event causes one of

the component stocks to move sharply (either up or down).

The trade is exposed to a practical danger. If the stocks and index move from the previous strikes, even if the correlation returns to normal levels, the option positions will not be exposed to vega in any significant way. The change in implied volatility will not have a substantial impact on the price of the options. This may mean that the trader will continuously have to adjust the straddle positions to remain close to ATM. It is precisely because of this difficulty that some market participants use the variance swap volatility for these trades.

Another danger is that of correlation ‘explosion’. This trade is a bet that implied correlation will come back to its normal levels. However, if correlation continues to increase, the trade may, in fact, lose money.

An additional difficulty occurs in some capital weighted (‘cap weighted’) indexes. As the market moves and stock prices change, so do the w_i s, the relative weights of the companies within the index.

An index with many stocks

Real-life indexes have more than two stocks. In that case, there is an entire correlation matrix. This matrix cannot be uniquely determined from the implied volatility numbers.

It is possible to use a historical correlation matrix in the portfolio equation. Of course, one has to choose carefully the

historical period for which correlation should be computed, the averaging method used etc. See Nelken⁴ for a discussion on volatility and correlation measurement. In any case, it is well known that the resulting correlation matrix must be symmetrical and positive semi-definite.

Using the historical correlation matrix in conjunction with the implied volatility numbers will cause the portfolio equation to be inexact. In practice, however, it is still usable as a signal of when to enter or exit a trade.

As it is impossible to determine the implied correlation matrix precisely, many traders attempt to estimate it.

- Begin with a historical correlation matrix.
- One now has two volatility numbers:
 - (a) the implied volatility of the index ('index volatility');
 - (b) the volatility of the index as computed using the implied volatility of the components and the historical correlation matrix ('stock volatility').
- These volatility numbers will not match. The market has its own view on future correlations (the so-called 'implied correlation' matrix), and it adjusts the historical correlation matrix in some undeterminable fashion. For example, when market participants are worried about a crash, the historical correlation matrix is adjusted upwards.
- Traders will typically adjust the historical correlation matrix until these two volatility numbers match. The resultant matrix may lose its features (symmetric,

positive semi-definiteness). It is possible to use techniques such as described in Higham⁵ to find the 'nearest' correlation matrix to the resultant matrix.

- The trader's goal is to estimate the 'implied correlation' matrix, that is, to find a symmetric, positive semi-definite matrix with unit diagonal that is as close as possible to the historical correlation matrix and makes the 'stock volatility' match the 'index volatility'.
- If the implied correlation matrix has large elements by comparison with the historical correlation matrix, it may be a signal to enter the trade and bet on declining correlations.

WHICH IMPLIED VOLATILITY?

If the trader uses ATM straddles to perform the trade, the market may 'run away' from the strikes. In that case, even if volatilities move in the expected directions, the trade will not make money, as the options that have become OTM have very little volatility exposure.

Some market participants are therefore using the variance swap volatility to create their volatility dispersion strategies.

VARIANCE SWAP VOLATILITY: A SHORT HISTORY

In 1996, Neuberger⁶ described the log contract. This is an option whose payout is tied to the logarithm of the stock price. Thus a call option would have a pay out of

$$\max[\ln(S) - X, 0]$$

where S is the stock price on expiration, X is the strike and $\ln$ is the natural logarithm function. The log contract has a unique feature in that it has stable ‘Greeks’. For example, a contour plot of the vega of such a contract versus the stock price results in a straight line. By contrast, a similar contour plot for vanilla European options results in a non-linear curve. The main conclusion was that ‘the log contract provides a much easier and more reliable way of betting on volatility’.

In 1998, in the aftermath of the Long Term Capital Management (LTCM) debacle, implied volatility figures rose to unprecedented levels. As Gatheral⁷ explains

‘Variance swaps took off as a product in the aftermath of the LTCM meltdown in late 1998 when implied stock index volatility levels rose to unprecedented levels. Hedge funds took advantage of this by paying variance in swaps (selling the realised volatility at high implied levels). The key to their willingness to pay on a variance swap rather than sell options was that a variance swap is a pure play on realised volatility — no labour — intensive delta hedging or other path dependency is involved. Dealers were happy to buy vega at these high levels because they were structurally short vega (in the aggregate) through sales of guaranteed equity-linked investments to retail investors and were getting badly hurt by high implied volatility levels.’

In 1999, Demeterfi *et al.*⁸ constructed a portfolio of European options. Their portfolio has $\Delta k/k^2$ units of each option

with a strike of k . It turns out that such a portfolio has a constant ‘variance vega’. The dollar amount gained by such a portfolio for a unit change in the variance (volatility squared) is insensitive to the stock price across a wide range of stock prices. The paper also described how this portfolio is almost equivalent to a log contract.

In 2003, the Chicago Board Options Exchange (CBOE)⁹ announced a major change in the way that the volatility index (VIX) was computed. This was followed by the initiation of trading on the VIX in 2004. The new VIX is fashioned after the variance swap portfolio previously introduced by Demeterfi *et al.*⁸

The generalised formula used by the CBOE is

$$\sigma^2 = \frac{2}{T} \sum_i \frac{\Delta K}{K_i^2} e^{RT} Q(K_i) - \frac{1}{T} \left[\frac{F}{K_0} - 1 \right]^2$$

where T is the time to expiration; F is the forward level of the stock (or index) as derived from the options; K_i is the strike price of the i th OTM option, a call if $K > F$ and a put if $K < F$; ΔK_i is the interval between strike prices (for single equities it is typically \$5); K_0 is the first strike price below the forward price F ; R is the risk-free interest rate; and $Q(K_i)$ is the quote for the specific option (it is taken as the mid price).

One of the main advantages of the way the new VIX is calculated is that it uses an entire collection of options to come at the index rather than just a few options as previously. Also, contrary to the normal implied volatility which relies on one’s own choice of model (eg Black-Scholes for European options), the VIX volatility does not rely on a particular choice of model.

Figure 1: Value of Dow Jones Index (DJI), index volatility and stock volatility March 2003–March 2004

Figure 2: Value of Dow Jones Index and difference between index volatility and stock volatility, March 2003–March 2004

SHOULD ONE TRADE VARIANCE SWAP VOLATILITY DISPERSION?

This paper looks at the Dow Jones Index. The DJI is composed of 30 well-known companies whose stocks are liquid. In addition, it is a price-weighted index.

Figure 1 plots the index level on the left-hand scale. The right-hand scale plots the ATM 30-day implied volatility of the

DJI as given by the index options. To get to a 30-day running volatility, it is typical to use interpolation between the front month option and the second to front month option. Also plotted is the volatility which was obtained using the portfolio formula, the implied volatility of the component options and a long-term historical correlation matrix. In 2003, the

Figure 3: Value of Dow Jones Index (DJI), index volatility and stock volatility (from stock options using the variance swap methodology), March 2003–March 2004

Figure 4: Value of Dow Jones Index (DJI), and difference between index volatility and stock volatility (computed from stock options using variance swap methodology) March 2003–March 2004

index rose sharply and implied volatility numbers fell.

Figure 2 plots the ATM volatility difference. This is the index implied volatility minus the index volatility as computed from the component (stock) options.

The index implied volatility ('index

volatility') is different from the volatility computed from the portfolio formula ('stock volatility'). That formula used historical correlation numbers. As mentioned before, when market participants are worried about a crash, the historical correlation matrix is adjusted upwards. As the stock market trends upwards, fears of a

Figure 5: Volatility differences, March 2003–March 2004

crash are ‘forgotten’, and the upward adjustment to the correlation matrix decreases. This reduces the difference between the index volatility and the stock volatility. This is clearly visible in Figure 2.

Figures 3 and 4 repeat the same computations, but now using the implied volatility numbers that are computed using the variance swap methodology.

It is instructive to place both volatility difference charts next to each other. Figure 5 plots the difference that was computed using the ATM implied volatilities and the difference as computed by the volatility dispersion methodology (the VIX approach).

It appears that the volatility difference using the VIX methodology is much less volatile than the volatility difference using

the ATM methodology. The volatility of the VIX chart is approximately 200 per cent, while that of the ATM chart is about 338 per cent.

SUMMARY

The theoretical results show that it may be simpler to trade the volatility dispersion strategy using the variance swap methodology. This is for two main reasons:

- (1) It eliminates the constant rebalancing (‘chasing the spot’) requirement when using ATM options.
- (2) The volatility difference, which is the quantity being traded, appears more tame under the variance swap (VIX) methodology.

It appears that some market participants have already adopted the variance swap methodology for volatility dispersion trading.

Acknowledgments

The author wishes to thank Krag Gregory and Venkatesh Balasubramanian at Goldman Sachs.

References

- (1) Schneeweis, T. and Spurgin, R. (2001) 'The benefits of index option-based strategies for institutional investors', *Journal of Alternative Investments*, Spring, pp. 44–52.
- (2) Bollen, N. P. and Whaley, R. (2002) 'Does price pressure affect the shape of implied volatility functions?', Working Paper, Duke University, Durham, NC, USA.
- (3) For an excellent guide to this see Markowitz, H. M., Todd, P. G. and Sharpe, W. F. (2000) 'Mean-Variance Analysis in Portfolio Choice and Capital Markets', John Wiley, Chichester, UK.
- (4) Nelken, I. (Ed.) (1997) 'Volatility in the Capital Markets', Glenlake, Chicago, IL.
- (5) Higham, N. (2002) 'Computing the nearest correlation matrix — A problem from finance', *IMA Journal of Numerical Analysis*, Vol. 22, pp. 329–343.
- (6) Neuberger, A. (1995) 'The Log Contract and Other Power Contracts', in Nelken, I. (Ed.), 'Handbook of Exotic Options', Irwin Professional Publishing (now McGraw-Hill), Burr Ridge, IL.
- (7) Gatheral, J. (2003) 'Replication of Quadratic Variation-based Payoffs', Merrill Lynch, available at: http://www.math.nyu.edu/fellows_fin_math/gatheral/lecture7_2003.pdf.
- (8) Demeterfi, K., Derman, E., Kamal, M. and Zou, J. (1999) 'More Than You Ever Wanted to Know About Volatility Swaps', *Goldman Sachs Quantitative Strategies Research Notes*, March.
- (9) For more details, see <http://www.cboe.com/micro/vix/index.asp> and also <http://www.cboe.com/micro/vix/vixwhite.pdf>