FEDERAL RESERVE BANK
OF CHICAGO

CAPITAL MARKETS NEWS

March 2001

Volatility and Variance Swaps

Overview

Equity-index volatility and variance swaps offer an efficient way for traders to take synthetic positions in pure volatility. These swaps were first traded in late 1998, and over the past two years the product has grown into a sizeable market. Volatility and variance swaps were not the first volatility-based derivative products to be engineered and traded as volatility derivatives on foreign currencies have been traded on major exchanges for several years. These include short and long dated volatility options on major currencies such as the British pound, Japanese yen, etc. The Deutsche Terminborse (DTB) became the first exchange in the world to list volatility futures based on an underlying equity index of implied volatility when it launched the VOLAX future on January 19, 1998.[1] The VOLAX is based on the implied volatility of DAX index options, which is represented by the VDAX, a set of eight volatility indices introduced by the DTB in 1997. Similar volatility indices have been developed domestically. The VIX, an index of implied volatility on the S&P100[2], was created in 1993[3] for the CBOE, and the VXN, an index of implied volatility on the NASDAQ 100 was launched on January 23, 2001. Although exchange traded futures or options based on the VIX or VXN have not yet been developed, OTC derivative contracts on these volatility indexes may be feasible.

As with most new products, equity-index volatility and variance swaps began as mostly a one-sided market, with clients consisting mainly of hedge funds selling volatility to dealers. Today there is respectable two-way flow, [4] with trades consisting of directional views on volatility as well as relative value spread trades. According to Spyros Papadakis, an exotic options trader at Lehman Brothers in New York, bid-ask spreads on the S&P 500 range from 2.5 volatility points for a 1 year contract to 5.0 volatility points for a 5 year contract.

Despite its name, a volatility or variance swap is actually a *forward contract* whose payoff is based on the realized volatility of a stated equity index, such as the S&P 500 or the NASDAQ. Unlike option-based strategies, these products do not provide exposure to the direction of price movements of the underlying asset. Rather, they synthetically alter the trader's exposure to pure volatility of the underlying index and, if used correctly, can be efficient products for trading and risk managing volatility exposures. Further, although these products are technically forward contracts, they can provide option-like payoffs depending on the structure. Indeed, volatility and variance swaps were engineered to be a substitute for traditional option-based volatility strategies such as straddles or hedged puts or calls. A major drawback to using traditional option-based strategies is as follows: Once the underlying stock or index moves, a delta neutral trade is not delta neutral any more. Rehedging becomes necessary to maintain a delta neutral position as the market moves. Since transaction and operational costs generally prohibit continuous rehedging, residual exposure (as well as time

Capital Markets News **is published quarterly by the Capital Markets Group of the Supervision and Regulation Department. Its primary intention is to further examiners' understanding of topical issues pertaining to derivatives and other capital markets subjects. Articles are not intended as exhaustive commentaries of the subject matter; rather, they are summaries meant to convey a basic understanding of the issue and to serve as a foundation for further analysis. Readers who would like further information on any of the articles may contact the author directly. For additional copies of back issues of the newsletter, please contact Joe Cilia at 312-322-2368.**

Any opinions expressed are the authors' alone and do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.

Volatility and Variance Swaps *continued*

decay in the case of long positions) of the underlying ultimately occurs from option-based volatility strategies.

Thus far, variance and volatility swaps have been discussed as if they were identical products. While they are used for the same purpose, there are subtle but important differences between the two. Specifically, volatility swaps have linear payoffs, while variance swaps have nonlinear (curvilinear) payoffs. Moreover, volatility swaps are harder to price and risk-manage. A technical discussion of both products follows; as volatility swaps are more intuitive, they are presented first.

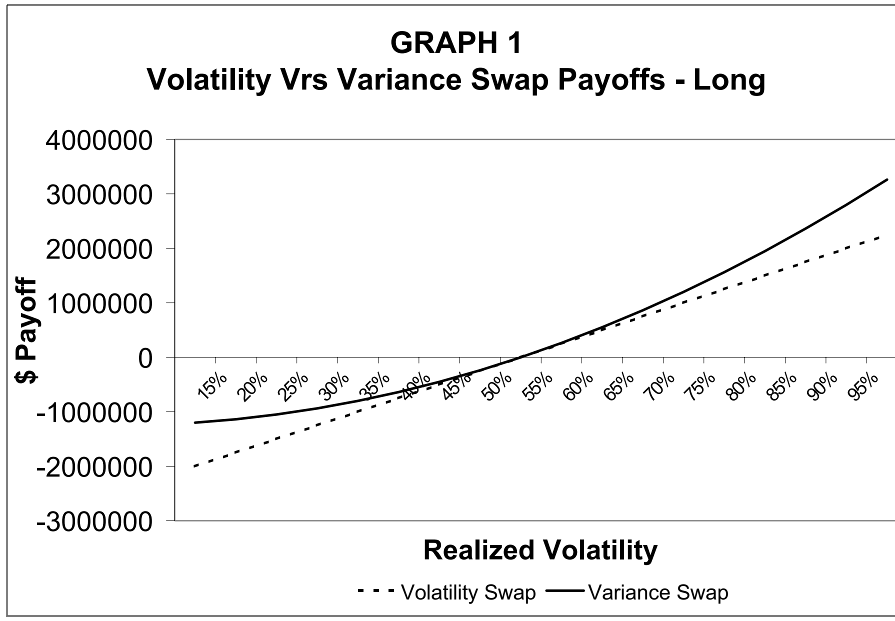

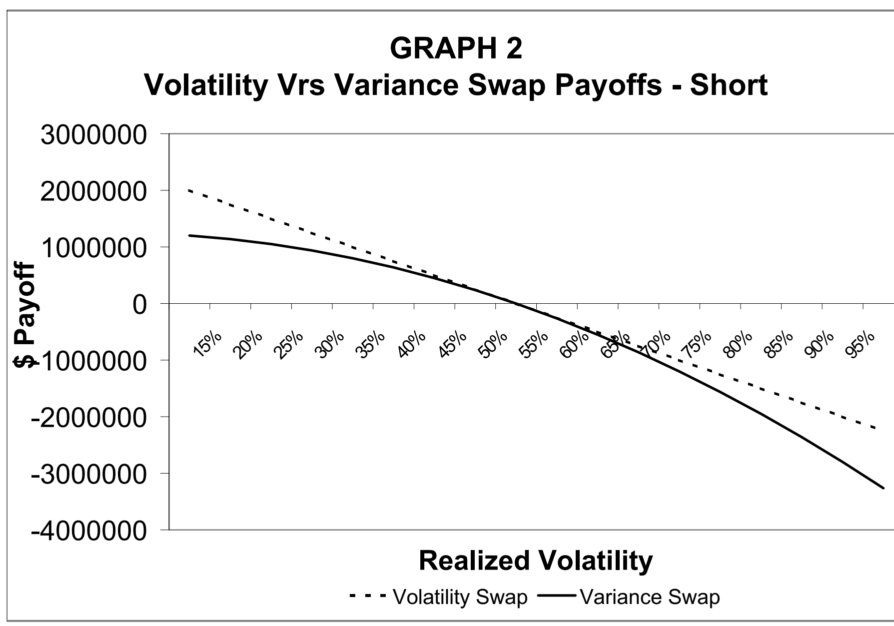

Volatility Swaps

As noted, a volatility swap is a forward contract on realized historical volatility of the underlying equity index. In such a contract, the buyer receives a payout from the counterparty selling the swap if the volatility of the stock index realized over the life of swap contract exceeds the implied volatility swap rate noted at the inception of the contract. The amount paid at expiration is based on a notional amount times the difference between the realized volatility and implied volatility, as specified here:

Payoff = $US notional x (realized volatility – implied volatility)

The notional is generally stated as multiple of $1000 per volatility point (i.e. 1%), and most contracts range between $100,000 and $500,000 per volatility point. The realized volatility is the floating volatility of the underlying asset realized over the life of the swap contract and is not known until contract expiry. The realized volatility is calculated as follows:

$$\sigma = \sqrt{A\left[\frac{1}{T}\right]\sum_{t=1}^{T}(\ln(P_{t+1}/P_t))^2}$$

Where ln(P) is the log of the asset price at time t and t+1, T is the length of the contract, and A is the annualization factor. The implied volatility is the fixed "swap rate", and is established by the writer of the swap at the inception of the contract. Essentially, the swap rate is the volatility implied by a replicating portfolio of puts and calls on the index. This portfolio of puts and calls is constructed such that its value is independent of moves in the stock price. Theoretically, the combination of options that achieves this state is a weighted combination of puts and calls across all strikes (i.e. from zero to infinity), with the weights consisting of the inverse of the square of the strike level[5]. For example, an option having a strike of $50 would have a weight of .0004 or .04% (i.e. $1/(50^2)$). However, in reality, option prices for all strikes don't routinely exist. This is especially true for options far out of the money. As a proxy, prices of less liquid or nontraded options are estimated via interpolation and extrapolation (or other skew estimation techniques) from the prices of liquid options. All the options within this "portfolio" have the same expiry as that of the volatility swap contract. Once the "portfolio" is constructed, its market value is determined. The volatility rate implied from the market value of this portfolio (under Black Scholes option pricing methodology) becomes the swap rate of the volatility derivative.

At expiration, the owner of the volatility swap receives "N" notional dollars for every point the stock index's realized volatility has exceeded the implied volatility swap rate that was established at the inception of the contract. Conversely, if the index's realized volatility is below the swap rate, then the swap holder makes a payment to the swap writer. The profit payoff of a long volatility swap is linear and is depicted as the dashed line in Graph 1 (the graph assumes a notional amount of $5 million and an implied volatility of 50%).

For larger notional amounts (all else equal), the line is more steeply sloped, resulting in a larger payoff. The payoff line pivots off of the breakeven point where implied volatility equals realized volatility. For higher (lower) levels of implied volatility (all else equal), this breakeven point of the line (zero intercept) shifts to the right (left).

The above example was discussed from the perspective of the swap owner; in this case, the trader is long the swap and therefore the higher the realized volatility, the higher the payoff. The payoff would be opposite from the perspective of a volatility swap seller; the payoff function would be downward sloping, intercepting at a volatility of 50%. See graph 2.

Variance Swaps

The overall structure and mechanics of a variance swap is similar to that of a volatility swap, except that realized variance (i.e. realized volatility-squared) and implied variance (implied volatility-squared) are used to calculate the payoff. Specifically,

$$\text{Payoff} = \$US \text{ notional x (realized volatility}^2 - \text{implied volatility}^2)$$

The use of variance instead of volatility swaps results in a nonlinear payoff. Refer to the solid line in graphs 1 and 2. Specifically, for long variance positions (graph 1), the payoff function exhibits positive convexity. This means that losses and gains are not symmetric: there is a larger payoff to the swap owner when realized variance exceeds implied variance, compared to the losses incurred when implied variance exceeds realized variance by the same volatility point magnitude. For example, using the long variance swap depicted in graph 1, a 30 point increase in realized variance over implied variance results in a net payoff to the swap owner of $1,950,000; however, a 30 point decrease results in net loss of only ($1,050,000.) A long variance swap allows the owner to enjoy gains, but buffers the owner from losses. This is particularly noticeable for large differences in implied and realized volatility.

Conversely, for short variance positions (graph 2), the payoff function exhibits negative convexity. As with positive convexity, gains and losses not symmetric; however, here there is a larger loss to the swap seller when realized variance exceeds implied variance, compared to the gains generated when implied variance exceeds realized variance by the same volatility point magnitude. Using the same example as above, a 30 point increase in realized variance over implied variance results in a net loss to the swap seller of ($1,950,000); however, a 30 point decrease results in net gain of only $1,050,000. A short variance swap benefits when realized volatility remains low during the contract period.

It is interesting to note that the payoff functions for long and short variance swaps loosely resemble those of long and short call options. As noted, a long variance swap allows the owner to enjoy gains, but buffers the owner from losses. Conversely, for short variance swaps upside potential appears capped, while downside risk is unlimited.

Strategies and Applications

There are several ways in which volatility or variance swaps can be used. It has been shown empirically that implied volatility is often higher than the volatility realized over the life of the option.[6] Given the structure of these derivatives, short variance swaps can be used to capture the difference between historical and implied volatility. Specifically, a trader can sell a variance swap, and earn profits as the contract expires. This, of course, assumes that there are no material spikes in realized volatility. However, sizeable increases in realized volatility do occur – especially during stock market declines. Therefore, if a trader expected volatility to increase, he could use a long variance swap to take a position on that view or, if his portfolio is naturally short vega, he could use the swap to hedge the loss in option book. In fact, a swap would provide a cleaner vega hedge than would offsetting options positions. The following example will help illustrate how the variance swap could be used to hedge.

Suppose a trader, whose option portfolio is short volatility with a vega[7] position of –$55,000, suspects an impending market correction and simultaneous increase in volatility sometime over the next six months. Suppose further that six-month implied volatility is 25%, and that the trader predicts that volatility will increase 35 percentage points (from 25% to 60%, for example) over this time frame. If this prediction comes to fruition, the option book could stand to lose $1,925,000 ($55,000 x 35) in market value as a result of changes in volatility alone. To hedge this volatility exposure, the trader could purchase options to neutralize vega. However, the problem with this strategy is that as soon as the underlying (i.e. S&P500) index

Volatility and Variance Swaps *continued*

moves, the position is no longer perfectly hedged. In other words, the trader would need to readjust the hedge every time the market moves – and this could be cumbersome and expensive in a volatile market. Rather than hedge with offsetting options, the trader could enter into a long variance swap. Using the swap rate of 25% (squared) above, a swap having a notional value of $6.5 million would effectively hedge the option book. If the actual volatility realized during this period was 60% (squared) as the trader projected, then he would receive at the end of the contract $1,933,750 ($6.5million x (60% squared less 25% squared), which slightly more than offsets the loss that would have been incurred on the option book.

There are other strategies that can be effected as well. For example, variance swaps may be used to execute stock index spread trading strategies. In these strategies, a short variance swap on an equity index such as the S&P 500 is partially hedged by a long swap on a different index, such as the NASDAQ 100. This spread trade has a payoff based on the difference between the realized volatility (or variance) in these two indices.[8] For example, suppose a trader shorted a variance swap on the S&P 500 with a swap rate of 30% and purchased a variance swap on the NASDAQ 100 with a swap rate of 35%. With this strategy, the trader is betting that the realized volatility of the NASDAQ 100 will exceed the realized volatility of the S&P 500 by more than 5 volatility points.

Market Risk[9] of a Variance Swap

At inception, the variance swap contract will have a zero market value, but throughout the life of the contract the market value of the swap will fluctuate. The mark-to-market value of a variance swap is primarily influenced by changes in the volatility surface for options of similar maturities based on the remaining life of the variance swap. Changes in the shape of the volatility surface, such as flattening or steepening, reflect changes in trader sentiment of market conditions. For example, if the demand for put options increases, the volatility surface steepens as put options become more costly. When these shifts occur during the life of the variance swap, the value of the swap fluctuates. An added layer of complexity (hence risk) relates to the fact that market prices for illiquid options are a matter of opinion and often differ from dealer to dealer. As such, dealers often employ different surface estimation techniques (linear extrapolation, cubic spline, etc) to derive the implied volatility of illiquid options.

Summary

The above article describes the structure, uses, and risks of volatility and variance swaps. These products may be useful in synthetically altering exposure to volatility, although care must be taken to understand the inherent risks in using such products. For additional information, please consult the referenced articles.

——*Cheryl L. Sulima, CFA*

References:

"The Investor Fear Gauge", Robert E. Whaley, last updated February 4, 2000.

"Derivatives on Market Volatility; Hedging Tools Long Overdue", Robert E. Whaley, Journal of Derivatives, 1, 1993, pages 71 – 84.

"A Guide to Volatility and Variance Swaps", Kresimir Demeterfi, Emanuel Derman, Michael Kamal, and Joseph Zou, The Journal of Derivatives, Summer 1999, pages 9 – 32.

"Variance Swaps", Paul Koo, Carl Mason, Ph.D., Scott Mixon, Ph.D., Warburg Dillon Read Global Equity Research, February 2000.

"Market Risk of Variance Swaps", Neil Chriss and William Morokoff, Risk Magazine, October 1999, pages 55 – 59.

"Equity Vol Swaps Grow Up", Nina Mehta, Derivatives Strategy Magazine, July 1999.

"The Art of the Variance Swap", Dean Curnutt, Derivatives Strategy Magazine, February 2000.

"The Quality of Market Forecasts Implied by S&P100 Option Prices", Jeff Fleming, Journal of Empirical Finance, 1998, 5, 317 – 345.

"Option Volatility & Pricing", Sheldon Natenberg, 1994.

[1] "Prepackaged Volatility Plays", DerivativesStrategy.com, April 1998.

[2] VIX is computed on a minute by minute basis from the implied volatilities of the eight near the money, nearby, and second nearby OEX option series. These implied volatilities are then weighted in such a manner that the VIX represents the implied volatility of a 30-calendar day (22 trading day) at the money OEX option.

[3] "The Investor Fear Gauge", Robert E. Whaley, February 4, 2000. See also: "Derivatives on Market Volatility; Hedging Tools Long Overdue", Robert E. Whaley, Journal of Derivatives, 1, 1993, pages 71 – 84.

[4] "Equity Vol Swaps Grow Up", Nina Mehta, Derivatives Strategy Magazine, July 1999.

[5] "A Guide to Volatility and Variance Swaps", Kresimir Demeterfi, Emanual Derman, Michael Kamal, and Joseph Zou, The Journal of Derivatives, Summer 1999, page 12.

[6] See Fleming's 1998 Journal of Empirical Finance article.

[7] The reader will recall that the vega of an option book is the dollar change in theoretical value for each one percentage point change in volatility. In this case, the portfolio would lose $55,000 in value for every percentage point increase in volatility.

[8] "Variance Swaps", Paul Koo, Carl Mason, Ph.D., Scott Mixon, Ph.D., Warburg Dillon Read Global Equity Research, February 2000.

[9] "Market Risk of Variance Swaps", Neil Chriss and William Morokoff, Risk Magazine, October 1999, pages 55 – 59.