Market Risk for Volatility and Variance Swaps

Neil Chriss * William Morokoff †

Submitted to Risk, July 1999

1 Introduction

Volatility swaps are swaps for which one counterparty agrees to the other a notional amount times the difference between a fixed level and a floating level of volatility. The fixed level is set by the writer of the swap and is determined by a variety of factors, implied volatilities of the underlying asset. The floating level is determined by a formula for the average variability of the asset over its life. The resulting product is therefore a pure bet on the level of volatility that institutional users are attracted to as an alternative to using options as a means to take on or hedge volatility exposure. These swaps, however, are newer and less well understood than options, and in particular very little has been written concerning their risk management.

The market for volatility swaps at the time of this writing is dominated by longer dated instruments with maturities in the one to five year range (for an overview of the market, see Mehta (1999)). Consequently, risk management is largely a matter of understanding fluctuations in the mark-to-market value of the swap. Recently a number of articles focusing on the pricing and hedging of volatility swaps (see Carr and Madan (1998), Demeterfi, Derman, Kamal and Zou (1999)) have appeared. These articles demonstrate that it is possible to hedge the payout risk of a variance swap using a combination of a static position in options and a dynamic stock strategy, but say nothing of mark-to-market risk. This article exclusively studies mark-to-market risk. We classify the types of risks the holder of a volatility swap faces, and argue that some of these risks are modelable and while others depend exclusively on the valuation of out-of-the-money options whose values are not available in the market.

*Asset Management, Goldman, Sachs & Co., neil.chriss@gs.com
†Firmwide Risk, Goldman, Sachs & Co., william.morokoff@gs.com

2 Definitions

We begin by defining the basic components of volatility and variance swaps. Let $S(t)$ be the value of a security at time t. A volatility swap on S is a contract that is traded at a time (the trade date) t_0 and that matures at a later time T (the maturity date) with a strike price K_{vol} that is determined at time t_0. We refer to this as the *fair value of volatility.* The payout on the volatility swap is given by the formula:

$$P_{vol} = N \cdot (\sigma_R - K_{vol}) \tag{1}$$

where σ_R is the realized volatility over the life of the swap and N is the swap notional. The realized volatility is defined by the writer of the swap contract in agreement with the counterparty. A typical definition is as follows:

$$\sigma_R = \sqrt{\frac{1}{T - t_0} \sum_{i=1}^{M} \left(\frac{S_i - S_{i-1}}{S_{i-1}} \right)^2} \tag{2}$$

for a swap covering M return observations, with S_i being the closing price of the asset on the i^{th} day. If the notional N is positive, the position is long; otherwise it is short. For practical purposes, the fair value K_{vol} is stated as an annualized percentage, often referred to as *volatility points.* The payout on a variance swap is similarly defined as

$$P_{var} = N \cdot (\sigma_R^2 - K_{var}) \tag{3}$$

where now σ_R^2 is the realized variance over the life of the contract, and K_{var} is the variance strike quoted in *volpoints*^{2}.

We now derive the mark-to-market value of a variance swap for a time t that falls between t_0 and T. Write $V(t_0, t)$ for the realized variance from time t_0 to t. The variable V satisfies the following additivity property:

$$V(t_0, T)(T - t_0) = V(t_0, t)(t - t_0) + V(t, T)(T - t) \tag{4}$$

It is this additivity that makes realized variance much easier to model than realized volatility. From the decomposition (4) it follows easily that

$$V(t_0, T) = \lambda V(t_0, t) + (1 - \lambda) V(t, T)\,, \tag{5}$$

where λ is the proportion of time already elapsed on the swap by time t defined by

$$\lambda(t) = \frac{t - t_0}{T - t_0}\,. \tag{6}$$

From Equation (3) we see that at time t the holder of a variance swap has learned something about the swap's final payout: it consists of a known, fixed amount (the realized variance up to time t) and an unknown amount, the variance for the remaining life of the swap. This unknown variance can be completely hedged by entering into an offsetting swap with notional $(1-\lambda)N$ over the period (t, T) with variance strike K_t. Therefore, the mark-to-market value of the variance swap at time t must be

$$\text{Var. Swap Value} = Ne^{-r(T-t)} \left(\lambda(V(t_0,t) - K_{var}) + (1-\lambda)(K_t - K_{var})\right) \quad (7)$$

where r is the continuously compounded risk free discount rate. We have arranged the equation to be intuitively easy to understand: it states that the mark-to-market value of the variance swap is equal to a time-weighted average of the realized payout on the variance swap through t and the change in the fair value of variance (where by fair value of variance we mean the variance strike). We will see later that the value K_t is a key factor in the risk management of variance swaps. Some dealers call this quantity the *unrealized volatility* of the underlying. In this paper we will simply refer to it as the variance strike of the offsetting swap, or when there is no possibility of confusion, simply the variance strike.

Unfortunately the lack of additivity for realized volatility means that there is no simple formula for the mark-to-market value of a volatility swap. It is not possible to hedge a partially matured volatility swaps with new volatility or variance swaps or any other commonly traded instruments. For this reason the exact pricing and marking-to-market depends on the choice of a stochastic volatility model for the underlying, any number of which may be consistent with observable market instruments. This is in contrast to the maturing variance swap which can be hedged by entering into an off-setting new variance swap (as described above), and for which the fair value of variance can be replicated by a portfolio of puts and calls.

While it is not possible to determine the mark-to-market value of a volatility swap without specifying a stochastic volatility model, it is possible to bound this value. As described in Morokoff, Akesson and Zhou (1999), the bounds can be obtained from an arbitrage argument involving an optimal hedge in a new volatility or variance swap. For example, a lower bound can be established by entering into α units of an off-setting volatility swap. For a given α, $0 < \alpha < \sqrt{1-\lambda}$, it can be shown that there is a worst case realized volatility scenario for which the hedged position has a minimum value at expiry. The hedge parameter α can then be chosen to maximize the value of this worst case. As the hedge swap has zero value at time t, the mark-to-market value of the original volatility swap must be at least as large this maximized worst case value (otherwise there is an arbitrage opportunity). Similarly, a variance swap hedge can be used to obtained an upper bound. This approach leads to the inequalities

$$Ne^{-r(T-t)} \left(\sigma^-(t) - K_{vol}\right) \le \text{Vol. Swap Value } \le Ne^{-r(T-t)} \left(\sigma^+(t) - K_{vol}\right) \quad (8)$$

where

$$\sigma^{-}(t) = \sqrt{\lambda V(t_0, t) + (1 - \lambda) K_{vol}^2(t)} \tag{9}$$

$$\sigma^{+}(t) = \sqrt{\lambda V(t_0, t) + (1 - \lambda) K_{var}(t)}\,. \tag{10}$$

Here $K_{vol}(t)$ and $K_{var}(t)$ are the respective volatility and variance strikes for swaps covering the period (t, T).

The same result is obtained if it is assumed that $K_{vol}(t)$, $K_{var}(t)$ and the mark-to-market volatility swap value are all expected values of their associated undetermined random components under some risk neutral measure through an application of Jensen's inequality. The choice of the risk neutral measure is equivalent to the choice of a stochastic volatility model, but this choice is expressed only through the value of $K_{vol}(t)$ in the bounds.

With these formulas in hand, we move on to the central concern of this paper: what are the chief sources of mark-to-market risk for a volatility swap and a variance swap? In the following section we identify and quantify these sources, and later we build a model for measuring this risk. In what follows we will primarily focus modeling the risks of variance swaps, while bearing in mind that Equation (8) implies that our results for variance swaps apply at least qualitatively for volatility swaps as well.

3 The Sources of Risk for a Variance Swap

The mark-to-market value of a variance swap has four important sources of risk:

1. **Movements of fair value of volatility (strike risk)**: the fair value of volatility K_t for an offsetting swap at time t (see Equation (7)) is the dominant source of mark-to-market risk for a variance swap. Its movements dominate continuous price movements of the underlying asset, so that while when the swap matures its final value depends solely on the intra-life movements of the underlying, mark-to-market risk is paradoxically determined almost exclusively by movements in fair value. The movements of fair value of volatility are themselves driven by the movements in implied volatility (to be explained below), and consequently the risk associated with movements in fair value of volatility is called *vega* risk. This risk is modelable and large. We will discuss how to model it later.

2. **Continuous underlying price movements (delta risk)**: if the process governing the underlying's price movements consists of a continuous component and a jump component, then under a large class of processes, the continuous component is a negligible consideration in measuring the risk of a volatility or variance swap. This risk as modelable and low.

3. **Asset price jumps (jump risk)**: jumps in the underlying's price can cause sharp movements in the value of a variance swap. The associated risk can be modeled with fat-tailed distributions for its risk factors or jump diffusion processes for the underlying. However, the pricing and hedging models that implicitly determine the mark-to-market value of a volatility swap are based on the assumption of no jumps, leading to additional model error (see Demeterfi, et al (1999) for a discussion of the impact of jumps on variance swap hedging). Such jumps may also have a significant impact on a dealer's risk appetite, leading to significant changes in their assessment of the fair value of volatility (see above). In addition the data available for parameterizing any such model are insufficient for creating reliable forecasts of jump frequencies. As a consequence, we regard this risk as high and unmodelable.

4. **Illiquid implied volatility (illiquidity risk)**: the theoretical value of the variance swap strike depends on the implied volatilities of options of all strikes, although only a small fraction of those options have observable prices. Put another way, movement in fair value of volatility are determined by movements in both observable and unobservable option values. Unobservable option values are, of course, the implied volatilities of deep out of the money calls and puts representing the far end of the volatility skew. Therefore a given move in fair value of volatility implies a particular movement in tail of the volatility skew. This sensitivity is in fact quite significant. However, the level of this implied volatility and the daily changes in the volatility skew in this unobservable region are essentially a matter of opinion and are governed as much by tastes and preferences as by any theory. This is a particularly thorny issue because a dealer's stance on the riskiness of out-of-the-money puts may change due to a variety of factors including an increased risk aversion. This steepening of the skew can have significant impact on the fair value of volatility and hence the value of any volatility swap on the books. Should this change in view be reflected in the value of the volatility swaps on the books? This risk is therefore unmodelable.

4 Basic Mark-to-Market Risk

In this section we will define and discuss the mark-to-market risk and derive some of the basic properties. As mentioned above we will focus exclusively on the risk of a variance swap for which we define the mark-to-market risk as the uncertainty in swap's value from one period to the next. More formally, at a time t, the mark-to-market risk for a period of time Δt is related to the uncertainty in

$$\Delta M = M(t + \Delta t) - M(t) . \tag{11}$$

This is a formal definition. To produce a more quantitatively useful definition, we make use of the mark-to-market formula (7). For modeling purposes, it is convenient to start from a slightly different definition of realized variance than the one above. If $\sigma(S(t), t)$ is the instantaneous volatility of the underlying asset with price $S(t)$, then the realized variance over a period (t_1, t_2) can be defined as

$$V(t_1, t_2) = \frac{1}{t_2 - t_1} \int_{t_1}^{t_2} \sigma^2(S(\tau), \tau)\, d\tau, \tag{12}$$

which is the continuous time analogue of the square of equation (2). This definition of realized variance clearly satisfies the additivity property of equation (4). We then take the risk-neutral expectation at time t of the variance swap payout (3) and find that the mark-to-market value of the swap $M(t)$ is given by Equation (7), where K_t is the risk neutral expectation of $V(t, T)$ conditioned on information available at time t.

Assume for the moment that the risk-free rate is identically zero. Then it is easy to derive the formula

$$\Delta M = N \cdot \left(\frac{1}{T - t_0} V(t, t + \Delta t)\Delta t + (1 - \lambda_{t+\Delta t})\Delta K - \frac{1}{T - t_0} K_t \Delta t \right) , \tag{13}$$

where $V(t, t+\Delta t)$ is the realized (annualized) variance over the time interval Δt, and ΔK is the change in variance strike over the same period, $K_{t+\Delta t} - K_t$. Put more directly, the change in value of a variance swap over a period is determined by realized variance, change in strike and the length of the time period itself.

As we will discuss in the next section, the variance strike K_t depends only on the volatility skew for options with expiry T, and therefore the quantity ΔK is determined by changes in the implied volatility surface over the interval Δt. If it is assumed that this surface evolves according to a continuous diffusion process, then $\Delta K = \mathcal{O}(\sqrt{\Delta t})$; that is, changes in K over the period Δt are on the order of size $\sqrt{\Delta t}$.

As the realized variance over Δt is bounded and K_t is known, the first and third terms of (13) are of size Δt. Thus movement of K_t dominate the mark-to-market risk. A related observation is that the sensitivity of the market value of a variance swap to changes in the underlying asset price (the "delta" of the swap) is small and goes to zero as $\Delta t \to 0$. This follows from the usual approximation for realized variance over a short interval (from Equation (2))

$$V(t, t + \Delta t)\Delta t \approx \left(\frac{\Delta S}{S(t)} \right)^2 . \tag{14}$$

Note that from this formula we see that continuous asset price movements are negligible over short horizons, but the swap's value is still sensitive to price jumps. The sensitivity of $M(t)$ to time decay, under constant asset and volatility levels (the θ of the swap), can be seen from (13) to be

$$\theta = -N \frac{K_t}{T - t_0} . \tag{15}$$

For the case of a non-zero interest rate that is constant over Δt, let $\beta_t = \exp(-r(T-t))$. The change in the market value of the swap over Δt is essentially the same as (13) with the exceptions that the time decay term becomes

$$\theta = rM_t - N\beta_t \frac{K_t}{T - t_0}, \tag{16}$$

and N is replaced by $N\beta_t$.

The upshot of all this is that if the $\mathcal{O}(\Delta t)$ terms are dropped from (13), then the resulting volatility of ΔM is equal to:

$$\sigma(\Delta M) \approx (1 - \lambda)\, N\beta_t\, \sigma(K_t). \tag{17}$$

Therefore the problem of defining the mark-to-market risk of a variance swap is reduced to that of finding the volatility of the variance strike K_t. We re-emphasize that in this context K_t is the variance strike for a swap with time to maturity $T - t$. This implies that as swaps evolve, the holder is exposed to the volatility of variance strike for all possible maturities less than the maturity of the swap itself. Thus, studying variance swap risk requires that we understand the volatility of variance strikes for all possible maturities. We will explicitly discuss strategies for computing $\sigma(K_t)$ at a later point in this paper, but first we digress to discuss in more detail the formation of the variance strike K_t.

5 Mathematical Modeling

To model the volatility of the variance strike and its movements, it is helpful to consider a formula for the variance strike in terms of the implied volatility surface. As described in recent articles (Carr and Madan (1998) and Demeterfi, et al (1999)), the fair strike level is obtained by essentially taking a weighted average of option prices (where "options" refers to options on the swap's underlying asset) across all strikes for options expiring on the swap's expiration. This average can be reformulated in terms of the implied volatility surface as follows:

$$K_{var}(t) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} f(\sigma(z, T-t), z, T-t)\, \exp(-z^2/2)\, dz \tag{18}$$

where

$$f(\sigma, z, T-t) = \frac{2}{T-t}\left(\sigma(z, T-t) \cdot z \cdot \sqrt{T-t} + \frac{1}{2}\sigma^2 \cdot (T-t) - \frac{\partial \sigma}{\partial z} \cdot \sqrt{T-t}\right) \tag{19}$$

(see Morokoff et al (1999) for details of this derivation). The function $\sigma(z, T-t)$ is a representation of the entire implied volatility surface in terms of time and a parameter z related to strike price that is defined below.

The importance of Equations (18) and (19) lies in the explicit representation of the variance strike as a function of the implied volatility surface. Understanding the relationship between the implied volatility skew, particularly for the unobservable far out-of-the-money regions, and the variance strike provides insight into the pricing as well as the risks associated with these swaps. Equations (18) and (19) also provide a convenient representation for numerical valuation of K_{var} as well as the fluctuations of K_{var} due to fluctuations in the implied volatility surface. These issues are explored in detail in the next sections.

6 Strike Risk: Models From Time Series

As the dominant factor influencing mark-to-market risk is movements in the variance strike of an offsetting swap, we will discuss at some length various techniques for modeling the variance strike. There are two basic methods. One is to directly model the time series behavior of the variance strike in order to make precise statements about its conditional distribution at a give time. The other is to build a model for movements of the implied volatility surface, and then use that model combined with equation (18) to make precise statements about the conditional distribution of the variance strike at a given time.

If time series for variance strikes with specific fixed maturities are available, then by assuming a time series model for ΔK (e.g., GARCH or exponential decay with normal or fat-tailed innovations, etc.) one can calibrate the model to the available data to yield a conditional distribution for the variance swap strike.

If an adequate time series for the variance strike is not available, an alternate approach is to use the movements in the square of at-the-money implied volatilities as a proxy for movements of the fair value of variance. If IV_{T-t} represents a time series of at-the-money implied volatility for options with time $T-t$ to expire, then the implied time series

$$\Delta K_{T-t} = IV_{T-t}^2 - L\, IV_{T-t}^2 \tag{20}$$

(where L is the one-period lag operator) represents an approximation for historical changes in variance strike. If this series can be constructed for a variety of maturities, then through judicious use of interpolation one can form a complete set of time series models for the K_{T-t} for various maturities. These in turn can be used to estimate $\sigma(\Delta K_{T-t})$ for any maturity $T-t$. As volatility of at-the-money implied variance for various maturities is relatively easy to obtain for most equity indices (at least for maturities of one year or less), this provides a simple estimate of the market risk of a variance swap in cases for which time series for the real variance swaps are not available.

7 Strike Risk From Implied Volatility Models

The simple approximation of modeling fluctuations of variance strike as the volatility of implied at-the-money variance can be made more rigorous by introducing an implied volatility surface model $\sigma(z, T-t)$ for use with Equation (18). We begin by considering the effect of the volatility skew on the variance strike, then continue on to consider how fluctuations of the volatility surface influence the changes in the variance strike.

7.1 Dependence of Variance Strike on Skew

The variable z is related to volatility and strike price κ by

$$z(x) = \frac{-\log x - .5\sigma^2(T-t)}{\sigma\sqrt{T-t}}\,. \tag{21}$$

Here x is the ratio of the forward price of the underlying asset F (for a forward with expiry T) to the strike price, so that $x = F/\kappa$.

The implied volatility surface σ can be modeled in a number of ways (as in, e.g. Brown and Randall (1999), Morokoff et al (1999)). If the volatility skew for options with expiry T is modeled as a function of x, so that $\sigma = \sigma(x)$, then Equation (21) implicitly defines σ as a function of z. This assumes that both $\sigma(x)$ and $z(x)$ are monotonic functions, which for equity indices is reasonable.

The z of Equation (21) is more commonly known as the "d_2" variable that appears in the Black-Scholes option pricing formula. The integration over z is equivalent to the integration over all strikes that appears in the usual variance swap strike formulas.

Alternatively, z can be related to a put option's delta through the formula Black-Scholes expression

$$\delta = N(z + \sigma\sqrt{T-t}) - 1, \tag{22}$$

where δ is the put's delta and $N(\cdot)$ is the cumulative normal distribution function. Another common model for implied volatility assumes that σ is a function of δ in which case we have

$$z = N^{-1}(\delta + 1) - \sigma(\delta)\sqrt{T-t}\,. \tag{23}$$

Again assuming that $z(\delta)$ and $\sigma(\delta)$ are monotonic, this relationship can be inverted to give $\sigma(z)$.

Once a suitable model for the implied volatility has been chosen, either as a function of strike or delta, and the associated function $\sigma(z)$ has been determined, the variance strike K_{var} can then be computed.

The function σ will also depend on a number of parameters that may be associated with quantities like the implied volatility level, the slope of the skew, etc. As a simple illustration, we consider the model

$$\sigma(\delta) = \sigma_0 + bf(\delta) \tag{24}$$

where $f(\delta)$ describes the shape of the skew (e.g., linear) and has the property $f(-.5) = 0$ so that the at-the-money volatility is σ_0. For this model, dependence on time to expiry will be suppressed. We now illustrate how the choice of the

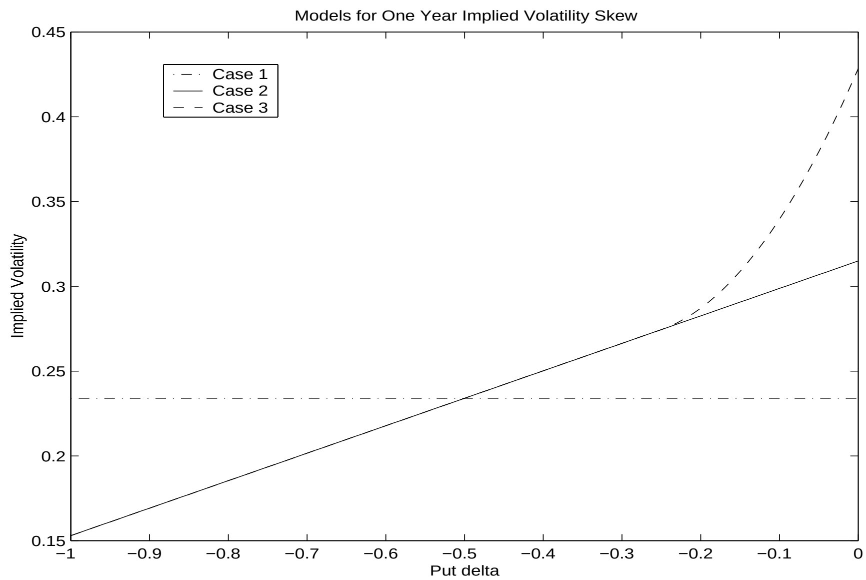

Figure 1: Implied volatility modeled as a function of put delta

parameters and the skew shape can influence the variance strike. We consider the following three cases:

Case 1	$\sigma_0 = .234$	$b = 0$		
Case 2	$\sigma_0 = .234$	$b = .162$	$f(\delta) = .5 + \delta$	
Case 3	$\sigma_0 = .234$	$b = .162$	$f(\delta) = .5 + \delta$	$\delta \leq -.25$
			$f(\delta) = 1.2 + 6.6\delta + 11.2\delta^2$	$\delta > -.25$

The first case is constant volatility, while the second case is a linear skew model considered in Demeterfi, et al (1999). The third case matches the second case over the usual range of observable options, but sets a substantially higher volatility for far out-of-the-money puts ($\delta > -.25$), with a maximum implied volatility of around 43%. The functions $\sigma(\delta)$ are plotted in Figure 1. for these three cases, while the corresponding $\sigma(z)$ are plotted in Figure 2. For these calculations $T-t$ is taken to be 1 year.

The plots in Figure 2 clearly show the difference that volatility skew makes in determining the variance swap strike, and in particular the effect of the implied volatility for far out-of-the-money puts. Although Cases 2 and 3 are essentially the same over the range of commonly traded deltas, the model that allows greater steepness in skew outer ranges of delta (a reasonable possibility) leads to significantly different results. For the case of constant volatility, it is easy to see that Equation (18) integrates to give $K_{var} = \sigma_0^2 = 0.0548$. For Case 2, $K_{var} = .0623$, a 14% increase attributable to the skew. For Case 3, $K_{var} = .0735$, an additional 18% larger than Case 2.

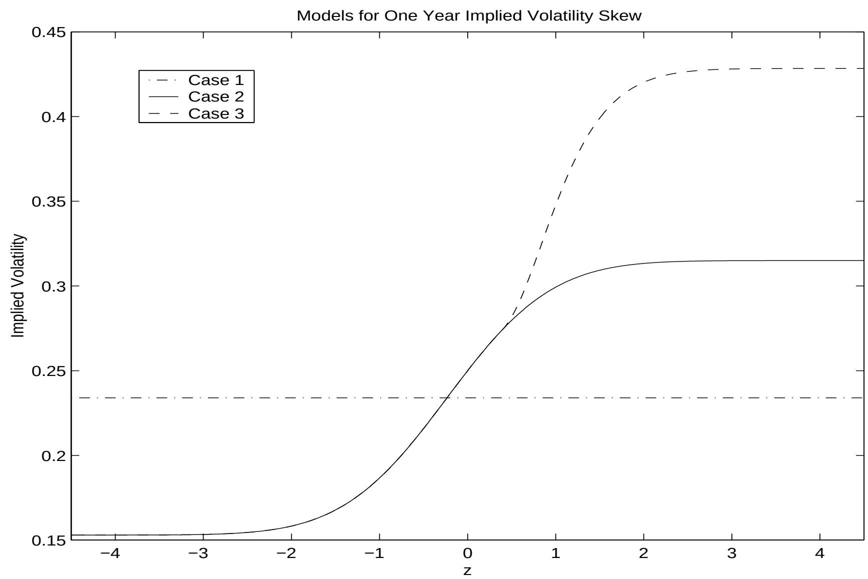

Figure 2: Implied volatility modeled as a function of z (Black-Scholes d_2)

The choice of a variance (or volatility) strike is in effect a statement of the dealer's views on the unobservable implied volatility of far out-of-the-money puts, or equivalently, a measure of the dealer's risk aversion to low probability events. A major source of market risk for a variance swap therefore lies in the possibility of changes in the dealer's level of risk aversion, an event that cannot be effectively modeled.

7.2 Strike Risk Arising From Implied Volatility

Equation (18) shows that the variance strike depends only on the implied volatility surface. Therefore the risk associated with K_t can be modeled by describing the evolution of the volatility surface. In general, if the surface can be parameterized by a set of P parameters α, as well as the strike variable z and time to expiry $T-t$, then under a linear approximation

$$\Delta\sigma = \sum_{i=1}^{P} \frac{\partial \sigma}{\partial \alpha_i} \Delta\alpha_i \,. \tag{25}$$

The associated change in K_t is then given by

$$\Delta K_t = \sum_{i=1}^{P} \left(\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \frac{\partial f(\sigma(z, T-t), z, T-t)}{\partial \alpha_i} \exp(-z^2/2)\, dz \right) \Delta\alpha_i \tag{26}$$

and

$$\frac{\partial f(\sigma, z, T-t)}{\partial \alpha_i} = \frac{2}{T-t}\left(\left(z \cdot \sqrt{T-t} + \sigma \cdot (T-t)\right) \frac{\partial \sigma}{\partial \alpha_i} - \frac{\partial}{\partial z}\left(\frac{\partial \sigma}{\partial \alpha_i}\right)\sqrt{T-t}\right). \tag{27}$$

Once a parameterized model for $\sigma(\delta, \alpha)$ or $\sigma(K/F, \alpha)$ is selected, it can be calibrated to the observable option prices, as well as to any views related to the unobservable strikes, to obtain the current α. From this, the coefficients of the $\Delta\alpha_i$ that appear in Equation (26) can then be evaluated. A historical time series for the $\Delta\alpha_i$ can be generated from the historical implied volatility data, leading to a model for the distribution of the stochastic variable $\Delta\alpha_i$. As ΔK_t is linear in the $\Delta\alpha$, the distribution of ΔK_t is then easily computed. For example, if $\Delta\alpha$ is modeled as multi-variate normal, then ΔK_t is also normally distributed.

As an example, consider the simply linear model described above for implied volatility as a function of option delta (Case 2). Here the parameters are $\alpha_1 = \sigma_0$ and $\alpha_2 = b$. Setting the current values of the parameters to those used above ($\alpha_1 = .234$, $\alpha_2 = .162$ and $T - t = 1$) and numerically evaluating the integrals in (26) leads to the model

$$\Delta K_t = 0.49\Delta\alpha_1 + 0.061\Delta\alpha_2. \tag{28}$$

From daily implied volatility data for $\delta = (-.75, -.5, -.25)$ on 12 month S&P options, the time series for $\Delta\alpha_1$ and $\Delta\alpha_2$ can be generated, leading to the covariance matrix (based on data from June 1998 to July 1999 weighted with 20% monthly decay)

$$\mathrm{Cov}(\Delta\alpha_1, \Delta\alpha_2) = 10^{-4}\begin{pmatrix} .385 & .364 \\ .364 & .908 \end{pmatrix} \tag{29}$$

Under the assumption of normality, then ΔK_t is normally distributed with mean zero and volatility 0.0034, which is very close to the volatility of the at-the-money implied variance. For a two year variance swap with one year left to expiry, a notional of 10 million dollars and a risk free rate of 5%, this puts the daily fluctuation of market value at around .016 million.

8 Jump Risk

Finally we consider the effects of jumps in the underlying asset and volatility processes. There are three ways in which such jumps affect the market risk of a variance swap. First, a large daily return for the asset can spike the realized variance, causing an increase in the market value of the swap. This effect is mitigated by the fact the each daily return carries a weight of only Δt. To a certain extent, the existence of some larger returns has already been priced into the swap. An analysis of the replication strategy (see Demeterfi, et. al.) shows that for a variance swap hedged by this strategy the P&L error associated with a daily return J is $\mathcal{O}(J^3)$. For typical diffusion returns, $J = \mathcal{O}(\sqrt{\Delta t})$ so that the P&L error is negligible. However, when a large jump J occurs, errors of size J^3 can throw off the balance between the realized variance and the variance strike. An interesting observation is that the sign of the J^3 error term depends on whether realized variance is measured as a percent return or a log return. A related hedging problem associated with jumps in the asset level is that this level may jump out of the range of validity of the approximation to the theoretical static hedge. More precisely, the finite number of options over a limited strike range initially available may not be adequate to replicate the variance if the asset moves out of this range.

A potentially larger source of jump risk arises from jumps in the implied volatility surface, and therefore unusually large fluctuations in K_t. Such jumps are generally correlated with jumps in the asset level. This risk can be captured to a certain extent by modeling the distribution of K_t or of the underlying factors α of $\sigma(\alpha, T, z)$ with fat-tailed distributions such as a mixture of normals. Such an approach should adequately model the risks associated with changes in implied volatility in the typically observable range of strikes.

Finally, jumps in asset level or observable implied volatility can induce large changes in the level of risk tolerance or aversion of a firm, and thus change the variance strike substantially beyond what the observable implied volatility fluctuations indicate. This effect is hard to measure as it depends more on the current risk appetite of a broker than a historical time series of σ or K_t may indicate.

9 Bibliography

Brown G and C Randall, 1999 *If the skew fits* Risk, April, pages 62-65

Carr P and D Madan, 1998 *Towards a theory of volatility trading* In Volatility: New Estimation Techniques for Pricing Derivatives, edited by R. Jarry, pages 417-427

Demeterfi K, E Derman, M Kamal and J Zou, 1999 *A guide to volatility and variance swaps* Journal of Derivatives 6(4), pages 9-32

Mehta, N (1999) *Equity Vol Swaps Grow Up,* Derivatives Strategy, July.

Morokoff W, F Akesson and Y Zhou, 1999 *Risk management of volatility and variance swaps* Firmwide Risk Quantitative Modeling Notes, Goldman, Sachs & Co.