Pricing Methods and Hedging Strategies
for Volatility Derivatives

H. Windcliff[*]   P.A. Forsyth[†]   K.R. Vetzal[‡]

May 4, 2003

Abstract

In this paper we investigate the behaviour and hedging of discretely observed volatility derivatives. We begin by comparing the effects of variations in the contract design, such as the differences between specifying log returns or actual returns, taking into consideration the impact of possible jumps in the underlying asset. We then focus on the difficulties associated with hedging these products. Naive delta-hedging strategies are ineffective for hedging volatility derivatives since they require very frequent rebalancing and have limited ability to protect the writer against possible jumps in the underlying asset. We investigate the performance of a hedging strategy for volatility swaps that establishes small, fixed positions in straddles and out-of-the-money strangles at each volatility observation.

1   Introduction

Recently there has been some interest in developing derivative products where the underlying variable is the realized volatility or variance of a traded financial asset over the life of the contract. The motivation behind introducing volatility derivative products is that they could be used to hedge vega exposure or to hedge against implicit exposure to volatility, such as expenses due to more frequent trades and larger spreads in a volatile market. In addition, these products could be used to speculate on future volatility levels or to trade the spread between the realized and implied volatility levels.

The simplest such contracts are volatility and variance swaps. For example, the payoff of a volatility swap is given by:

$$\text{volatility swap payoff} = (\sigma_R - K_{vol}) \times B \ , \tag{1}$$

[*]School of Computer Science, University of Waterloo, Waterloo ON, Canada N2L 3G1, hawindcliff@elora.math.uwaterloo.ca

[†]School of Computer Science, University of Waterloo,Waterloo ON, Canada N2L 3G1, paforsyt@elora.math.uwaterloo.ca

[‡]Centre for Advanced Studies in Finance, University of Waterloo, Waterloo ON, Canada N2L 3G1, kvetzal@watarts.uwaterloo.ca

where σ_R is the realized annualized volatility of the underlying asset, K_{vol} is the annualized volatility delivery price and B is the notional amount of the swap in dollars per annualized volatility point. More complex derivative contracts are also possible, such as volatility options and products which cap the sizes of the discretely sampled returns.

The analysis of variance is inherently easier than the analysis of volatility and consequently a lot of work in this area [4, 7, 5, 11] has focused on variance products. There are two commonly proposed hedging models for variance. The first involves hedging with a log contract [17], which can be approximated by trading in a large number of vanilla instruments [4, 8]. A second hedging approach involves direct delta-hedging of the variance product [11]. Interestingly, the proponents of each method indicate that the other method is likely to fail in the presence of transaction costs, a point we will investigate in this paper. Further, most analytic work [4, 8, 12] specifies continuously realized variance, whereas in practice the variance is discretely monitored.

Another collection of papers has focused on volatility derivative products, considering them to be a square root derivative of variance as discussed in [8]. In [2] the authors provide a volatility convexity correction relating variance and volatility products. One problem with hedging volatility products is that they require a dynamic position in the log contract, which will result in a large amount of trading in far out-of-the-money vanilla instruments. Due to the difficulties with hedging these products, some authors have even suggested pricing these products via expectation in the real physical measure [13].

In this paper we develop pricing and hedging methods for discretely sampled volatility derivatives. We focus on the structure that is imposed by the design of the contract rather than on a specific model for the stochastic process followed by the underlying asset. We will find that the contract structure will affect the feasibility of various hedging methods when applied to these products. Even in a constant volatility Black-Scholes setting, delta hedging strategies must be rebalanced so frequently that they are not a practical method for hedging discretely observed volatility. Further, if there are possible jumps in the underlying asset price then even if the delta hedge is rebalanced very frequently it does not effectively manage downside tail events. As an alternative, we will investigate the performance of a delta-gamma hedging strategy with an appropriate selection of vanilla hedging instruments. This strategy can be viewed as an approximation of the log contract hedge, while avoiding rebalancing a large number of positions in far out-of-the-money vanilla instruments. Simulation experiments provided in this paper demonstrate that this technique can provide very effective downside risk management. We will conclude by investigating the impact of transaction costs on the various proposed hedging strategies.

2 Volatility Derivative Products

In the introduction, we discussed a very simple volatility derivative product, the volatility swap. Even restricting ourselves to volatility swaps, there are many possible contract variations. For example, there are many possible ways that the volatility derivative contract may define the realized volatility and many ways that the discretely sampled returns can be calculated. In this section we discuss some common volatility

and variance derivative contracts.

2.1 Calculation of Returns

If we sample the underlying asset price at the times:

$$\{t_{obs,i} \mid i = 0, \ldots, N\} , \tag{2}$$

then there are two common contractual definitions of the return during the interval $[t_{obs,i-1}, t_{obs,i}]$. If we define $\Delta t_{obs,i} = t_{obs,i} - t_{obs,i-1}$ then the *actual return* is defined to be:

$$R_{actual,i} = \frac{S(t_{obs,i}) - S(t_{obs,i-1})}{S(t_{obs,i-1})} . \tag{3}$$

We define the *log return* to be:

$$R_{log,i} = \log\left(\frac{S(t_{obs,i})}{S(t_{obs,i-1})}\right) . \tag{4}$$

Both of these definitions of the return involve dividing by the previous asset level and the contract would need to define how the payoff is calculated in the event that the asset price becomes zero.

2.2 Calculation of Volatility

In addition to specifying how the discretely sampled returns are measured, the contract must also specify how the volatility or variance is calculated. From a discrete sample of N returns, the annualized realized volatility, $\sigma_{R,stat}$, can be measured by:

$$\sigma_{R,stat} = \sqrt{\frac{AN}{N-1}\left[\frac{1}{N}\left(\sum_{i=1}^{N} R_i^2\right) - \left(\frac{1}{N}\sum_{i=1}^{N} R_i\right)^2\right]} . \tag{5}$$

The annualization factor, A, converts this expression to an annualized volatility and for equally spaced discrete observations is given by $A = 1/\Delta t_{obs}$. In order to convert units of volatility into volatility points we would multiply by 100.

Although this is how one would statistically define an estimate for the standard deviation of returns from a sample, volatility derivatives often define a simpler approximation for the volatility. Since many volatility derivative products are sampled at market closing each day and the mean daily returns are typically quite small, often the contract defines the realized volatility, $\sigma_{R,std}$, to be:

$$\sigma_{R,std} = \sqrt{\frac{A}{N}\sum_{i=1}^{N} R_i^2} , \tag{6}$$

the average of the squared returns. Notice that the factor $N/(N-1)$ has been removed from the definition of $\sigma_{R,std}$ since it was used to account for the fact that there is a loss of one degree of freedom used to determine the mean return in (5). In this paper we will refer to (5) as the *statistical realized volatility*, whereas we will say that (6) is the *standard realized volatility*.

2.3 Contractual Payoffs

Once the contract has defined how the volatility is to be calculated, the derivative payoff can be specified. As mentioned above, the payoff of a volatility swap is given by:

$$\text{volatility swap payoff} = (\sigma_R - K_{vol}) \times B \ . \tag{7}$$

There are two objectives that are of interest when pricing volatility swaps. Since there is no cost to enter into a swap, one objective is to determine the fair delivery price K_{vol}, which makes the no-arbitrage value of the swap initially zero. The volatility delivery price can be found by computing the value of a swap with zero delivery price and multiplying by e^{rT}, where r is the risk-free rate and T is the maturity date of the contract.

A second objective is to determine the fair value of the volatility swap at some time during the contract's life given the initially specified delivery price. Because of the simplicity of the payoff of the swap contract, it is sufficient to be able to find the no-arbitrage value of a contract which pays σ_R at maturity.

In some markets severe volatility spikes are occasionally observed. In order to protect the short volatility position some contracts cap the maximum realized volatility. For example, a capped volatility swap would have a payoff given by:

$$\text{capped volatility swap payoff} = (\min(\sigma_R, \sigma_{R,max}) - K_{vol}) \times B \ . \tag{8}$$

In the variance swap market the maximum realized volatility is typically set to be 2.5 times the variance delivery price. The payoffs for variance based derivative products can be obtained by substituting in σ_R^2 in place of σ_R in the above definitions. Although we will not focus on more complex contract structures in this work, some institutions do offer products such as corridor swaps [3].

3 A Computational Model for Pricing Volatility Derivatives

In this section we describe two computational frameworks, one based on a numerical PDE approach and the other based on Monte Carlo simulation methods, that can be used to price volatility and variance based derivative products. In this paper we focus on our ability to hedge a volatility derivative product with value $V = V(S, t; \ldots)$. We utilize numerical PDE methods to obtain accurate delta, $\Delta = V_S$, and gamma, $\Gamma = V_{SS}$, hedging parameters. We then simulate the performance of various hedging strategies by simulating their performance under the real-world (physical) measure and compare the resulting distributions of profits and losses.

The numerical experiments provided in this paper assume a jump-diffusion model. The sizes of the jumps, J, are drawn from a lognormal distribution with:

$$\log J \sim N(\mu_J, \gamma_J^2) \ . \tag{9}$$

and the underlying asset price, S, which follows the SDE:

$$dS = (\mu - \lambda m)S \, dt + \sigma(S, t)S \, dW + (J - 1)S \, dq \ , \tag{10}$$

4

where $m = E[J - 1] = \exp(\mu_J + \frac{1}{2}\gamma_J^2) - 1$ and $E[\,\cdot\,]$ is the expectation operator. Also, μ is the drift rate of the underlying asset in the physical measure, $\sigma = \sigma(S, t)$ is the (state dependent) volatility function, and dW is an increment from a Wiener process. Jumps in the underlying asset price are modelled by the last term with dq being a Poisson process with arrival intensity λ:

$$dq = \begin{cases} 1 & \text{with probability } \lambda dt \\ 0 & \text{with probability } 1 - \lambda dt \,. \end{cases} \tag{11}$$

The situation where the underlying asset price evolves continuously without jumps can be modelled by setting the arrival intensity $\lambda = 0$. For simplicity, we assume that no dividends are paid by the underlying asset although it is straightforward to incorporate either a continuous dividend yield or discrete dollar dividends in our numerical framework. We also point out that in our numerical approach it is possible to use any jump size distribution in place of (9).

3.1   Risk-Neutral Valuation

Some of the numerical results provided in this paper were obtained using Monte Carlo simulation. Further, we will use the risk-neutral valuation ideas presented here to analyze the asymptotic behaviour of volatility derivative contracts.

Under the risk-neutral measure, Q, the underlying asset follows the SDE:

$$dS = (r - \lambda m)S \, dt + \sigma(S, t)S \, dW + (J - 1)S \, dq \,. \tag{12}$$

The local volatility surface, $\sigma(S, t)$, and the jump parameters λ, μ_J and γ_J, have been selected so that the model correctly prices existing options in the market. The no-arbitrage value is found by approximating the expectation:

$$V(S(0), 0) = e^{-rT} E^Q[V(S, T; \, \sigma_R)] \,, \tag{13}$$

by averaging over many sample asset paths and computing the realized quantity σ_R along each of these paths. Although this technique is very straightforward to implement, it is difficult to obtain accurate estimates of the delta and gamma derivatives throughout the life of the contract, which are necessary when we simulate the performance of various hedging strategies. When a general volatility surface is used we cannot integrate (12) analytically, although we can generate the risk-neutral random walks numerically using, for example, an Euler timestepping method.

3.2   Numerical PDE Framework

Many of the results provided in this paper were obtained using a numerical partial differential equation (PDE) framework. This allows us to efficiently compute the delta and gamma derivatives used later in this paper to simulate the performance of various hedging strategies for these contracts.

In [16] the authors provide an efficient computational model for pricing discretely sampled variance swaps in a Black-Scholes setting. The efficiency of their method comes from exploiting the linear structure of variance products and cannot be extended to volatility derivative products, which have matters complicated by the coupling of the realized returns through the square root function.

3.2.1 State Variables and Updating Rules

In order to price a general volatility derivative product we introduce two additional state variables. Let P represent the stock price at the previous volatility observation time and let Z be the average of the squared returns observed to date:

$$Z_i = \frac{1}{i} \sum_{j=1}^{i} R_j^2 \ . \tag{14}$$

In some situations it is possible to use a similarity reduction in the variable $\xi = S/P$. However, for a general volatility function, $\sigma(S, t)$, this dimensionality reduction is not possible.

Initially the state variables are set to:

$$P(0) = S(0) \tag{15}$$
$$Z(0) = 0 \ . \tag{16}$$

These variables are changed only at the discrete volatility sampling times, $t_{obs,i}$, $i = 1, \ldots, N$ according to the following jump conditions. If $t_{obs,i}^{-}$ and $t_{obs,i}^{+}$ represent the instants immediately before and after the i^{th} observation date then:

$$P(t_{obs,i}^{+}) = S(t_{obs,i}^{-}) \ , \tag{17}$$

$$Z(t_{obs,i}^{+}) = Z(t_{obs,i}^{-}) + \frac{R_i^2 - Z(t_{obs,i}^{-})}{i} \ . \tag{18}$$

The return, R_i, can be computed from the state variables contained in the computational model. For example, if the contract specifies that log returns are used then:

$$R_i = \log \left(\frac{S(t_{obs,i}^{-})}{P(t_{obs,i}^{-})} \right) \ . \tag{19}$$

The updating rules for the state variables are implicitly defined by the volatility derivative contract and are independent of any assumptions regarding the behaviour of the underlying asset. We will find that this structure has important ramifications when we consider the hedging of these products.

3.2.2 Evolution Equations Between Volatility Observations

Between the discrete volatility sampling times the state variables do not change. Consequently, between observations we can think of the value of the volatility derivative product as being a function of the underlying asset price S and time t, parameterized by the state variables:

$$V = V(S, t; P, Z) \ . \tag{20}$$

So far in this section our discussion has been independent of any assumptions regarding the behaviour of the underlying asset. In order to model the behaviour of the contract between volatility observations we need to make some assumptions. In this paper we will work with a one factor model that utilizes a local volatility surface. In some examples we allow the possibility of jumps in the underlying asset price. It could

6

be argued that it would also be useful to consider a stochastic volatility model as in [12, 13, 11]. However, our focus in this paper is to investigate hedging results that are independent of the assumptions about the evolution of the underlying asset. The simple one factor, jump-diffusion model is sufficient to illustrate our point that delta hedging strategies are ineffective for managing the risk associated with these products.

In the jump-diffusion model, the value of the volatility derivative satisfies the partial integro-differential equation (PIDE):

$$V_t + (r - \lambda m)SV_S + \frac{1}{2}\sigma^2(S,t)S^2V_{SS} - rV + \lambda E[\Delta V] = 0 \ , \tag{21}$$

where:

$$E[\Delta V] = E[V(JS,t)] - V(S,t) \tag{22}$$

$$= \int V(JS,t)p(J)dJ - V(S,t) \ , \tag{23}$$

and $p(\,\cdot\,)$ is the probability density function for the jump size. This equation is solved backwards from maturity, $t = T$, to the present time, $t = 0$, to determine the current fair value for the contract. For a description of the computational methods used to solve this PIDE the reader is referred to [9].

3.2.3 Maturity Conditions

If the volatility is defined without the mean according to (6) then it is straightforward to specify the value of the volatility derivative as a function of the state variables. For example, from the contractual payoff we see that the appropriate terminal condition for a volatility swap would be:

$$V(S,T;\ P,Z)_{\text{volatility swap}} = (100\sqrt{AZ} - K_{vol}) \times B \ , \tag{24}$$

where K_{vol} is the volatility delivery price, A is the annualization factor and B is the notional amount. The terminal condition for a variance swap would be:

$$V(S,T;\ P,Z)_{\text{variance swap}} = (100AZ - K_{var}) \times B \ , \tag{25}$$

where K_{var} is the variance delivery price. More exotic volatility payoffs are also possible in this framework. For example the terminal condition for a capped volatility swap would be:

$$V(S,T;\ P,Z)_{\text{capped volatility swap}} = (\min(100\sqrt{AZ}, \sigma_{R,max}) - K_{vol}) \times B \ . \tag{26}$$

In summary, the value of the volatility derivative product is a time-dependent function of three space-like variables. After applying the terminal condition at maturity we solve a collection of independent backward equations (21) between the discrete observation times. At the discrete volatility sampling times we apply the jump conditions (17)-(18). When we reach the date of sale of the contract, the no-arbitrage value of the volatility derivative is given by:

$$V(S = S(0), t = 0;\ P = S(0), Z = 0) \ . \tag{27}$$

An example of this technique applied to a different type of path-dependent option is given in [23].

3.2.4 Asymptotic Boundary Conditions

In order to complete the numerical problem, we determine appropriate conditions at the boundary of the computational domain, $S = S_{min}$ and $S = S_{max}$. Although it is possible to reduce the boundary truncation error in the region of interest near $S = S(0)$, $t = 0$, to an arbitrary tolerance by sufficiently extending the computational domain [14], it is of practical interest to accurately specify the boundary behaviour in order to reduce the number of nodes in the grid.

The payoff of a volatility option or swap (capped or otherwise) is linear in σ_R. Thus, it suffices to analyze the asymptotic behaviour of a contract that pays off the realized volatility at maturity, $V(S,T) = \sigma_R$. To determine appropriate boundary conditions we look at the asymptotic form of the jump conditions. Notice that these can be thought of as specifying initial data over a given volatility observation period.

We begin by analyzing the value of the volatility derivative at the instant immediately preceding the j^{th} volatility observation, $t_{obs,j}^{-}$. If we let $\mathcal{F}(t)$ represent the information available at time t then, assuming that σ_R is defined according to (6), using risk-neutral valuation we find:

$$V(S, t_{obs,j}^{-}) = e^{-r(T-t_{obs,j})} E^Q[\sigma_R | \mathcal{F}(t_{obs,j}^{-})] \tag{28}$$

$$= e^{-r(T-t_{obs,j})} \sqrt{\frac{A}{N}} \, E^Q \left[\sqrt{\sum_{i=1}^{j-1} R_i^2 + R_j^2 + \sum_{i=j+1}^{N} R_i^2} \, \middle| \mathcal{F}(t_{obs,j}^{-}) \right] . \tag{29}$$

The first term in the square root is a constant, independent of S, as it represents the past volatility observations. The second term, R_j^2, represents the current volatility observation. It depends on S in a the way specified by the contractual definition of the observed returns. At time $t_{obs,j}^{-}$ the last term is random, corresponding to the level of future volatility samples. This decomposition is illustrated in Figure 1.

Suppressing the explicit reference to the time, $t_{obs,j}^{-}$, for S far away from the previous asset level P, if the volatility function is suitably well behaved, the current volatility observation, R_j^2, will dominate in (29). Thus the value is approximately a linear function of the current return at the boundaries. For actual returns we have:

$$R_j = \frac{S - P}{P}, \qquad \frac{dR_j}{dS} = 1/P, \qquad \frac{d^2 R_j}{dS^2} = 0 . \tag{30}$$

This indicates that $V_{SS} \to 0$ at both boundaries when actual returns are specified. For log returns we have:

$$R_j = \log(S/P), \qquad \frac{dR_j}{dS} = 1/S, \qquad \frac{d^2 R_j}{dS^2} = -1/S^2 , \tag{31}$$

which indicates that $V_{SS} \to 0$ as the asset level becomes large. In practice, the volatility derivative contract would need to specify how future returns would be computed in the event that the asset price became zero. However, the lower boundary is an outflow boundary [22] and using the approximation $V_{SS} = 0$ at $S = S_1$ will not affect the solution near $S = S(0)$, $t = 0$, assuming that the computational domain is sufficiently wide [14].
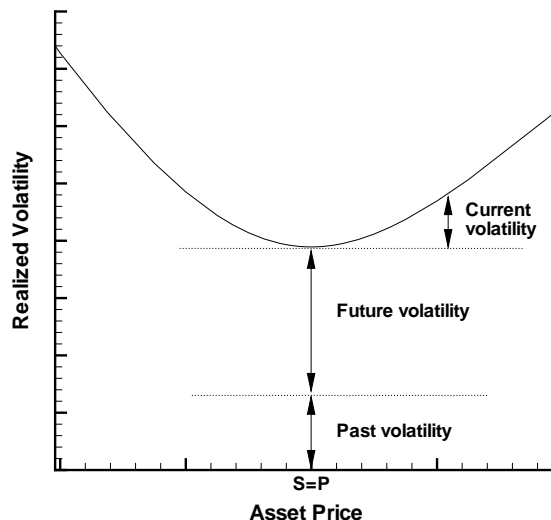
FIGURE 1: *Heuristic decomposition of the realized volatility in terms of the past, current and future volatility samples, at a time immediately preceeding a volatility observation.*

4 Pricing Volatility Exposure

Now that we have described numerical methods for pricing these contracts we can investigate the impact of various modelling assumptions and contractual designs on the fair value of these products. Specifically, we would like to determine how robust the pricing and hedging results are against changes in our assumptions regarding the modelling of the underlying asset price movements. Also, we would like to understand the effect that variations in the contract design will have on the pricing of these products.

4.1 Effect of the Underlying Asset Price Model

In this section we compare the value of the volatility swap assuming a jump-diffusion model, a local volatility function model with no jumps, and a constant volatility model with no jumps. We consider a market where the underlying asset price contains possible jumps and that these jumps are priced into a market of available options. The options market consists of European call and put options with strikes spaced by $\Delta K = \$10$ and maturities spaced by $\Delta T = .1$ year, or approximately one month. We assume that the writers of options in this market use the risk-adjusted parameters $\lambda = .1$, $\mu_J = -.9$, $\gamma_J = .45$ and $\sigma = .2$ to price these instruments and charge the fair value.[1] This defines a market consistent with the jump-diffusion model parameters given above. In order to facilitate comparisons between the various models of the underlying asset, we calibrated a local volatility function[2] as described in [6], and a constant implied

---

[1]These parameters are approximately the values reported in [1], which the authors found were implied in a certain set of S&P options market prices.

[2]Source: the local volatility function was computed using the `Calcvol` volatility surface calibration program developed at Cornell University.
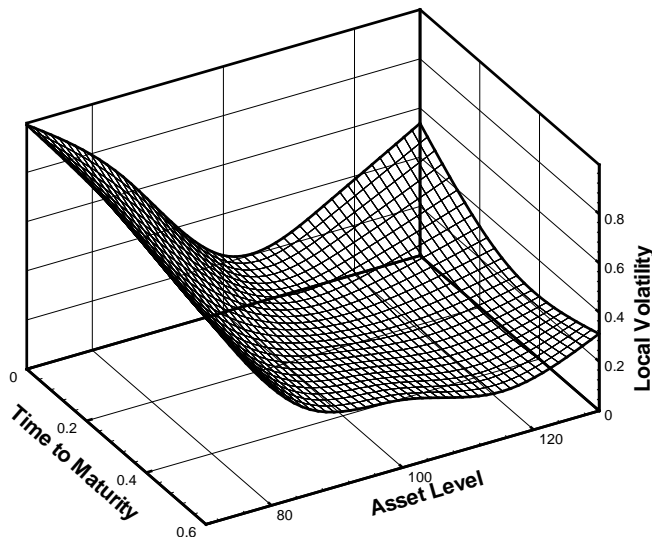
9

FIGURE 2: *Local volatility function computed to match the prices of call and put options in a synthetically generated market. The options were priced assuming $r = .05$, $\sigma = .2$ with jump parameters $\lambda = .1$, $\mu_J = -.9$ and $\gamma_J = .45$, $S(0) = \$100$.*

volatility to these market prices of vanilla options.

Although jump-diffusion models have recently been gaining popularity, solving the PIDE (21) for exotic options requires advanced numerical software and is somewhat more complex than the techniques required to value exotic options in the standard Black-Scholes framework without jumps. As a result it is common to use local volatility surfaces in order to price exotics consistently with observed market prices. If we calibrate a local volatility function consistent with the option prices observed in our synthetic market, the resulting local volatility function is as shown in Figure 2. The local volatility function exhibits the skewed smile that is often observed in options markets, which flattens off for longer maturities.

Even simpler than using a local volatility function, we can consider matching a single constant implied volatility, σ_{imp}, using an at-the-money option with the same maturity as the volatility swap we are pricing. We find that an implied volatility of $\sigma_{imp} = 0.25046$ matches the price of an at-the-money option in our synthetic market.

We now have three possible models for the underlying asset price that are all plausible given currently observed market prices. In practice, the person hedging the volatility derivative would not know which of these (or other) models truly generates the underlying price process and would need to choose among them. Here we briefly discuss some of the similarities and differences that can occur in the valuation and hedging of volatility products under these different models.

In Figure 3(a) we see that there are some qualitative properties that hold for all of the models for the underlying asset process. All models have a minimum occurring near the initial asset level (corresponding to the previous asset price during the first

10
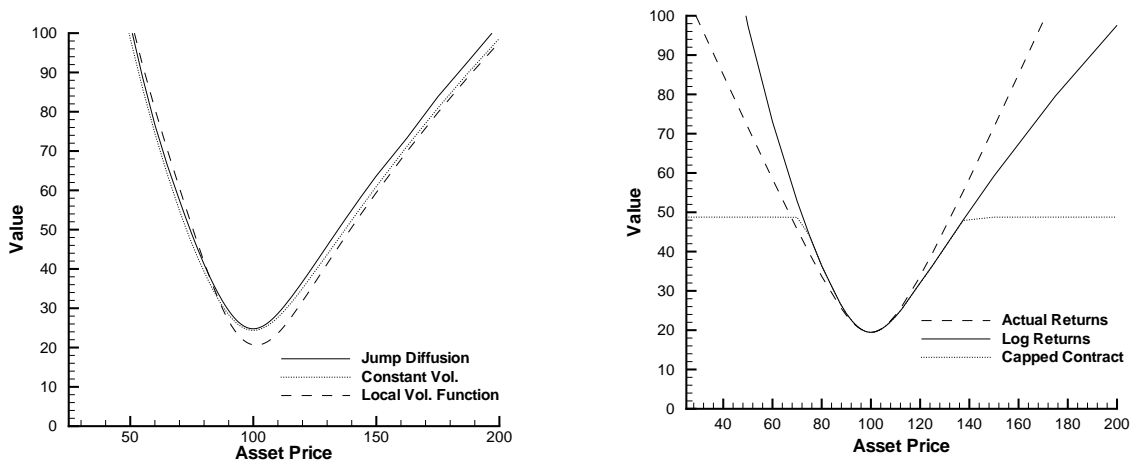
(a) *The effect of assumptions regarding the underlying asset price process. The local volatility function and constant implied volatilities were chosen to be consistent with the pricing of vanilla call and put options under the jump-diffusion process, which utilized $\sigma = .20$, $\mu_J = -.9$, $\gamma_J = .45$ and $\lambda = .1$.*

(b) *The effect of the contractually defined return on the fair value of a volatility swap contract. The capped contract used log returns with a maximum realized volatility of $\sigma_{R,max} = .50$. The underlying asset price followed geometric Brownian motion with $\sigma = .20$, $S(0) = \$100$, and no jumps.*

FIGURE 3: *The volatility swap payoff was calculated using standard realized volatility, $T = .5$, $K_{vol} = 0$ and $B = 1$ with daily observations, $\Delta t_{obs} = .004$. The initial asset price was $S(0) = \$100$ and the risk-free rate was $r = .05$.*

volatility sample). As one moves away from the previous asset level, the value of the volatility swap increases because more volatility will accrue during the current volatility sample. Looking at the slope, which corresponds to the delta hedging parameter, we see that a delta hedging strategy will hold a long position in the underlying asset if $S > P$ to protect against further increases in the asset price. Similarly, a delta hedging strategy will hold a short position in the underlying asset if $S < P$ to protect against the volatility accrued if the asset value decreases further.

As we would expect, there are some quantitative differences between the valuations obtained using the different models for the underlying asset. Although the constant volatility model and the jump-diffusion model give very similar solutions, the local volatility function model gives somewhat different results. This is because the local volatility function behaves as if the volatility is state dependent, and from Figure 2 we see that the local volatility function imposes a higher volatility when the asset price is either well below or well above $S(0) = \$100$. Although the valuations, and hence the implied hedging positions, differ slightly for the various underlying asset models, the qualitative properties, and the general hedging results given in Section 5 based on these qualitative properties, continue hold for different models of the underlying asset.

4.2 The Influence of Product Design on Pricing

In this section, we investigate the impact of variation in the design of the contract on the fair volatility delivery price. Specifically, we investigate the differences caused by

11

Jumps	Sampling frequency	Return type	K_{vol} (volatility points)
No	Daily	Log	19.961
		Actual	19.961
		Capped	19.961
No	Weekly	Log	19.806
		Actual	19.835
		Capped	19.806
Yes	Daily	Log	25.440
		Actual	23.052
		Capped	21.354
No	Daily	Log ($\sigma_{R,stat}$)	19.961
	Weekly	Log ($\sigma_{R,stat}$)	19.794

TABLE 1: *The impact of variations in the contract definition on the fair forward delivery price. The capped contracts specified a maximum realized volatility of $\sigma_{R,max} = .50$ with log returns. Unless mentioned otherwise, the volatility swap specified standard calculation of realized volatility, $T = .5$, and $B = 1$. For daily observations $\Delta t_{obs} = .004$ while for weekly observations $\Delta t_{obs} = .02$. The risk-free rate is $r = .05$ and a constant volatility of $\sigma = .20$ was used. The experiments that included jumps in the underlying asset price specified $\lambda = .1$, $\mu_J = -.9$ and $\gamma_J = .45$.*

the definition of return, the frequency of observation and the impact of whether or not the mean is included in the calculation of volatility. The numerical computations given in this section were performed using a sufficiently fine discretization that the solutions are accurate to within approximately $\pm.001$.

In Table 1 we see that when the volatility is sampled very frequently (i.e. daily or weekly) and the asset price evolves continuously, the definition of the return has very little impact on the fair value of the contract. The capped contract used log returns and the total realized volatility was limited to a maximum of $\sigma_{R,max} = .50$. In the simulations that were carried out, with daily and weekly sampling the cap was sufficiently large that it never affected the payoff when there were no jumps in the underlying asset price. Figure 3(b) illustrates the differences between the realized volatility when the contract specifies log returns, actual returns and a cap on the fair value of these contracts. The sampling frequency has a larger impact on the fair value of these contracts, with differences between weekly and daily sampling occurring in the third digit. As a result, hedging strategies based on a continuously observed volatility may become less effective for longer sampling intervals.

If the underlying asset price jumps then the differences between log, actual and capped returns become more noticeable. In Table 1 we see that the capped contract is less affected when we introduce a jump component to our simulation model. In this case we have introduced jumps according to a Poisson process with intensity, $\lambda = .1$. If a jump occurs, the size of the jump is drawn from a lognormal distribution with mean, $\mu_J = -.9$, and standard deviation, $\gamma_J = .45$. Notice in Figure 3(b) that the value of the contract using log returns increases more quickly than the value of the contract using actual returns when $S \ll P$. Since on average the jumps are downward,

contracts defined using log returns are the most dramatically impacted by the jump component.

At the bottom of Table 1 we investigate the impact of statistically defining the realized volatility as in equation (5) compared with the more common standard definition of the realized volatility given by equation (6). We find that for daily sampling the differences are minimal, affecting the fifth digit. As the sampling becomes less frequent, there is more difference between the fair values of the contracts depending on whether or not the mean is included in the calculation of the volatility. For example with weekly sampling, the effects of whether or not the mean in included lie in the fourth digit.

5   Hedging Volatility Exposure

We will see that hedging volatility swap contracts is more difficult than hedging simple vanilla call and put options. There are two standard dynamic approaches that we can use to hedge these contracts; delta hedging and delta-gamma hedging. In this section we look at the relative merits of each of these approaches and investigate the performance of these hedging methods considering the effects of transaction costs and jumps in the underlying asset price.

In this paper we computed the delta and gamma hedging parameters using a sufficiently fine mesh during the numerical PDE computations such that further refinements did not appreciably affect the hedging results provided in this section. We then performed simulations where the asset path evolved according to (10) using non-risk-adjusted parameters (i.e. in the physical measure) in order to investigate the performances of the various hedging strategies. The profit and loss (P&L) is the value of the hedging portfolio less the value of the payout obligation for the short volatility swap at the maturity of the contract. For each simulation study we provide the expected profit (or loss if negative), the standard deviation of the P&L distribution and the 95% conditional value at risk (CVaR) which is the average of the worst 5% of the outcomes in the P&L distribution. The CVaR measure satisfies certain axiomatic properties [18] that are consistent with the notion of risk. It has also been recognized as a more robust measure downside risk than standard value at risk (VaR) when the profit and loss distribution has fat tails [18].

5.1   Model-Independent Hedging Results

There are two main model-independent results that we focus on in this section. These are imposed on us by the structure of the volatility contract and consequently hold for general models of the price movements by the underlying asset. First, we demonstrate that discretely observed volatility derivative products require very frequent rebalancing. Second, we offer suggestions as to appropriate hedging instruments based on the profile of the realized volatility during the current observation.

5.1.1   Frequency of Rebalancing

Consider the situation of the investor who is short the floating leg of the volatility swap. In theory, one can delta hedge risk exposure to a short position in a derivative contract

| | | Profit and Loss Distribution | | |
Hedge type	Δt_{hedge}	Mean	Std. dev.	95% CVaR
No hedge	None	-.005	1.27	-2.66
Delta hedge	Δt_{obs}	-.002	1.26	-2.65
	$\Delta t_{obs}/2$	-.003	.89	-1.86
	$\Delta t_{obs}/4$	-.002	.63	-1.32
Delta-Gamma hedge	Δt_{obs}	.004	.03	-.05

TABLE 2: *Statistics of the profit and loss distribution of a discretely hedged, short volatility swap position. The volatility swap specified log returns, no mean, $T = .5$, $K_{vol} = 19.961$, $\Delta t_{obs} = .004$, and $B = 1$. It was assumed that $r = .05$, $\mu = .1$ and $\sigma = .2$. The numerical computations were obtained from Monte Carlo experiments using 1,000,000 simulations.*

with value $V(S,t)$ by holding $V_S(S,t)$ shares in the underlying asset at all times. This strategy can be viewed as setting up a local tangent line approximation to the value of the volatility swap. In practice, we define a regular hedging interval, Δt_{hedge}, and adjust the hedging position at $t_h = h\Delta t_{hedge}$, $h = 1, 2, \ldots, n_h$, where $n_h = \lfloor T/\Delta t_{hedge} \rfloor$. In order to delta hedge over the time interval $[t_h, t_{h+1})$, the investor holds $V_S(S(t_h), t_h)$ shares of the underlying asset. In order for the discrete delta hedging strategy to be accurate we need to choose Δt_{hedge} sufficiently small so that:

$$V_S(S(t_h), t_h; P(t_h), Z(t_h)) \approx V_S(S(t), t; P(t), Z(t)) \tag{32}$$

for all $t \in [t_h, t_{h+1})$, where we have explicitly written the dependence of the underlying asset price and state variables on time. Since the state variables change at the volatility sampling times, we require that the delta hedging interval cannot be longer than the volatility sampling period, $\Delta t_{hedge} \leq \Delta t_{obs}$.

To illustrate the fact that very frequent rebalancing is required for discretely observed volatility derivative contracts, we consider a very simple Black-Scholes setting with a constant volatility model. Focusing on the middle section of Table 2 we see that the delta hedging strategy must be rebalanced four times per observation in order to substantially reduce the risk when compared with the unhedged position. This excessive rebalancing makes delta hedging appear to be inappropriate for these contracts. In a more realistic non-constant volatility model we would need to delta hedge the current volatility exposure as well as manage changes in the level of volatility, making this hedging approach even less viable. In the next section we will consider a more flexible delta-gamma hedging strategy. We will find that this hedging strategy can provide good performance even if we only rebalance our hedging positions at the volatility observation times.

5.1.2 Appropriate Hedging Instruments

We have seen that delta hedging strategies must be rebalanced much more frequently than the volatility sampling frequency. We now investigate the structure of the updating rules for the state variables in order to gain insight as to why the underlying asset is not an appropriate hedging instrument. One reason for this is illustrated in Figure 4. In this figure we see that when $S = P$ the tangent to the curve denoting the
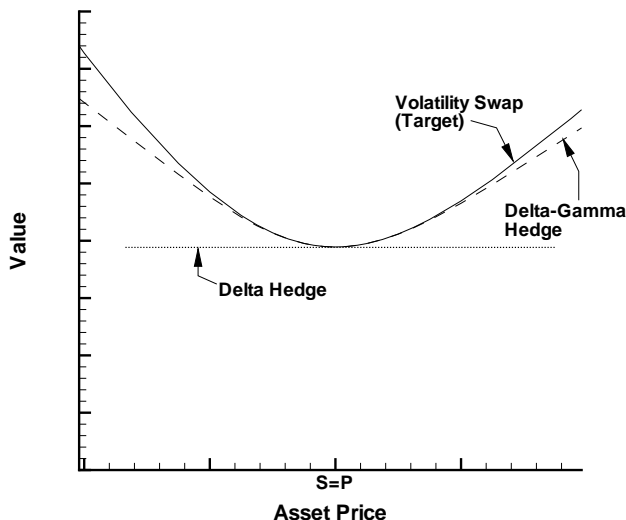
FIGURE 4: *Demonstration of the ability of a delta-gamma hedging strategy to match the value profile of a daily sampled volatility swap. The delta-gamma hedge was constructed using an at-the-money straddle position as the secondary hedging instrument. The delta hedge takes no position in the underlying asset and is unable to hedge against price movements in either direction.*

value as a function of underlying asset price is horizontal. This indicates that $V_S \approx 0$ and that the delta hedge does not take a position in the underlying asset at this time. Unfortunately, most of the time $S \approx P$ since the previous asset level is set to the current asset level at each volatility observation date. This is evident in Table 2 where we see that delta hedging only at the volatility sampling times yields almost identical results to the situation where the writer elects not to hedge the volatility product at all. The underlying asset is not flexible enough to *simultaneously* hedge the volatility that would be accrued if the asset price moved in *either* direction. As a result, in order to delta hedge our volatility exposure we will need to adjust our hedging positions much more frequently than the volatility sampling frequency.

Looking at Figure 4, we see that the value of the volatility swap attributed to the current sample is quite similar to the payoff of a straddle position struck at the previous asset level. If the underlying asset price moves away from the previous asset level in either direction, then this sample will accrue a positive amount towards the final realized volatility. As a result, we suggest constructing straddle or out-of-the-money strangle positions at each volatility observation. Although this will still involve rebalancing at each volatility observation, the positions taken will be quite small since we are only hedging the volatility that accrues over the current volatility sampling period.

In order to hedge a short position in the volatility derivative with price given by V, a delta-gamma hedging strategy holds positions x_1 in the underlying asset and x_2

in appropriate short-term options according to:

$$x_1 = V_S - x_2 I_S \ , \tag{33}$$

$$x_2 = \frac{V_{SS}}{I_{SS}} \ , \tag{34}$$

where I_S, I_{SS} are the delta and gamma respectively of the secondary instruments. We will choose the secondary instruments so that I_{SS} is large enough so that the position in the secondary instruments given by (34) does not become too large. Although the weights of the hedging instruments have been chosen to locally match the delta and gamma of the product we are hedging, we have also chosen the secondary instruments to be consistent with the far-field behaviour. As a result, if there are large asset price swings, the proposed hedging strategy qualitatively matches the target profile. In Section 5.4 we will see that this strategy is closely related to a hedging strategy for variance swaps that utilizes a log contract.

We assume that the writers set up their hedging positions using short term, exchange traded options. Exchange traded options tend to have a fixed range of available strike prices. In our experiments, it was assumed that the strike prices of available options used as secondary instruments were spaced by $\Delta K = \$10$ and the initial asset price was $S(0) = \$100$. The delta-gamma hedging strategy constructs either straddle or out-of-the-money strangle positions at each volatility observation using the strike prices nearest the current asset price while attempting to maintain a roughly symmetric risk exposure to large price movements. Specifically, if $K_i \leq S \leq K_{i+1}$ at time $t_{obs,j}$, then we construct:

- A straddle position with strike K_i if $S - K_i < .2\Delta K$.

- A straddle position with strike K_{i+1} if $K_{i+1} - S < .2\Delta K$.

- An out-of-the-money strangle position using put options with strike K_i and call options with strike K_{i+1} otherwise.

In order to avoid excessive transaction costs, once we establish an out-of-the-money strangle position, we will only change the secondary hedging instruments if the asset level moves beyond the strike prices of either the call or put option.

In our experiments we assume that the options in the market mature at approximately monthly intervals where $\Delta T = .1$ year. In general we use the shortest term options whose maturity date is later than the next volatility observation since short term options have a higher ratio of gamma to value, which will be useful in reducing transaction costs. However, as we near the maturity date of the secondary options their gammas become too localized around the strike prices and we choose to restrict ourselves to using options with a minimum remaining time to maturity of half of a month, i.e. $T - t \geq .05$ years.

In Table 2 we see that the delta-gamma hedging strategy performs very well relative to the delta hedging strategy. If we only adjust the delta-gamma hedge at the volatility observations, the standard deviation of the profit and loss distribution is reduced by a factor of over 20 when compared with a delta hedging strategy that is re-balanced four times per volatility sampling period. We refer to this delta-gamma hedging strategy as a semi-static hedge because it constructs small, fixed positions at each volatility observation which are not adjusted until the next volatility sampling date.

16

		Profit and Loss Distribution		
Hedge type	Δt_{hedge}	Mean	Std. dev.	95% CVaR
No hedge	None	4.897	7.59	-8.33
Delta hedge	$\Delta t_{obs}/4$	4.896	7.52	-7.21
Delta-Gamma hedge	Δt_{obs}	2.364	4.82	-2.25
MVO hedge (underlying)	$\Delta t_{obs}/4$	4.622	7.29	-8.70
MVO hedge (underlying,puts,calls)	Δt_{obs}	.548	3.31	-2.21

TABLE 3: *Statistics of the profit and loss distribution of a discretely hedged, short volatility swap position when there are jumps in the asset price. The volatility swap specified log returns, no mean, $T = .5$, $K_{vol} = 25.440$, $\Delta t_{obs} = .004$, and $B = 1$. In the hedging simulations it was assumed that the non-risk-adjusted parameters were $r = .05$, $\mu = .1$, $\sigma = .2$, $\lambda = .02$, $\mu_J = -.45$ and $\gamma_J = .45$. The numerical computations were obtained from Monte Carlo experiments using 100,000 simulations.*

5.2   Hedging in a Jump-Diffusion Setting

We now imagine that the asset price occasionally jumps discontinuously and investigate the impact of jumps on our ability to hedge these contracts. In [1] the authors found that the risk-adjusted jump parameters, $\lambda^{ra} = .1$, $\mu_J^{ra} = -.9$ and $\gamma_J^{ra} = .45$ were implied in a particular set of S&P option prices. Valuing our volatility contract consistently with these vanilla instruments gives a volatility delivery price of $K_{vol} = 24.440$. The arrival intensity and typical jump sizes given by these implied parameters are much larger than those given by historical time series data. In [1] the authors argue that $\lambda^h = .02$, $\mu_J^h = -.45$, $\gamma_J^h = .45$ are more appropriate estimates of the jump parameters under the physical measure.

In Table 3 we compare the performances of various hedging strategies under the physical measure. We see that the variability of the hedged position measured by the standard deviation of the profit and loss distribution is much larger when there are possible jumps in the underlying asset price. The fair volatility delivery price is such that on average the profit and loss of an unhedged position in a risk-neutral setting has mean zero. In the physical measure the expected P&L is positive because of the risk aversion built into the implied jump parameters. However, the CVaR indicates that occasionally the writer experiences a very large loss in the relatively rare event that a jump occurs.

It is interesting to notice that delta hedging, even with very frequent rebalancing, does very little to reduce the downside risk associated with hedging these contracts when there are jumps. In fact, looking at the CVaR we see that the worst case outcomes when delta hedging are only marginally better than the worst case outcomes when the volatility swap is not hedged. To see why this is the case, consider the following situation:

- Suppose the asset price at the previous volatility observation was $100.

- After the volatility observation, the asset price rises and the delta hedging strategy takes on a positive position in the underlying asset to hedge against further increases in the asset level before the next volatility observation.

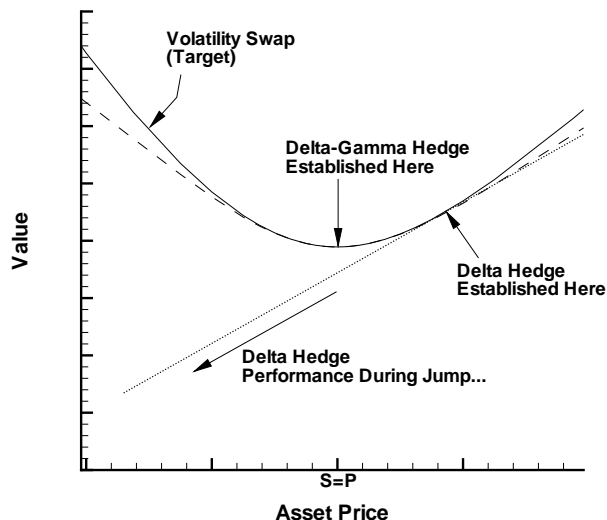- There is a large downward jump in the asset price.

FIGURE 5: *Illustration of the performance of a delta and a delta-gamma hedging strategy when a jump occurs in the underlying asset price.*

In this case, the writer will face a hit in the short realized volatility position due to the large downward jump and, to make matters worse, the attempted hedging position (consisting of a long position in the underlying asset) will have also decreased in value. This situation is illustrated in Figure 5.

We contrast this situation with a delta-gamma hedging strategy which sets up straddle/strangle positions at the beginning of the volatility observation as described in the previous section. In Table 3, the delta-gamma hedged position still offers significant risk reductions over not hedging these contracts at all. In Figure 5 we see that the possible jumps in the underlying asset have much less negative impact when a delta-gamma hedging strategy is implemented because the secondary instruments have been chosen to qualitatively match the far-field behaviour of the volatility derivative contract.

The delta and delta-gamma hedging strategies try to reduce risk by locally matching the sensitivities of the hedging and target portfolio to changes in the underlying asset price. In an attempt to improve the performance of the hedges we can choose our positions in the hedging assets in order to minimize the variance of the partially hedged position as described in [10, 20, 19]. Each time the hedge is rebalanced, we solve an optimization problem and select our hedging positions so that the variance of the hedged position is minimized. In Table 3 we see that the minimum variance optimal (MVO) hedge using only the underlying asset as a hedging instrument still offers very little risk improvement over the unhedged position. This provides a further demonstration that hedging using only the underlying asset is inappropriate for managing the risk associated with writing volatility derivatives. The MVO hedge using the instruments used in our delta-gamma hedge offers somewhat better risk reduction compared with the local delta-gamma strategy. The MVO hedge can select the weightings in the puts and calls that comprise the strangle position independently in order to better match

18

the target profile of the contract that we are hedging.

5.3   Hedging with a Bid-Ask Spread

We now investigate the impact of transaction costs on the valuation and hedging of these contracts. Specifically, we assume that the hedger incurs transaction costs due to a bid-ask spread. We define the one-way transaction cost loss due to trading in the underlying asset to be:

$$\kappa = \left(\frac{1}{2}\right)\frac{S_{ask} - S_{bid}}{S_{ask}} \ . \tag{35}$$

Typically, the bid-ask spread for liquidly traded assets is quite small and in our experiments we use $\kappa = .001$ or 10 basis points. On the other hand, the bid-ask spread for exchange traded options can be quite large, and typical values for the transaction cost parameter for the secondary instruments would be around $\kappa_{SI} = .05$.[3]

When there are transaction costs [21] we replace (21) with:

$$V_t + rSV_S + \frac{1}{2}\sigma^2 S^2 V_{SS} - \kappa\sigma S^2\left(\sqrt{\frac{2}{\pi\Delta t_{hedge}}}\right)|V_{SS}| - rV = 0 \ . \tag{36}$$

The new term containing $|V_{SS}|$ estimates the expected costs of changing the delta hedged position at the end of the hedging interval. In general this equation is non-linear and must be solved numerically. However, when actual returns are specified the gamma, V_{SS}, is always positive and we can simply use the Leland volatility correction [15]:

$$\sigma_{Leland} = \sigma\left(1 + \left(\sqrt{\frac{8}{\pi\Delta t_{hedge}}}\right)\frac{\kappa}{\sigma}\right)^{1/2} \ . \tag{37}$$

Even when log returns were specified, the regions where the gamma changes sign are so far away from the region of interest (see Figure 3(b)) that we could not notice any differences between the solution computed using (36) compared with the solution computed using (21) with (37).

Assuming that $\kappa = .001$ and re-balancing the delta hedged position four times per volatility observation, $\Delta t_{hedge} = .001$, the cost of hedging the realized volatility is \$21.786 at time $t = 0$, giving a fair delivery price at maturity of $K_{vol} = 22.338$. Comparing with the fair delivery price without considering transaction costs given in Table 1, we see that the expected transaction costs are \$2.377. In other words, approximately 10% of the value of the delivery price is lost through hedging transaction costs. If we do choose not to hedge the product, then at maturity we expect to have approximately \$2.377 in profit, at the expense of the additional risk we take on by not hedging. This is very close to the expected excess of the unhedged position in Table 4.

---

[3]An alternative to exogenously specifying the option bid-ask spread would be to determine the inferred spread in terms of the bid-ask spread for the underlying asset using a transaction cost model. However, the rebalancing interval used to hedge the vanilla instruments would probably be much longer than the rebalancing interval used to hedge the volatility derivative contract, making the implied bid-ask spread somewhat arbitrary. Instead, we choose κ_{SI} to be representative a typical options market.

| | | | Profit and Loss Distribution | | |
Hedge type	ΔK	κ_{SI}	Mean	Std. dev.	95% CVaR
No hedge	N/A	N/A	2.375	1.26	-.27
Delta hedge	N/A	N/A	-.427	.68	-1.84
Delta-Gamma hedge	$10	.05	-3.564	.85	-5.35
	$10	.001	.843	.15	.50
	$20	.001	.991	.23	.55

TABLE 4: *Statistics of the profit and loss distribution of a discretely hedged, short volatility swap position when there are transaction costs. The volatility swap specified log returns, no mean, $T = .5$, $K_{vol} = 22.338$, $\Delta t_{obs} = .004$, and $B = 1$. It was assumed that $r = .05$, $\mu = .1$, $\sigma = .2$, $\kappa = .001$ and $S(0) = \$100$. The numerical computations were obtained from Monte Carlo experiments using 1,000,000 simulations.*

If we attempt to delta hedge the contract we find that the profit and loss distribution has a negative mean value, indicating that equation (36) has underestimated the expected transaction costs. To explain this, recall that the Leland transaction cost model uses a gamma approximation in order to determine the expected change in the position of the hedging asset. In Figure 4 we see that this second-order approximation underestimates the curvature, and hence the expected transaction costs, for larger asset price movements. The risk associated with the delta hedged position is similar to that found in Table 2 in the absence of a bid-ask spread.

To this point, the delta-gamma hedging strategy that involves setting up straddle-strangle positions at each volatility observation has performed very well. However, the bid-ask spread on exchange traded options is quite large and typical values might be around $\kappa_{SI} = .05$. We see in the bottom row of Table 4 that when faced with these large transaction costs, the delta-gamma hedging strategy is no longer feasible.

On the other hand, it might be plausible to assume that the institution hedging the volatility swap has clients who can act as natural counterparties to the positions in these vanilla options. In this situation we assume a much smaller transaction cost associated with trading in the secondary instruments, $\kappa_{SI} = .001$. Now the delta-gamma hedging strategy incurs fewer transaction costs and has less variability than the delta hedged position with transaction costs. In fact, now approximately half of the additional mark-up due to the transaction costs is now retained as profit and the delta-gamma hedging is so effective that even the 95% CVaR is positive.

When there are transaction costs there are tradeoffs that must be considered in the selection of the hedging instruments. At the bottom of Table 4 we compare the performance of the delta-gamma hedging strategy when $\Delta K = \$10$ and when $\Delta K = \$20$. We see that by choosing wider strangle positions we expect to incur less transaction costs since we will change the structure of the positions less frequently. This is at the expense of a moderate increase in variability because the wider strangle positions are less precise.

5.4 Hedging Using a Log Contract

Much of the existing work on volatility based derivatives [8, 16, 4] has focused on variance swaps. We now summarize one of the important results given in [8], which can be used to price and hedge variance swaps. Assuming that the underlying asset evolves continuously according to:

$$dS = \mu S \, dt + \sigma(S, t) S \, dW \tag{38}$$

where μ is the drift rate of the underlying asset in the physical measure, $\sigma(S, t)$ is the volatility function and dW is an increment from a Wiener process. Using Ito's lemma we see that:

$$d(\log(S)) = \frac{dS}{S} - \frac{1}{2}\sigma^2 dt \ . \tag{39}$$

Integrating both sides with respect to time we find that:

$$\log\left(\frac{S(T)}{S(0)}\right) = \int_0^T \frac{dS}{S} - \frac{1}{2}\int_0^T \sigma^2 dt \ . \tag{40}$$

Rearranging we find that the continuously observed variance is given by:

$$\sigma_{R,cts}^2 = \frac{1}{T}\int_0^T \sigma^2 dt \tag{41}$$

$$= \frac{2}{T}\left[\int_0^T \frac{dS}{S} - \log\left(\frac{S(T)}{S(0)}\right)\right] \ . \tag{42}$$

We refer to this hedging strategy as a semi-static hedge, since it consists of a static short position in a log contract[4], and a dynamic position in the underlying asset that is always instantaneously long $1/S(t)$ units worth \$1, with a scaling factor of $2/T$. In fact, we will see that this semi-static hedging strategy is very closely related to the delta-gamma hedging strategy that we discussed earlier for volatility derivative products using out-of-the-money strangle positions. It is interesting to note that the pricing and hedging of a variance swap given by (42) is independent of the form of the volatility function $\sigma(S, t)$, in the sense that this has been appropriately priced into the log contract.

There are several issues that one encounters when we consider hedging volatility derivative products using this log contract formulation.

- This pricing and hedging strategy will only be accurate if $\sigma_{R,cts}^2$ is a good estimation of the discretely sampled variance defined in the contract. We would expect this to be the case if Δt_{obs} is sufficiently small, such as daily sampling, but it may not hold for longer sampling intervals, such as weekly samples. We will look at how the hedging strategy using the log contract should be extended to handle discretely observed variance sampling.

- Log contracts are not traded and must be synthetically created. In [8] the authors demonstrate how to approximate a log contract using a static hedge in traded options with a discrete set of available strikes. In order to hedge volatility derivative products we need to hedge a square root derivative on the variance swaps, which

---

[4]A log contract [17] is simply a derivative product whose payoff at time T is given by $\log(S(T)/S(0))$.

will involve dynamic trading in these log contracts. Since synthetically replicating the log contract requires trading in a large number of exchange traded options, this may result in excessive transaction costs. We will find that the delta-gamma hedging strategy discussed previously in this paper can be viewed as an approximation of the log contract hedging strategy that reduces the number of positions taken in the exchange traded options.

- Finally, the derivation of this hedging strategy for variance swaps assumes that there are no jumps and that the underlying asset price evolves in a continuous manner. Consistent with ideas in [8], we expect that the log contract hedge will offer reasonable performance even when there are possible jumps in the asset price due to its connection with the delta-gamma hedging strategy.

We now investigate the connection between the hedging strategy for variance swaps using a log contract with the delta-gamma hedging strategies discussed in Section 5.1.2. We will then look at how we can efficiently modify these techniques to handle the direct hedging of volatility derivative products.

5.4.1 Connection to the Delta-Gamma Hedging Strategy

In order to see the connection between the hedging strategy using the log contract and the delta-gamma hedging strategy discussed earlier in this paper we can write a discrete representation of (42) as:

$$\sigma_{R,cts}^2 \approx \frac{2}{T} \sum_{h=1}^{N} \left[\frac{S(t_h) - S(t_{h-1})}{S(t_{h-1})} - \log\left(\frac{S(t_h)}{S(t_{h-1})}\right) \right] \tag{43}$$

where $t_h = h\Delta t_{hedge}$, $h = 1, 2, \ldots, n_h$, where $n_h = \lfloor T/\Delta t_{hedge} \rfloor$. In the limit as $\Delta t_{hedge} \to 0$ this becomes identical to (42) where we have conceptually expanded the log contract into an equivalent sequence of log contract positions.

In order to replicate the accrued variance during the time interval $[t_{h-1}, t_h)$, we construct the portfolio, consisting of $2/(TS(t_{h-1}))$ shares of the underlying asset and a short position of $2/T$ in a contract whose payoff at time t_h is given by $\log[S(t_h)/S(t_{h-1})]$. If we let $d\Pi_{h-1}$ represent the value of the payoff received at time t_h of this replicating portfolio established at time t_{h-1}, then:

$$d\Pi_{h-1} = \frac{2}{T} \left[\frac{S(t_h) - S(t_{h-1})}{S(t_{h-1})} - \log\left(\frac{S(t_h)}{S(t_{h-1})}\right) \right] . \tag{44}$$

We find [8] that the position in the underlying asset exactly neutralizes the delta of the log position and that a second-order approximation gives:

$$E^Q \left[d\Pi_{h-1} \big| \mathcal{F}(t_{h-1}) \right] = \frac{\sigma^2 \Delta t_{hedge}}{T} . \tag{45}$$

As $\Delta t_{hedge} \to 0$ this becomes the continuous variance over the sampling period $[t_{h-1}, t_h]$.

We can think of the log contract hedge as being a delta-gamma hedge with a clever choice of hedging instruments. Because the gamma of the log contract scales with $1/S^2$, and because the payoff of the variance swap contract depends linearly on the individual accrued variances, our position in the log contract does not need to change

22

throughout the life of the contract. As the asset price process evolves, we simply need to adjust our position in the underlying asset so that the hedging position is delta neutral and so that the instantaneous variance is captured.

If the contract specifies a *discretely* observed variance, then as suggested in [8] we can imagine using the same short position in the log contract to hedge our gamma exposure, except that now we would only adjust our position in the underlying asset to make our position delta neutral at the sampling times. In other words, we would use the hedging strategy implied by (43) where we set $\Delta t_{hedge} = \Delta t_{obs}$. Although this hedging strategy is no longer exact, we saw in Table 2 that this second-order approximation can provide very significant risk reductions if the sampling occurs quite frequently.

5.4.2 Using Log Contracts to Hedge Volatility Derivative Products

So far, we have discussed ways to use log contracts to hedge variance exposure. When hedging volatility exposure we need to hedge a square root contract on the underlying variance. This will involve adjusting our positions in the log contract as the realized variance to date fluctuates. In [8] the authors describe how to replicate a log contract using out-of-the-money options and many of these contracts are far away from the current asset level. Since our holdings in the log contract are uncertain, if there are transaction costs it makes sense to avoid holding these far out-of-the-money options until they actually have an impact on the performance of the hedging strategy. Of course there will be some tradeoff between the possibility of facing transaction costs to reacquire positions versus the transaction costs of acquiring positions that are never utilized. The delta-gamma hedging strategies described in this paper using straddle and out-of-the-money strangle positions can be thought of as constructing only the portion of the log contract that is near the current asset price.

6 Conclusions

This paper has focused on several issues concerning the pricing and hedging of volatility derivative products. First, we described a computational framework for pricing volatility products using numerical PDE methods that can be extended to handle a variety of modelling assumptions including local volatility models, jump-diffusion models, and transaction cost models. Using this framework we investigated the effects of assumptions regarding the underlying asset price movements and effects of the contractual design on the pricing of volatility derivatives. We then studied our ability to hedge these products using delta and delta-gamma hedging strategies in a variety of settings.

When we began investigating the hedging of these products, we found it convenient to think of the volatility as decomposing into three parts: the past, current and future realized volatility. We then focused on our ability to actively hedge the volatility that will accrue over the current observation. Dynamic delta hedging is not effective for hedging volatility exposure because it requires extremely frequent rebalancing. Also, if there are jumps in the underlying asset price, then there are situations where delta hedging can increase the risk of the net position. As a result we consider delta-gamma hedging strategies that have been constructed using instruments that have similar

profiles as the volatility product we are replicating; namely straddle and strangle positions using exchange traded options with discretely spaced strikes. These delta-gamma hedging strategies provided excellent risk reduction and were still effective when the underlying asset price contained jumps since the hedging instruments were chosen to qualitatively match the far-field behaviour of the volatility product. We also discussed the close connections between the proposed delta-gamma hedging strategies described in this paper with the log contract hedging strategies for variance swaps.

If there is a large bid-ask spread in the market prices of the exchange traded options then transaction costs may make the proposed delta-gamma hedging strategy infeasible. However, if the institution writing the volatility swap has natural counterparties for their positions in the vanilla options then the delta-gamma hedging strategies proposed here using straddle/strangle positions can be very effective for managing downside risk.

References

[1] L. Andersen and J. Andreasen. Jump-diffusion processes: Volatility smile fitting and numerical methods for option pricing. *Review of Derivatives Research*, 4:231–262, 2000.

[2] O. Brockhaus and D. Long. Volatility Swaps Made Simple. *Risk*, 13(1):92–95, 2000.

[3] P. Carr and K. Lewis. Corridor variance swaps. Working Paper, 2002.

[4] P. Carr and D. Madan. Towards a theory of volatility trading. In *Volatility: New Estimation Techniques for Pricing Derivatives*. Risk Publications, London, 1998.

[5] N. Chriss and W. Morokoff. Market Risk of Variance Swaps. *Risk*, 12(10):55–59, 1999.

[6] T. F. Coleman, Y. Li, and A. Verma. Reconstructing the unknown local volatility function. *Journal of Computational Finance*, 2(3):77–102, 1999.

[7] K. Demeterfi, E. Derman, M. Kamal, and J. Zou. A Guide to Variance Swaps. *Risk*, 12(6):54–59, 1999.

[8] K. Demeterfi, E. Derman, M. Kamal, and J. Zou. A guide to volatility and variance swaps. *Journal of Derivatives*, 6(4):9–32, 1999.

[9] Y. d'Halluin, P. A. Forsyth, and K. R. Vetzal. Robust numerical methods for contingent claims under jump diffusion processes. Working paper, University of Waterloo, School of Computer Science, 2003.

[10] D. Duffie and H. R. Richardson. Mean-variance hedging in continuous time. *The Annals of Applied Probability*, 1:1–15, 1991.

[11] S. L. Heston and S. Nandi. Derivatives on Volatility: Some Simple Solutions Based on Observables. Technical report, Federal Reserve Bank of Atlanta, November 2000. Working paper 2000-20.

[12] S. D. Howison, H. O. Rasmussen, and A. Rafailidis. A Note on the Pricing and Hedging of Volatility Derivatives. Technical report, Financial Mathematics Research Group, King's College, 2002. Working paper.

[13] A. Javaheri, P. Wilmott, and E. G. Haug. GARCH and Volatility Swaps. Working paper, 2003.

[14] R. Kangro and R. Nicolaides. Far field boundary conditions for Black-Scholes equations. *SIAM Journal on Numerical Analysis*, 38(4):1357–1368, 2000.

[15] H. Leland. Option pricing and replication with transaction costs. *Journal of Finance*, 40:1283–1301, 1985.

[16] T. Little and V. Pant. A finite-difference method for the valuation of variance swaps. *Journal of Computational Finance*, 5(1):81–101, 2001.

[17] A. Neuberger. The Log Contract: A new instrument to hedge volatility. *Journal of Portfolio Management*, pages 74–80, Winter 1994.

[18] R. T. Rockafellar and S. Uryasev. Conditional Value-at-Risk for General Loss Distributions. *Journal of Banking and Finance*, 26(7):1443–1471, 2002.

[19] M. Schal. On quadratic cost criteria for option hedging. *Mathematics of Operations Research*, 19(1):121–131, 1994.

[20] M. Schweizer. Mean-variance hedging for general claims. *The Annals of Applied Probability*, 2:171–179, 1992.

[21] P. Wilmott. *Derivatives.* John Wiley and Sons, West Sussex, England, 1998.

[22] H. Windcliff, P. A. Forsyth, and K. R. Vetzal. Analysis of the stability of the linear boundary condition for the Black-Scholes equation. Working paper, University of Waterloo, School of Computer Science, 2003.

[23] H. Windcliff, P. A. Forsyth, K. R. Vetzal, A. Verma, and T. F. Coleman. An object oriented framework for valuing shout options on high-performance computing architectures. *Journal of Economic Dynamics and Control*, 27(6):1133–1161, 2003.