NEW YORK UNIVERSITY
STERN SCHOOL OF BUSINESS
FINANCE DEPARTMENT

Working Paper Series, 1996

Testing the Volatility Term Structure Using Option Hedging Criteria

Engle, Robert F. and Joshua Rosenberg

FIN-96-24

Testing the volatility term structure using option hedging criteria

February 1997

Robert F. Engle
Department of Economics
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0508
(619) 534-1054

Joshua V. Rosenberg
Department of Finance
NYU - Stern School of Business
44 West 4th Street, Suite 9-190
New York, New York 10012-1126
(212) 998-0311

Abstract

The volatility term structure (VTS) reflects market expectations of average asset volatility over different time horizons. Various stochastic volatility models provide forecasts of the VTS and how it shifts in response to changes in market conditions. This paper develops a methodology for testing VTS forecasts using option hedging performance. An innovative feature of the hedging approach is its increased sensitivity to several important forms of model misspecification relative to previous testing methods.

Hedging tests using S&P500 index options indicate that the GARCH components with leverage VTS estimate is most accurate. The poorer hedging performance of the alternative models suggests that volatility term structure shifts are related to the magnitude and level of recent returns. Strong evidence is obtained for mean-reversion in volatility.

The volatility term structure (VTS) reflects market expectations of average asset volatility over a range of time horizons. An accurate model of the VTS and its dynamics is essential in pricing and hedging options. Since the VTS is generated by a particular volatility process, VTS models may be compared by estimating different volatility models and comparing their implied term structures with an objective measure. This is the approach taken by Heynen, Kemna, and Vorst (HKV, 1994), who develop a volatility term structure test that considers a range of candidate volatility models. Their approach measures the closeness of a forecast VTS from a volatility model to a realized VTS based on implied volatilities.

Several other papers including Stein (1989), Diz and Finucane (1993), and Xu and Taylor (1994) directly estimate the volatility term structure using option implied volatilities. These papers do not consider volatility models estimated from the behavior of the underlying asset, such as GARCH models, and thus potentially exclude an important predictor of future volatility.

In contrast to the HKV approach, our paper focuses on testing volatility term structure dynamics, i.e. day-to-day changes in the VTS due to changes in market conditions, using option hedging performance criteria. In particular, VTS tests are developed based on success in predicting relative changes in medium and short term straddle prices due to shifts in the VTS. Since a VTS model provides a forecast of the effect of a volatility shock on short and medium term average volatility, this may be used along with option price sensitivity to a volatility shock, to generate testable implications about relative option price changes. These hedging tests may be applied to a wide variety of stochastic volatility models.

Hedging tests provide several benefits over comparing a forecast VTS to a realized VTS. First, hedging tests focus on the relative levels of forecast average variance over different horizons and are less sensitive to unconditional variance forecasts. So, hedging tests may be able to distinguish among alternative models with different term structure shapes but similar levels of unconditional variance. HKV (1994) find that differences in unconditional volatility forecasts drive their test results.

Second, hedging tests evaluate predicted term structure dynamics rather than the closeness of fit to the term structure levels. In tests of term structure levels, it may be difficult to distinguish among models with similar rates of mean reversion but different explanatory variables. In other words, two volatility models with different information variables may perform similarly in pricing options, but quite differently in hedging options. Hedging tests may be superior at identifying omitted variables or interrelationships in the volatility model, because hedging performance depends on eliminating sensitivity to all of sources of volatility changes.

Third, hedging tests offer interpretable metrics for measuring forecast accuracy, such as value-at-risk, that relate more directly to the objectives of agents using the forecasts. This contrasts with volatility forecast tests that compare realized and forecast variance using a mean-squared-error or similar fit criterion.

In addition to introducing an improved methodology for comparing VTS models, this paper provides a comprehensive empirical comparison of the efficacy of techniques for hedging medium term at-the-money straddles with short term at-the-money straddles. GARCH delta-gamma hedging is found to be one of the most effective option hedging techniques. However, the methodology presented in this paper is only applicable to at-the-money options and may not be directly used to hedge an options book which has options with a variety of moneynesses. Engle and Rosenberg (1995) provide an alternative GARCH option hedging methodology, based on Monte-Carlo simulation, which is applicable to options of any moneyness.

Previous papers concerned with hedging performance are limited by their focus on a single type of stochastic volatility model and reliance on interpolated prices or option values at expiration for empirical hedging results. For instance, Hull and White (1987a) present simulation results that show that under a continuous-time integrated variance process, constant volatility (CV) delta-gamma hedging works better when implied volatility is relatively stable and CV delta-vega hedging works best otherwise. The HW empirical hedging results depend on interpolating OTC option prices using Black-Scholes evaluated at the market implied volatility. Melino and Turnbull (1995) present a method for option hedging under a particular continuous-time stochastic volatility by numerically obtaining deltas and vegas from the pricing partial differential equation. They find evidence that a constant volatility model performs poorly in pricing and hedging long-term currency options under stochastic volatility. However, instead of direct hedging tests involving the day-to-day changes in the hedge portfolio value, only a single option price is used. Model performance is ranked based on the replicating error of a hedge portfolio held until the target option expiration.

In this paper, volatility hedge ratios are developed based on an approximate option pricing formula that holds for at-the-money options for a variety of volatility processes. VTS models for the S&P500 are compared based on performance in hedging the effects of changes in volatility on medium term S&P500 index straddles. All price changes are based on market prices rather than interpolated prices. The evidence presented in this paper indicates that the S&P500 volatility process is mean-reverting with asymmetric (leverage) effects, and that the volatility term structure is sensitive to the magnitude of underlying asset returns. The volatility model with the overall best hedging performance is the GARCH components with leverage model, which incorporates these characteristics. This is a new type of GARCH volatility model in which volatility has both a long-run and a short-run factor that have different degrees of mean-reversion and there is an asymmetric effect that applies to the short-run component.

The importance of asymmetries in the variance process is also found by Amin and Ng (1994) who compare market prices and estimated prices of individual equity options under several GARCH specifications. Amin and Ng (1994) find that the two GARCH models with asymmetric effects outperform a GARCH model with no asymmetries. In tests of Philips and the EOE index, HKV (1994) find that an EGARCH model is preferred to two alternative volatility models with no asymmetries.

This paper is structured as follows. Section 1 presents an approximate option pricing formula applicable to at-the-money options under stochastic volatility, Section 2 details the derivation of stochastic volatility hedging parameters and the stochastic volatility hedge ratios. Section 3 describes the estimation of the S&P500 volatility models and volatility hedge ratios. In Section 4, the five candidate volatility models are tested based on their ability to hedge the effects of volatility changes on medium term S&P500 index straddles. Section 5 concludes the study.

1. Approximate option pricing in a stochastic volatility environment

The Black-Scholes (1973) option pricing model assumes that the underlying asset variance is constant over the life of the option. Substantial empirical evidence has been presented, see e.g. Bollerslev, Chou, and Kroner (1992), that time-dependent stochastic variance characterizes many financial return time-series. A number of option pricing models, such as Hull and White (1987b), Melino and Turnbull (1990, 1995), Amin and Ng (1993, 1994), and Duan (1995) have been developed that allow for stochastic volatility. However, each model provides a different specification for the variance process. The method for selecting the appropriate variance process and thus the correct model for option pricing and hedging is left open.

Since pricing options in a generic stochastic volatility environment is not a solved problem, we utilize an approximate option pricing formula for *at-the-money* options that may be applied to a variety of volatility models. The use of a single approximate option pricing formula for different volatility models facilitates the derivation of the option hedge parameters using a consistent methodology. The five volatility models evaluated in this paper are a constant volatility model (CV), an autoregressive implied volatility model (ARIV), a GARCH(1,1) model, a GARCH(1,1) with leverage model (GJR), and a GARCH components with leverage model (GCOMP).

This paper uses the following approximate pricing formula, which will be referred to as the Black-Scholes-plug-in formula or BSP.

$$(1) \qquad P_t \cong BSP(E_t[\bar{\sigma}_{t,T}(S_t)], S_t, T) \qquad\qquad C_t \cong BSP(E_t[\bar{\sigma}_{t,T}(S_t)], S_t, T)$$

In equation (1), BSP is the Black-Scholes pricing formula for a put or call. P_t and C_t are the call and put premia, S_t is the current underlying asset price, and T is the number of days until option expiration. Dependence on the riskless rate and strike price are suppressed. In this approximate pricing formula, expected average volatility $\bar{\sigma}_{t,T}(S_t)$ is "plugged into" the standard Black-Scholes formula to obtain the stochastic volatility option price. When volatility is constant, BSP simplifies to the Black-Scholes formula with the constant volatility plugged-in.

The potential dependence of expected average volatility on the current underlying price is reflected in the notation. In fact, future volatilities may also depend on the most recent return magnitude which is a function of S_t and S_{t-1}. For simplicity, dependence on S_{t-1} is suppressed in equation (1).

The accuracy of the BSP formula depends on two factors. First, to ensure accuracy, the options to be priced should be at-the-money. This is because the BSP approximation relies on the linearity of the Black-Scholes formula in the volatility parameter for at-the-money options. Second, the effect of volatility risk premia must be small, since average expected volatility rather than average volatility under the risk-neutral measure is used. In addition, conditional log-normality under the risk-neutral measure is necessary to obtain a BSP type option pricing formula. The accuracy the hedge ratios derived under BSP and used in this paper will be verified using simulations under the appropriate risk-neutral measure in section 2.

2. Hedging options in a stochastic volatility environment

This paper develops a methodology for selecting the most accurate volatility term structure based on option hedging performance. Since option prices and option price changes depend on volatility forecasts and volatility forecast changes, the correct underlying asset volatility model and option pricing formula should generate accurate hedge parameters. Thus, success at hedging options is a natural criterion for evaluating alternative volatility models.

In this section, the hedge parameters and hedge ratios are derived that are used in the VTS hedging performance tests. In addition to their use in the hedging tests, the hedge parameters provide insight into how volatility news is incorporated into option prices. Since the methodology for deriving the hedge parameters is general, hedge parameters may be derived for volatility models not considered in this paper.

Hedge parameters measure the sensitivity of option prices to changes in the state variables, and are obtained by differentiating the Black-Scholes plug-in formula with respect to its state-variables. The approximate hedge parameters developed in this section are appropriate for hedging over a short time period, such as one day, and may not perform well for hedging over a longer horizon.

To derive the hedge parameters, it is necessary to identify the sources of random changes in the option price. When a day passes, the option price will change in part because the underlying asset price changes and in part because the average volatility forecast changes. Changes due solely to the passage of time or changes in interest rates will be ignored. The change in the option price due to a change in volatility will depend on the parameters of the volatility process. Thus, tests of success in hedging changes in volatility are indirect tests of the estimated volatility process and its forecast term structure.

The change in the option value due to changes in the state variables may be approximated using a Taylor series expansion. In this case, it is natural to think of expanding the end-of-day option price as a function of the end-of-day state variables. Just as Black-Scholes delta and gamma are derived by taking the first-derivative of the Black-Scholes formula with respect to the end-of-day underlying price (S_t), the BSP delta and gamma are obtained by taking derivatives of BSP with respect to S_t. If volatility (σ) is considered to be a separate stochastic state variable, a volatility hedge parameter may be obtained using a partial derivative of BSP with respect to volatility.

Evaluating the derivatives of BSP under constant volatility (CV) at current values of the state variables gives the familiar Black-Scholes delta, gamma, and vega hedge parameters. Typically, these hedge parameters are used to hedge option price changes in response to the first and second-order effect of changes in the underlying asset price and the first-order effects of changes in the underlying asset variance.

(2) $$\Delta_{CV} = \frac{\partial P_t}{\partial S_t} \qquad \Gamma_{CV} = \frac{\partial^2 P_t}{\partial S_t^2} \qquad \Lambda_{CV} = \frac{\partial P_t}{\partial \sigma}$$

To derive the hedge parameters in a stochastic volatility environment, it is necessary to be precise about the current state variables and their interrelationships. The chain rule may be used to develop extended hedge parameters which capture these interrelationships. The potential dependence of expected average volatility on the underlying price suggests that additional chain rule terms may appear in the BSP hedge parameters. In fact, the stochastic volatility hedge parameters will be combinations of CV delta, CV gamma, CV vega, and derivatives of the volatility term structure defined by the volatility process.

For example, the stochastic volatility delta is obtained by differentiating BSP in (1) with respect to the underlying price. Delta (Δ) measures the option price change due to a small (first-order) change in the current underlying price (S_t) at the current level of volatility.

(3) $$\Delta = \Delta_{CV} + \Lambda_{CV} * VM \qquad VM = \frac{\partial \bar{\sigma}_{t,T}}{\partial S_t} \qquad \bar{\sigma}^2_{t,T} = \frac{1}{T} E_t\left[\sum_{i=1}^{T} \sigma^2_{t+i}\right]$$

The stochastic volatility delta indicates that a change in the underlying price affects the option price directly through CV delta and indirectly through CV vega and a change in volatility. The vega multiplier (VM) in equation (3) measures the change in the forecast average volatility, $\bar{\sigma}_{t,T}$, due to a change in the current underlying price. The forecast average volatility over the next T days is a single point on the VTS, but all forecast average volatilities are potentially affected by a change in the current underlying price.

The differences in deltas among the volatility models are due to different vega multipliers and different levels of forecast average volatility. For the CV and ARIV models, the vega multiplier in the delta equation is zero. This is because the underlying price and volatility are independent, so the derivative of forecast average volatility with respect to today's underlying price is zero. Setting the vega multiplier equal to zero simplifies equation (3) to CV delta evaluated at the forecast average volatility.

When evaluated at the point of no return surprise, the vega multiplier in equation (3) is also zero for the GARCH volatility processes considered in this paper. This is because a small (first-order) change in the underlying price does not affect forecast volatility when it conveys no volatility news. In GARCH models, volatility changes due to underlying price changes (including the leverage effect) are captured by second and higher order terms when evaluated at the no return surprise point.

To emulate the impact of additional news on the option price under GARCH, rather than additional news cumulated with the already observed news of the current day's return, consider the following methodology used in this paper. First, obtain the most general vega multiplier formula for the GARCH models by differentiating average volatility with respect to the current underlying price. The vega multiplier may then be substantially simplified by eliminating terms that disappear when S_t equals its expected value at date t-1, $E[S_t|S_{t-1}]$. This corresponds to a no return surprise situation, and results in a vega multiplier of zero. Second, for evaluating the expected average volatility and the CV delta formula, use the current day's return and price. This ensures that current information available to market participants is included in the hedge parameters.

In contrast to the deltas which are fairly similar for all the volatility models, the volatility hedge parameters are potentially quite different across stochastic volatility environments. This means that hedging tests that use the volatility hedge parameters should be able to distinguish the relative accuracy of the volatility models. So, delta-vega or delta-gamma hedging tests are strongly preferred to delta hedging tests for ranking volatility models.

Consider hedging a change in volatility when volatility is independent of the underlying price such as in the CV and ARIV models. While the CV model is not a stochastic volatility model, one might consider the effects on the option price of continual updating of volatility estimates based on an investor's expanding information set. If volatility is constant but is estimated with error, then an investor should update his volatility estimate each day which will result in a change in the option price.

In this case, the volatility hedge parameter is CV vega as given in (2), which measures the option price change due to a change in σ. The CV volatility term structure is flat, since average volatilities over all horizons are equal to the estimated unconditional variance. In addition, volatility shocks cause parallel shifts in the term structure, since a revision to the unconditional volatility estimate affects all average variance forecasts equally.

For the ARIV model, the volatility hedge parameter is ARIV vega. This is obtained by differentiating the BSP formula with respect to the volatility news using the chain rule. The first term is the derivative of BSP with respect to volatility, which is CV vega. The first term is multiplied by the derivative of average volatility with respect to the volatility news, which is the vega multiplier (VM). ARIV vega is evaluated at the current forecast average volatility ($\overline{\sigma}_{t,T}$) proxied by the implied volatility, and the current stock price (S_t). Thus, ARIV vega is:

(4) $$\Lambda_{ARIV} = \Lambda_{CV} * VM \qquad VM = \frac{\partial \overline{\sigma}_{t,T}}{\partial \sigma^2_{t+1}}$$

In an ARIV model, news arrives each day about the one-day ahead squared volatility. This news results in updates of all n-step ahead one-day forecast volatilities. The relationship between a change in the one-day ahead volatility and the forecast average volatility used in the BSP formula is given by the VM.

A mean-reverting ARIV model has the property that volatility news has the greatest effect on the one-day ahead volatility, and the effect on future volatilities decays with time. If the volatility process were integrated, the volatility news would affect all future volatilities equally. While the volatility term structure at any given time might be upward or downward sloping, volatility shocks always decay at a rate defined by the autoregressive parameter of the ARIV process, which will be less than one if the process is mean-reverting. This decay rate also determines the term structure shape.

Now, consider hedging changes in volatility when volatility forecasts depend on the magnitude of the current underlying asset return such as in the GARCH models. In GARCH models, each day the underlying asset return is observed. The magnitude (and possibly the level) of this return is used to predict future volatilities. Thus, a change in the current underlying price, which changes the magnitude of the current return, changes the forecast volatility.

The volatility hedge parameter for GARCH models is based on a derivative of BSP with respect to the underlying price, which captures the impact of volatility news. As noted earlier, the first-order volatility effects of underlying price change in GARCH models are zero. In contrast, the second derivative of the GARCH process with respect to the underlying price captures the volatility effect. This second derivative is the GARCH volatility hedge parameter: GARCH gamma.

As with the GARCH delta formula, the GARCH gamma formula is derived by eliminating the terms in the vega multiplier that are zero at the point of no return surprise. This eliminates many cross-product terms that would appear in the formula that do not measure the incremental effect of volatility news. GARCH gamma is then estimated using current information including the realized return for date t, and the current underlying price. GARCH gamma is given by:

$$(5) \qquad \Gamma_G = \Gamma_{CV} + \Lambda_{CV} * VM \qquad\qquad VM = \frac{\partial^2 \overline{\sigma}_{t,T}}{\partial S_t^2}$$

The GARCH gamma in (5) incorporates both a volatility hedge and a hedge against non-linear price response, since it includes both CV vega and CV gamma. The vega multiplier (VM) measures the second-order change in the average volatility forecast due to a change in the underlying price. This is also a measure of a shift in one point on the VTS due to an underlying price change. As was noted previously, the vega multiplier will incorporate all the parameters of the GARCH process including leverage terms.

Using the volatility exposures measured by the hedge parameters, it is possible to set up volatility hedging tests for medium term at-the-money straddles. An accurate volatility model should be able to estimate the correct number of short term straddles to purchase per medium term straddles written to neutralize effects of volatility changes. The number of short term straddles to purchase per medium term straddle will be referred to as the volatility hedge ratio and is calculated as the ratio of the volatility hedge parameters. A hedging test involving straddles is appropriate because straddles are especially sensitive to changes in volatility.

In continuous-time hedging, the hedge ratios are measured each instant, and perfect instantaneous hedge performance is expected from the correct VTS model. However, our tests are conducted in discrete time with a rebalancing period of one day, so hedge portfolios will have some risk. In fact, the standard deviation minimizing discrete time hedge ratios will not necessarily be equal to the hedge ratios derived above. Robins and Schachter (1994) show this for the case of CV delta hedges. As long as the error terms are uncorrelated with the risk factors and each other, the hedge ratios given in this paper will generate factor-neutral hedge portfolios. To address this general problem, we use the continuous time hedge ratios as an approximation to a dispersion minimizing hedge but measure hedging performance based on multiple criteria. The four hedging performance measures used in this paper are discussed in detail in section 4.

The final step before hedge ratio estimation is evaluation of the accuracy of the approximate hedge ratios. Since the BSP formula is approximate for the ARIV and GARCH models, the ARIV and GARCH hedge ratios may deviate from their "true values." Using Monte-Carlo simulation and the appropriate risk-neutralized processes for the ARIV and GARCH models, the "true" hedge ratios may be calculated under the risk-neutral measure and compared with the BSP approximate hedge ratios.

The true hedge ratios are calculated using centered finite difference approximations of option pricing formula derivatives. The finite difference approximations for GARCH delta and GARCH gamma are described in Rosenberg and Engle (1995). For this study, Amin and Ng's (1994) risk-neutralization for

GARCH processes is used, and derivatives are evaluated at a one-tenth standard deviation shock centered around the expected asset return. The finite difference approximation for ARIV vega is calculated by first taking the difference between a simulated straddle price evaluated at the initial level of volatility plus and minus a one-tenth standard deviation shock. This option price difference is divided by the difference in the initial volatilities, giving an estimate of ARIV vega. ARIV delta is calculated in an analogous manner using small changes in the initial underlying price.

All "true" hedge ratios for at-the-money straddles are calculated with 50,000 simulation replications for maturities from 5 to 90 days representing the range of maturities of options in this study. In each simulation experiment, the level of volatility is set to the unconditional risk-neutral volatility for the given process. The gamma and vega hedge ratios are based on hedging a medium term straddle with 25 to 90 days left with a short term straddle with 20 or fewer days until maturity. The delta hedge ratios are generated for at-the-money straddles with 5 to 90 days left.

The "true" hedge ratios are compared with the approximate BSP hedge ratios evaluated at the unconditional level of volatility for the original process. The difference between the BSP hedge ratio and the fully simulated alternative is defined to be the approximation error. The BSP delta, gamma, and vega hedge ratios are based on equations (3), (11), and (19) with details given in the next section.

We expect that the particular method for delta calculation will have little effect on the hedging tests for two reasons. First, Table 1.1 shows that average at-the-money simulated straddle deltas across maturities are close to zero, ranging from .0288 to .0468. This indicates that changes in the underlying asset price have a relatively small contribution to straddle variance. Empirically, a delta hedge provides a negligible amount of hedging benefit for a straddle. The straddle delta hedging experiments described in Table 6 show that delta hedging a straddle using any of the volatility models provides only a marginal or no improvement over a position that is not delta hedged, for all but the value-at-risk criteria.

Second, Table 1.1 also shows that the BSP at-the-money straddle deltas are a close approximation of the simulated deltas. For example, the simulated average at-the-money GCOMP put and call deltas are .5238 and -.4771, netting out to an average straddle delta of -.0466. The GCOMP BSP approximation preserves the nearly offsetting character of the "true" call and put deltas with an average approximation error of -.0261 and an error standard deviation of .0035.

The ARIV and GARCH volatility hedge ratio approximation errors are also acceptably small as shown in Table 1.2. The average error across maturities and moneynesses for each volatility model may be calculated by averaging the each model's first row in Table 1.2. The average errors for the ARIV, GARCH(1,1), GJR, and GCOMP models are .0045, .0021, .0017, and .0048 respectively with average error standard deviations ranging from .0047 to .0067. Using the GCOMP model as a representative example, this level of approximation error represents an error of about .6% relative to the average hedge ratio of .8. This also corresponds to an error of about 5% relative to the time-series standard deviation of hedge ratios for the sample options data. These results suggest that the contribution of approximation error to the hedging results is insignificant.

3. Estimating the volatility models and the volatility hedge ratios

In order to implement hedging tests, it is necessary to estimate the volatility models which provide the sensitivity of the volatility term structure to volatility news. The empirical volatility hedge ratios depend on the parameters of the estimated stochastic volatility models, and thus hedging performance is a function of the accuracy of the estimated models. Using the methods described in this section, the GARCH, ARIV or CV hedge ratios could be estimated for European options traded on other indices, commodities, or individual stocks.

The five volatility models considered in this paper differ in structure and in the type of shocks that drive the volatility process. In some sense, the most basic distinction between the GARCH, ARIV, and CV models is whether the current return provides information about future average volatility. In the GARCH models, the magnitude of the current return is a predictor of future volatility, while the ARIV and CV models do not incorporate this information. The three GARCH models differ among themselves in their specification of the dependence of conditional volatility on returns. The differences between the models result in different hedge ratios, and provide a basis for comparison using hedging tests. For each model, the rate of mean reversion, which defines the term structure shape, is the central determinant of the sensitivity of the volatility term structure to volatility news.

First, consider the CV model. We estimate σ, the constant volatility parameter, as the sample standard deviation of S&P500 index total daily log-returns from January 1986 to February 1992. This is the period over which hedging performance for S&P500 index options is tested. The estimated CV model in Table 2.1 indicates that σ is about 1.2% per day, which corresponds to an annualized S&P500 return standard deviation of 19.05%. Hedge parameters and hedge ratios for the CV model are defined in the previous section. A trailing 20-day standard deviation of returns and standard deviation based on an expanding window of trailing returns are also computed for comparison in out-of-sample hedging tests. For the expanding window, the first estimation begins with one-thousand prior returns.

In this paper, the autoregressive volatility model (ARIV) is based on the specification given by Heynen, Kemna, and Vorst (1994). The ARIV model may be viewed as a reduced form of a stochastic autoregressive volatility model (SARV) model in which the factors that drive the volatility process are unobservable, but average volatility is observable. For further discussion of SARV models, see Andersen (1994) or Taylor (1994).

Consider a first-order autoregressive volatility model in which σ_t is the volatility on day t, and σ_{t-1} is the one day lag. The ARIV model may be written in variances as:

(6) $$\sigma_t^2 = \omega + \rho\sigma_{t-1}^2 + \varepsilon_t$$

or equivalently,

(7) $$(\sigma_t^2 - \bar{\sigma}^2) = \rho(\sigma_{t-1}^2 - \bar{\sigma}^2) + \varepsilon_t$$

The volatility term-structure estimated at date t describes the relationship between average daily volatility over the next T days and the forecast horizon T. The VTS and the vega multiplier for the ARIV model are defined by:

$$\bar{\sigma}_{t,T} = \sqrt{\bar{\sigma}^2 + \frac{1}{T}\left[\frac{1-\rho^T}{1-\rho}\right](\sigma^2_{t+1} - \bar{\sigma}^2)} \tag{8}$$

$$VM = \frac{\partial \bar{\sigma}_{t,T}}{\partial \sigma_{t+1}} = \frac{1}{2\bar{\sigma}_{t,T}T}\left[\frac{1-\rho^T}{1-\rho}\right]$$

In this case, the term structure of volatility is upward or downward sloping depending on the level of tomorrow's volatility, σ_{t+1} compared to the long-term average volatility, $\bar{\sigma}$. The shape of the term structure and the sensitivity of the term structure to volatility shocks are determined by the mean reversion parameter ρ.

Since our purpose is to relate changes in average short and medium term implied variance, we use (6) and the definition of average variance to get the restriction:

$$(\bar{\sigma}^2_{t,T_m} - \bar{\sigma}^2) = \frac{T_s}{T_m}\frac{(1-\rho^{T_m})}{(1-\rho^{T_s})}(\bar{\sigma}^2_{t,T_s} - \bar{\sigma}^2) \tag{9}$$

Taking the derivative of medium term implied variance with respect to short term implied variance, the following is obtained:

$$\frac{\partial \bar{\sigma}^2_{t,T_m}}{\partial \bar{\sigma}^2_{t,T_s}} = \frac{T_s}{T_m}\frac{(1-\rho^{T_m})}{(1-\rho^{T_s})} \tag{10}$$

A change in short term implied variance results in a change in medium term implied variance defined by the ratio in (10). When the ARIV process is mean-reverting, ρ is less than one, implying that the ratio is less than one. So, a shock to short term implied variance results in a less than one-for-one change in medium term implied variance. Thus, the shorter maturity end of the volatility term structure is more volatile than the longer maturity end to volatility news.

We estimate the ARIV process in equation (6) using a regression of the short term implied variances on first lags. The closest-to-the-money, nearest maturity S&P500 call and put implied variance are averaged to obtain the short term implied variance. The option contracts used for implied volatility estimation are rolled over on the first business day of the short term contract expiration week. Implied variances are extracted using the Black-Scholes formula with an adjustment for the dividends to be paid over the remaining life of the option. The risk-free rate used is the yield, based on the average bid and ask price, for a three-month Treasury bill with one month remaining. Notice that there is no historical data on underlying asset returns used in this estimation procedure.

From the time series of implied volatilities, there is strong evidence for mean reversion in volatility as indicated by the estimated ARIV model in Table 2.1. The estimated ρ is substantially lower than one with an estimated value of .88 and a standard error of 0.02. For a unit change in 20 day implied variance, 40 day implied variance is expected to change by .54 based on the estimated ρ and equation (10).

As expected from a mean-reverting model, short term implied variances and their first differences are more volatile than for medium term variances. The Table 2.1 summary statistics for implied variances show that the volatility of volatility is declining with maturity. Mean reversion is also confirmed by the raw data for option returns in Table 4, which indicates that short term returns are substantially more volatile than medium term returns.

The ARIV volatility hedge ratio relates the price change of a medium term straddle due to a volatility shock to the price change in a short term option straddle due to the same shock. As noted previously, it is also the number of short term straddles to purchase per medium term straddle sold to ensure neutrality to volatility news. As given in section 2, the ARIV volatility hedge ratio is the ratio of ARIV vegas:

$$(11) \qquad \frac{\Lambda_{CV}(\overline{\sigma}_{t,T_m}, S_t, T_m)}{\Lambda_{CV}(\overline{\sigma}_{t,T_s}, S_t, T_s)} \frac{\overline{\sigma}_{t,T_s}}{\overline{\sigma}_{t,T_m}} \frac{T_s}{T_m} \frac{(1-\rho^{T_m})}{(1-\rho^{T_s})}$$

The medium and short term average volatilities used in (18) are given by the implied standard deviations from the appropriate maturity options, and ρ is given by the ARIV estimate based on equation (6). For out-of-sample hedging tests, a time-varying ρ is also reestimated daily based on an expanding window of implied variances.

In contrast to the ARIV model, the GARCH models use historical data from the underlying asset to estimate the volatility process and the volatility term structure. For these models, the current return magnitude predicts future volatility. Bollerslev's (1986) GARCH(p,q) model was developed as an modification to the Engle's (1982) ARCH(p) formulation to incorporate long memory effects in a more parsimonious manner. The GARCH(1,1) model with a constant risk premium may be written as:

$$(12) \qquad \ln(S_t / S_{t-1}) - r = \mu - \frac{1}{2}\sigma_t^2 + \varepsilon_t, \quad \varepsilon_t \sim N(0, \sigma_t^2)$$

$$(13) \qquad \sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2$$

where $\ln(S_t/S_{t-1})$-r is the excess log return, μ is a constant risk premium, σ_t^2 is the one day conditional variance, and r is the risk-free rate. The parameters α and β reflect the relative importance of the one day lagged squared return and the prior day's conditional variance on today's conditional variance.

In high frequency data, the GARCH-in-mean model developed by Engle, Lilien, and Robins (1987) which allows for a time varying risk-premium is frequently rejected. We find this to be the case for S&P500 log excess returns, so a constant risk premium is used in the GARCH models.

The GJR model (Glosten, Jagannathan, and Runkle, 1993) or TARCH model (Zakoian, 1994) generalizes the GARCH(1,1) model to allow negative return shocks to disproportionately increase volatility. This asymmetric effect is frequently called leverage, reflecting the increase in the debt-equity ratio that follows a reduction in a firm's market capitalization. The GJR model with a constant risk premium may be written as:

$$(14) \qquad \ln(S_t / S_{t-1}) - r = \mu - \frac{1}{2}\sigma_t^2 + \varepsilon_t, \quad \varepsilon_t \sim N(0, \sigma_t^2)$$

$$(15) \qquad \sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \gamma Max[0, -\varepsilon_{t-1}]^2 + \beta\sigma_{t-1}^2$$

where γ measures the relationship between "bad news," i.e. a negative return shock, and conditional variance. The other terms are defined as above.

The GARCH components with leverage model (GCOMP) developed by Engle and Lee (1993) allows for greater volatility dynamics and a leverage effect. In this model, volatility shocks have different effects on a long run and short run volatility component. Each component has a different level of mean reversion.

$$(16) \quad \ln(S_t / S_{t-1}) - r = \mu + \varepsilon_t, \ \varepsilon_t \sim N(0, \sigma_t^2)$$

$$(17) \quad \sigma_t^2 = q_t^2 + \alpha(\varepsilon_{t-1}^2 - q_{t-1}^2) + \gamma(Max[0, -\varepsilon_{t-1}]^2 - .5q_{t-1}^2) + \beta(\sigma_{t-1}^2 - q_{t-1}^2)$$

$$(18) \quad q_t^2 = \omega + \rho q_{t-1}^2 + \phi(\varepsilon_{t-1}^2 - \sigma_{t-1}^2)$$

Of the parameters, α reflects the effect of a shock on the short run component of volatility, γ captures the short run asymmetric effect of "bad news" on volatility, β reflects the influence of the prior day's volatility forecast, ρ measures the persistence of the long run component, and φ represents the effect of a shock on the long term component.

The GARCH(1,1) and GJR models, like the ARIV model, imply a monotonic upward or downwards sloping volatility term structure. The GARCH(1,1) and GJR term structures are of the same form as equation (8) with ρ replaced by the sum of α and β for the GARCH(1,1) model and by the sum of α, β, and one-half γ for the GJR model. These sums correspond to the volatility decay rate and are central determinants of shape of the volatility term structure. The GCOMP model allows for non-monotonic term structure shapes. In the GCOMP model, the term-structure shape is dominated by to the sum of α, β, and one-half γ over the short run and by ρ over the long run. The GARCH variance term-structures are given by the first equality in equations (20), (22), and (24).

The GARCH models are estimated for the full sample period using maximum likelihood with the week of the October 1987 crash downweighted in the likelihood function. Table 2.2 provides the results of the full sample estimation. The GARCH models are also reestimated daily using an expanding window of trailing returns for out-of-sample hedging tests. In this case, the first estimation begins with one-thousand historical returns.

For the full sample estimation, decay rates are around .96 to .98, which implies mean reversion in volatility and a declining volatility of volatility. However, persistence is substantially higher than in the ARIV model, indicating that the GARCH models have greater relative sensitivity of medium term average volatility to volatility news and flatter term structures. The leverage effect is significant in both the GJR and GCOMP models suggesting that negative S&P500 returns increase volatility disproportionately. This effect is ignored in the ARIV model.

The GARCH volatility hedge ratios are defined in section 2 as the of the ratio of the GARCH gamma of the medium term straddle and the GARCH gamma of the short term straddle. GARCH gamma is given by (5), so the GARCH volatility hedge ratio is:

$$(19)\qquad \frac{\Gamma_m(\bar{\sigma}_{t,T_m}, S_t, T_m) + \Lambda_m(\bar{\sigma}_{t,T_m}, S_t, T_m) VM(T_m)}{\Gamma_s(\bar{\sigma}_{t,T_s}, S_t, T_s) + \Lambda_s(\bar{\sigma}_{t,T_s}, S_t, T_s) VM(T_s)}$$

where $VM(T_m)$ and $VM(T_s)$ are the medium and short term vega multipliers, which are derivatives of average volatility with respect to a volatility shock. As before, the vega multiplier represents the sensitivity of the volatility term structure to a volatility shock. The differences in GARCH volatility hedge ratios among the GARCH models are due to their different estimates of average volatility and their different estimates of the vega multipliers. S_{t-1} appears in the denominator of all of the GARCH volatility multipliers as a result of taking the second derivative with respect to S_t of the current squared return $(r_t)^2$ in the volatility equation, which is a function of S_{t-1}.

For the GARCH(1,1) model the following are used in equation (19):

$$(20)\qquad \bar{\sigma}_{t,T}^2 = \bar{\sigma}^2 + \frac{1}{T}\left(\frac{1-(\alpha+\beta)^T}{1-(\alpha+\beta)}\right)(\sigma_{t+1}^2 - \bar{\sigma}^2),\ \ \bar{\sigma}^2 = \frac{\omega}{1-(\alpha+\beta)}$$

$$(21)\qquad VM(T) = \frac{\alpha}{T\bar{\sigma}_{t,T} S_{t-1}^2}\left(\frac{1-(\alpha+\beta)^T}{1-(\alpha+\beta)}\right)$$

For models with leverage effects as formulated by GJR, the second derivative of volatility with respect to the current underlying price does not exist. From the left, the second derivative is zero, and from the right, it is γ, so we use one-half γ as an approximation. The GJR volatility hedge ratio is given by equation (19) with the following substitutions:

$$(22)\qquad \bar{\sigma}_{t,T}^2 = \bar{\sigma}^2 + \frac{1}{T}\left(\frac{1-(\alpha+\beta+.5\gamma)^T}{1-(\alpha+\beta+.5\gamma)}\right)(\sigma_{t+1}^2 - \bar{\sigma}^2),\ \ \bar{\sigma}^2 = \frac{\omega}{1-(\alpha+\beta+.5\gamma)}$$

$$(23)\qquad VM(T) = \frac{\alpha+.5\gamma}{T\bar{\sigma}_{t,T} S_{t-1}^2}\left(\frac{1-(\alpha+\beta+.5\gamma)^T}{1-(\alpha+\beta+.5\gamma)}\right)$$

Notice that the GCOMP model has both long and short run mean reversion. The GCOMP hedge ratio is defined by equation (19) with the following substitutions:

(24) $$\bar{\sigma}^2_{t,T} = \bar{\sigma}^2 + \frac{1}{T}\left(\frac{1-(\alpha+\beta+.5\gamma)^T}{1-(\alpha+\beta+.5\gamma)}\right)(\sigma^2_{t+1} - q_{t+1}) + \frac{1}{T}\left(\frac{1-\rho^T}{1-\rho}\right)(q_{t+1} - \bar{\sigma}^2)$$

$$\bar{\sigma}^2 = \frac{\omega}{1-\rho}$$

(25) $$VM(T) = \frac{1}{T\bar{\sigma}_{t,T}S^2_{t-1}}\left((\alpha+.5\gamma)\frac{1-(\alpha+\beta+.5\gamma)^T}{1-(\alpha+\beta+.5\gamma)} + \phi\frac{1-\rho^T}{1-\rho}\right)$$

To better understand these volatility hedging ratios, consider hedging the volatility sensitivity of a medium term at-the-money straddle with 30 days to maturity with a short term at-the-money straddle with 10 days to maturity. Table 3 indicates that a CV volatility or vega hedge will require the purchase of 1.73 short term straddles for every medium term straddle sold. A hedge ratio greater than one reflects the increase in vega with maturity. This corresponds to an experiment where volatility is changed once and for all, and therefore has a larger impact on longer-lived options. Figure 1 illustrates that the CV vega is increasing as the straddle's time to maturity increases.

In the ARIV model, the ARIV vega hedge ratio is used to eliminate the effect of changes in volatility on the option portfolio. The vega multiplier incorporates mean reversion in volatility, which counteracts the rise in vega with maturity. For this example and the estimated ARIV model, the vega hedge ratio is .68. Figure 1 shows that ARIV vega declines with time to maturity for straddles with greater than two weeks until expiration resulting in a hedge ratio less than one.

Using GARCH models, the volatility hedge is formed using GARCH gamma hedge ratios, since volatility changes respond to underlying price changes. In this example, the GARCH(1,1) gamma hedge ratio is .66, the GJR gamma hedge ratio is .65, and the GCOMP gamma hedge ratio is .62. Since these hedge ratios are approximately linear combinations of the BS gamma and BS vega, they lie between the CV gamma and CV vega hedge ratios.

The term structure shape determines how CV gamma and CV vega are weighted in GARCH gamma. In general, if only short term variances are sensitive to volatility shocks, then the weights will give more emphasis to the BS gamma. If the process is IGARCH, then more weight will be given to BS vega. GARCH gamma is always higher than CV gamma since GARCH gamma adds the volatility effect of a price shock to the non-linear effect of a price shock which is measured by CV gamma. This is illustrated in Figure 1.

Using the estimated volatility models, it is straightforward to calculate hedge ratios over the sample period. This gives a another picture of the model differences. For at-the-money straddle positions, all average hedge ratios are less than one except for the CV vega hedges. This is due to the increase in CV vegas with time to maturity. The CV gamma hedge and the GCOMP gamma hedge have the lowest hedge ratios and are closest to the ex-post variance minimizing hedge.

The ex-post hedge portfolio is obtained by regressing medium term straddle price changes on short term straddle price changes and the underlying price change. In this case, the hedge ratio is forced to be constant over the sample period with portfolio weights given by the regression coefficients. This portfolio indicates reasonable average values for the hedge parameters, although there is no reason why it should not be dominated by a time-varying hedge ratio.

4. Hedging tests

The hedging tests developed in this paper measure the accuracy of the volatility hedge ratios derived from each volatility model. Since these hedge ratios reflect predicted volatility term structure dynamics as well as the predicted term structure shape, the hedging tests may be interpreted as tests of the VTS implied by each volatility model. We expect that the most effective hedges will be generated by the volatility model with best term structure estimate.

The hedging tests are sensitive to omitted variables and other forms of volatility model misspecification. Omitted variables in the volatility equation will result in poorer hedging performance, since shifts in the volatility term-structure related to these variables will be overlooked. The hedging tests also focus on the term structure shape and rate of mean-reversion, which are central to the hedge ratio estimates. Thus, the hedging tests reduce the influence of unconditional volatilities on the results which HKV (1994) find to be a dominant differentiating factor in their comparisons using predicted and realized volatility term structures.

In this section, the volatility term structure estimates of the five models are compared by constructing volatility hedges for $100 medium term S&P500 index straddle positions using short term S&P500 index straddles. Each straddle position consists of an equal number of calls and puts with identical strike prices and times to expiration. The positions are also delta hedged to minimize the influence of correlation with the underlying asset on the tests. All option prices are market prices from an options data set provided by the Chicago Board Options Exchange and for the period January 1986 through February 1992. In contrast to some previous papers measuring hedging performance, no interpolated option prices are used in this paper.

Daily closing prices for the nearest-to-the-money Standard and Poor's 500 index put and call options with closest and next-closest maturities are used in this analysis. Only the 552 data points for which daily price changes are available for the medium and short term straddles are used. Options that are further than one percent from the money are excluded. Table 4 summarizes the data used.

The underlying price used in the analysis is the closing price of the S&P500 index as reported on the CRSP tapes. Hedge parameters are adjusted for dividends by discounting the index level by the present value of dividends to be paid over the life of the option. The S&P500 daily dividend series is also from CRSP. The risk-free rate used is the yield, based on the average bid and ask price, for a three-month Treasury bill with one month remaining from the Fama file.

The hedging tests are implemented as follows. Each trading day, a medium term straddle position worth $100 is sold. The five volatility and delta hedge ratios corresponding to the volatility models are then calculated. The number of short term straddles to purchase is given by the volatility hedge ratio times the number of medium term straddles sold. The number of shares of the underlying to purchase or sell per medium term straddle is given by the delta hedge ratio. These transactions are made, creating the volatility hedge portfolio. This portfolio is held for one day and then sold. Each day, portfolio price differences are calculated, and new positions established.

At the end of the analysis, hedging effectiveness is analyzed using the standard deviation of realized price changes, interquartile range, value-at-risk, and factor sensitivity regressions. Each of these measures provides important information about hedging effectiveness, and different measures may be important in different contexts. For the purpose of the hedging tests in this paper, the best VTS model is selected based on the number of criteria for which it is superior to all the other models. In practice, agents may have different attitudes towards the types of risk quantified in these four measures, and might choose a different rule to select a subjectively optimal VTS model.

Consider the characteristics of the standard deviation and interquartile range (IQR) as risk measures. Both quantify average hedge portfolio variability and are estimated using sample statistics taken from the realizations of the 552 hedge portfolio price changes. The interquartile range may be a more robust measure than the standard deviation, since IQR is less sensitive to outliers.

Value-at-risk is measured as the fifth percentile of realized hedge portfolio returns. The idea of the value-at-risk measure is that it captures the magnitude of a low probability, large negative outcome while not being influenced by the shape of the distribution of positive outcomes. In practice, the value-at-risk or "5th percentile" is calculated by sorting the realized hedge portfolio price changes from lowest to highest and selecting the 28th smallest price change, which corresponds to five percent of a sample of 552 observations. Value-at-risk is useful in measuring "downside risk" or a worst-case scenario rather than average variability.

The F-probability from factor sensitivity regressions is an important risk measure because a low value (below .01) indicates exposure of the "hedge" portfolio to systematic sources of risk, namely risk due to changes in the level of the index. This systematic risk would not necessarily be reduced by holding a diversified options portfolio. The reported F-probability is based on an OLS regression of the hedge portfolio price changes on changes and squared changes in the level of the S&P500 index. The F-probability is the probability of being able to explain by chance larger amount of the variance in hedge portfolio price changes, under the hypothesis that the portfolio is factor-neutral.

The hedging performance results are summarized in Table 5. We also examined results hedging performance using unexpected price changes, that is changes net of theta, and the outcomes were quite similar. We do not expect that sensitivity to interest rate changes will have a substantial impact on test results.

All of the models generate quite similar results in delta hedge performance as reported in Table 6 under the "Hedge using underlying" label. This is not surprising, since the primary impact of hedging differences using the volatility models is through volatility hedging, and straddle deltas are small. None of the delta hedges reduce the standard deviation of the non-delta hedged portfolio. Value-at-risk is somewhat improved using delta-hedging, and interquartile range improves marginally.

The volatility hedging performance detailed in Table 5 indicates that the GCOMP estimates of the term structure shape are most accurate, although the performance of all the GARCH models is similar. It should be noted that the GCOMP model is the most flexible of the GARCH models, encompassing both the GJR and GARCH(1,1) formulations. It appears that long and short run mean-reversion as well asymmetric effects are important aspects of S&P500 volatility process. These are reflected in the volatility term structure shape and the sensitivity of the term structure to volatility news.

GCOMP delta-gamma volatility hedges reduce the option portfolio standard deviation by about 10% and value-at-risk by about 36% compared to the unhedged medium term straddle. These are the largest improvements for all delta-volatility hedges based on these two criteria. The GCOMP model reduces option portfolio interquartile range by about 11%, which is slightly inferior to the GARCH(1,1) based on this criterion. The results of the factor sensitivity regressions in Table 5 show that all of the GARCH hedge portfolios eliminate sensitivity to underlying price changes and their squares, while neither the CV or ARIV models is able to accomplish this.

The delta-volatility hedges for the CV and ARIV models, in which underlying price changes are assumed to be unrelated to volatility changes, are inferior in dispersion reduction compared to the GARCH hedges. This indicates that squared price changes provide important information about shifts in volatility term structure, and should be incorporated in the conditional volatility model. It is notable that the constant volatility delta-vega hedge substantially increases portfolio variance compared to a no-hedge alternative. Clearly, treating a volatility shock as a one-and-for-all change that affects all parts of the volatility term structure equally is unrealistic based on this result. The fact that all of the GARCH models are superior to the ARIV model indicates that the addition of mean-reversion to the volatility process, without a realistic specification of the relationship between return magnitudes and future volatility, is inadequate for modeling volatility term-structure dynamics and hedging changes in volatility. Thus, the hedging performance tests highlight a particular form of model misspecification: an omitted variable in the volatility equation.

The second part of Table 5 provides an out-of-sample evaluation of hedging performance. For the out-of-sample tests, all volatility models and volatility hedge ratios are estimated using only data available at the time of portfolio construction. These results are similar to those in the first part of Table 5. The GCOMP model provides superior performance based on the interquartile range, value-at-risk, and factor sensitivity regressions compared to all the other models. The GARCH(1,1) model is a slight improvement over the GCOMP model based on the standard deviation criterion.

The out-of-sample hedging performance of all the GARCH models is somewhat lessened compared to in-sample performance, but is still substantially better than the unhedged alternative. Interestingly, the out-of-sample ARIV model performs better than the in-sample ARIV model based on all hedging criteria, possibly suggesting a time-varying rate of mean-reversion.

Table 6 reports hedging performance results for some "benchmark" portfolios providing information about the optimality of GCOMP gamma hedging as a dispersion reducing strategy. This table may be used to evaluate GCOMP delta-gamma hedging as a practical hedging strategy for ATM straddles, rather than as a measure of the superiority of the GCOMP volatility term structure forecasts. In this case, alternative hedging techniques are considered that are not necessarily derived from volatility models, but that might be effective for other reasons.

As noted previously, the impact of delta hedging on medium term straddles is limited to minor improvement in interquartile range and value-at-risk. All delta hedging strategies perform fairly similarly. Several new methods for hedging a medium term straddle using a short term straddle and the underlying asset are presented in Table 6. Two methods of CV delta-vega hedging are considered that use alternative estimates of volatility. The first hedging method, "trailing 20 day historical volatility delta-vega hedge," uses the standard deviation of returns over the past twenty days as a proxy for the current level of volatility. The second hedging method, "implied volatility delta-vega hedge," uses the average of current short term put and call implied volatilities as a proxy for current volatility. Both of these hedging methods have substantially inferior performance compared to GCOMP gamma hedging and even to a no-hedge alternative.

In addition, two methods of delta-gamma hedging are considered that use alternative estimates of volatility. The first hedging method, "constant volatility delta-gamma hedge," uses the historical standard deviation of daily returns from January 1985 until February 1992 as a volatility estimate. This is an in-sample hedging method. The second hedging method, "trailing 20 day historical volatility delta-gamma hedge," uses the uses the standard deviation of returns over the twenty days prior to hedge portfolio construction as a volatility estimate. The "constant volatility delta-gamma hedge" strategy improves on the in-sample GCOMP hedge based on a single criterion: standard deviation. The "trailing historical volatility 20 day delta-gamma hedge" outperforms the out-of-sample GCOMP hedge based on two criteria: standard deviation and factor sensitivity. Using the simple decision rule, this would result in a tie in hedging effectiveness of the trailing 20 day method with the GCOMP method for the purposes of implementing hedges.

It is somewhat surprising that the simple delta-gamma hedges work so well based on the standard deviation criterion. One explanation for the volatility hedging success is that CV or historical volatility hedging is a good proxy for the GARCH volatility hedge on quiet days in the market. In fact, it is possible that on quiet days GARCH gammas add noise to hedge portfolios, and that it is only on larger return days that the price-volatility relationship is crucial for volatility hedging. This suggests that GARCH-type hedges may be most effective when they are likely to be most needed, that is when there is a large market move.

A comparison of the performance of the constant volatility and GCOMP delta-gamma hedges given in Table 7 supports this explanation. GARCH hedge ratios work best when the price volatility link is strongest, namely when there are large returns or large negative returns. However, the poorer performance of GARCH hedges over the large positive returns indicates that there may be a weaker price-volatility relationship for small positive returns than specified by the GARCH models, even the asymmetric models.

Another explanation for the effectiveness of the CV delta-gamma hedge compared to stochastic volatility alternatives is mispricing in the S&P500 index options market. A systematic overreaction of short maturity options to news would make them excessively variable and would make the low hedge ratio observed in our sample optimal. Mispricing may be in the market, or it may be due to data anomalies such as non-synchronous prices or bid-ask bounce. In either case, the optimal hedge ratios would be different than those derived from BSP.

Diz and Finucane (1993) find some evidence for short term overreaction of options to new information, while Stein (1989) finds evidence of short term underreaction. If the options market is not efficient in that it does not incorporate all available volatility information into current prices, then it may be possible to earn excess profits by using alternative volatility forecasts. Noh, Engle, and Kane (1994) present evidence that GARCH models can profitably forecast when to be short or long at-the-money S&P500 index straddles. If the mispricing is substantial and of unknown form, it could have unpredictable effects on hedging performance.

It is also possible that profitability in options trading reflects risk premia rather than mispricing. For example, Table 4 shows that average put returns are substantially below expected put returns derived from the CAPM and instantaneous option betas. It is plausible that put sellers receive a risk premium due to decreasing relative risk aversion. When the S&P500 falls, aggregate wealth falls, and put options have a positive payoff. Since puts payoff when the marginal utility of consumption is highest, they might require a premium for this characteristic. A constant put risk premium would have no effect on hedging ratios.

Two additional regression based hedges used in the Table 6 comparisons are the "ex-post minimum variance hedge" and the "trailing 20 day minimum variance hedge." For the ex-post hedge, the medium term straddle portfolio price changes are regressed (with no intercept) on the short term straddle price changes and the changes in the level of the index for the entire sample. The regression coefficients provide the gamma and delta hedge ratios. The ex-post method, which is an in-sample method since it depends on all the realized price changes over the sample, does well based on the standard deviation criterion against the GCOMP model but fails based on the other criteria.

A serious competitor to GCOMP delta-gamma hedging is the "trailing 20 day minimum variance hedge." This is an out-of-sample method that uses the same regression as the ex-post method but with only the last 20 days of realized straddle price changes. This method, which generates a daily updated hedge ratio, is a significant improvement over the out-of-sample GCOMP alternative in standard deviation and interquartile range, but inferior in terms of value-at-risk and eliminating factor sensitivity. Regression type hedges often work well in practice, and it is interesting that GCOMP method is able to tie the regression method.

Overall, there are two hedging methods that tie the hedging effectiveness of the GCOMP method in out-of-sample tests. These are "trailing historical volatility 20 day delta-gamma hedge" and the "trailing 20 day minimum variance hedge." Each is superior to the GCOMP model based on two criteria. The GCOMP method is superior to both based on the value-at-risk criteria. A ranking of the relative hedging effectiveness of these methods would depend on an agents relative weighting of the four types of risk. If reducing average hedge portfolio variability were of primary concern then one of these two alternatives to the GCOMP method might be preferred. If reducing the magnitude of large negative outcomes were primary, then the GCOMP method might be preferred.

5. Conclusions

This paper provides a methodology for testing the volatility term structure using option hedging performance, and implements it to analyze the volatility term structure of S&P500 index returns. As part of the testing methodology, a technique is developed to derive approximate at-the-money option hedge parameters for a wide variety of volatility model specifications. These hedge parameters may be used in VTS tests or to implement hedges of at-the-money options. The hedging technique is innovative in that it is particularly sensitive to several important forms of volatility model misspecification. In particular, hedging tests measure the accuracy of predicted VTS dynamics, the predicted VTS shape, and the predicted relationships between the underlying price and volatility process.

Five models of the S&P500 volatility term structure are compared based on hedging performance for medium term S&P500 index straddles using short term straddles. The volatility models considered are a constant volatility model (CV), an autoregressive implied volatility model (ARIV), a GARCH(1,1) model, a GARCH(1,1) with leverage model (GJR), and a GARCH components with leverage model (GCOMP).

The hedging test results indicate that the GARCH components with leverage estimate of the volatility term structure is most accurate. This suggests that long and short run mean reversion as well as asymmetric effects are important components of an accurate volatility term structure model. The poorer hedging performance of the CV and ARIV models implies that the magnitude of underlying asset returns provide important information about shifts in the volatility term structure. There is strong evidence against a flat volatility term structure, and for mean-reverting volatility, based on the inferior performance of the CV volatility hedge.

Bibliography

Amin, K. I. and V. K. Ng, 1994, "A Comparison of Predictable Volatility Models using Option Data," Manuscript.

Amin, K. I. and V. K. Ng, 1993, "Option Valuation with Systematic Stochastic Volatility," *Journal of Finance*, 48, 881-910.

Andersen, T. G., 1994, "Stochastic Autoregressive Volatility: A Framework for Volatility Modeling," *Mathematical Finance*, 4, 75-102.

Black, F. and M. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637-654.

Bollerslev, T., 1986, "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, 31, 307-327.

Bollerslev, T., R. Y. Chou and K. F. Kroner, 1992, "ARCH Modeling in Finance - A Review of the Theory and Empirical Evidence," *Journal of Econometrics*, 52, 5-59.

Diz, F. and T. J. Finucane, 1993, "Do the Options Markets Really Overreact?," *Journal of Futures Markets*, 13, 299-312.

Duan, J. C., 1995, "The GARCH Option Pricing Model," *Mathematical Finance*, 5, 13-32.

Engle, R. F., 1982, "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica*, 50, 987-1007.

Engle, R. F. and G. Lee, 1993, "A Permanent and Transitory Component Model of Stock Return Volatility," University of California - San Diego, UCSD discussion paper, 92-44R.

Engle, R. F., D. M. Lilien and R. P. Robins, 1987, "Estimating Time Varying Risk Premia in the Term Structure: The ARCH-M model," *Econometrica*, 55, 391-407.

Engle, R. F. and J. V. Rosenberg, 1995, "GARCH Gamma," *Journal of Derivatives*, 2, 47-59.

Glosten, L. R., R. Jagannathan and D. E. Runkle, 1993, "On the Relation Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks," *Journal of Finance*, 48, 1779-1801.

Harrison, J. M. and D. M. Kreps, 1979, "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory*, 20, 381-408.

Heynen, R., A. Kemna and T. Vorst, 1994, "Analysis of the Term Structure of Implied Volatilities," *Journal of Financial and Quantitative Analysis*, 29, 31-56.

Hull, J. and A. White, 1987a, "Hedging the Risks from Writing Foreign Currency Options," *Journal of International Money and Finance*, 6, 131-152.

Hull, J. and A. White, 1987b, "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance*, 42, 281-301.

Melino, A. and S. M. Turnbull, 1995, "Misspecification and the Pricing and Hedging of Long Term Foreign Currency Options," *Journal of International Money and Finance*, 14, 373-393.

Melino, A. and S. M. Turnbull, 1990, "Pricing Foreign Currency Options with Stochastic Volatility," *Journal of Econometrics*, 45, 239-265.

Noh, J., R. F. Engle and A. Kane, 1994, "A Test of Efficiency for the S&P500 Index Options Market using Variance Forecasts," *Journal of Derivatives*, 2, 17-30.

Robins, R. P. and B. Schachter, 1994, "An Analysis of the Risk in Discretely Rebalanced Option Hedges and Delta-based Techniques," *Management Science*, 40, 798-808.

Rubinstein, M., 1976, "The Valuation of Uncertain Income Streams and the Pricing of Options," *Bell Journal of Economics and Management Science*, 7, 407-425.

Stein, J., 1989, "Overreactions in the Options Market," *Journal of Finance*, 44, 1011-1023.

Taylor, S. J., 1994, "Modeling Stochastic Volatility: A Review and Comparative Study," *Mathematical Finance*, 4, 183-204.

Zakoian, J.-M., 1994, "Threshold Heteroskedastic Models," *Journal of Economic Dynamics and Control*, 18, 931-955.

Table 1.1 - Comparison of simulated and BSP approximated hedge ratios

BSP hedge ratios are calculated according to equations 3, 11, and 19.
Simulated hedge ratios are calculated for 1:1 straddles using estimated GARCH/ARIV model as the risk-neutral measure, finite-difference approximations, and simulation (50,000 replications).
Simulated gammas and vegas are smoothed using the method in Engle and Rosenberg (1995).

Delta hedge ratio errors (BSP-simulated)
For straddles with maturities 5-90 days.

ARIV delta hedge ratio error

	Moneyness (S/K)				
	0.990	0.995	1.000	1.005	1.010
Mean error	-0.0015	0.0003	0.0018	0.0035	0.0050
Std. error	0.0028	0.0022	0.0019	0.0027	0.0034
Min. error	-0.0097	-0.0063	-0.0019	-0.0019	-0.0010
Max. error	0.0033	0.0049	0.0058	0.0096	0.0135
Avg. sim. call delta	0.5745	0.5448	0.5143	0.4839	0.4534
Avg. sim. put delta	-0.4254	-0.4552	-0.4856	-0.5160	-0.5465
Avg. sim. strd. delta	0.1491	0.0896	0.0288	-0.0320	-0.0931

GARCH (1,1) delta hedge ratio error

	Moneyness (S/K)				
	0.990	0.995	1.000	1.005	1.010
Mean error	-0.0144	-0.0127	-0.0104	-0.0080	-0.0055
Std. error	0.0042	0.0044	0.0042	0.0042	0.0042
Min. error	-0.0272	-0.0266	-0.0234	-0.0201	-0.0163
Max. error	-0.0078	-0.0072	-0.0049	-0.0013	0.0019
Avg. sim. call delta	0.5927	0.5554	0.5167	0.4780	0.4392
Avg. sim. put delta	-0.4073	-0.4445	-0.4832	-0.5219	-0.5608
Avg. sim. strd. delta	0.1854	0.1108	0.0335	-0.0439	-0.1216

GARCH with leverage (GJR) delta hedge ratio error

	Moneyness (S/K)				
	0.990	0.995	1.000	1.005	1.010
Mean error	-0.0315	-0.0285	-0.0242	-0.0192	-0.0135
Std. error	0.0039	0.0037	0.0025	0.0020	0.0028
Min. error	-0.0423	-0.0386	-0.0296	-0.0240	-0.0201
Max. error	-0.0247	-0.0229	-0.0194	-0.0149	-0.0034
Avg. sim. call delta	0.5990	0.5622	0.5240	0.4855	0.4467
Avg. sim. put delta	-0.3956	-0.4354	-0.4771	-0.5191	-0.5616
Avg. sim. strd. delta	0.2034	0.1268	0.0468	-0.0336	-0.1149

GARCH components (GCOMP) delta hedge ratio error

	Moneyness (S/K)				
	0.990	0.995	1.000	1.005	1.010
Mean error	-0.0227	-0.0249	-0.0261	-0.0265	-0.0259
Std. error	0.0025	0.0027	0.0035	0.0042	0.0042
Min. error	-0.0280	-0.0301	-0.0316	-0.0344	-0.0366
Max. error	-0.0172	-0.0200	-0.0192	-0.0194	-0.0185
Avg. sim. call delta	0.5993	0.5623	0.5238	0.4851	0.4458
Avg. sim. put delta	-0.3956	-0.4354	-0.4771	-0.5191	-0.5616
Avg. sim. strd. delta	0.2037	0.1269	0.0466	-0.0341	-0.1157

Table 1.2 - Comparison of simulated and BSP approximated hedge ratios

Vega and gamma hedge ratio errors (BSP-simulated):
For short and medium term straddles with maturities 20 days apart.
Short term maturities (5-70 days), medium term maturities (25-90 days).

ARIV vega hedge ratio error

	Moneyness (S/K)				
	0.990	0.995	1.000	1.005	1.010
Mean error	0.0041	0.0045	0.0050	0.0048	0.0042
Std. error	0.0051	0.0055	0.0050	0.0046	0.0035
Min. error	-0.0092	-0.0095	-0.0078	-0.0072	-0.0047
Max. error	0.0110	0.0126	0.0127	0.0119	0.0101
Avg. sim. hedge rat.	0.8104	0.8059	0.8044	0.8059	0.8105

GARCH (1,1) gamma hedge ratio error

	Moneyness (S/K)				
	0.990	0.995	1.000	1.005	1.010
Mean	-0.0041	0.0028	0.0031	0.0028	0.0057
Std	0.0061	0.0012	0.0098	0.0029	0.0094
Min	-0.0109	-0.0008	-0.0138	-0.0030	-0.0057
Max	0.0086	0.0057	0.0170	0.0071	0.0241
Avg. sim. hedge rat.	0.8249	0.8181	0.8159	0.8181	0.8250

GARCH with leverage (GJR) gamma hedge ratio error

	Moneyness (S/K)				
	0.990	0.995	1.000	1.005	1.010
Mean error	-0.0032	0.0011	0.0042	0.0032	0.0031
Std. error	0.0073	0.0044	0.0079	0.0066	0.0042
Min. error	-0.0165	-0.0044	-0.0149	-0.0083	-0.0066
Max. error	0.0068	0.0096	0.0105	0.0117	0.0064
Avg. sim. hedge rat.	0.8143	0.8075	0.8053	0.8075	0.8145

GARCH components (GCOMP) gamma hedge ratio error

	Moneyness (S/K)				
	0.990	0.995	1.000	1.005	1.010
Mean error	-0.0015	0.0037	0.0077	0.0068	0.0074
Std. error	0.0039	0.0070	0.0146	0.0031	0.0050
Min. error	-0.0085	-0.0113	-0.0115	0.0011	-0.0035
Max. error	0.0034	0.0097	0.0340	0.0097	0.0127
Avg. sim. hedge rat.	0.8089	0.8014	0.7989	0.8014	0.8091

Table 2.1- Estimation of volatility models

Constant volatility model (CV)
Standard deviation of log total S&P500 returns (Jan. 1986 - Feb. 1992)

Estimated CV model

	Number obs.	Sample standard deviation
σ	1557	0.0116

Autoregressive implied volatility model (ARIV)

Implied variance process is AR(1), see Heynen, Kemna, and Vorst (1994).
Daily implied variance used is average of short term put and call implied variances.
Contracts rolled over on first trading day of short term contract expiration week.
930 observations based on data availability from Jan. 1986-Feb. 1992.

Estimated ARIV model

	Coefficient	t-stat	Prob > |t|
α	1.7432E-05	6.77	0.0001
ρ	0.8829	51.11	0.0001

RMSE	3.09E-05
First order autocorr.	-0.1532
Ljung-Box stat., 6 lags	25.59

Summary statistics for implied variances

	Number obs.	Mean	Std. dev.	Skewness	Excess kurtosis	First order autocorr.
Short term implied variance	930	1.38E-04	6.03E-05	1.45	2.55	0.86
Medium term implied	930	1.44E-04	5.85E-05	1.31	1.79	0.82
Change in short term implied variance	930	1.39E-06	3.32E-05	0.66	7.75	-0.25
Change in medium term implied	930	-6.17E-07	2.54E-05	0.50	6.49	-0.36

Table 2.2 - Estimation of volatility models
(continued)

GARCH(1,1)

Daily log excess returns for S&P500 index from CRSP, with dividends.
1557 observations (Jan. 1986 - Feb. 1992)
Maximum likelihood estimation with normal as the underlying density.
Week of Oct. 1987 crash down-weighted.

	Coefficient	Std Error	t-stat	Robust Std Err	Robust t-stat	Ljung-Box(15)
μ	0.0006	0.0002	2.38	0.0002	2.38	11.95
ω	2.67E-06	3.72E-07	7.17	1.39E-06	1.92	
α	0.0151	0.0031	4.86	0.0666	0.23	
β	0.9538	0.0063	150.46	0.0786	12.14	

GJR - GARCH(1,1)

Daily log excess returns for S&P500 index from CRSP, with dividends.
1557 observations (Jan. 1986 - Feb. 1992)
Maximum likelihood estimation with normal as the underlying density.
Week of Oct. 1987 crash down-weighted.

	Coefficient	Std Error	t-stat	Robust Std Err	Robust t-stat	Ljung-Box(15)
μ	0.0006	0.0002	2.36	0.0003	2.22	12.43
ω	3.04E-06	3.35E-07	9.08	1.04E-06	2.83	
α	1.00E-08	0.0073	0.00	0.0521	0.00	
β	0.9501	0.0063	151.74	0.0695	13.71	
γ	0.0273	0.0099	2.75	0.0207	1.16	

GARCH components

Daily log excess returns for S&P500 index from CRSP, with dividends.
1557 observations (Jan. 1986 - Feb. 1992)
Maximum likelihood estimation with normal as the underlying density.
Week of Oct. 1987 crash down-weighted.

	Coefficient	Std Error	t-stat	Robust Std Err	Robust t-stat	Ljung-Box(15)
μ	0.0004	0.0002	1.66	0.0003	1.26	8.62
ω	1.08E-06	2.60E-07	4.13	5.07E-07	2.12	
α	1.00E-08	1.55E-02	0.00	9.50E-02	0.00	
β	0.7824	0.0565	13.84	0.1307	5.99	
γ	0.0843	0.0253	3.33	0.0362	2.33	
ϕ	0.0045	0.0013	3.52	0.0040	1.14	
ρ	0.9854	0.0034	293.57	0.0077	128.65	

Summary statistics for daily log S&P500 index excess returns, with dividends

	Number obs.	Mean	Standard deviation	Skewness	Kurtosis	Kolmogorov normality test p-value	Ljung-Box (15) on log returns	Ljung-Box (15) on squared log returns
S&P500 index daily log total return	1557	0.0004	0.0122	-4.63	88.65	0.001	30.10	93.75

Table 3 - Estimated volatility hedge ratios

Volatility hedging a medium term at-the-money straddle with a short term at-the-money straddle.
The hedge ratio is the number of short term positions to purchase per medium term position sold.

Expected average volatility=.01 daily, risk-free rate=0, yield=0, S=K=100

Model	Medium T=30 / Short T=10	Medium T=35.3 / Short T=12.7	Medium T=40 / Short T=20	Average hedge ratio over sample
CV vega hedge ratio	1.73	1.67	1.41	1.84
ARIV vega hedge ratio	0.68	0.66	0.72	0.80
GARCH(1,1) gamma hedge ratio	0.66	0.69	0.78	0.65
GJR gamma hedge ratio	0.65	0.67	0.76	0.65
GARCH comp w/lev. gamma hedge ratio	0.62	0.64	0.74	0.62
Benchmark: CV gamma hedge ratio	0.58	0.60	0.71	0.57
Benchmark: Ex-post minimum variance constant hedge ratio	0.45	0.45	0.45	0.45

Table 4 - Summary of option data

Data gathered by the Chicago Board Options Exchange, end of day price quotes, Jan. 1986-Feb. 1992. Nearest-to-the-money options for which current and next day's price are available are used in the study. Prices are in points ($100 units), moneyness is underlying price (S) / Strike price (K).

	Short term call options	Medium term call options	Short term put options	Medium term put options
Num. obs.	552	552	552	552
Average price	5.42	9.74	4.53	7.97
Std. price	1.93	2.75	1.73	2.28
Average price change	0.04	-0.01	-0.18	-0.18
Std. dev. price change	1.72	1.73	1.55	1.54
Skewness of price change	0.88	0.23	1.30	0.90
Kurtosis of price change	2.71	1.22	6.20	4.72
Average return	0.0209	0.0092	-0.0399	-0.0213
Average expected return (Rubinstein,	0.0211	0.0112	-0.0208	-0.0120
Std. dev. return	0.3451	0.1787	0.4228	0.2056
Skewness of return	1.3531	0.4569	4.9555	2.7029
Kurtosis of return	4.5733	1.1097	56.9712	24.8478
Average time to matur.	12.0	35.3	12.0	37.2
Std. time to matur.	6.3	12.7	6.3	13.7
Average moneyness	1.002	1.002	1.002	1.002
Std. moneyness	0.005	0.005	0.005	0.005

Year	1986	1987*	1988	1989	1990	1991	1992
Data points available	33	75	51	88	123	152	30

*Only 2 points are used from Oct. 1987: Oct. 6, Oct. 7.

Table 5 - Hedging tests of the volatility term-structure

Hedging the volatility sensitivity of a $100 medium term at-the-money straddle position with short term at-the-money straddles. Volatility hedges based on estimated volatility mo

The hedge portfolio is also delta-neutralized using the underlying asset.
Options on S&P500 index, Jan. 1986- Feb. 1992 (552 observations with all data).

Analysis of daily total change in hedge portfolio price

Volatility hedge portfolios: (using in-sample forecasts)	*Standard deviation*	*Interquartile Range*	*5th percentile*	*F prob. from factor sensitivity regression*
Constant volatility delta-vega hedge	12.04	8.89	-14.16	0.0001
ARIV delta-vega hedge	6.81	6.65	-9.73	0.0001
GARCH(1,1) delta-GARCH gamma hedge	6.04	6.06	-8.63	0.1327
GJR - GARCH(1,1) delta-GARCH gamma hedge	6.04	6.16	-8.70	0.1303
GARCH comp w/lev. delta-GARCH gamma hedge	6.01	6.08	-8.39	0.2213

Volatility hedge portfolios: (out-of-sample forecasts)	*Standard deviation*	*Interquartile Range*	*5th percentile*	*F prob. from factor sensitivity regression*
Trailing historical volatility delta-vega hedge	12.68	8.98	-14.19	0.0001
ARIV delta-vega hedge	6.29	6.58	-9.38	0.0153
GARCH(1,1) delta-GARCH gamma hedge	6.18	6.40	-8.92	0.0120
GJR - GARCH(1,1) delta-GARCH gamma hedge	6.23	6.36	-8.98	0.0031
GARCH comp w/lev. delta-GARCH gamma hedge	6.19	6.24	-8.84	0.0157

Table 6 - Hedging results for benchmark portfolios

Options on S&P500 index, Jan. 1986- Feb. 1992 (552 observations with all data).
Hedging a $100 medium term straddle using other instruments.

Benchmarks:	*Standard deviation*	*Interquartile Range*	*5th percentile*	*F prob. from factor sensitivity regression*
No hedge:				
100$ Medium term straddles	6.62	6.73	-11.43	.0001
Hedge using underlying:				
Constant volatility delta hedge	6.64	6.64	-10.69	.0001
ARIV delta hedge	6.65	6.66	-10.59	.0001
GARCH(1,1) delta hedge	6.67	6.68	-10.45	.0001
GJR delta hedge	6.67	6.67	-10.45	.0001
GARCH comp w/lev. delta hedge	6.68	6.69	-10.45	.0001
Hedge using underlying and short term straddle:				
Trailing 20 day historical volatility delta-vega hedge	12.64	8.93	-15.40	.0001
Implied volatility delta-vega hedge	11.80	8.98	-13.95	.0001
Constant volatility delta-gamma hedge	5.98	6.23	-9.06	.2895
Trailing 20 day historical volatility delta-gamma hedge	5.98	6.32	-9.06	.3291
Ex-post minimum variance hedge	5.96	6.20	-10.07	.8820
Trailing 20 day minimum variance hedge	5.97	5.86	-9.33	.0001

Table 7 - Comparison of GARCH gamma and CV gamma hedging
(Subsample analysis)

The GARCH components model and CV model are estimated over the full sample period.
All options data is used.

	Standard deviation	*Interquartile Range*	*5th percentile*
GARCH comp w/lev. delta-GARCH gamma hedge	6.01	6.08	-8.39
Constant volatility delta-gamma hedge	5.98	6.23	-9.06

100 largest absolute S&P500 returns

	Standard deviation	*Interquartile Range*	*5th percentile*
GARCH comp w/lev. delta-GARCH gamma hedge	7.59	7.95	-7.46
Constant volatility delta-gamma hedge	7.63	8.11	-9.37

100 largest negative S&P500 returns

	Standard deviation	*Interquartile Range*	*5th percentile*
GARCH comp w/lev. delta-GARCH gamma hedge	6.78	6.41	-10.18
Constant volatility delta-gamma hedge	6.82	6.69	-11.06

100 largest positive S&P500 returns

	Standard deviation	*Interquartile Range*	*5th percentile*
GARCH comp w/lev. delta-GARCH gamma hedge	6.53	7.46	-7.57
Constant volatility delta-gamma hedge	6.46	7.32	-7.63

Figure 1 - Comparison of hedge parameters
At-the-money S&P500 index straddles
sigmabar=.01, risk-free rate=0, yield=0, S=K=100