Pricing and Hedging Index Options under Stochastic Volatility: An Empirical Examination

Saikat Nandi

Federal Reserve Bank of Atlanta
Working Paper 96-9
August 1996

Abstract: An empirical examination of the pricing and hedging performance of a stochastic volatility (SV) model with closed form solution (Heston 1993) is provided for options on the S&P 500 index in which the unobservable time varying volatility is jointly estimated with the time invariant parameters of the model. Although, out-of-sample, the mean absolute pricing error in the SV model is always lower than in the Black-Scholes model, still substantial mispricings are observed for deep out-of-the-money options. The degree of mispricing in different options classes is related to bid-ask spreads on options and options trading volume after controlling for moneyness and maturity biases. Taking into account the transactions costs (bid-ask spreads) in the options market and using S&P 500 futures to hedge, it is found that the stochastic volatility model yields lower variance for a minimum variance hedge portfolio than the Black-Scholes model for most classes of options and the differences in variances are statistically significant.

JEL classification: G13

This paper is from the second chapter of the author's dissertation (1995) at Washington University in St. Louis. The author is grateful to his committee members, Kerry Back (chairperson), Phil Dybvig, and especially Steve Heston for various suggestions and comments. He thanks Mike McIntyre for pointing out a missing term in one of the equations and Dan Waggoner for excellent assistance with processing data sets. He also thanks the participants of the sixth annual Cornell-Queens Derivative Securities Conference and the Atlanta Finance Workshop. The views expressed here are those of the author and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Any remaining errors are the author's responsibility.

Please address questions of substance to Saikat Nandi, Research Department, Federal Reserve Bank of Atlanta, 104 Marietta Street, N.W., Atlanta, Georgia 30303-2713, 404/614-7094, 404/521-8810 (fax), snandi@frbatlanta.org.

Questions regarding subscriptions to the Federal Reserve Bank of Atlanta working paper series should be addressed to the Public Affairs Department, Federal Reserve Bank of Atlanta, 104 Marietta Street, N.W., Atlanta, Georgia 30303-2713, 404/521-8020, http://www.frbatlanta.org.

Pricing and Hedging Index Options under Stochastic Volatility: An Empirical Examination

I Introduction

It is well known that volatility of the underlying asset price (as measured by its standard deviation or variance) is one of the most determinants of option prices and hedge ratios. The famous Black-Scholes (henceforth, BS) formula of option pricing is predicated on the assumption of a constant or a deterministically evolving volatility process. Rubinstein (85, 92) document various option pricing anomalies in a BS framework in the equity and index options market, the most notable of which is the volatility smile evident in option prices.[1] Several others have documented similar biases and anomalies in other financial markets, a comprehensive discussion of which can be found in Bates (95). Although the presence of these anomalies/biases does not immediately imply that volatility evolves randomly, a random volatility process that has non-zero correlation with the asset returns process can yield leptokurtosis and skewness in asset returns that have the promise to explain many of these biases. A random volatility process, however, takes us away from the pure arbitrage foundations of option pricing, since volatility is another state variable that needs to be considered in the pricing and hedging of options and volatility or a known function of it is not a traded asset.

The issue of stochastic volatility (henceforth SV) and its effects on option prices has been studied extensively.[2] Some of the earlier work in this area include Johnson and Shanno (87), Hull and White (87), Scott (87), Wiggins (87), Chesney and Scott (89) and Melino and Turnbull (90). The models developed by these authors require either the use of Monte-Carlo simulation or the numerical solution of a two dimensional parabolic partial differential equation (PDE) to get option prices and hedge ratios and are computationally demanding. Some of these models have the questionable assumption of zero correlation between spot asset returns and the volatility and/or zero risk premium on volatility. Stein and Stein (90)

[1]The fact that as the strike price of an option increases, its implied volatility (from the BS model) decreases is often referred to as the volatility smile.

[2]In this paper stochastic volatility models will refer to models in which volatility is driven by innovations that are different from the innovations driving the asset returns, thus precluding the GARCH models.

provide a solution for option prices that can be gotten via bivariate numerical integration. However their model is also developed under the assumption of zero correlation between spot asset returns and the volatility process.

Heston (93) has developed a tractable closed-form solution for European option pricing with stochastic volatility (henceforth, SV) that also permits correlation between asset returns and volatility. Correlation between asset returns and volatility is an important factor for any option pricing model because the substantial negative skewness in the risk neutral distribution of the underlying asset's returns that can give rise to overpricing/underpricing of out-of-the-money calls/puts [see Bates (94)] can only be gotten via a negative correlation between returns and volatility and not just by making the tails of the distribution fatter that any SV model can accomplish. Furthermore, Heston's solution also permits closed form solutions for the option's delta and vega. Bates (96) and Scott (93) use the similar framework as that of Heston to develop models in which the asset returns follow a jump diffusion and Scott also allows for stochastic interest rates.

Although a stochastic volatility option pricing model may be more appropriate than the BS model for many option markets, extant empirical tests of SV models have had to cope with the lack of a closed form solution making it difficult to imply out parameters that drive the volatility process (and whose values are needed to price and hedge options) and the level of volatility from option prices. Using the time series of BS implied volatilities to estimate the parameters and determine the level of volatility in a SV model is inconsistent and Kearns (92) has shown that the method of moments approach used in some papers [Scott (87), Chesney and Scott (89) and Melino and Turnbull (90)] to estimate the parameters of the volatility process is sensitive to the choice of moments.

Furthermore, formally testing the hedging properties of SV models that lack closed-form solutions is very difficult because computing hedge ratios by numerically solving the required PDE or in a Monte Carlo framework is very computationally demanding. Hedging is a very important issue (especially in the real world) and empirical tests of SV models that do not throw light on this issue miss an important dimension. Heston's model (if implemented on European options) provides one with an opportunity to test both the pricing and hedging effectiveness of a SV model without being encumbered by the additional complications of

other SV models and the possible clouding of inference therein.

Bates (95) provides an empirical examination of Heston's SV framework in the currency options market as a subset of his jump diffusion – stochastic volatility model. The currency options are American and the add-on approximation for the early exercise feature is derived from a deterministic volatility setup. There are no out-of-sample tests thus making it difficult to judge whether the improved fit between model and market prices in-sample results form the presence of extra parameters or the model truly does better. Also there is no examination of the very important issue of the hedging performance of the model. Knoch (92) has used Heston's model to explain option prices in the foreign currency options market. However, he estimates the parameters of the volatility process from the time series of exogenous implied volatilities (as given by the Black-Scholes model) thus introducing a potential inconsistency in the methodology and sets the risk premium of volatility to zero.

This paper provides an empirical test of the pricing and hedging effectiveness of the SV model developed by Heston (93) in the S&P 500 index options market. It distinguishes itself from the above mentioned papers in using European options for the empirical test, avoids the inconsistency of using BS volatilities in a SV framework by estimating volatility jointly with other parameters, provides out-of-sample forecasts of option prices of the SV model, relates the out-of-sample model mispricings to bid-ask spreads and trading activity of options and checks, out-of-sample, the hedging performance of the SV model after taking into account the transactions costs (bid-ask spreads) in the index options market. Bakshi et al. (96) also test for the pricing and hedging effectiveness of Heston's SV model (a subset of their stochastic volatility – stochastic interest rates model) in the S&P 500 index options market and in spirit is closest to this paper. However, in their estimation, Bakshi et al. treat the time invariant parameters of the model to be time varying (they estimate these parameters daily) introducing an inconsistency with the prescriptions of the theoretical model. In contrast, this paper treats the time invariant parameters of the SV option pricing model to be constant during the period of estimation and these parameters are estimated by implicitly concentrating out the daily variance from a criterion function that uses both cross-sectional and time-series information and allows for autocorrelation in the residuals. The hedging test is performed only by using one option both for BS and SV model, thus

enabling one to make controlled comparisons between the two models.[3] Also, unlike Bakshi et al., S&P 500 futures are used instead of the spot value of the index available from the data set (that requires transacting in the 500 stocks constituting the index) in the hedge portfolio and round trip transaction costs in the options market are taken into account, resulting in a more realistic hedging test.

The main results of this paper are – Out-of-sample, the SV model results in a decrease of the mean absolute pricing error (where price is measured as the mid point of bid-ask spread) by 24 cents for calls and 15 cents for puts as compared to the BS model. Out-of-sample, both models significantly underprice out-of-the-money puts and overprice out-of-the-money calls, though the degree of mispricing is much lower in the SV model. In addition to strike price and time to maturity biases, the degree of mispricing in both models is related to bid-ask spreads (a form of transactions cost) on options and options trading volume, though trading volume does not seem to have any economic significance. In terms of hedging, where both models are compared using only one option, the SV model always results in lower root mean squared errors and lower variance for a minimum variance hedge portfolio than the BS model for most classes of options and the differences in variances between the two models are statistically significant. This shows that not taking into account the risk from volatility shocks and therefore the shocks to the asset returns process through the correlation between returns and volatility can result in poor hedges in the index options market.

Section II describes the model, Section III describes the methodology and data, Section IV describes the results and Section V concludes.

II Model

Let the asset price S_t and volatility v_t follow the following two processes

$$dS(t) = \mu S dt + \sqrt{v(t)} S \, dz_1(t) \tag{1}$$

$$dv(t) = \kappa[(\theta - v(t)]dt + \sigma\sqrt{v(t)} \, dz_2(t) \tag{2}$$

[3]The hedging criterion used is that of minimum variance hedging that requires only one option even for a SV model.

where z_1 and z_2 are two Wiener processes with correlation ρ. The asset price evolves as a geometric Brownian motion and the variance of the instantaneous returns v_t evolves as a square root process with long-term mean θ.[4] The speed with which v_t reverts towards θ is measured by κ. It can be shown (see the appendix) that the half life of the volatility process is $\frac{\ln(2)}{\kappa}$.

Given this, Heston (93) shows that the price at time t of a European call option that matures at T ($\tau \equiv T - t$) and has an exercise price of K is given by

$$C(S, v, t) = SP_1 - Ke^{-r_f \tau} P_2 \tag{3}$$

where

$$P_j(x, v, \tau; ln(K)) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \mathrm{Re}\left(\frac{e^{-i\phi \ln(K)} f_j(x, v, \tau; \phi)}{i\phi} \right) d\phi \tag{4}$$

where Re() denotes the real part of a complex variable and $x \equiv \ln(S)$.

$$f_j(x, v, \tau; \phi) = e^{C_j(\tau;\phi) + D_j(\tau;\phi)v + i\phi x} \tag{5}$$

where $C_j()$ and $D_j()$ are functions of $\kappa\theta$, $\kappa + \lambda$, ρ, σ, r_f and τ, and λ is a constant such that λv_t measures the volatility risk premium.

Differentiating with respect to S, the delta of the option is given by

$$\frac{\partial C}{\partial S} = \frac{1}{\pi} h_1 + P_1 - \frac{Ke^{-r\tau}}{\pi} h_2 \tag{6}$$

where

$$h_j = \int_0^\infty \mathrm{Re}\left(e^{-i\phi \ln(K) + C_j() + D_j()v + i\phi x} \right) d\phi$$

[4]The variance, v_t follows a square root process and therefore is precluded from being negative [Cox, Ingersoll and Ross, (85)].

Differentiating with respect to v, the vega of the option is given by

$$\frac{\partial C}{\partial v} = \frac{S}{\pi} v_1 - \frac{Ke^{-r\tau}}{\pi} v_2 \tag{7}$$

where

$$v_j = \int_0^\infty \mathrm{Re}\left(\frac{D_j()e^{-i\phi \ln(K)+C_j()+D_j()v+i\phi x}}{i\phi}\right) d\phi$$

The option price and hedge ratios are functions not only of S_t, v_t, r, K, τ (that appear in the Black-Scholes formula of constant or deterministic volatility) but also of $\kappa+\lambda$, $\kappa\theta$, ρ and σ that are unobservable and have to be estimated. $\kappa+\lambda$ and $\frac{\kappa\theta}{\kappa+\lambda}$ are the mean reversion parameter and the long run mean of the risk neutralized volatility process respectively.

III Data and Methodology

III-A Data

Intra day data on S&P 500 index options (ticker symbol SPX) that are traded on the Chicago Board Options Exchange (CBOE) are used to test the model. The market for S&P 500 index options is the second most active index options market in the United States and, in terms of open interest in options, it is the largest. The options are European and settled for a cash amount equal to 100 times the difference between the index level and strike price. Unlike options on the S&P 100 index, there are no wild card features that can complicate valuation. Also there is a very active market for the S&P 500 futures, making the replication of the index much easier. Thus, according to Rubinstein (94) it is one of the best markets for testing a European option pricing model.

The minimum tick for series that trades below \$3 is 1/16 and for all other series the tick is 1/8. Strike prices are spaced 5 points apart for near months and 25 points for far away months. The options expire in the three near-term months in addition to the months from the quarterly cycle of March, June, September, December.

The following rules are applied to filter data needed for the empirical test.

(1) For each day, a bid-ask quote or a transaction has to occur between 9.00 A.M.

(central-time) and 2.30 P.M. (central-time). This criterion excludes quotes and transactions that may be influenced by artificial pricing at the beginning or close of the trading day at the CBOE (Rubinstein, 85).

(2) A transaction has to satisfy the no-arbitrage relationships as outlined in Propositions 1 and 2 of Cox and Rubinstein (85, p. 129-133).[5]

(3) A transaction price has to be at least 50 cents to reduce problems arising from minimum bid-ask spreads (that can constitute a big percentage of price) in low priced options.

(4) All transactions that were executed electronically are excluded.[6]

(5) A transaction cannot be a part of a straddle or a part of trade that involves the simultaneous execution of two or more trades.

Criteria (4) and (5) are used to maintain as much homogeneity in the sample as possible.

Stale index levels of the S&P 500 (stemming from infrequent trading of the constituent securities) could potentially be a problem. One must note that the S&P 500 is a value weighted index and the bigger stocks that trade more frequently constitute the bulk of the index level. Since intra day data and not the end-of-the-day option prices is used, the staleness problem may not be severe enough to undermine the estimation procedure. In theory, one could possibly overcome this problem by using implied index levels from the put-call parity equation. However, this is conditional on put call parity holding as an equality and in the presence of transactions costs (bid-ask spreads that are non-negligible), the equality becomes an inequality. Thus the implied index levels from the put-call parity equation may not equal the true index level. Also, even if one assumes away transactions costs, it is very difficult to create a sample of sufficient size by creating matched pairs of puts and calls because the level of the S&P 500 index changes quite frequently through the day. As a result, the index levels as reported in the data set are used in parameter estimation. However, in the hedging test, S&P 500 futures data is used instead of the spot level of the index. The daily dividend yield on the S&P 500 index are taken from Data Stream. For the risk free rate, I use a series of monthly yields from the daily treasury yield

[5]The propositions are adapted from Merton (73).

[6]Only a small fraction of the trades is electronically executed.

curves prepared by Salomon Brothers and interpolate to match the yield to the expiration date of the option.[7] The filtered data set on which the estimation of the parameters of the SV model is carried out consists of 22,134 transactions, from 01/21/91 to 04/10/92, a total of 310 trading days. Of these 7,853 are calls and 14,281 are puts.

III-B Methodology

Central to the empirical test is the estimation of the unobservable parameters, $\kappa + \lambda$, $\kappa\theta$, ρ, σ and v_t for the stochastic volatility model that appear in the pricing and hedging formulae.[8] Since the variance v_t is unobservable, the estimation of these parameters from the time series of v_t is difficult. Although one has high frequency intra day data, the construction of a daily v_t series from the intra day data is subject to question due to the bid-ask bounce present in the intra day data. For the SV model, the moments of the logarithm of the spot price in the original probability measure are unknown to this author. Although the moments of the log spot price in the risk neutral measure can be computed in this model using Heston's (93) formula for the characteristic function of the log spot price, one cannot use log spot price data to estimate the parameters for the SV model, unless it is assumed that the risk premium of the asset price process is zero. Therefore it does not seem to be feasible to estimate the parameters from the index returns data by using the Generalized Method of Moments of Hansen (82) (henceforth GMM) unless one makes some questionable assumptions about the risk premium of the asset price process. Even if the risk premium is assumed to be zero, Kearns (92) has shown that using GMM to estimate parameters from the returns data is sensitive to the choice of moments. Foster and Nelson (94) describe a procedure to estimate the parameters (i.e., κ, θ, ρ, σ) of the unobservable volatility process from the returns data. Unfortunately, their procedure to estimate the parameters from the daily data of S&P 500 returns did not work on this data set. The procedure involves drawing standard normals and numerical overflow occurred during the unconstrained optimization procedure with some standard normals being drawn from the extreme tail of the distribution. Further, this procedure is not able to identify λ that has

[7] I thank Sanjiv Das for letting me use this data set.

[8] Note that one does not have to separately estimate κ and θ because the pricing and hedging formulae are in terms of $\kappa + \lambda$ and $\kappa\theta$.

to be estimated from the options data conditional on the estimates of other parameters.

Since closed-form solutions are available for the option price, a natural candidate for the estimation of the parameters that enter into the pricing and hedging formulae is a non-linear least squares procedure involving minimization of the sum of squared errors between the model and market prices. Let $O_{i,t}$ denote the market price of option i on day t and $f_{i,t}(\Psi, v_t)$ denote the model price of the option i at time t as a function of the set of parameters Ψ, v_t, where $\Psi \equiv (\kappa+\lambda, \kappa\theta, \rho, \sigma)$ for the SV model. It is assumed that the variance v_t for day t does not vary intra day. Admittedly, this assumption is somewhat inconsistent with the spirit of the model because v_t and S_t are correlated. However, v_t is a parameter to be estimated and allowing v_t to change every time t changes will increase the number of parameters by tens of thousands given the size of the data set. Also the variation in v_t through a typical trading day is rather small (with some of the variation resulting from spurious bid-ask bounce) making this assumption not very unrealistic. Let,

$$e_{i,t}(\Psi, v_t) = O_{i,t} - f_{i,t}(\Psi, v_t) \tag{8}$$

If e_{it} are i.i.d., the criterion function is

$$\min_{\Psi, v_t} \sum_{t=1}^{T} \sum_{i=1}^{N_t} e_{i,t}^2(\Psi, v_t) \tag{9}$$

where T is the number of trading days in the sample and N_t is the number of options on a particular day t in the sample. However, (9) cannot be used if the errors are autocorrelated. Time series correlation or autocorrelation in the residuals (difference between model and market prices) can arise from a misspecification of the model, such as the omission of a relevant state variable or from bid-ask bounce. Since all models are simplifications of the real world, autocorrelation in the residuals is an important issue that ought to be addressed in our testing. A preliminary estimation of the model using (9) indicates that the errors are first order serially correlated. Therefore it is assumed that the residuals evolve as follows

$$e_{i,t} = \rho^* e_{i,t-1} + \nu_t$$

where ν_t is i.i.d.

Given this, it can be shown that the appropriate sum of squared errors (SSE) that one needs to minimize over Ψ, ρ^* and v_t is (refer to Judge et al. (85), pp 283–290)

$$
\begin{aligned}
S(\Psi, \rho^*, v_t) &= \sum_{i=1}^{N_1}(1-\rho^{*2})[O_{i,1} - f_{i,1}(\Psi, v_1)]^2 \\
&\quad + \sum_{t=2}^{T}\sum_{i=1}^{N_t}[O_{i,t} - \rho^* O_{i,t-1} - f_{i,t}(\Psi, v_t) + \rho^* f_{i,t-1}(\Psi, v_{t-1})]^2 \qquad (10)
\end{aligned}
$$

Note that if $\rho^* = 0$, then the above criterion function reduces to the one with i.i.d. residuals. (10) is the total SSE for the sample and is the sum of the SSE of the individual days. The first expression in the RHS representing SSE for day 1 and the next the sum of all SSE $\forall\ t \geq 2$. The SSE for $t \geq 2$ is given by $\sum_{i=1}^{N_t}[O_{i,t} - \rho^* O_{i,t-1} - f_{i,t}(\Psi, v_t) + \rho^* f_{i,t-1}(\Psi, v_{t-1})]^2$. To see why the above criterion function is appropriate, note that the inefficiency of the estimates based on the i.i.d. assumption result from the fact that $E[ee'] \neq \sigma_e^2 I_{N^*}$ where I is the $N^* \times N^*$ identity matrix (N^* is the total number of observations) and σ_e^2 is the variance of the residuals. This implies that we need to transform the residuals such that the transformed residuals are i.i.d. If Ω denotes the covariance matrix of the residuals e_{it}, then we can find a matrix Q such that $Q\Omega Q' = I_{N^*}$. Since Ω is positive definite, such a matrix always exists [Strang, (88)]. Now if we consider the transformed residuals $e^* = Qe$, then it is the case that the covariance matrix of the transformed residuals is

$$
\begin{aligned}
E\left[e^* e^{*\prime}\right] &= E\left[Qee'Q'\right] \\
&= QE\left[ee'\right]Q' \\
&= \sigma_e^2 Q\Omega Q' \\
&= \sigma_e^2 I_{N^*}
\end{aligned}
$$

If the residuals are first order autocorrelated, then the appropriate Q is given in Judge et al. (78, p. 285) and one minimizes the sum of squares of the transformed residuals which is equation (10).

As mentioned before, v_t is a parameter to be estimated. Since it is assumed that v_t is fixed for day t (i.e., does not vary through the day), the number of parameters to be estimated, other than (Ψ, ρ^*), increases by the number of days in the sample. One could possibly overcome this problem by estimating v_t and Ψ (treating ρ^* to be zero) every

day from the daily minimization of the sum of squared errors between model and market option prices as done in Bakshi et al. (96). In this way one has a different Ψ and v_t every day. However, it must be noted that according to the model, Ψ is constant and only v_t varies and estimating Ψ along with v_t daily runs contrary to the model. This paper treats (Ψ, ρ^*) to be constant through the sample (i.e., the same Ψ is used $\forall\ t$) and v_t is estimated simultaneously with (Ψ, ρ^*) under the criterion function (10). The procedure is as follows: Given a particular value of (Ψ, ρ^*), one uses the SSE for day t to estimate v_t, conditional on v_{t-1} (note that the sum of squared errors for day t, $\forall\ t \geq 2$, contains both v_t and v_{t-1}). Since the sum of squared errors for the first day contains only v_1, at the beginning, given (Ψ, ρ^*), v_1 is estimated from the SSE expression for day 1. Thereafter, for each $t \geq 2$, v_t is estimated, conditional on the given Ψ, ρ^* and the estimated v_{t-1}, from the SSE expression for day t. Having estimated v_t for each day t for the given (Ψ, ρ^*), one goes back and estimates (Ψ, ρ^*) conditional on the estimated v_t by minimizing the entire SSE for the sample (i.e., the sum of SSE for the individual days) which is (10). Thereafter, given the new (ψ, ρ^*), re-estimate v_t as described above and keep on iterating between (Ψ, ρ^*) and v_t until convergence according to some particular criterion is achieved.[9] For the estimation of v_t (each day given (Ψ, ρ^*)), we have on an average 50 options transactions that should be sufficient for the reliable estimation of one parameter. Note that this procedure yields estimates for Ψ conditional on the estimates for v_t and similarly yield estimates for v_t conditional on Ψ. Thus the standard errors obtained are conditional standard errors. Implicitly, the variance v_t is being concentrated out of the criterion function [conditional on (Ψ, ρ^*)] where as ideally one would want to integrate out v_t from the criterion function.

(10) assumes that Ψ is constant over the time period of estimation and this criterion function enables one to utilize both the cross-sectional and time series information in the sample. Heteroscedasticity in the residuals is taken into account in constructing the asymptotic variance-covariance matrix using White's (84) procedure. $\kappa + \lambda$, $\kappa\theta$ and v_t are constrained to be non-negative by an exponential transformation.[10] ρ and ρ^* are constrained between -1 and 1 by a sine/cosine transformation. The SV model price for the

[9]The basic NLS routine used is Marquardt's method as given in Press et al. (92).

[10]$\kappa + \lambda$ has to be non-negative for stationarity of the risk neutralized variance processes and $\kappa\theta$ has to be positive for the long-run mean of the risk neutralized variance process to be non negative.

option is evaluated by Gaussian quadrature (with hermite polynomials) and the number of quadrature points used is 80.[11] Derivatives of the model option price with respect to Ψ are evaluated numerically.

Ideally parameter estimation would use all the information in the index values and option prices. For example, in the SV model, one could maximize the full likelihood function which recognizes innovations in the index values, innovations in the volatility and errors in the option prices. Because volatility is a latent variable, it should be integrated out to arrive at the full likelihood function. In contrast to our procedure, full information likelihood would identify all the parameters κ, θ and λ. Unfortunately, maximizing the full likelihood function seems infeasible given current technology. Even the Gibbs sampling procedure [Gefland and Smith (90)] is difficult to apply because the non-linearity of the option prices in the index values and variance makes the required conditional densities intractable.

The results of the estimation for call and put options are discussed in section IV.

III-C Minimum Variance Hedging

A complete test of an option pricing model should take into account not only how well the model prices match observed option prices, but also, and equally importantly, the ability of the model to hedge against changes in the underlying state variables. The hedging issue is of interest not only to academics, but also to actual traders in the option market who often have to hedge their positions in the underlying asset or options market.

In the BS model of option pricing, if one creates a portfolio consisting of a long position in a certain number of units of the underlying asset (as given by the option's delta) and short one unit of the call option, then the portfolio has no risk over the next instant. In this case, there is only one state variable (the asset price) and one can create a hedge portfolio by trading only in the underlying asset and one option. However, if there are two state variables as in the SV model, then one more option is needed to hedge the additional volatility risk and three assets are needed to form an instantaneously risk-free portfolio.

[11] For most option records, we need a maximum of 20 quadrature points to arrive at a reliable option price. However, experimentation with artificial data indicates that a greater number of quadrature points is needed for out-of-the-money options. The quadrature routine that we use is "gauher" in Press et al.

In theory, one can compare the hedging performance of the two models by constructing the respective instantaneous risk free portfolios. However, the SV model does not specify a unique three asset hedge. As a result a controlled comparison of the hedging performance of the two models is difficult using this type of criterion. So this paper takes a different approach, using one option in each case to compare the hedging performance of the SV and BS models

Let us suppose that we create a portfolio consisting of a long position in 1 unit of the underlying asset and a short position in h units of an option. Then the variance of the change in the portfolio value over the next instant is $\text{var}(dS - hdC)$. One can choose h to minimize the variance of this portfolio over the next instant. It can be shown that for the BS model, the minimum variance hedge ratio h is $\frac{1}{C_S}$ and for the SV model, the minimum variance hedge ratio is $\frac{C_S S_t^2 + C_v \rho \sigma S_t}{C_S^2 S_t^2 + C_v^2 \sigma^2 + 2C_S C_v \rho \sigma S_t}$ (the proofs are in the appendix). In the BS model, the minimum variance hedge ratio is the same as the reciprocal of the delta of the option.[12] On the other hand, the minimum variance hedge ratio for the SV model not only takes into account the sensitivity of the option price to the underlying asset price and the variance, but also the correlation between the asset price and variance as well as the volatility of variance. Since S&P 500 futures data are used instead of the spot level of the S&P 500 index (that is available in the options data set), one has to take a long position in $exp(r_{f,t} - d_t)$ units of futures in the hedge portfolio, where $r_{f,t}$ is the risk free yield on day t corresponding to the selected futures and d_t is the dividend yield on day t for the S&P 500 index.

In order to perform the out-of-sample hedging analysis, v_t is estimated each day by minimizing the sum of squared errors between model and market prices, conditional on Ψ that has been estimated from the past data. Note that according to our convention, if h is negative, one has a long position in that option. The way a time series of hedging returns is created is as follows: On day t, (given Ψ and $v_{t-1}$for the SV model and the previous day's implied volatility for the BS model), an option in a particular moneyness-maturity

[12]Note that for the DV model the minimum variance hedge ratio is the reciprocal of the delta of the option because a portfolio short in $\frac{1}{\text{delta}}$ units of the option and long one unit of the asset is riskless over the next instant with zero variance.

class is selected for the first bid-ask quote in that option after 9 A.M. by buying/selling h units at the ask/bid. Immediately after that, one takes a long position in S&P 500 futures by selecting the next futures transaction (from the futures data set) that occurs after the selected bid-ask quote in the options data set. Thus one has a portfolio of a long position in a given number of units of the S&P 500 index futures and a long/short position in h units of the option. Then on day $t+1$, or $t+4$ (two rebalancing intervals are used) the value of the portfolio is computed from the first bid-ask quote in that option after 9 A.M. and the next available futures transaction. The value of the option position on day $t+1/t+4$ is calculated by liquidating the previous long/short position at the bid/ask i.e., if one bought/sold options on the day t, these options are sold/bought on day $t+1/t+4$, thus recognizing the round trip transactions costs in the options market. The normalized change in the hedge portfolio is computed by dividing the change in the portfolio value by the value of the portfolio on day t.[13] This portfolio is formed for all options belonging to a particular maturity/moneyness class and the normalized change of the hedge portfolio of a particular option class is the sum of normalized changes of the individual portfolios in that class. Repeating this procedure for all days in our sample, time series of normalized changes of the minimum variance hedge portfolios of options belonging to various maturity and moneyness classes are created with one and three day rebalancing intervals.

Trades in the S&P index options are for a minimum of one contract (100 options) and multiples thereof and the bid-ask quotes on index options is firm only for ten contracts (i.e., 1000 options). The required number of options to be traded (as given by the models) at the quoted bid/ask may be in excess of ten contracts. However, in order to get around this apparent problem, one has to model how market makers revise their extant quotes in response to trades that exceed the minimum trade size, an issue outside the scope of this paper.

The minimum variance portfolio is being rebalanced at discrete intervals where as the minimum variance formula is based on the notion of continuous rebalancing. The discrete rebalancing interval introduces an additional error term that has the effect of making

[13]With futures as part of the portfolio, returns of the hedge portfolios are not well defined, hence the normalized change

the returns of the minimum variance portfolio heteroscedastic [Boyle and Emanuel, (80)] and possible model misspecification will introduce autocorrelation in the portfolio returns. Therefore, a GMM statistic is used to compare the normalized changes of the minimum variance portfolio of the two models and is discussed below.

In order to assess the hedging performance of the two models, one has to compare how the variances of the daily normalized changes in the hedge portfolios from the two models compare with each other. Let $r_{bs,t}$ and $r_{sv,t}$ denote the normalized changes in the hedge portfolio minus the population mean (estimated by the sample mean) of the hedge portfolios corresponding to the BS and the SV models respectively. We want to test whether

$$\begin{aligned} E(r_{bs,t})^2 - E(r_{sv,t})^2 &= 0 \\ or, E[(r_{bs,t})^2 - (r_{sv,t})^2] &= 0 \end{aligned} \tag{11}$$

Let $Z_t \equiv {r_{bs,t}}^2 - {r_{sv,t}}^2$. Then we want to test whether $E(Z_t) = 0$. A GMM statistic (in the Newey-West framework) for testing this is given by

$$V = \sqrt{T}\frac{\sum_{t=1}^{T}(Z_t - \mu_z)}{\sqrt{\sum_{t=1}^{T}(Z_t - \mu_z)^2 + 2\sum_{j=1}^{m}\sum_{t=j+1}^{T}(Z_t - \mu_z)(Z_{t-j} - \mu_z)}} \tag{12}$$

where μ_z represents the sample mean of Z_t and V is distributed asymptotically as a standard normal. The choice of m is somewhat ad hoc. However, different values of m are tested with and the results discussed in a following section remain essentially unchanged.

IV Results

IV-A Parameter Estimates

The estimates of the various parameters of the two models from are reported in Table 1 for calls and puts separately. All the reported estimates are from using transaction prices. Although one may want to use the mid-point of bid-ask as the option price in the estimation, the number of bid-ask quotes in the data set is enormous as compared to the number of transactions, making it extremely time consuming to estimate the parameters of the SV

model for the entire sample period using bid-ask quotes.[14] However, for a smaller sample, the mid-point of bid-ask quotes was used for estimating the parameters of the SV model and there were not any discernible changes in the parameter estimates relative to using transaction prices.

For the SV model, the long-run volatility (standard deviation) of the risk-neutralized variance process is 16.73% (as measured by $\sqrt{\frac{\kappa\theta}{\kappa+\lambda}}$) . The half-life of the variance process is 2.47 months for calls and 3.01 months for puts (as measured by $\left((\frac{\ln(2)}{\kappa+\lambda}\right))$ under the risk neutral distribution. As noted before, the estimation procedure is unable to identify λ and, therefore, cannot identify κ and θ which are the mean reversion coefficient and the long run mean under the empirical measure. However, one does not need the values of λ, κ and θ separately to price and hedge options as only the risk neutralized parameters enter the pricing and hedging formulae. For the SV model, σ is around 0.19 for both calls and puts. ρ is -0.82 for calls and -0.88 for puts indicating that volatility and the S&P 500 index level are highly negatively correlated.

A rough estimate of σ, given a discrete time series of v_t is $\sqrt{\frac{(\sum_{t=1}^{t=T}(v_t-v_{t-1})^2}{\sum_{t=1}^{t=T} v_t}}$.[15]

The estimated values of σ seems unusually high given the estimated v_t series in Figure 1 and Figure 2 for calls and puts, thus driving a wedge between σ estimated from option prices and a rough estimate of σ from the estimated time series of v_t. A similar result for the volatility of the square root variance process has been reported by Bates (96) for currency options using Heston's SV model. This may indicate that the square root variance process is misspecified. In this context, one must note that for index options a deterministic volatility model i.e., BS is known to severely underprice out-of-the-money puts and overprice out-of-the-money calls [Dumas et. al (96)] that indicates extreme left skewness in the risk neutral distribution of the index returns. The SV model, while minimizing the sum of squared errors, perhaps tries to account for this skewness through an increased σ and a highly negative ρ.

[14]The estimation of the parameters of the SV model takes around 3.5 days on a non-dedicated SPARC 10 workstation, using transaction prices for the entire sample period.

[15]The rough estimate of σ is computed from the fact that $dv_t{}^2 = v_t\sigma dt$ and approximating the integration in both sides by summation.

IV-B Out-of-Sample Pricing and Pricing Biases

The results of out-of-sample pricing from 04/13/92 to 11/05/92 (132 days) are shown in Table 2 for the SV and BS models. For the BS model a daily implied volatility is computed by minimizing the sum of squared errors (SSE) between model and market prices. For the SV model, conditional on the estimates of Ψ obtained from the previous period, a daily variance v_t is computed by minimizing the SSE between model and market prices. The implied volatility on day t is then used to price the options on day $t+1$ for all options under both models. Note that in doing this the BS model is getting an unfair advantage because the BS model is based on the notion of a constant volatility and one is estimating the volatility daily instead of estimating a single volatility from the previous period (over which Ψ was estimated for the SV model) and using this volatility to price options in the second period. However, since volatility is indeed time varying, doing such an exercise would yield unacceptably high SSE in-sample and perhaps even more in out-of-sample for the BS model. What one is interested in here is that if someone had carried out the estimation of the SV model and priced options in a subsequent period on the basis of that, would it have resulted in prices that were closer to the market than BS prices, where the BS price is computed (as is normally done in the industry) by estimating a daily implied volatility on the previous day? The option prices that are used now are the mid-points of bid-ask because in computing the actual out-of-sample mispricing, one would want to avoid the spurious mispricing from the bid-ask bounce that results from using transaction prices. In addition, similar to Longstaff (95), one would want to check the effects of bid-ask spreads on options mispricing which is possible only if one is using the bid-ask records. This may raise the question as to whether using estimates obtained from using transaction prices should be used to compute mispricings with respect to the mid-point of bid-ask. However, as noted before this was done to avoid computational burden and the use of a smaller sample for the SV model did not show any particular change in the parameter estimates from using either of the two price series.

The average absolute dollar mispricing for calls (with 318,197 observations and an average bid-ask spread of 63.4 cents) under BS is 67 cents and under SV is 43 cents Thus one is able to improve mispricing by around 24 cents using the SV model for calls and bring the

average option price within the bid-ask spread. For puts (with 333,390 observations and an bid-ask spread of 62.4 cents), the average absolute dollar mispricing is 77 cents using BS and 62 cents using SV, an improvement of 15 cents. Looking at Table 2, one can see that the pricing differences between the BS and the SV model are especially significant for deep out-of-the-money puts and out-of-the-money calls. For example the average absolute percentage mispricing for deep out-of-money calls ($\frac{K}{S} - 1 > 0.06$) in the 0-3 months maturity range is 149% for BS and 74% for SV. In fact out-of-the-money calls are significantly mispriced by both models for all maturities. Similarly, for deep out-of-the-money puts in the 3-9 months range, the average absolute mispricing is 52% for BS and 17% for SV. If one computes the raw (not absolute) mispricings (market price - model price) for out-of-the-money options, one finds that both models significantly underprice out-of-the-money puts and overprice out-of-the-money calls, although the amount of mispricing is much lower for the SV model. The directions of mispricing for out-of-the-money options are in contrast to those for individual equity options reported in Rubinstein (85) who found that out-of-the-money calls are always underpriced by the BS model indicating a thicker right tail in the risk-neural distribution of spot equity returns. The direction of mispricing for out-of-the-money options is consistent with an unusual high concentration of probability mass in the left tail of the risk-neutral distribution of the index returns that could have resulted from a crash-phobia regarding the market and have also been reported by Dumas et. al (96) for S&P 500 index options and Bates (94) for options on S&P 500 futures. Although the SV model can put extra mass in the left tail of the risk-neutral distribution by increasing σ that makes the tails fatter and a negative ρ that makes the distribution skewed to the left, it seems that market priced these options under a risk-neutral distribution that seems implausible under this SV model. It is also to be noted that the SV model does not yield better prices than the BS model for all moneyness and maturity categories, though it does for most. For example, for near-the-money puts in the (-0.01,0.01) moneyness range (moneyness is defined to be $\frac{K}{S} - 1$) and maturity range of 0-3 months, the absolute percentage mispricing under BS is 8% as compared o 16% for the SV model. Also for call options, under both models, the mispricing seems to decrease with an increase in maturity.

Following Longstaff (95), the effects of market frictions like bid-ask spread and measure

of trading activity like trading volume are checked on the amount of mispricing for both models by regressing the out-of-sample raw mispricing on these variables after controlling for moneyness and maturity biases i.e., by including them as additional independent variables. However instead of using all options like Longstaff (95), one checks for the mispricing biases by looking at various classes of calls and puts (classified by moneyness and time to maturity) separately. The trading volume used is the sum total of options contracts that were transacted for a particular class of options on each day. The results of these regressions are reported in Table 3(a) and Table 3(b). For the BS model the coefficient on the bid-ask spread is significant for all ten classes of options and for the SV model, this coefficient is significant for eight out of ten classes of options (out-of-the-money calls and long-maturity puts being the exceptions). For the BS model, the coefficient on the bid-ask spread is always negative. This suggests that if the option is underpriced/overpriced by BS, then an increase in the bid-ask spread increases/decreases the underpricing/overpricing. The effects of the bid-ask bias on the SV model are mixed, with some option classes reflect the same bias as the BS model and some other the opposite. For both models, the coefficient on trading volume is very small and though statistically significant for many classes of options, does not seem to have any economic significance. The coefficient on the moneyness is always significant for both models, indicating that both models have strike price biases. The coefficient on time to maturity is also always significant for both models indicating time to maturity biases, although the size of the coefficient is lower than the one on moneyness.

Thus, although the SV model yields lower pricing errors than the BS model, it still suffers from bid-ask, strike price and time to maturity biases. Overall, the adjusted R^2 in these regressions tend to be somewhat lower for the SV model i.e., the degree of variation in the mispricings that can be explained by the variations in these variables is lower for the SV model than the BS model. The bid-ask bias suggests that the bid-ask is proxying for variables that are not captured by any of the two models and it is interesting to note in this context that in options market the bid-ask spread could be related to the information asymmetry about the underlying asset and/or volatility of the underlying asset [see Back (93) and Nandi (95)], features that are not captured by the extant option pricing models that assume symmetric information among investors about the variables that affect the

price of options.

IV-C Hedging Results

The minimum variance hedging results are reported in Table 4(a) and 4(b) for one day and three day rebalancing intervals respectively. For the one day rebalancing interval, the SV model yields a lower variance of the hedge portfolio than the BS model for eight of the ten classes of options that were considered. Taking into account the autocorrelation and heteroscedasticity in the normalized changes in the hedge portfolios, the differences in variances of the hedge portfolios are statistically significant at the 5% level. Although the variances of the hedge portfolios under the BS model are smaller for near-the-money calls and long-maturity puts, the differences are not statistically significant according to the GMM statistic. Similar results are obtained for the three day rebalancing interval, except that an increase in the rebalancing interval is accompanied with an increase in the variance of the hedge portfolios for both models. For all classes of options, with a three day rebalancing period, the volatility of the hedge portfolio under the BS model is greater than the daily volatility of the S&P 500 index over this time period (0.06%). For the SV model, such an increase is observed only for short-maturity puts and in-the-money puts. The differences between the two models in terms of hedging are especially pronounced for out-of-the-money options, a class where there are substantial differences in terms of pricing too. Also for both models, the calls are better hedged than the puts. In terms of root mean squared error (RMSE) of the hedge portfolios, where the RMSE is defined to be the root mean squared returns (it is assumed that a perfect hedge portfolio would earn a zero return), the SV model also tends to dominate the BS model under both rebalancing intervals, with the RMSE also increasing with an increase in the rebalancing interval.

The hedging results suggest that failure to take into account the risk of a randomly changing volatility can result in poor hedges in the index options market for most options classes. This result is encouraging in the context of results reported in Dumas et al. (96) who found that models based on implied binomial trees [Rubinstein (94)] where volatility is implicitly modeled to be a deterministic function of the level of stock price and time performed poorly than the simple BS model in an out-of-sample hedging scenario. An

important feature of the SV model is that not only does this model allow for volatility shocks that are different from the shocks to the returns process, but volatility and spot index returns are negatively correlated (an often documented feature of equity indices) allowing this model to capture the changes in the index returns resulting from volatility shocks.

V Conclusion

The stochastic volatility option pricing model developed by Heston (93) that gives closed form solutions for option prices and hedge ratios was tested for its pricing and hedging effectiveness in the S&P 500 index options market. In order to avoid the potential inconsistency of using either Black-Scholes implied volatilities in a stochastic volatility model or estimating the time invariant parameters of the model daily, the unobservable time varying volatility was estimated jointly with the other time invariant parameters, using both cross-sectional and time series data.

Out-of-sample, the SV model yields a lower mean absolute pricing error than the BS model. However, both models consistently underprice out-of-the-money puts and overprice out-of-the-money calls, although the amount of mispricing is much lower in the stochastic volatility model. This type of mispricing indicates an unusual concentration of probability mass in the left tail of the risk neutral distribution of the index returns, part of which can result from negative correlation between index returns and volatility. The degree of options mispricing is related to the bid-ask spreads on options and options trading volume after controlling for moneyness and maturity biases, which shows that transaction costs and options trading activity are important in explaining intra day options prices. However, it is possible that bid-ask spreads are proxying for other effects such as information asymmetry about the index and/or volatility of the index as well as discrete rebalancing risk.

In terms of minimum variance hedging, after taking into account the transactions costs (bid-ask spreads) in the index options market and using S&P 500 futures to hedge, the SV model results in lower variance of the hedge portfolio than the BS model for eight out of ten classes of options and the differences in variances are statistically significant. This

shows that option pricing models that do not incorporate shocks to the volatility process and its subsequent effect on returns through the correlation between returns and volatility can result in poor hedges in the index options market. Future research can be directed towards developing and estimating SV models that can account for the transaction costs biases in option prices.

Appendix

The purpose of this appendix is to derive a formula for the half-life of the square root volatility process and the minimum variance hedge ratios for the SV and the BS model.

If variance of asset returns follows the process prescribed in (2) and $t - s \equiv \tau$ then (please refer to [Cox,Ingersoll, and Ross(85)])

$$E\left[v_s|v_t\right] = v_t e^{-\kappa\tau} + \theta(1 - e^{-\kappa\tau}) \tag{13}$$

Now if τ measures the half life of the volatility process, it means that half of the volatility shock of time t has been dissipated at time s. This implies

$$\begin{aligned} \frac{v_t + \theta}{2} &= v_t e^{-\kappa\tau} + \theta(1 - e^{-\kappa\tau}) \\ \Rightarrow v_t(\frac{1}{2} - e^{-\kappa\tau}) &= \theta(\frac{1}{2} - e^{-\kappa\tau}) \end{aligned} \tag{14}$$

If the above equation is to hold for all values of v_t it is required that $\frac{1}{2} - e^{-\kappa\tau} = 0$. This in turn yields, $\tau = \frac{\ln(2)}{\kappa}$. Next, we derive the minimum variance hedge portfolios in the two setups. If the asset price S_t and volatility v_t follow the processes as in (1) and (2) then, the price of an option $C_t = f(S_t, v_t, t)$ follows the process

$$\begin{aligned} dC_t &= f_t dt + f_S\left[\mu S dt + \sqrt{v} S dz_1\right] + f_v\left[\kappa(\theta - v)dt + \sigma\sqrt{v}\, dz_2\right] \\ &\quad + \frac{1}{2} f_{SS} v S dt + \frac{1}{2} f_{vv}\sigma^2 v dt + f_{Sv}\rho\sigma v dt \end{aligned} \tag{15}$$

Hence the variance of a portfolio which is long 1 unit of the asset and short h units of the option is

$$\begin{aligned} \mathrm{var}(dS - hdC) &= \mathrm{var}(dS) + h^2\mathrm{var}(dC) - 2h\mathrm{Cov}(dS, dC) \\ &= S^2 v dt + h^2(f_S^2 S^2 v + f_v^2 v\sigma^2 + 2 f_S f_v S v\rho\sigma)dt \\ &\quad - 2h(f_S v S^2 + f_v\rho\sigma v S)dt \end{aligned} \tag{16}$$

Differentiating with respect to h we get the required result for the minimum variance hedge ratio in the SV model. A similar calculation for the BS model yields the minimum variance

hedge ratio for that model.

REFERENCES

Back, K. , 1993, "Asymmetric Information and Options", *Review of Financial Studies*, 6, 435-472.

Bailey, W., and R. Stultz, 1989, "The Pricing of Stock Index Options in a General Equilibrium Model", *Journal of Financial and Quantitative Analysis*, 24, 1-12.

Bakshi, G., C. Cao and Z. Chen, 1996, "Option Pricing and Hedging Performance under Stochastic Volatility and Stochastic Interest Rates", Working Paper, Ohio State University.

Bates, D., 1996 , "Jumps & Stochastic Volatility: Exchange Rate Processes Implicit in Deutschemark Options", *Review of Financial Studies*, 9, 69-107.

Bates, D., 1995, "Testing Option Pricing Models", Working Paper, The Wharton School, University of Pennsylvania.

Bates, D., 1994 , "The Skewness Premium: Option Pricing Under Asymmetric Processes", Working Paper, University of Pennsylvania.

Black, F. And M. Scholes, 1973, "The Valuation of Options & Corporate Liabilities", *Journal of Political Economy*, 81, 399-417.

Box, G. and G. Jenkins, 1976, "Time Series Analysis: Forecasting and Control", Holden-Day.

Boyle, P., and D. Emanuel, 1980, "Discretely Adjusted Option Hedges", *Journal of Financial Economics*, 8, 259-282.

Brown, S., and P. Dybvig, 1986, "The Empirical Implications of the Cox, Ingersoll, Ross Theory of the Term Structure of Interest Rates", *Journal of Finance*, 41, 617-630.

Chesney, M., and L. Scott, 1989, "Pricing European Curreny Options: A Comparison of the Modified Balck-Scholes Model and a Random Variance Model", *Journal of Financial and Quantitative Analysis*, 24, 267-284.

Cox, J., J. Ingersoll, and S. Ross, 1985, "A Theory of the Term Structure of Interest Rates", *Econometrica*, 53, 385-408.

Cox, J., and M. Rubinstein, 1985, "Options Markets", Prentice - Hall, Inc., New Jersey.

Dumas, B., J. Fleming, and R. Whaley, 1996, "Implied Volatility Functions: Empirical Tests", Working Paper, Jones Graduate School of Administration, Rice University.

Gefland, A., and A. Smith, 1990, "Sampling Based Approaches to Calculating Marginal Densities", *Journal of the American Statistical Association*, 85, 398-409.

Hansen , L.P., 1982, "Large Sample Properties of Generalized Methods of Moments Estimation", *Econometrica*, 50, 1029-1054.

Heston, S., 1993, "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options", *Review of Financial Studies*, 6, 327-343.

Hull, J.C. & A. White, 1987, "The Pricing of Options on Assets with Stochastic Volatilities", *Journal of Finance*, 42, 281-300.

Johnson, H. and D. Shanno, 1987, "Option Pricing when the Variance is Changing", *Journal of Financial and Quantitative Analysis*, 22, 143-151.

Judge, G., W. Griffiths, R. Hill, H. Lutkepohl & T. Lee, 1985, "The Theory & Practice of Econometrics", John Wiley & Sons.

Kearns, P., 1992, "Pricing Interest Rate Derivative Securities When Volatility is Stochastic", Working Paper, University of Rochester.

Knoch, H.J., 1992, "The Pricing of Foreign Currency Options with Stochastic Volatility", Ph.D. Dissertation, Yale School of Organization and Management.

Longstaff, F., 1995, "Option Pricing and the Martingale Restriction", *Review of Financial Studies*, 8, 1091-1124.

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

Merton, R., 1973, "Theory of Rational Option Pricing", *Bell Journal of Economics and Management Science*", 4, 141-183.

Nandi, S., 1995 , "Asymmetric Information about Volatility and Options", Working Paper 95-19, Federal Reserve Bank of Atlanta.

Press, W., S. Teukolsky, W. Vetterling & B. Flannery, 1992, " Numerical Recipes in C - The Art of Scientific Computing", Cambridge University Press.

Rubinstein, M., 1994, "Implied Binomial Trees", *Journal of Finance*, 69, 771-818.

Rubinstein, M., 1985, "Nonparametric Tests of Alternative Option Pricing Models Using All Reported Trades and Quotes on the 30 Most Active Option Classes from August 23, 1976 through August 31, 1978", *Journal of Finance*, 40, 455-480.

Scott, L.O., 1987, "Option Pricing When the Variance Changes Randomly: Theory, Estimation & Application", *Journal of Financial and Quantitative Analysis*, 22, 419-438.

Scott, L.O., 1993, "Pricing Stock Options in a Jump-Diffusion Model with Stochastic Volatility and Interest Rates: Applications of Fourier Inversion Methods", Working Paper, University of Georgia.

Stein, E.M. & J.C Stein, 1991, "Stock Price Distributions with Stochastic Volatility: An Analytic Approach", *Review of Financial Studies*, 4, 727-752.

H. Stoll & R. Whaley, 1993, "Futures & Options, Theory & Applications", Southwestern Publishing Company.

Strang, G., 1988, "Linear Algebra and Its Applications", Harcourt Brace Jovanovich.

Wiggins, J.B., 1987, "Option Values under Stochastic Volatilities", *Journal of Financial Economics*, 19, 351-372.

Table 1
Parameter Estimates

Non-Linear Least Squares parameter estimates of the stochastic volatility model using transaction prices from 01/21/91 to 04/10/92 (310 days). Total number of observations = 22,134. Number of calls = 7853 and number of puts = 14,281. Half Life of the variance process is given by $\frac{\ln(2)}{\kappa+\lambda}$ and is measured in months. Asymptotic standard errors appear in parentheses. These standard errors are conditional standard errors (conditional on the estimated v_t series). The in-sample sum of squared errors (SSE) reported is for the entire sample and is the sum of the daily SSEs.

Data Set	SSE	ρ^*	$\kappa\theta$	$\kappa+\lambda$	σ	ρ	Half-Life
Calls	2316.53	0.31	0.094	3.371	0.194	-0.82	2.47
		(0.19)	(0.0015)	(0.225)	(0.013)	(0.093)	
Puts	5,183.14	0.48	0.077	2.758	0.191	-0.88	3.01
		(0.26)	(0.0008)	(0.143)	(0.019)	(0.096)	

Table 2
Out-of-Sample Pricing Errors, 04/13/92 - 11/05/92

Table for average out-of-sample absolute mispricing of different classes of options by moneyness and maturity. Moneyness is defined as (K/S)-1, where K is the strike price and S is the spot price. Mispricing for a single option is defined to be | Model Price - Market Price |/Market Price and the market price is measured as the mid point of bid-ask. The average mispricing in one particular class is the sum of the individual mispricings in that class divided by the total number of observations in that class. Also shown is the average absolute dollar mispricng for calls/puts which is the sum of the absolute mispricings (in $) for all calls/puts divided by the total number of calls/puts. Number of calls = 318,197 (average bid-ask spread of 63.4 cents) and number of puts = 333,390 (average bid-ask spread of 62.4 cents). SV is the stochastic volatility model and BS is the Black-Scholes model.

	Moneyness	BS	SV		Moneyness	BS	SV
Short-Term Calls (0-3 months)	< -0.06	0.03	0.024	Short-Term Puts (0-3 months)	< −0.06	0.60	0.42
	(-0.06,-0.01)	0.03	0.04		(-0.06,0.01)	0.21	0.28
	(-0.01,0.01)	0.18	0.11		(-0.01,0.01)	0.08	0.16
	(0.01,0.06)	0.55	0.23		(0.01,0.06)	0.03	0.03
	> 0.06	1.49	0.74		> 0.06	0.02	0.01
Medium-Term Calls (3-9 months)	< -0.06	0.03	0.01	Medium Term Puts (3-9 months)	< −0.06	0.52	0.17
	(-0.06,-0.01)	0.02	0.02		(-0.06,-0.01)	0.21	0.06
	(-0.01,0.01)	0.08	0.06		(-0.01,0.01)	0.07	0.04
	(0.01,0.06)	0.18	0.11		(0.01,0.06)	0.03	0.03
	> 0.06	0.85	0.45		(> 0.06)	0.03	0.02
Long-Term Calls (9 months & above)	< -0.06	0.03	0.01	Long-Term Puts (9 months & above)	< -0.06	0.40	0.11
	(-0.06,-0.01)	0.04	0.02		(-0.06,-0.01)	0.17	0.03
	(-0.01,0.01)	0.03	0.06		(-0.01,0.01)	0.11	0.04
	(0.01,0.06)	0.10	0.09		(0.01,0.06)	0.04	0.04
	> 0.06	0.35	0.27		> 0.06	0.02	0.03
Average Absolute $ Mispricing		$0.67	$0.43			$0.77	$0.62

Table 3(a)
Mispricing Biases for the Black-Scholes Model

This table shows the OLS regression of the mispricing of different options classes on the bid-ask spread, moneyness, maturity and trading volume. Bid-Ask spread is the relative bid-ask spread i.e, $\frac{ask-bid}{0.5*(bid+ask)}$ and trading volume is the total number of contracts transacted in that option category on a particular day. The coefficient on trading volume is multiplied by 10^9. T statistics apear in parentheses. Mispricing is the actual mispricing defined as (Model price - Market price)/Market price where the market price is measured as the mid point of bid-ask. An option is near-the-money if |(S-K)/S| <0.01, a call is out-of-the money if (K-S)/S >0.01 and a call is in-the-money if (S-K)/S >0.01. Similarly the moneyness for puts is defined. Short-maturity calls are those that have 0-3 months matrity and long-maturity calls are those that have a time to maturity greater than 3 months.

Option Type	No. of Obs	Intercept	Bid-Ask	Moneyness	Maturity	Volume	Adjusted R^2
In-the-money calls	284,666	-0.92	-0.001	0.97	-0.00001	-9.32	0.36
		(-56.15)	(-3.42)	(81.97)	(-14.45)	(-1.96)	
Near-the-money calls	25,076	-6.69	-0.27	6.97	-0.0004	-384.09	0.75
		(-33.56)	(-21.58)	(35.03)	(-11.30)	(-4.52)	
Out-of-the-money calls	8,449	-6.96	-0.35	7.39	-0.001	-233.80	0.84
		(-18.24)	(-8.13)	(19.88)	(-14.03)	(-1.69)	
Short-Maturity calls	173,538	-1.93	-0.06	2.09	-0.0003	-2.17	0.73
		(-98.56)	(-23.09)	(106.08)	(-22.48)	(0.43)	
Long-Maturity calls	144,654	-1.26	-0.02	1.35	-0.0003	-197.07	0.94
		(-55.92)	(-7.48)	(61.10)	(-4.40)	(-3.92)	
In-the-money puts	297,462	-0.28	-0.02	0.29	-0.00004	253.21	0.95
		(-23.97)	(-11.48)	(25.52)	(-12.10)	(4.540)	
Near-the-money puts	18,786	-3.91	-0.25	3.95	-0.00004	169.55	0.98
		(-5.31)	(-5.87)	(5.36)	(-0.35)	(0.90)	
Out-of-the-money puts	17,136	-6.15	-0.11	6.11	0.0003	263.81	0.98
		(-27.51)	(-2.53)	(26.29)	(4.32)	(1.39)	
Short-Maturity puts	191,208	-0.92	-0.03	0.88	0.0002	879.86	0.91
		(-42.93)	(-8.70)	(41.40)	(12.88)	(5.92)	
Long-Maturity puts	142,178	-1.88	-0.01	1.83	-0.0002	340.76	0.96
		(-92.28)	(-3.92)	(89.45)	(-24.53)	(2.04)	

Table 3(b)
Mispricing Biases for the Stochastic Volatility Model

This table shows the OLS regression of the mispricing of different options classes on the bid-ask spread, moneyness, maturity and trading volume. Bid-Ask spread is the relative bid-ask spread i.e, $\frac{ask-bid}{0.5*(bid+ask)}$ and trading volume is the total number of contracts transacted in that option category on a particular day. The coefficient on trading volume is multiplied by 10^9. T statistics apear in parentheses. Mispricing is the actual mispricing defined as (Model price - Market price)/Market price, where the market price is measured as the mid point of bid-ask. An option is near-the-money if |(S-K)/S| <0.01, a call is out-of-the money if (K-S)/S >0.01 and a call is in-the-money if (S-K)/S >0.01. Similarly the moneyness for puts is defined. Short-maturity calls are those that have 0-3 months matrity and long-maturity calls are those that have a time to maturity greater than 3 months.

Option Type	No of Obs.	Intercept	Bid-Ask	Moneyness	Maturity	Volume	Adjusted R^2
In-the-money calls	284,666	-0.28	0.50	0.25	0.0002	-62.87	0.11
		(-69.69)	(32.15)	(66.61)	(92.13)	(-9.18)	
Near-the-money calls	25,076	-4.00	-2.39	4.01	0.0008	-608.62	0.63
		(-16.96)	(-10.25)	(16.97)	(38.91)	(-5.42)	
Out-of-the-money calls	8,449	-4.62	-0.10	4.72	-0.0004	-172.06	0.91
		(-15.45)	(-0.16)	(16.29)	(-11.54)	(-1.33)	
Short-Maturity calls	173,538	-0.42	0.15	0.36	0.001	-13.70	0.61
		(-37.09)	(2.77)	(28.81)	(88.08)	(-3.34)	
Long-Maturity calls	144,654	-0.14	-0.12	0.12	0.0002	481.82	0.91
		(-77.27)	(-5.91)	(75.09)	(5.41)	(-5.05)	
In-the-money puts	297,462	-0.14	-0.12	0.13	0.0002	481.82	0.91
		(-15.03)	(-3.72)	(14.65)	(74.18)	(10.97)	
Near-the-money puts	18,786	-5.75	-2.17	5.58	0.001	-114.96	0.87
		(-9.87)	(3.98)	(9.60)	(27.13)	(-0.72)	
Out-of-the-money puts	17,136	-1.81	-1.80	1.65	0.0009	34.87	0.89
		(-10.43)	(-3.15)	(9.19)	(23.92)	(0.20)	
Short-Maturity puts	191,208	-1.78	1.63	1.58	0.0009	483.87	0.81
		(-98.90)	(22.19)	(98.53)	(48.61)	(3.55)	
Long-Maturity puts	142,178	-0.47	0.13	0.45	0.00003	807.12	0.97
		(-27.52)	(1.41)	(29.53)	(4.02)	(5.17)	

Table 4 (a)
Out-of-Sample Minimum Variance Hedging – Daily Rebalancing.

This table shows the mean and standard deviation (in absolute numbers, not percentages) of the daily normalized changes of the minimum variance hedge portfolios from 04/13/92 - 11/05/92 (132 days) of different classes of options for daily rebalancing. Also shown is the root mean squared error (RMSE) of the hedge portfolios. An option is near-the-money if |(S-K)/S| <0.01, a call is out-of-the money if (K-S)/S >0.01 and a call is in-the-money if (S-K)/S >0.01. Similarly the moneyness for puts is defined. Short-maturity calls are those that have 0-3 months matrity and long-maturity calls are those that have a time to maturity greater than 3 months. SV is the stochastic volatility model and BS is the Black-Scholes model. V is the GMM test statistic shown in the text for testing the equality of the variances of the minimum variance hedge portfolios of the two models and is asymptotically distributed as a standard normal.

Option Type	Mean		RMSE		Std-Dev		$\mid V_{bs,sv} \mid$
	BS	SV	BS	SV	BS	SV	
In-the-money calls	-0.0027	0.00049	0.004	0.001	0.0024	0.0013	3.14
Near-the-money calls	-0.0019	0.00046	0.002	0.002	0.0014	0.0021	0.34
Out-of-the-money calls	-0.0019	0.00052	0.003	0.0009	0.0025	0.0007	4.22
Short-maturity calls	-0.0021	0.00046	0.003	0.001	0.0023	0.0012	2.19
Long-Maturity calls	-0.0028	0.0006	0.003	0.001	0.0016	0.0011	2.05
In-the-money puts	-0.003	-0.022	0.004	0.002	0.0024	0.0015	2.06
Near-the-money puts	-0.0026	0.0008	0.004	0.004	0.0036	0.0002	3.05
Out-of-the-money puts	-0.00276	0.0003	0.0038	0.0005	0.0042	0.0026	7.06
Short-Maturity puts	-0.00252	0.000722	0.004	0.008	0.0038	0.0011	1.95
Long-maturity puts	-0.00332	-0.0248	0.003	0.004	0.0013	0.0037	1.05

Daily mean return of the S&P 500 index = 0.0002

Standard deviation of the daily S&P 500 index return = 0.006

Table 4 (b)
Out-of-Sample Minimum Variance Hedging – Three Day Rebalancing.

This table shows the mean and standard deviation (in absolute numbers, not percentages)of the normalized changes of the minimum variance hedge portfolios from 04/13/92 - 11/02/92 (129 days) of different classes of options for three day rebalancing. Also shown is the root mean squared error (RMSE) of the hedge portfolio. An option is near-the-money if |(S-K)/S| <0.01, a call is out-of-the money if (K-S)/S >0.01 and a call is in-the-money if (S-K)/S >0.01. Similarly the moneyness for puts is defined. Short-maturity calls are those that have 0-3 months matrity and long-maturity calls are those that have a time to maturity greater than 3 months. SV is the stochastic volatility model and BS is the Black-Scholes model. V is the GMM test statistic shown in the text for testing the equality of the variances of the minimum variance hedge portfolios of the two models and is asymptotically distributed as a standard normal.

Option Type	Mean		RMSE		Std-Dev		$\mid V_{bs,sv} \mid$
	BS	SV	BS	SV	BS	SV	
In-the-money calls	-0.0021	0.0005	0.01	0.002	0.012	0.002	4.94
Near-the-money calls	-0.0008	0.0005	0.009	0.001	0.01	0.001	4.45
Out-of-the-money calls	-0.001	0.001	0.006	0.002	0.0062	0.0017	4.10
Short-Maturity Calls	-0.0011	0.0005	0.001	0.001	0.01	0.0015	4.67
Long-Maturity Calls	-0.0021	0.0007	0.01	0.002	0.01	0.002	4.73
In-the-money puts	-0.004	0.005	0.013	0.04	0.013	0.043	1.06
Near-the-money puts	-0.004	0.0003	0.01	0.002	0.011	0.0022	4.49
Out-of-the-money puts	-0.0041	0.0004	0.008	0.0008	0.007	0.0007	4.11
Short-Maturity puts	-0.0042	0.003	0.02	0.03	0.011	0.034	1.11
Long-maturity puts	-0.0043	0.0012	0.02	0.006	0.009	0.005	3.24

Daily mean return of the S&P 500 index = 0.0002

Standard deviation of the daily S&P 500 index return = 0.006

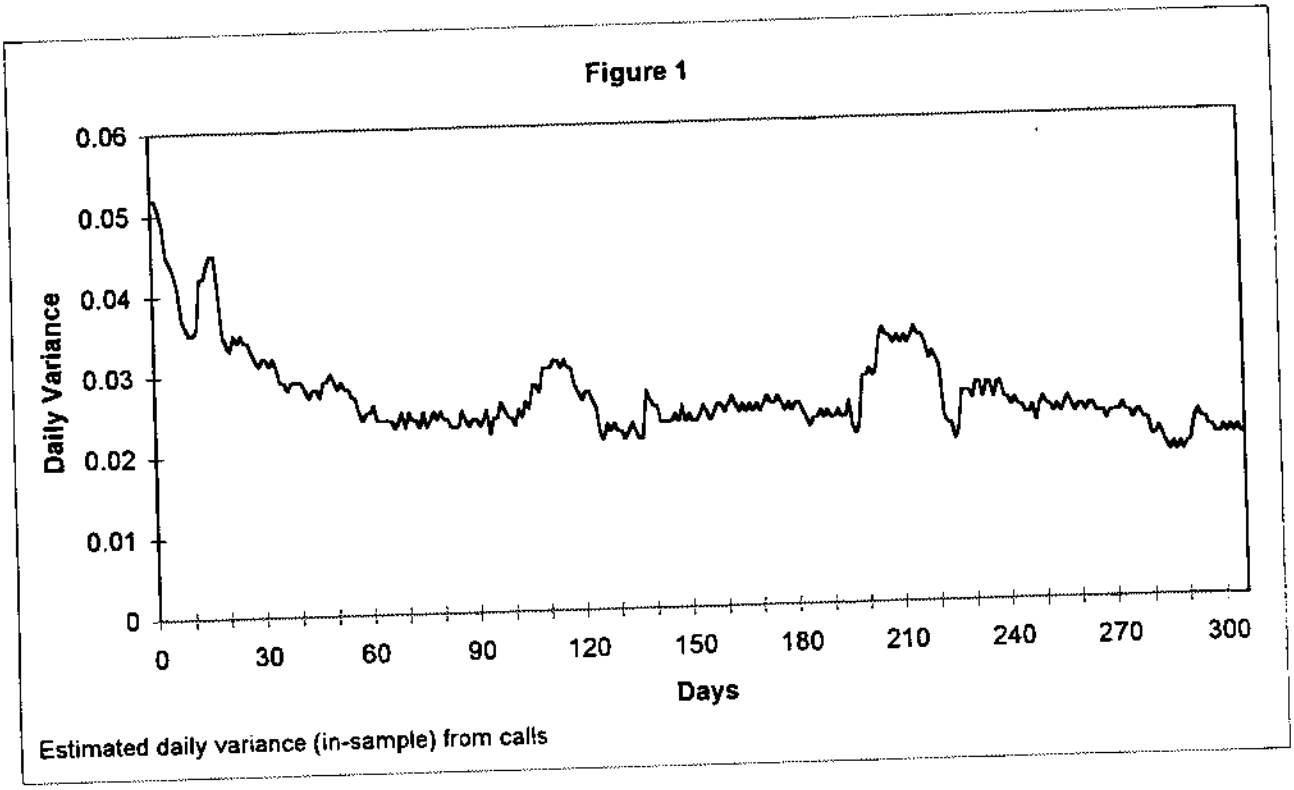
Figure 1
0.06
0.05
0.04
0.03
0.02
0.01
0
Daily Variance
0
30
60
90
120
150
180
210
240
270
300
Days
Estimated daily variance (in-sample) from calls

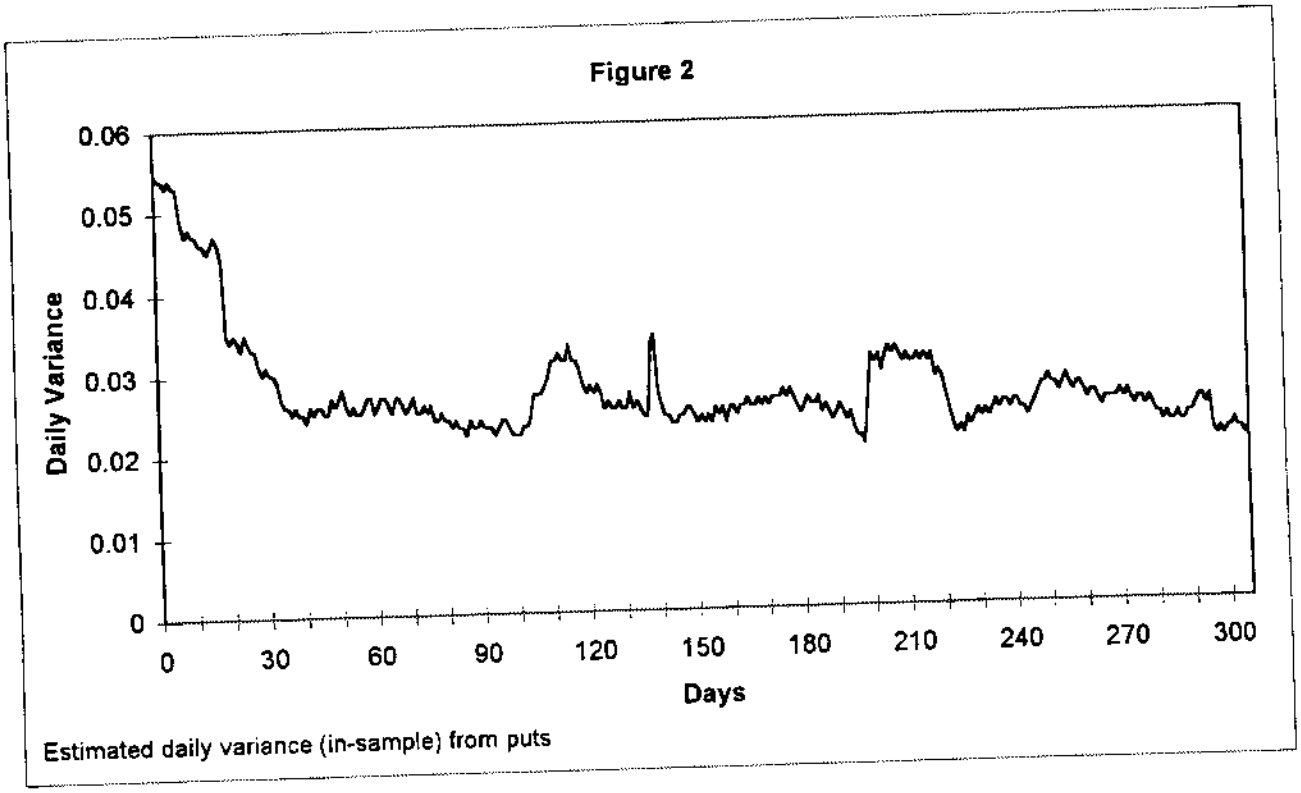
Figure 2
0.06
0.05
0.04
0.03
0.02
0.01
0
Daily Variance
0
30
60
90
120
150
180
210
240
270
300
Days
Estimated daily variance (in-sample) from puts