Jump Starting GARCH:

Pricing and Hedging Options with Jumps in Returns and Volatilities

J. Duan * P. Ritchken † Z. Sun ‡

September 15, 2004

*Rotman School of Management, University of Toronto, Ontario, Canada, M5S 3E6. Email jcduan@rotman.utoronto.ca.

†Weatherhead School of Management, Case Western Reserve University, Cleveland, Ohio, 44106. Email phr@po.cwru.edu.

‡Weatherhead School of Management, Case Western Reserve University, Cleveland, Ohio, 44106. Email zxs9@po.cwru.edu

Jump Starting GARCH:
Pricing and Hedging Options with Jumps in Returns and Volatilities

ABSTRACT

This paper considers the pricing of options when there are jumps in the pricing kernel and correlated jumps in asset returns and volatilities. Our model nests Duan's GARCH option model, as well as generalized versions of Merton's model where jump risk is priced. Empirical analysis on the $S\&P$ 500 index reveals that the incorporation of jumps in returns and volatilities adds significantly to the description of the time series process. Since the state variables are fully determined by the path of prices, once the parameters have been estimated, option prices can readily be computed. We find that option prices, even 50 weeks after the parameters are estimated, are fairly precise. In addition to pricing tests, we examine hedging effectiveness, and provide evidence that the hedges can be maintained very well over time.

(*GARCH option models, stochastic volatility models with jumps, pricing and hedging options*)

In this paper we introduce a new family of GARCH-Jump models and derive the corresponding option pricing theory. These discrete time processes are of interest since the conditional returns of the underlying asset allow levels of skewness and kurtosis to be matched to the data and option prices can readily be established that are influenced by changing volatility and jumps in both returns and volatilities. When the GARCH process is curtailed, but jumps allowed, our limiting model nests the jump-diffusion model of Merton (1976), or the more general model of Naik and Lee (1990). When the jumps are suppressed, both in returns and volatilities, our model reduces to a GARCH model, and option prices are given by to Duan's GARCH option model. Just as Duan's model provides a discrete time filter for an array of continuous time stochastic volatility models, including Heston (1993), Hull and White (1987) and Scott (1987), among others, our more general models can be viewed as filters for continuous-time stochastic volatility models with jumps in returns and volatility. For example, in the more general case, our limiting models nest the time series models in Eraker, Johannes and Polson (2003) as well as option models along the lines of Duffie, Singleton and Pan (1999).

Our option pricing theory is developed along the line of Duan (1995) via an equilibrium argument. The modeling approach can be viewed as providing a pair of corresponding discrete-time systems for describing the underlying asset price and for pricing options under the physical and risk-neutral probability measures.

Why is it important to incorporate jumps in volatility? Empirical research has shown that models which describe returns by a jump-diffusion process with volatility being characterized by a correlated diffusive stochastic process are incapable of capturing empirical features of equity index returns or option prices. For example, both Bates (2000) and Pan (2002) examine such models, and are unable to remove systematic option pricing biases that remain.[1] While jumps in the return process can explain large daily shocks, these return shocks are transient and have no lasting effect on future returns. At the same time, with volatility being diffusive, changes occur gradually and with high persistence. These models are unlikely to generate clustering of large returns associated with temporarily high levels of volatility, a feature that is displayed by the data. Both of the above authors recommend considering models with jumps in volatility.

In response to these findings, researchers have begun to investigate models that incorporate jumps in both returns and volatility. In general, estimating the parameters of continuous-time processes when the volatility is not only unobservable, but the return and volatility processes both contain diffusive and jump elements, is difficult. While in the last decade, significant

[1]Stochastic volatility option models have been considered by Hull and White (1987), Heston (1993), Nandi (1998), Scott (1987), among others. Bakshi, Cao and Chen (1997) provide empirical tests of alternative option models, none of which contain jumps in volatility. Naik (1993) considers a regime switching model where volatility can jump. For additional regime switching models, see Duan, Popova and Ritchken (2002). More recently Bakshi and Cao (2003) provide empirical support for some stochastic volatility models with jumps in returns and volatility.

advances in econometric methodology have been made, these estimation problems are still fairly delicate.[2] In a very convincing study, Eraker, Johannes and Polson (2003), for example, examine the jump in volatility models proposed by Duffie, Singleton and Pan (1999), and show that the addition of jumps in volatility provide a significant improvement to explaining the returns data on the S&P 500 and Nasdaq 100 index returns. The econometric technique in their study was based on returns data alone. In contrast, Eraker (2001) estimated parameters using the time series of returns together with the panel of option data, using methodology similar to Chernov and Ghysels (2000) and Pan (2002). He also confirmed that the time series of returns was better described with a jump in volatility. Surprisingly, however, the model did not provide significantly better fits to option prices beyond the basic stochastic volatility model.

Our empirical analysis focuses on five nested models. At one extreme, we consider models where volatility does not jump, but returns can jump. A Merton-like model is considered, where jump risk is not priced, and a generalized version of this model is also considered where jump risk is priced. At the other extreme we consider Duan's NGARCH option pricing model which serves as a proxy for a stochastic volatility model with no jumps in returns or volatility. Finally, we consider two models that contain jumps in both returns and volatilities.

Our empirical analysis follows a different path to most studies in option models. In particular, if our models are good, then estimates of the parameters, based on time series of the underlying alone, should be sufficient to price options, and eliminate all biases. Unfortunately, for two out of our five models, it is not possible to estimate all the parameter values necessary for pricing options, solely from the time series of the underlying asset value. In these two cases, we augment the asset time series with a panel of at-the-money option contracts. With this set of information all parameter values can be estimated. Once all the parameters are estimated, we observe the future path of the asset, and based on the path, we can update all the state variables, and compute prices of all options. We do this daily for up to one year after the parameters are estimated. Given the "out-of-sample" estimates of option prices, we conduct tests to examine option biases and to evaluate whether incorporating jumps adds value to the model. In addition we conduct hedging tests. Our results for the $S\&P$ 500 demonstrate that incorporating jumps in volatility adds significantly to explaining the time series properties of the index and the patterns in option prices.

The paper proceeds as follows. In section 1 we provide the basic setup for the pricing kernel and the dynamics of the underlying asset. We also identify the risk neutral measure, and

[2]Eraker, Johannes and Polson (2003) provide an excellent review of the difficulties in adopting standard MLE or GMM approaches. Singleton (2001) discusses an approach using characteristic functions. An alternative approach based on simulation methods using Efficient Method of Moments, and Monte Carlo Markov Chains does resolve some of these issues. For an overview on econometric techniques to estimate continuous-time models see Renault (1997), Jacquier, Polson, and Rossi (1994), Eraker, Johannes and Polson (2003), and the references therein.

establish our 5 nested models which represent interesting special cases of the model. In section 2 we discuss time series estimation, option pricing, and hedge construction issues in the discrete GARCH Jump framework. In section 3 we examine our five nested GARCH with jumps models and present empirical evidence from time series of the $S\&P500$ index, adjusted for dividends. We also examine the ability of these models to price European options. We investigate how theoretical option prices, computed up to 50 weeks after the parameters are estimated, perform and examine the hedging effectiveness of these models. Section 5 concludes.

1 The Basic Setup

We consider a discrete-time economy for a period of $[0, T]$ where uncertainty is defined on a complete filtered probability space $(\Omega, \mathcal{F}, \mathsf{P})$ with filtration $\mathbf{F} = (\mathcal{F}_t)_{t\in\{0,1,\ldots,T\}}$ where $\mathcal{F}_0$ contains all P-null sets in $\mathcal{F}$.

Let m_t be the marginal utility of consumption at date t. For pricing to proceed, the joint dynamics of the asset price, and the pricing kernel, $\frac{m_t}{m_{t-1}}$, needs to be specified. We have

$$S_{t-1} = E^{\mathsf{P}}\left[S_t \frac{m_t}{m_{t-1}} \middle| \mathcal{F}_{t-1}\right] \tag{1}$$

where S_t is the total payout, consisting of price and dividends. The expectation is taken under the data generating measure, P, conditional on the information up to date $t-1$.

We assume that the dynamics of this pricing kernel, m_t/m_{t-1}, is given by:

$$\frac{m_t}{m_{t-1}} = e^{a+bJ_t} \tag{2}$$

where J_t is a standard normal random variable plus a Poisson random sum of normally distributed variables. That is,

$$J_t = X_t^{(0)} + \sum_{j=1}^{N_t} X_t^{(j)} \tag{3}$$

where,

$$\begin{aligned} X_t^{(0)} &\sim N(0,1) \\ X_t^{(j)} &\sim N(\mu, \gamma^2) \text{ for } j = 1, 2, ... \end{aligned}$$

and N_t is distributed as a Poisson random variable with parameter λ. Although we have assumed a constant λ, our theoretical results remain valid if the Poisson parameter is stochastic but $\mathcal{F}_{t-1}$-measurable. The random variables $X_t^{(j)}$ are independent for $j = 0, 1, 2...$ and $t = 1, 2, ..., T$.

The asset price, S_t, is assumed to follow the process:

$$\frac{S_t}{S_{t-1}} = e^{\alpha_t + \sqrt{h_t}\bar{J}_t} \tag{4}$$

where $\bar{J}_t$ is a standard normal random variable plus a Poisson random sum of normal random variables. In particular:

$$\bar{J}_t = \bar{X}_t^{(0)} + \sum_{j=1}^{N_t} \bar{X}_t^{(j)} \tag{5}$$

where,

$$\begin{aligned} \bar{X}_t^{(0)} &\sim N(0,1) \\ \bar{X}_t^{(j)} &\sim N(\bar{\mu}, \bar{\gamma}^2) \text{ for } j = 1, 2, ... \end{aligned}$$

Furthermore, for $t = 1, 2, ..., T$:

$$Corr(X_t^{(i)}, \bar{X}_\tau^{(j)}) = \begin{cases} \rho & \text{if } i = j \text{ and } t = \tau \\ 0 & \text{otherwise,} \end{cases}$$

and N_t is the *same* Poisson random variable as in the pricing kernel.

The Poisson random variable provides shocks in period t. Given that the number of shocks in a particular period is some nonnegative integer k, say, the logarithm of the pricing kernel for that period consists of a drawing from the sum of $k+1$ normal distributions, while the logarithmic return of the asset also consists of a drawing from the sum of $k+1$ correlated normal random variables. In either case, the first normal random variable is standardized to have mean 0 and variance 1 because its location and scale have already been reflected in the model specification.

The local variance of the logarithmic returns for date t, viewed from date $t-1$ is $h_t Var(\bar{J}_t) = h_t(1 + \lambda\hat{\gamma}^2)$, where

$$\hat{\gamma}^2 = \bar{\mu}^2 + \bar{\gamma}^2.$$

We shall refer to h_t as the local scaling factor because it differs from local variance by a constant. In general, the local scaling factor h_t can be any predictable process. However, to make matters specific, we shall assume that h_t follows a simple NGARCH(1,1) process:

$$h_t = \beta_0 + \beta_1 h_{t-1} + \beta_2 h_{t-1} \left(\frac{\bar{J}_{t-1} - \lambda\bar{\mu}}{\sqrt{1 + \lambda\hat{\gamma}^2}} - c \right)^2 \tag{6}$$

where β_0 is positive, β_1 and β_2 are nonnegative to ensure that the local scaling process is positive. Here we normalize $\bar{J}_{t-1}$ in the last term to make this equation comparable to the NGARCH model which typically uses a random variable with mean 0 and variance 1. Notice that when $\lambda = 0$, the model reduces to the NGARCH-normal process used by Duan (1995). By permitting jumps in the pricing kernel and stock price, the dynamics of the model may better reflect the time series properties of asset returns, and the fat tails and negative skewness in the risk neutral distribution may fully reflect the information contained in away-from-the-money option prices. Note that by Duan (1997), the h_t process is strictly stationary if $\beta_1 + \beta_2(1 + c^2) \leq 1$. The unconditional mean of h_t is finite and equals $\beta_0 / \left[1 - \beta_1 - \beta_2(1 + c^2)\right]$ if $\beta_1 + \beta_2(1 + c^2) < 1$.

We assume that the single period continuously compounded interest rate is constant, say, r.[3] Thus, the following restrictions must hold:

$$E^{\mathsf{P}}\left[\frac{m_t}{m_{t-1}}\middle|\mathcal{F}_{t-1}\right] = e^{-r} \tag{7}$$

$$E^{\mathsf{P}}\left[\frac{m_t}{m_{t-1}}\frac{S_t}{S_{t-1}}\middle|\mathcal{F}_{t-1}\right] = 1 \tag{8}$$

These equilibrium conditions impose a specific form on α_t. The dynamics of the asset price can be rewritten as in the following proposition.

Proposition 1

Under measure P*, the dynamics of the asset price can be expressed as:*

$$\frac{S_t}{S_{t-1}} = e^{\alpha_t + \sqrt{h_t}\bar{J}_t} \tag{9}$$

where

$$\alpha_t = r - \frac{h_t}{2} - \sqrt{h_t}b\rho + \lambda\kappa\left(1 - K_t(1)\right) \tag{10}$$

$$\kappa = \exp\left(b\mu + \frac{1}{2}b^2\gamma^2\right) \tag{11}$$

$$K_t(q) = \exp\left(q\sqrt{h_t}(\bar{\mu} + b\rho\gamma\bar{\gamma}) + \frac{1}{2}q^2 h_t\bar{\gamma}^2\right). \tag{12}$$

Proof: See Appendix

Given these dynamics, we want to be able to price derivative claims in a risk neutral framework. Towards that goal we assume date T to be the terminal date that we are considering and define measure Q by

$$d\mathsf{Q} = e^{rT}\frac{m_T}{m_0}d\mathsf{P}. \tag{13}$$

Lemma 1

(i) Q *is a probability measure.*

(ii) For any $\mathcal{F}_t$ *measurable random variable,* Z_t*:*

$$Z_{t-1} = E^{\mathsf{P}}[Z_t\frac{m_t}{m_{t-1}}|\mathcal{F}_{t-1}] = e^{-r}E^{\mathsf{Q}}[Z_t|\mathcal{F}_{t-1}].$$

Proof: See Appendix.

[3]Note the constant interest rate assumption is not a necessity. We make this assumption so that there is no need to specify an additional stochastic process for the interest rate.

Given a specification for the dynamics of the pricing kernel and the state variable, all the information that is necessary for pricing contingent claims is provided. While pricing of all claims can proceed, the advantage of the Q measure is that pricing can proceed as if risk neutrality holds.

Proposition 2

(i) Under measure Q*, the dynamics of the asset price is distributionally equivalent to:*

$$\frac{S_t}{S_{t-1}} = e^{\tilde{\alpha}_t + \sqrt{h_t}\tilde{J}_t} \tag{14}$$

where

$$\begin{aligned}
&\tilde{\alpha}_t = r - \frac{h_t}{2} + \tilde{\lambda}\left(1 - K_t(1)\right) \\
&\tilde{J}_t = \tilde{X}_t^{(0)} + \sum_{j=1}^{\tilde{N}_t} \tilde{X}_t^{(j)} \\
&\tilde{X}_t^{(0)} \sim N(0,1) \;\; \textit{for } t = 1,2,...,T \\
&\tilde{X}_t^{(j)} \sim N(\bar{\mu} + b\rho\gamma\bar{\gamma}, \bar{\gamma}^2) \textit{ for } t = 1,2,...,T \textit{ and } j = 1,2,... \\
&\tilde{X}_t^{(j)} \textit{ are independent for } t = 1,2,...,T \textit{ and } j = 0,1,2,...
\end{aligned}$$

and $\tilde{N}_t$ *has a Poisson distribution with parameter* $\tilde{\lambda} \equiv \lambda\kappa$*.*

(ii) Further, when the updating equation for the local scaling factor, h_t*, is given by equation (6), then under measure* Q*, it has the form*

$$h_t = \beta_0 + \beta_1 h_{t-1} + \beta_2^* h_{t-1}\left(\frac{\tilde{J}_{t-1} - \tilde{\lambda}(\bar{\mu} + b\rho\gamma\bar{\gamma})}{\sqrt{1 + \tilde{\lambda}\tilde{\gamma}^2}} - c^*\right)^2 \tag{15}$$

where

$$\begin{aligned}
\beta_2^* &= \beta_2\left(\frac{1 + \tilde{\lambda}\tilde{\gamma}^2}{1 + \lambda\hat{\gamma}^2}\right) \\
c^* &= \frac{c\sqrt{1 + \lambda\hat{\gamma}^2} + \lambda\bar{\mu} - \tilde{\lambda}(\bar{\mu} + b\rho\gamma\bar{\gamma}) - b\rho}{\sqrt{1 + \tilde{\lambda}\tilde{\gamma}^2}} \\
\tilde{\gamma}^2 &= (\bar{\mu} + b\rho\gamma\bar{\gamma})^2 + \bar{\gamma}^2
\end{aligned}$$

Proof: See Appendix

Under measure Q, the overall dynamics of the asset price is similar in form to the dynamics under the data generating measure, P. In particular, the logarithmic return is still a random Poisson sum of normal random variables. However, under measure Q, the mean of each of the

normal random variables is shifted. Similarly, the random variable, N_t, distributed as a Poisson random variable under measure P, is still Poisson under measure Q but with a shifted parameter.

Notice that each normal random variable has the same variance under both measures. However, the local variance of the innovation under measure Q is $h_t(1+\tilde{\lambda}\tilde{\gamma}^2)$, which is not equal to the local variance under the original P measure unless $\kappa = 1$ and $b\rho\gamma = 0$.[4] In other words, one should not in general expect the local risk-neutral valuation principle to apply.

Finally, when the local scaling factor h_t follows a NGARCH process, as in equation (6) then under measure Q, the updating scheme translates into a similar NGARCH process as in equation (15). Specifically, as in equation (6), $\tilde{J}_t$ in equation (15) undergoes a mean-scale transformation so that the resulting variable has mean 0 and variance 1.

Proposition 2 allows us to easily compute derivative prices consistent with our underlying preferences and dynamics.

1.1 Decomposition of the Risk Premium

Under measure P, the expected total return on the stock can be expressed as:

$$E^{\mathsf{P}}\left[\frac{S_t}{S_{t-1}}\right] = e^{(r+\eta_t)}$$

where the risk premium η_t is given by:

$$\begin{aligned} \eta_t &= \lambda\kappa(1-K(1)) - \lambda(1-e^{\bar{\mu}\sqrt{h_t}+\frac{\bar{\gamma}^2 h_t}{2}}) - \sqrt{h_t}b\rho & (16)\\ &\approx [\lambda\bar{\mu}(1-\kappa) - b\rho(1+\lambda\kappa\gamma\bar{\gamma})]\sqrt{h_t} + \lambda\bar{\gamma}^2(1-\kappa)\frac{h_t}{2}, & (17) \end{aligned}$$

where the approximation is justified if h_t is small.[5]

To gain some insight into the model, first consider the case when $\kappa = 1$ and $\gamma = 0$. In this case, the risk premium, η_t reduces to $-b\rho\sqrt{h_t}$. With $\kappa \neq 1$ and $\gamma = 0$ in the pricing kernel, the sensitivity of the risk premium to $\bar{\gamma}$ is very small. That is, the randomness about the jump size adds minimally to the risk premium. Additionally, with $\kappa = 1$, the jump risk is fully diversifiable. This corresponds to the assumption used in Merton (1976). Naik and Lee (1990) extends Merton's model to the case where jump risk is not diversifiable. In our model this is accomplished by releasing κ from 1 and γ from 0.

[4] This result differs from the local risk-neutral valuation conclusion of Duan (1995) because the innovation term is generated by a Poisson random sum of normal random variables as opposed to the use of normally distributed innovations in Duan (1995). Of course, when the Poisson parameter is switched off, the local variance will remain unaltered with the measure change and the pricing result reduces to the pricing model of Duan (1995).

[5] In our empirical studies we obtain h_t is of order 10^{-6}

With $\kappa = 1$ and $\gamma > 0$, the risk premium is

$$\eta_t \approx -b\rho\sqrt{h_t} - b\rho\lambda\gamma\bar{\gamma}\sqrt{h_t}.$$

Here, the uncertainty of the jump size, as measured by $\bar{\gamma}$, adds to the risk premium as does the intensity. Finally, when κ is released from 1, the impact of the intensity of the process on the risk premium becomes more complex.

The expected value of the pricing kernel, fully determines interest rates, and is given by:

$$E^{\mathsf{P}}[\frac{m_t}{m_{t-1}}|\mathcal{F}_{t-1}] = e^{a+b^2/2+\lambda(\kappa-1)}.$$

For the case when $\kappa = 1$ (i.e., $\mu = -b\gamma^2/2$), the effects of the jump in the pricing kernel play no role on the interest rate. For all other values of κ, the jump process explicitly effects both the interest rate and asset price.

1.2 Five Nested Models

We shall refer to the general model as the NGARCH-Jump model, and the special case of this model when $\kappa = 1$ and $\gamma = 0$ as the Restricted NGARCH-Jump model or RNGARCH-Jump. If in addition, we shut down the GARCH effects, ($\beta_1 = \beta_2 = 0$), we obtain a model where in each period conditional on the number of jumps, the return distribution is normal, with the same variance. Since jump risk is diversifiable and the local scaling factor, h_t, is constant, we refer to this model as the discrete-time Merton model, or MERTON, for short. The Merton model, but with κ and γ released from 1 and 0, will be called the generalized Merton model (hereafter, G-MERTON). The final special case we consider is the case where we permit stochastic volatility, but shut down the jumps, ie. $\lambda = 0$. In this case, the general model reduces to the standard NGARCH-Normal model.

1.3 Continuous Time Limits

Duan, Ritchken and Sun (2004) have extended this GARCH model by introducing the element of interval length. Their model would reduce to our model when the length of the time increments is set to one day. The advantage of their framework is that it allows them to explore various limiting cases as the time increments tend to zero. They show that our discrete time model can be viewed as a good approximation for several limiting models that include jump and diffusive elements in returns and/or volatilities.

In light of their results, even if one prefers modeling in continuous time, the results here could be useful, since our model could serve as a useful proxy for continuous time processes.

Given the ease of estimating GARCH parameters from the time series of the underlying, this could be very beneficial.

We now explore which, if any, of the models nested in our family can simultaneously explain both the time series of the $S\&P$ 500 and the cross sectional variation of option prices over a broad array of strikes and maturities.

2 Experimental Design for Pricing and Hedging

In this section we consider the empirical performance of the NGARCH-Jump model using time series data on the $S\&P$ 500 index and dividends. Our main goal is to estimate models using time series data alone and then to evaluate the ability of these models to price and hedge options. We are particularly interested in evaluating the full model as well as its four nested special cases, namely the RNGARCH-Jump model, the NGARCH-Normal model, the G-MERTON model and the MERTON model.

2.1 Description of Data

The $S\&P$ 500 index options are European options that exist with maturities in the next six calendar months, and also for the time periods corresponding to the expiration dates of the futures. Our price data on the call options, covering the five year period from January 1991 to December 1995, comes from the Berkeley Option Database. We collected daily data and excluded contracts with maturities fewer than 10 days. We only used options with bid/ask price quotes during the last half hour of trading. For these contracts we also captured the reported concurrent stock index level associated with each option trade.

In order to price the call options we need to adjust the index level according to the dividends paid out over the time to expiration. We follow Harvey and Whaley (1992), and Bakshi, Cao and Chen (1997), and use the actual cash dividend payments made during the life of the option to proxy for the expected dividend payments. The present value of all the dividends is then subtracted from the reported index levels to obtain the contemporaneous adjusted index levels. This procedure assumes that the reported index level is not stale and reflects the actual price of the basket of stocks representing the index. Since intra day data and not the end of the day option prices are used, the problem with the index level being stale is not severe.[6] Since we used

[6]There are other methods for establishing the adjusted index level. The first is to compute the mid points of call and put options with the same strikes and then to use put-call parity to imply out the value of the underlying index. Of course, this method has its own problems, since with non negligible bid ask spreads, put call parity only holds as an inequality. An alternative approach is to use the stock index futures price to back out the implied dividend adjusted index level. This leads to one stock index adjusted value that is used for all option contracts.

the actual contemporaneous index level associated with each option trade that was reported in the data base, the actual adjusted index level would vary slightly among the individual contracts depending on their time of trade. Finally, we used the T-Bill term structure to extract the appropriate discount rates.

We have 1250 trading days in our time series, with 250 consecutive weeks of cross sectional option prices. We split the data up into an in-sample period of 200 weeks, and an out-of-sample period of the remaining 50 weeks. Over the first 200 weeks we use the daily time series on the index to estimate the parameters of some of the nested models. As we shall see, the parameters of some of the models cannot be fully identified from the time series alone. In these cases we complement the daily time series with weekly observations of the prices of the at-the-money call option price with maturity closest to 30 days. Once all models are estimated, we use the parameter estimates and the daily time series of the index to compute the full time series of the local scaling factor not only over the 200 week historical time period, but also for the successive days over the next 50 weeks.

Our first set of experiments are concerned with using the time series data on the S&P 500 index alone to compare the performance of some of the nested models, and to establish the importance of incorporating jumps and NGARCH effects. Our second set of experiments evaluates how well the fitted models from the time series are able to price options, conditional on the index, and on the computed local scaling factor, over the 50 weeks in the "out-of-sample" period. We also compare these models to the more general models which required that some parameters be estimated from the historical time series augmented with option prices.

An option model is viewed positively if the in-sample fits are precise and unbiased, and if, conditional on future state variables, the "out-of-sample" price predictions are also precise and unbiased. In addition to investigating pricing biases in the out-of-sample period, we also investigate the performance of delta hedging strategies and report the hedging effectiveness associated with the models.

2.2 Estimation from the Time Series

We use a maximum-likelihood approach to estimate the parameters from the models using the time series of historical asset returns.

Let $y_t = \ln(S_t/S_{t-1}) - \alpha_t$. We rewrite the GARCH process of return under measure P as

$$\begin{aligned} y_t &= \sqrt{h_t}\bar{J}_t \\ h_t &= v(h_{t-1}, \bar{J}_{t-1}) \end{aligned}$$

For a discussion of these approaches see Jackwerth and Rubinstein (1996).

where the function α_t is given in Proposition 1, and the function $v(\cdot)$ is given by (6). The initial value of the local scaling factor is determined by

$$h_1 = V/(1+\lambda\hat{\gamma}^2) \tag{18}$$

where V is the sample variance of the asset return and as defined earlier, $\hat{\gamma}^2 = \bar{\mu}^2 + \bar{\gamma}^2$.[7] Our model parameter set is

$$\theta = \{\beta_0, \beta_1, \beta_2, c, b\rho, \kappa, \gamma, \bar{\mu}, \bar{\gamma}, \lambda\}$$

The conditional probability density function, $l(y_t|h_t, y_{t-1})$, of y_t is:

$$l(y_t|h_t, y_{t-1}) = \sum_{i=0}^{\infty} \frac{\lambda^i}{i!} e^{-\lambda} f_{(\mu_i(t), \sigma_i^2(t))}(y_t)$$

where $f_{(\mu_i(t),\sigma_i^2(t))}(\cdot)$ is the normal density function with mean $\mu_i(t) = i\bar{\mu}\sqrt{h_t}$ and variance $\sigma_i^2(t) = h_t(1+i\bar{\gamma}^2)$.[8] The log-likelihood function of the sample is:

$$L(\theta; y_1, ..., y_T) = \sum_{t=2}^{T} \ln\left[l(y_t|h_t, y_{t-1})\right]. \tag{19}$$

The maximum likelihood estimator for θ is the solution of maximizing the above log-likelihood function. Given the asset return process, $\{\ln \frac{S_t}{S_{t-1}}\}_{1\le t\le T}$, we can write down the likelihood function recursively, and solve this optimization problem numerically.

In principle, the entire set of parameters can be identified by only using a time series of asset returns. In practice, however, two of them are hard to pin down empirically. To understand this assertion, recall that:

$$\begin{aligned} 1 - K_t(1) &= 1 - \exp(\sqrt{h_t}(\bar{\mu} + b\rho\gamma\bar{\gamma}) + \frac{1}{2}h_t\bar{\gamma}^2) \\ &\approx -\sqrt{h_t}(\bar{\mu} + b\rho\gamma\bar{\gamma}) - \frac{1}{2}h_t((\bar{\mu} + b\rho\gamma\bar{\gamma})^2 + \bar{\gamma}^2). \end{aligned}$$

Hence,

$$\alpha_t \approx r - \frac{1}{2}h_t(1 + \lambda\kappa(\bar{\gamma}^2 + (\bar{\mu} + b\rho\gamma\bar{\gamma})^2)) - \sqrt{h_t}(b\rho + \lambda\kappa\bar{\mu} + \lambda\kappa b\rho\gamma\bar{\gamma}) \tag{20}$$

First note that h_t is much smaller than $\sqrt{h_t}$ because $\sqrt{h_t}$ takes on small values already. The term with $\sqrt{h_t}$ effectively dominates the term with h_t. The above formula suggests that the coefficient of $\sqrt{h_t}$, i.e., $-(b\rho + \lambda\kappa\bar{\mu} + \lambda\kappa b\rho\gamma\bar{\gamma})$, practically acts as a single term, which makes it hard to separate $b\rho$, κ and γ. Note that parameters λ, $\bar{\mu}$ and $\bar{\gamma}$ directly enter into the density

[7] In fact, $Var(y_t) = E(h_t)Var(\bar{J}_t) = (1+\lambda\hat{\gamma}^2)E(h_t)$, and we assume that the initial scaling factor is the long-run average of h_t.

[8] Conditioning on $N_t = i$, the variance of $\bar{J}_t$ is $1 + i\bar{\gamma}^2$. Without conditioning, however, the variance becomes $1 + \lambda\hat{\gamma}^2$.

function. In contrast, parameters $b\rho$, κ and γ only appear through α_t in the equation for y_t. Since only the sum $-(b\rho + \lambda\kappa\bar{\mu} + \lambda\kappa b\rho\gamma\bar{\gamma})$ matters in the sample likelihood function, two of the three parameters - $b\rho$, κ and γ - are indeterminate. In the estimation, we thus introduce a composite parameter $\delta = -(b\rho + \lambda\kappa\bar{\mu} + \lambda\kappa b\rho\gamma\bar{\gamma})$. To deal with the indeterminacy we set $\kappa = 1$ and $\gamma = 0$ and view $b\rho$ as a function of δ. As a result, we actually estimate the parameter set

$$\theta^* = \{\beta_0, \beta_1, \beta_2, c, \delta, \bar{\mu}, \bar{\gamma}, \lambda\}$$

Notice that when $\kappa = 1$ and $\gamma = 0$ we obtain the RNGARCH-Jump model. This model, as well as the MERTON model (with $\beta_1 = \beta_2 = 0$) can be fully estimated using maximum likelihood estimates from the time series data alone. In addition, the NGARCH-Normal model can be readily estimated. Of course, if option data is available as well, then the parameters κ and γ can be released in both the MERTON and RNGARCH-Jump models to obtain estimates for the G-MERTON and the full NGARCH-Jump model.

Since there are no simple analytical expressions for the options, their prices are generated by Monte Carlo simulation. Hence, rather than doing a joint optimization, for the last two models we use the option prices only to estimate κ and γ. Specifically, we fix the other parameters, and then for a given κ and γ we generate the daily values of h_t over the last 20 weeks of the in-sample data. Every five days, we compute, via Monte Carlo simulations, the theoretical price of the short-dated (closest to 30 days) nearest the money contract. We then select the parameter values that result in the minimum sum of squared percentage errors.

2.3 Fitting of Option Prices

Given the parameter estimates from the time series data, extracted over the "in-sample" period, and given the time series of the index over the next 250 days, we can construct the time series for h_t over all the days in the out-of-sample period. Given the index and local scaling factor at any date, we can compute option prices using simulation. The option prices are computed using $10,000$ sample paths and antithetic variance reduction techniques. We refer to all the theoretical option prices computed after day 1000, as "out-of-sample" option prices. These prices, of course, are conditional on the index level being observed, and on the level of the local scaling factor h_t, that determines the local volatility. Over our 250 day out-of-sample period, we compute the theoretical prices of all option contracts, each week, for a total of 50 weeks, for our models. For the different models, the same stream of random variables are used. The residuals for each contract and model are stored. If a model is good, the fitted option prices in the out-of-sample period should be unbiased across maturities and strike prices. That is, the model should explain the volatility skew and the maturity bias inherent in the Black-Scholes model.

Investigating the fit of option contracts using models estimated from the time series of prices alone has been used in many studies. For example, Jaganathan, Kaplin and Sun (2001), estimate

several multifactor Cox-Ingersoll-Ross models, using time series data on swaps, and then assess how well the resulting calibrated models fit swaption contracts. Alternatively, parameters can be implied out from a set of derivatives in one market and then used to price claims in a related market. For example, Longstaff, Santa Clara and Schwartz (2001), calibrate models of the term structure using caps and floors, and then assess their models by considering the performance of the fitted models in the swaption market. In our analysis, we want to estimate our models using time series data as much as possible, and then assess the models not only on their time series fit, but also on their ability to price the panel of option contracts in the out-of-sample weeks.

Our analysis here stands in strong contrast to the common procedure of repeatedly re-estimating models based on cross sectional option prices and examining properties of the pricing residuals and implied parameters. Our purpose here is to place as much weight on the time series of prices as possible, to use the minimal amount of option information, and then to examine whether we are capable of pricing options, over an array of strikes and maturities, in out-of-sample tests. Specifically, for a particular model we only need to optimize once to obtain all the parameters. Then, since the future state variables are fully determined by the trajectory of the underlying price, as it evolves we can easily update option prices. In this regard, our "out-of-sample" residuals can be based off parameter values estimated up to 50 weeks earlier. Our goal is to demonstrate that from the time series of asset prices, we can fit option prices well and that the out-of-sample performance is fairly precise, even 50 weeks after our parameters were estimated.

2.4 Hedging Effectiveness

Our final tests will be to evaluate the hedging performance of our models using the out-of-sample period of 250 days. We compute the hedge ratios for our models and set up hedges for each contract for each day. The performance of each hedged position dynamically rebalanced over a 15 day interval is recorded. This allows us to compare the relative performance of the hedges.

Specifically, consider an option that is to be hedged over n successive periods (days) of length Δt, starting from date k. Define the discrete delta hedged gains, $\pi(n;k)$, over the n days as:

$$\pi(n;k) = (C_{n+k} - C_k) - \sum_{i=0}^{n-1} \Delta_{k+i}(S_{k+i+1} - S_{k+i}) - \sum_{i=0}^{n-1} r(C_k - \Delta_{k+i}S_{k+i})\Delta t$$

where Δ_{k+i} is the hedge ratio for the option at date $k+i$, and is given by the model.

The hedging tests are conducted over the last 250 days of data, using models, the parameters which are estimated using data from the first 1000 days. The "out-of-sample" hedging performance for our models is compared to the "in-sample" hedging performance of the Black-Scholes model, where the hedge for each contract is determined by its *own* concurrent implied volatility.

That is

$$\Delta_{k+i} = \Delta_{k+i}^{BS} = N(d_1(\sigma_{k+i}(X,T))$$

where $N()$ is the standard normal cumulative distribution function and

$$d_1 = \frac{\ln(S_{k+i}/X) + (r + \sigma_{k+i}(X,T)^2/2)T}{\sigma_{k+i}(X,T)\sqrt{T}}$$

where T is the time remaining to expiration, r is the yield to maturity over date T, and $\sigma_{k+i}(X,T)$ is the implied Black volatility at date $k+i$ that equates the theoretical price to the actual option price.

Notice that this benchmark against which our hedging is to be compared is difficult to beat. At each day the hedge is constructed so that every option matches its actual price. At any single date, this model has as many parameters as there are contracts, and over n successive days the number of parameters in this model is n times the number of contracts! In contrast, the models we test are based on parameters estimated using historical data alone and at any date, theoretical option prices will not exactly match actual option prices.

We now discuss how the hedge ratios for our GARCH models are established. Discrete-time GARCH models do not allow for hedging along the lines of Black-Scholes because markets are incomplete. Nevertheless, one can view the hedge ratio as the partial derivative of the option pricing function with respect to the stock price while holding the local volatility fixed. The hedge ratio naturally becomes

$$\Delta_t = e^{-r(T-t)} E^{\mathsf{Q}} \left[\frac{S_T}{S_t} 1_{[\frac{S_T}{S_t} > \frac{X}{S_t}]} \right] \tag{21}$$

a result first established in Duan (1995) and later collaborated by Garcia and Renault (1998) by applying the homogeneity of degree one property of the option pricing function. In our hedging analysis, we adopt equation (21) for computing the hedge ratio.

For each model, we compute the hedge ratios numerically, using Monte Carlo simulation with 10,000 paths, and antithetic variables. The same set of random numbers are used for the different models. The above analysis will reveal how effective the models are in their ability to hedge the full array of call options by moneyness and maturity.

3 Empirical Results

3.1 Time Series Estimation

Table 1 shows the parameter estimates based on the time series data over the first 1000 trading days for the 5 models. In particular, we report the point estimates and their standard deviations.

Table 1 Here

First, consider the three models for which no option data was used. It can be seen that λ is significantly different from 0 indicating that the incorporation of jumps is significant. In addition, the parameters β_1 and β_2 are significantly different from 0, indicating that GARCH effects are important. As is well documented the non-linear term c, capturing the so called leverage effect, is also significant in the NGARCH as well as in the RNGARCH-Jump model.

Table 1 also reports on the additional estimates for the G-MERTON and full NGARCH-Jump model, when option data was used to estimate κ and γ. The effects of γ were found to be very small and not significant at the 5% level of significance. Hence, the results that are reported are based on optimizations conducted with $\gamma = 0$. In both models, we see that κ is significantly different from 1. The option information therefore has been useful in extracting information on jump risk premia. In particular according to equation (17), the contribution of the diffusion risk premium, $b\rho\sqrt{h_t} = 0.06329\sqrt{h_t}$, while the jump risk premium, $\lambda\bar{\mu}(1-\kappa)\sqrt{h_t} = 0.009\sqrt{h_t}$. So the jump risk premium accounts for about 12.5% of the total risk premium.

Table 2 reports the skewness and kurtosis of the conditional daily return residual normalized by the square root of the local scaling factor, i.e., $y_t/\sqrt{h_t}$. We know that the NGARCH model assumes that the conditional distribution of daily return residuals is normal, which means the kurtosis of residuals should be 3. But Table 3 shows that the actual kurtosis is larger than 3.[9]

Table 2 Here

Eraker, Johannes and Polson (2003) find that jumps are infrequent events, occurring on average about twice every three years, tend to be negative, and are very large relative to normal day to day movements. In contrast, our average jump frequency is close to two a day. The jumps add conditional skewness and kurtosis to the daily innovations, rather than providing large shocks. Indeed, the mean and standard deviation of our jump size variable is not particular large compared to the standard normal innovation. By mixing a random number of normal distributions, the conditional distribution displays higher kurtosis. In our case $\bar{J}_t$ consists of one standard normal random variable together with a Poisson random sum of independent normal random variables with mean $\bar{\mu}$ and variance, $\bar{\gamma}^2$. The NGARCH-Jump model with the estimated parameter values has a predicted kurtosis of normalized daily return residuals equal to 4.12, close to the observed kurtosis of 4.02.

[9]Under measure P the conditional skewness and kurtosis of NGARCH-Jump innovation, $\bar{J}_t$, can be shown to be $skewness = \frac{\lambda(3\bar{\mu}\bar{\gamma}^2+\bar{\mu}^3)}{(1+\lambda(\bar{\gamma}^2+\bar{\mu}^2))^{3/2}}$ and $kurtosis = \frac{\lambda(3\bar{\gamma}^4+6\bar{\gamma}^2\bar{\mu}^2+\bar{\mu}^4)}{(1+\lambda(\bar{\gamma}^2+\bar{\mu}^2))^2}$.

3.2 Option Pricing Performance

Once the parameter estimates of the models have been obtained, the full time series of the local scaling factor can be established. Given the two state variables, (S_t, h_t), at any date t, we can compute the theoretical option prices in the out-of-sample period and compare them to actual option prices.

In our out-of-sample period, the stock market steadily increased. As a result, there are many more very deep in-the-money contracts. We define moneyness as $(S_t - X)/S_t$ where X is strike price. Our default moneyness buckets consisted of 8 bins set up around the at-the-money bucket of moneyness values in the range $(-0.01, 0.01)$ as follows: $1 = (< -0.04)$, $2 = (-0.04, -0.02)$, $3 = (-0.02, -0.01)$, $4 = (-0.01, 0.01)$, $5 = (0.01, 0.02)$, $6 = (0.02, 0.04)$, $7 = (0.04, 0.08)$, and $8 = (> 0.08)$. We separated out all contracts into 4 maturity buckets, of less than 30 days, $30 - 60$ days, $60 - 90$ days and greater than 90 days. When reporting information graphically, finer partitions of moneyness or maturity buckets were used. Our out-of-sample call option set consists of $17,891$ (3,710) contracts, when computed daily (weekly).

Figure 1 summarizes the pricing performance of all 5 models. The left panel provides a plot of the average percentage error in prices of contracts versus moneyness. The percentage pricing error is defined as the model price minus the market price and then divided by the market price. The four plots are for each maturity bucket. On each graph there are five lines, each representing a particular model. In all graphs, the topmost line refers to the MERTON model, the next line the G-MERTON model, followed by the NGARCH-Normal, RNGARCH-Jump and NGARCH-Jump.

Figure 1 Here

For all maturities, and for all models, there are systematic moneyness biases. In general, on average, all models overprice deep out-of-the-money contracts, and very slightly underprice in-the-money contracts. This can be seen in the skewness of the curves which are, in general, positive over the range of out-the-money contracts, then negative for in-the-money contracts and eventually converging to zero for very deep in-the-money contracts.

The MERTON model has very large biases and is clearly dominated by other models. G-MERTON improves upon MERTON in reducing the skewness of the curves. The NGARCH-Normal and RNGARCH-Jump appear to be fairly similar with the jump component providing small benefits. Finally, there appears to be a significant flattening out of the curve for the NGARCH-Jump model.

For short-dated contracts, and 30-60 day contracts, the NGARCH-Jump model produces very good results relative to the others. For the longest dated contracts, the model does the

best at fitting deep out-the-money contracts, but underprices contracts close to the money and in-the-money.

Since the big discrepancy among the performance of the models is for out-the-money and at-the-money contracts, we take a closer look at pricing errors for these contracts. The second panel in Figure 2 presents the box and whisker plots for the percentage errors of the five models plotted against moneyness. These plots clarify the results presented in the graphs and highlight the distribution of residuals. For each moneyness bucket, the five plots are ordered from left to right as MERTON, G-MERTON, NGARCH-Normal, RNGARCH-Jump and NGARCH-Jump.

Figure 2 shows the average percentage errors in prices for each model across expiration dates. Since the results depend on moneyness, five graphs are presented, each graph for a different moneyness category. If there were no bias in the results, then the plots should be horizontal lines near zero. For all models, and for all moneyness buckets, the overall trend of the lines is downward sloping. Average percentage errors in short-dated options are higher than average percentage errors for long-dated contracts. For all out-the-money and at-the-money contracts, the NGARCH-Jump model has the flattest curve closest to zero. However, the underpricing of in-the-money contracts is clearly revealed, with the problem becoming more extreme with longer-dated contracts.

Figure 2 Here

In our analysis of the NGARCH-Jump model, we used short-dated at-the-money options to estimate κ. As a result, it may not be surprising that short-dated, at-the-money option prices are well fit. If we re-estimate κ using longer-dated options, then the "out-of-sample" fit to longer-dated contracts does improve.

In general there is no unambiguously preferable metric for computing and presenting in-sample or out-of-sample fits.[10] Our estimate of κ was done by minimizing the sum of squared percentage errors for the set of short-dated, at-the-money options. Alternative approaches could use more cross sectional option prices, and/or use absolute errors, rather than percentage errors. Bates (2002) suggests that dividing dollar errors by the underlying asset price makes results more comparable and is more appropriate given that option prices are theoretically nonstationary. Since our goal is to compare model performance, using estimates extracted from time series information, as much as possible, we refrained from recalibrating our models using panel option data and alternative best fit criterion.

The next pricing analysis conducted on the "out-of-sample" residuals was to compare, pairwise, the residuals of the RNGARCH-Jump model, with the MERTON and the NGARCH-Normal model. The parameters for these three models were completely determined by the time

[10]For a good review of the alternative approaches, see Bates (2002).

series of returns alone. For completeness, we also included pairwise comparisons for the G-MERTON and NGARCH-Jump models. Table 3 presents the fraction of times the RNGARCH model produced a smaller absolute residual than the other models. The analysis is performed for all contracts in the fifty week period after the parameters of the models were computed. The top panel performs the tests by moneyness; the bottom panel repeats the tests by maturity.

Table 3 Here

The table shows that in aggregate, over all contracts, the RNGARCH-Jump model outperforms all other models. Further, for almost all maturity and moneyness buckets, this model outperformed MERTON and G-MERTON. The NGARCH-Normal model was a bit more competitive, but, at the 5% level of significance, the RNGARCH-Jump model is preferable. As we saw earlier, the unrestricted NGARCH-Jump model was significantly better than the restricted model in pricing out-the-money and at-the-money contracts.

The top panel of Table 4 shows the median absolute percentage pricing errors for the RNGARCH-Jump model, for all contracts in the 50 week out-of-sample period in each maturity-strike price bucket. The total number of contracts in each bucket are also provided. The median absolute pricing error over all $17,891$ contracts was 3.5%.

The errors reported here are somewhat similar to the errors reported in one-week "out-of-sample" tests conducted by Bakshi and Cao (2003) for their stochastic volatility model with correlated return and volatility jumps. For out-(at-)the-money calls their average absolute percentage errors ranged from 14% to 27% (5% to 12%) depending on maturity. Comparisons of our residuals with theirs are somewhat difficult to make for several reasons. First, in our study we used time series of the underlying index to estimate most of the parameters, while they fit their parameters based on out-the-money contracts. If we had used information on out-the-money contracts in the optimization, our fits of out-the-money options would improve, possibly at the expense of in-the-money contracts. However, our goal was to evaluate whether models estimated from the time series of returns would lead to good models of option prices. Second, our model is never re-estimated. In particular, included in our sample of residuals are contracts whose prices are computed up to 50 weeks after the model parameters were determined.[11] In spite of differing objectives, our error terms appear to be of similar magnitudes to their reported values.

Table 4 Here

[11] Bakshi and Cao's analysis is primarily geared towards examining volatility skews for stock options rather than index options. Bakshi, Kapadia and Madan (2003), relate individual security skewness to the skew of the market, and identify conditions where the skewness of the market is greater. This skewness directly relates to the volatility smile.

Recall, that our model is never recalibrated. As a result it may be the case that the errors propagate over time in an uneven way. The bottom panel of Table 4 shows the medians of the absolute pricing errors by moneyness, for each successive 10-week period. Interestingly, the performance of the model does not seem to deteriorate over the 10-week blocks. Indeed, the pricing errors, 40-50 weeks out-of-sample, are no worse than the errors in the first 10-week block.

3.3 Hedging Performance

Figure 3 shows the box and whisker plots of the raw hedging errors (discrete delta hedged gains) from dynamically hedging over 15 successive trading day periods, when the delta values are computed by the G-MERTON, NGARCH-Normal, RNGARCH-Jump, NGARCH-Jump and Black-Scholes models. The leftmost plot in each block of six plots indicates the change in the unhedged position. The plots are ordered by the original moneyness of the option at the start of the hedging period.

Figure 3 Here

From the box and whiskers plot of the unhedged residuals, we see that over the sample period, the S&P 500 steadily increased, and on average buying calls was profitable. The figure shows that all five hedges were remarkably effective, in spite of the fact that some of these models performed poorly in pricing. Indeed, at this aggregate level, there appears to be very minor differences between the five hedges. The amount of unhedged variability explained by the delta hedging strategies are similar for all models. Indeed, while the ordering of the R^2 values by model align with the results from the pricing, the differences are hardly significant.[12] To our knowledge, the results reported here are the first results that attempt to measure the effectiveness of hedges established using GARCH based models.

Notice that all the models produced hedge results as good as the *ad hoc* Black-Scholes model. Indeed, over all 706 hedges that were tested, the NGARCH-Jump model outperformed the *ad hoc* Black-Scholes model on 54% of occasions.

Recall that the *ad hoc* Black-Scholes model used the implied volatilities of each contract as the basis for the delta hedge. As a result, this model perfectly priced each contract each day. The fact that these "out-of-sample" hedges performed as well as the "in-sample" Black-Scholes model indicates that these models are useful for explaining option prices moves. This is quite remarkable, since the Black-Scholes equation as used here, is not really a model but serves only

[12]The R^2 values for G-MERTON, NGARCH-Normal, RNGARCH-Jump, NGARCH-Jump and the *ad hoc* (in-sample) Black-Scholes were 0.79, 0.85, 0.87, 0.90, and 0.89, respectively.

as a calibrating device.[13]

The average return from all the delta hedged strategies, across all moneyness categories is clearly negative. This result is consistent with that found by Bakshi and Kapadia (2001), who explained this finding by postulating a negative risk premium for volatility risk.

Eraker (2001) found that while the inclusion of jumps in volatility improved the time series fit of the S&P 500 time series of returns, the benefits of the jump in explaining option prices were surprisingly not significant. In our analysis, we find that our "jumps" are significant in the time series, and that the benefits of incorporating these jumps flow over into option pricing. However, from a hedging perspective, even the simplest NGARCH-Normal model does an outstanding job in producing hedge ratios that reduces the risk associated with selling naked calls. Indeed, while the hedges constructed from the NGARCH-Jump model were not worse, they added little to the explanatory power.

The hedging results indicate that even crude models might be very effective in hedging European call options. However, this may just be a property of European calls, and may not generalize to the hedging of exotic options, such as barrier options. Indeed, evidence that this is indeed the case is provided by Davydov and Linetsky (2001). The above hedging results can therefore be interpreted positively. The hedges are effective for Europeans, and, unlike the *ad hoc* Black-Scholes model, the methodology used to construct the hedges flows over very naturally to the hedging of exotic derivatives.

4 Conclusion

In this paper we have extended Duan's (1995) GARCH option model to incorporate jumps. A big advantage of this model is that it permits daily conditional logarithmic returns to be fat-tailed and/or skewed. The conditional return's skewness and kurtosis increase with the jump intensity. Our model, contains as special limiting cases, jump-diffusion models, like Merton (1976) and diffusive stochastic volatility models, like Heston (1993) or Hull and White (1987). Indeed, the exact linkages between our discrete time models and these continuous time models has recently been explored by Duan, Ritchken and Sun (2004).

Using data on the S&P 500 index and the set of European options, we have provided empirical tests of the ability of GARCH-Jump models to price and hedge options. We show that introducing jumps that allow for fat tails and higher kurtosis adds significantly to explaining the time series behavior of the S&P 500 and to pricing options. Further, the model is able to hedge European options very effectively.

[13]For example, American or exotic options cannot be priced using the ad hoc model information, while American or exotic GARCH option prices can easily be computed.

This paper is the first to provide empirical tests on the hedging performance of any GARCH option model. More important, we evaluate the performance of the NGARCH-Jump model and some of its nested special cases. From a time series perspective the NGARCH-Jump model outperforms the NGARCH-Normal model. Further, option prices generated using the NGARCH-Jump model are significantly better than NGARCH-Normal option prices with the primary explanation of its benefits coming from the fact that the conditional daily return is no longer normally distributed. However, while improving upon the NGARCH-Normal model, the NGARCH-Jump model, estimated using time series data alone, was not able to fully explain the volatility smile effect.

With regard to hedging over a fifteen-day period, the results of both the NGARCH and NGARCH-Jump model were impressive. The hedges, balanced daily could explain about 90% of the variability of the unhedged position.

In our analysis, we used historical information on asset returns to estimate most of the parameters required for option pricing. Given the parameter estimates and the path of the underlying index for daily data over the next year, we could update the state variables and price options in consecutive weeks. Our "out-of-sample" analysis over 50 weeks revealed that even out-the-money contracts could be well priced.

We illustrated how incorporating short-term at-the-money contracts into the analysis improved the fit of out-the-money contracts. If our sole goal was only to price and hedge options, then we should be able to improve our results by incorporating more historical information provided by the time series of all the options.

In general, distinguishing between stochastic volatility and jumps is difficult. Our empirical results showed that the jumps were small and frequent and their primary role was to generate short-term fat-tailed return distributions that better reflected the kurtosis. It is possible that introducing more dependence in the dynamics of the pricing kernel will have the effect of allowing for greater skewness and kurtosis in return distributions over longer time horizons, leading to a better explanation of the volatility skew. Finally, all our results related to a local volatility updating equation of the form given in equation (6). Härdle and Hafner (2000) have compared this volatility updating mechanism with a threshold, or TGARCH updating mechanism, and, by using simulations, they conclude that this scheme might be preferable. Since our theory is not limited to updates as in equation (6), other updating schemes can be easily implemented. These issues are left for future research.

References

Bakshi, G., and C. Cao (2003) Risk Neutral Kurtosis, Jumps, and Option Pricing: Evidence from 100 Most Actively Traded Firms on the CBOE, *working paper,* University of Maryland.

Bakshi, G., C. Cao and Z. Chen (1997) Empirical Performance of Alternative Option Pricing Models, *The Journal of Finance*, 53, 499-547.

Bakshi, G., and N. Kapadia (2003) Delta-Hedged Gains and the Negative Market Volatility Risk Premium, *Review of Financial Studies,* 527-566.

Bakshi, G., N. Kapadia and D. Madan (2003) Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options, *Review of Financial Studies,* 101-143.

Bates, D. (2000) Post-'87 Crash Fears in S&P 500 Futures Options, *Journal of Econometrics*, 94, 181-238.

Bates, D. (2002) Empirical Option Pricing: A Retrospection, *forthcoming, Journal of Econometrics.*

Chernov, M. and E. Ghysels (2000) Towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for the Purpose of Option Valuation, *Journal of Financial Economics*, 56, 407-458

Corradi V. (2000) Reconsidering the Continuous Time Limit of the GARCH(1,1) Process, *Journal of Econometrics*, 96, 145-153.

Das, S., and R. Sunderam (1999) Of Smiles and Smirks: A Term Structure Perspective, *Journal of Financial and Quantitative Analysis,* 34, 211-240.

Davydov, D., and V. Linetsky (2001) The Valuation and Hedging of Barrier and Lookback Options under the CEV Process, *Management Science,* 47, 949-965.

Duan, J. (1995) The GARCH Option Pricing Model, *Mathematical Finance*, 5, 13–32.

Duan, J. (1997) Augmented GARCH(p,q) Process and its Diffusion Limit, *Journal of Econometrics* 79, 97-127.

Duan, J., I. Popova and P. Ritchken (2002) Option Pricing Under Regime Switching, *Quantitative Finance*, 2, 116-132.

Duan,J., P. Ritchken and Z. Sun (2004) Approximating GARCH-Jumo Models, Jump-Diffusion Processes, and Option Pricing, *Working paper, University of Toronto and Case Western Reserve University.*

Dumas, B., J. Fleming and B. Whaley (1998) Implied Volatility Functions: Empirical Tests, *The Journal of Finance*, 53 2059-2106.

Duffie, D., K. Singleton and J. Pan (1999) Transform Analysis and Asset pricing for Affine Jump Diffusions, *Econometrica*, 68, 1343-1376.

Eraker, B. (2001) Do Equity Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices, *working paper*, University of Chicago.

Eraker B., M. Johannes and N. Polson (2003) The Impact of Jumps in Volatility and Returns, *Journal of Finance*, 3, 1269-1300.

Garcia, R. and E. Renault (1998) A Note on Hedging in ARCH and Stochastic Volatility Option Pricing Models, *Mathematical Finance* 8, 153-161.

Härdle W., and C. Hafner (2000) Discrete Time Option Pricing with Flexible Volatility Estimation, *Finance and Stochastics*, 4, 189-207.

Harvey, C. and R. Whaley (1992) Dividends and $S\&P$ 100 Index Option Valuation, *Journal of Futures Markets*, 12, 123-137.

Heston, S. (1993) A Closed-Form Solution for Options with Stochastic Volatility, *Review of Financial Studies* 6, 327-344.

Heston, S., and S. Nandi (2000) A Closed Form GARCH Option Pricing Model, *The Review of Financial Studies,* 13, 585-625.

Hull, J. and A. White (1987) The Pricing of Options on Assets with Stochastic Volatility, *The Journal of Finance*, 42, 281-300.

Jackwerth and M. Rubinstein (1996) Recovering Probability Distributions from Option Prices, *The Journal of Finance*, 51, 1611-1631.

Jacquier, E., N. Polson, and P. Rossi (1994) Bayesian Analysis of Stochastic Volatility Models, *Journal of Business and Economic Statistics*, 12, 1-19.

Jagannathan, R., A. Kaplin, and S. Sun (2001) An Evaluation of Multi-Factor CIR Models Using LIBOR, Swap Rates, and Cap and Swaption Prices, forthcoming, *Journal of Econometrics.*

Kurtz, T. and P. Protter (1991) Weak Limit Theorems for Stochastic Integrals and Stochastic Differential Equations, Annals of Probability, 19, 1035-1070.

Longstaff, F., P. Santa-Clara, and E. Schwartz (2001) The Relative Valuation of Caps and Swaptions: Theory and Empirical Evidence, *The Journal of Finance*, 56, 2069-2109.

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

Merton, R. (1976) Option Pricing when the Underlying Stock Returns are Discontinuous, *Journal of Financial Economics* 3, 125-144.

Maheu, J. and T. McCurdy (2002) News Arrivals, Jump Dynamics and Volatility Components for Individual Stock Returns, *working paper*, University of Toronto.

Naik, V. (1993) Option Valuation and Hedging Strategies with Jumps in the Volatility of Asset Returns, *The Journal of Finance*, 48, 1969-1984.

Naik, V. and M. Lee (1990) General Equilibrium Pricing of Options on the Market Portfolio with Discontinuous Returns, *The Review of Financial Studies*, 3, 493-521.

Nandi, S. (1998) How Important is the Correlation Between Returns and Volatility in a Stochastic Volatility Model? Empirical Evidence from Pricing and Hedging in the S&P 500 Index Option Market,*Journal of Banking and Finance,* 22, 589-610.

Nelson, D. (1990) ARCH Models as Diffusion Approximations, *Journal of Econometrics*, 45, 7-38

Pan, J. (2002) The Jump-risk Premia Implicit in Options: Evidence from an Integrated Time-series Study, *Journal of Financial Economics*, 63, 3-50.

Renault, E. (1997), Econometric Models of Option Pricing Errors, *Advances in Economics and Econometrics: Theory and Applications*, 3 (Kreps, D. and K. Wallis, eds), Cambridge University Press, 223-278.

Singleton, K. (2001) Estimation of Affine Asset Pricing Models using the Empirical Characteristic Function, *Journal of Econometrics*, 102, 111-141.

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

Appendix

Proof of Proposition 1

Substituting for the dynamics of the pricing kernel, we compute the following expectation:

$$E^{\mathsf{P}}\left[\frac{m_t}{m_{t-1}}|\mathcal{F}_{t-1}\right] = \exp\left[a + b^2/2 + \lambda(\kappa - 1)\right].$$

Since this value is the price of a one period discount bond with face value \$1, we have:

$$r = -\left(a + b^2/2 + \lambda(\kappa - 1)\right).$$

This equation uniquely identifies a in terms of the other parameters.

Now consider the pricing equation for the asset. We have, from equation (8),

$$E^{\mathsf{P}}\left[\frac{m_t}{m_{t-1}}\frac{S_t}{S_{t-1}}|\mathcal{F}_{t-1}\right] = 1.$$

Substituting for the dynamics of the pricing kernel and the asset price, the equation can be reexpressed as

$$E^{\mathsf{P}}\left[e^{\alpha_t + a + \tilde{X}_t^{(0)} + \sum_{j=1}^{N_t}\tilde{X}_t^{(j)}}\right] = 1$$

where

$$\begin{aligned}\tilde{X}_t^{(0)} &\sim N(0, \sigma_{0t}^2)\\ \tilde{X}_t^{(j)} &\sim N(b\mu + \sqrt{h_t}\bar{\mu}, \sigma_t^2)\end{aligned}$$

with

$$\begin{aligned}\sigma_{0t}^2 &= h_t + b^2 + 2\sqrt{h_t}b\rho\\ \sigma_t^2 &= h_t\bar{\gamma}^2 + b^2\gamma^2 + 2\sqrt{h_t}b\rho\gamma\bar{\gamma}\end{aligned}$$

Computing this expectation, the equation leads to:

$$\alpha_t + a + \sigma_{0t}^2/2 - \lambda + \lambda e^{b\mu + \sqrt{h_t}\bar{\mu} + \sigma_t^2/2} = 0$$

Finally, substituting the expression for a into the above equation leads to:

$$\alpha_t = r - \frac{h_t}{2} - \sqrt{h_t}b\rho + \lambda\kappa\left[1 - \exp\left(\sqrt{h_t}\left(\bar{\mu} + b\rho\gamma\bar{\gamma}\right) + \frac{h_t\bar{\gamma}^2}{2}\right)\right],$$

and the result follows.

Proof of Lemma 1

The proof follows along the line of Duan (1995).

(i) Q is a probability measure since:

$$\begin{aligned}\int 1 d\mathsf{Q} &= \int e^{rT}\frac{m_T}{m_0}d\mathsf{P} \\ &= E^{\mathsf{P}}\left[e^{rT}\frac{m_T}{m_0}\middle|\mathcal{F}_0\right] \\ &= E^{\mathsf{P}}\left[e^{rT}\frac{m_{T-1}}{m_0}E^{\mathsf{P}}\left(\frac{m_T}{m_{T-1}}\middle|\mathcal{F}_{T-1}\right)\middle|\mathcal{F}_0\right] \\ &= E^{\mathsf{P}}\left[e^{r(T-1)}\frac{m_{T-1}}{m_0}\middle|\mathcal{F}_0\right]\end{aligned}$$

where the last equality follows from the fact that:

$$E^{\mathsf{P}}\left[\frac{m_T}{m_{T-1}}\middle|\mathcal{F}_{T-1}\right] = e^{-r}.$$

Continuing this process we obtain

$$\int 1 d\mathsf{Q} = 1.$$

(ii) Now, for any $t < T$, we have:

$$\begin{aligned}E^{\mathsf{Q}}[Z_t|\mathcal{F}_{t-1}] &= E^{\mathsf{P}}\left[Z_t e^{r(T-t+1)}\frac{m_T}{m_{t-1}}\middle|\mathcal{F}_{t-1}\right] \\ &= E^{\mathsf{P}}\left[Z_t e^{r(T-t+1)}\frac{m_{T-1}}{m_{t-1}}E^{\mathsf{P}}\left(\frac{m_T}{m_{T-1}}\middle|\mathcal{F}_{T-1}\right)\middle|\mathcal{F}_{t-1}\right] \\ &= E^{\mathsf{P}}\left[Z_t e^{r(T-t)}\frac{m_{T-1}}{m_{t-1}}\middle|\mathcal{F}_{t-1}\right]\end{aligned}$$

Continuing this process, we obtain:

$$E^{\mathsf{Q}}[Z_t|\mathcal{F}_{t-1}] = e^r E^{\mathsf{P}}\left[Z_t\frac{m_t}{m_{t-1}}\middle|\mathcal{F}_{t-1}\right] = e^r Z_{t-1}.$$

So, Q is a local risk neutral probability measure.

Proof of Proposition 2

The first part of the proof follows along the line of Duan (1995). Let W_t represent the logarithmic return over period $[t-1, t]$. Then,

$$W_t = \alpha_t + \sqrt{h_t}\bar{J}_t.$$

We now consider the moment generating function of W_t under Q:

$$\begin{aligned}E^{\mathsf{Q}}[e^{cW_t}|\mathcal{F}_{t-1}] &= E^{\mathsf{P}}\left[e^{cW_t+r}\frac{m_t}{m_{t-1}}\middle|\mathcal{F}_{t-1}\right] \\ &= E^{\mathsf{P}}\left[e^{c\alpha_t+c\sqrt{h_t}\bar{J}_t+r+a+bJ_t}\middle|\mathcal{F}_{t-1}\right] \\ &= e^{c\alpha_t+r+a}E^{\mathsf{P}}\left[e^{c\sqrt{h_t}\bar{X}_t^{(0)}+bX_t^{(0)}+\sum_{j=1}^{N_t}(c\sqrt{h_t}\bar{X}_t^{(j)}+bX_t^{(j)})}\middle|\mathcal{F}_{t-1}\right]\end{aligned}$$

We know that

$$
\begin{aligned}
&E^{\mathsf{P}}(c\sqrt{h_t}\bar{X}_t^{(0)} + bX_t^{(0)}) = 0 \\
&E^{\mathsf{P}}(c\sqrt{h_t}\bar{X}_t^{(j)} + bX_t^{(j)}) = b\mu + c\sqrt{h_t}\bar{\mu}, \text{ for } j = 1, 2, ... \\
&Var^{\mathsf{P}}(c\sqrt{h_t}\bar{X}_t^{(0)} + bX_t^{(0)}) = c^2 h_t + b^2 + 2c\sqrt{h_t}b\rho \\
&Var^{\mathsf{P}}(c\sqrt{h_t}\bar{X}_t^{(j)} + bX_t^{(j)}) = c^2 h_t\bar{\gamma}^2 + b^2\gamma^2 + 2c\sqrt{h_t}b\rho\gamma\bar{\gamma}, \text{ for } j = 1, 2, ...
\end{aligned}
$$

Using these results, we obtain

$$
E^{\mathsf{Q}}[e^{cW_t}|\mathcal{F}_{t-1}] = \exp\left(c\alpha_t + r + a + \frac{1}{2}(c^2 h_t + b^2 + 2c\sqrt{h_t}b\rho) - \lambda(1 - \kappa K_t(c))\right) \tag{22}
$$

where $K_t(c)$ has been defined in Proposition 1.

Now, let $c = 0$. Then,

$$
1 = \exp\left(r + a + \frac{b^2}{2} - \lambda(1 - \kappa)\right)
$$

or, equivalently,

$$
r + a + \frac{b^2}{2} = \lambda(1 - \kappa)
$$

Substituting this expression into equation (22), we obtain

$$
E^{\mathsf{Q}}[e^{cW_t}|\mathcal{F}_{t-1}] = \exp\left(c\alpha_t + \frac{1}{2}(c^2 h_t + 2c\sqrt{h_t}b\rho) - \lambda\kappa(1 - K_t(c))\right) \tag{23}
$$

Now let $c = 1$. Then $E^{\mathsf{Q}}[e^{W_t}|\mathcal{F}_{t-1}] = e^r$. Hence:

$$
r = \alpha_t + \frac{1}{2}h_t + \sqrt{h_t}b\rho - \lambda\kappa(1 - K_t(1)),
$$

from which:

$$
\alpha_t + \sqrt{h_t}b\rho = r - \frac{1}{2}h_t + \lambda\kappa(1 - K_t(1)).
$$

Hence:

$$
E^{\mathsf{Q}}[e^{cW_t}|\mathcal{F}_{t-1}] = \exp\left[c\left(r - \frac{1}{2}h_t + \lambda\kappa(1 - K_t(1))\right) + \frac{1}{2}c^2 h_t - \lambda\kappa(1 - K_t(c))\right]
$$

Let

$$
\begin{aligned}
\tilde{\alpha_t} &= r - \frac{1}{2}h_t + \tilde{\lambda}\,(1 - K_t(1)) \\
\tilde{\lambda_t} &= \lambda\kappa
\end{aligned}
$$

We can write:

$$
E^{\mathsf{Q}}[e^{cW_t}|\mathcal{F}_{t-1}] = \exp\left[c\tilde{\alpha}_t + \frac{1}{2}c^2 h_t - \tilde{\lambda}\,(1 - K_t(c))\right] \tag{24}
$$

Now consider the following system:

$$\tilde{W}_t = \tilde{\alpha}_t + \sqrt{h_t}\tilde{J}_t$$

where

$$\begin{aligned}
\tilde{J}_t &= \tilde{X}_t^{(0)} + \sum_{j=1}^{\tilde{N}_t} \tilde{X}_t^{(j)} \\
\tilde{N}_t &\sim Poisson\left(\tilde{\lambda}\right) \\
\tilde{X}_t^{(0)} &\sim N(0,1) \\
\tilde{X}_t^{(j)} &\sim N(\bar{\mu} + b\rho\gamma\bar{\gamma}, \bar{\gamma}^2)
\end{aligned}$$

It is straightforward to verify that the moment generating function of $\tilde{W}_t$ is the same as that in equation (24). Thus, under measure $\mathtt{Q}$, W_t is distributionally equivalent to $\tilde{W}_t$.

The volatility dynamic can be expressed in terms of $\tilde{J}_t$ using $\bar{J}_t = \tilde{J}_t + b\rho$, which can be obtained via the return definition. The new innovation $\tilde{J}_t$ has mean $\tilde{\lambda}(\bar{\mu} + b\rho\gamma\bar{\gamma})$ and variance $(1 + \tilde{\lambda}\tilde{\gamma}^2)$, and thus requires the appropriate standardization in the expression.

Table 1
Estimates for the Five Models

The table shows the point estimates and standard deviations for the parameters of the three models that can be identified using the time series of daily data on the S&P 500 Index. The G-Merton and NGARCH-Jump parameter estimates in the bottom rows were obtained by incorporating the time series of the closest to 30 day at-the-money contracts. In particular the δ value is replaced by κ, $b\rho$ and γ and these parameters are estimated using least squares as explained in the text. Since γ was not significantly different from zero, the optimizations reported are based under the constraint that it was zero.

Parameter	Merton	G-Merton	NGARCH Normal	RNGARCH Jump	NGARCH Jump
β_0	6.41E-06	6.41E-06	1.83E-06	1.65E-07	1.65E-07
	(2.96E-07)	(2.96E-07)	(8.52E-06)	(6.63E-09)	(6.63E-09)
β_1	-	-	0.84795	0.84431	0.84431
			(0.0040)	(0.0062)	(0.0062)
β_2	-	-	0.07962	0.07560	0.07560
			(0.0035)	(0.0041)	(0.0041)
c	-	-	0.66425	0.77139	0.77139
			(0.0412)	(0.0008)	(0.0008)
λ	1.4365	1.4365	-	2.20226	2.20226
	(0.0471)	(0.0471)		(0.0004)	(0.0004)
$\widehat{\gamma}$	2.0705	2.0705	-	2.09608	2.09608
	(0.0451)	(0.0451)		(0.0014)	(0.0014)
$\widehat{\mu}$	0.12941	0.12941	-	0.0332	0.0332
	(0.0202)	(0.0202)		(0.0161)	(0.0161)
δ	0.081681	0.081681	-	8.48E-04	8.48E-04
	(0.0293)	(0.0293)		(0.0064)	(0.0064)
γ	0	0	-	0	0
κ	1	0.7252	-	1	0.8766
		(0.0127)			(0.0053)
$b\rho$	-0.10422	-0.05310	-0.02249	-0.07233	-0.0633
			(0.0085)		
ML Value	3605.9	3605.9	3616.57	3635.1	3635.1

Table 2
Skewness and Kurtosis of J_t

The table presents the actual and theoretical conditional skewness and kurtosis values for the three models estimated from the time series of S&P 500 Index values alone.

		NGARCH-Normal	MERTON	RNGARCH-Jump
Skewness	Actual	-0.030548	0.20467	-0.023474
	Theoretical	0	0.12437	0.027643
Kurtosis	Actual	4.008	4.9103	4.0233
	Theoretical	3	4.5474	4.119

Table 3
Pairwise Relative Performance of the RNGARCH-Jump Model

This table presents the fraction of times the RNGARCH-Jump model produces a smaller absolute pricing error than that produced by the model in the second row. The analysis is performed for all contracts in the middle of each week (Wednesday) over the fifty week periods after the parameters were estimated. The top panel performs the tests by moneyness; the bottom panel repeats the test by maturity.

Moneyness	Number Contracts	RNGARCH-Jump vs. Merton	G-Merton	NGARCH Normal	NGARCH Jump
(<-0.04)	341	0.95	0.38	0.45	0.19
(-0.04, -0.02)	288	0.85	0.54	0.54	0.30
(-0.02, -0.01)	168	0.82	0.55	0.54	0.43
(-0.01, -0.01)	371	0.70	0.60	0.49	0.62
(0.01, 0.02)	577	0.54	0.81	0.53	0.81
(0.02, 0.04)	669	0.71	0.87	0.56	0.80
(0.04, 0.08)	501	0.73	0.83	0.60	0.73
(>0.08)	795	0.74	0.69	0.63	0.68
Total	3710	0.73	0.71	0.55	0.64

Maturity	Number Contracts	RNGARCH-Jump vs. Merton	G-Merton	NGARCH Normal	NGARCH Jump
<30	510	0.75	0.60	0.56	0.55
31-60	978	0.76	0.71	0.63	0.62
61-90	831	0.76	0.71	0.55	0.62
>90	1391	0.68	0.74	0.51	0.70
Total	3710	0.73	0.71	0.55	0.64

Table 4
Median Absolute Percentage Errors

The parameter values for the RNGARCH-Jump model are estimated using the time series of asset prices up to week 200, as well as at-the-money option prices. In the following 50 weeks, the parameter values are not updated and the theoretical RNGARCH-Jump option prices are updated solely based on the path followed by the stock index. The top panel shows the median (and number of contracts) of the absolute percentage errors over all contracts for each maturity-moneyness bucket for all contracts for which we had actual prices over all days in the 50 week period. All option prices are computed conditional on the index value and the percentage error is defined as the theoretical price less the actual price divided by the actual price. The second panel shows the median of the absolute percentage errors for each moneyness bucket and for each 10-week time block in the out of sample period.

Moneyness	Maturity				
	<30	**30-60**	**60-90**	**>90**	**Total**
<-0.04	16.8% (15)	22.7% (273)	18.7% (466)	14.0% (922)	16.3% (1,676)
(-0.04, -0.02)	13.6% (185)	11.6% (419)	9.5% (365)	7.8% (444)	9.8% (1,413)
(-0.02, -0.01)	10.2% (155)	8.1% (235)	7.0% (218)	7.0% (245)	7.7% (853)
(-0.01, 0.01)	6.4% (338)	5.3% (497)	5.8% (451)	7.7% (555)	6.5% (1,841)
(-0.01, 0.02)	4.4% (176)	4.8% (258)	5.3% (243)	7.5% (297)	5.8% (974)
(0.02, 0.04)	3.6% (311)	4.0% (533)	4.8% (463)	6.7% (592)	5.0% (1,899)
(0.04, 0.08)	2.3% (410)	2.5% (938)	3.7% (799)	4.9% (1,160)	3.6% (3,307)
>0.08	1.0% (749)	1.1% (1,317)	1.4% (1,174)	1.8% (2,688)	1.4% (5,928)
Total	2.3% (2,339)	2.7% (4,470)	3.7% (4,179)	4.2% (6,903)	3.5% (17,891)

Moneyness	Out-of-Sample (Weeks)					
	1-10	**11-20**	**21-30**	**31-40**	**41-50**	**Total**
<-0.04	28.6%	17.1%	9.7%	16.1%	15.0%	16.3%
(-0.04, -0.02)	11.9%	7.7%	8.5%	11.2%	11.9%	9.8%
(-0.02, -0.01)	6.7%	5.0%	6.7%	6.3%	10.6%	7.7%
(-0.01, 0.01)	4.6%	4.5%	5.2%	10.2%	9.1%	6.5%
(-0.01, 0.02)	5.2%	4.7%	4.0%	7.6%	7.7%	5.8%
(0.02, 0.04)	5.1%	4.5%	3.5%	6.0%	5.9%	5.0%
(0.04, 0.08)	3.9%	3.3%	2.8%	3.9%	4.3%	3.6%
>0.08	1.8%	1.5%	1.3%	1.3%	1.2%	1.4%
Total	4.1%	3.3%	3.0%	3.0%	4.1%	3.5%

Figure 1
Box and Whisker Plots of Option Pricing Errors

The left column presents the mean percentage errors by moneyness for the four maturity buckets. In all cases, the ordering of the lines is MERTON, G-MERTON, NGARCH-Normal, RNGARCH-JUMP and NGARCH-JUMP, with the highest errors being for the first model and the lowest errors for the last model. The right column presents the box and whisker plots for the percentage pricing errors for each of the five models, plotted against moneyness over the range (-0.05 to +0.01) where there are significant differences. In each figure, the leftmost plot is MERTON, followed by G-MERTON, NGARCH-Normal, RNGARCH-Jump, and finally the NGARCH-Jump model.

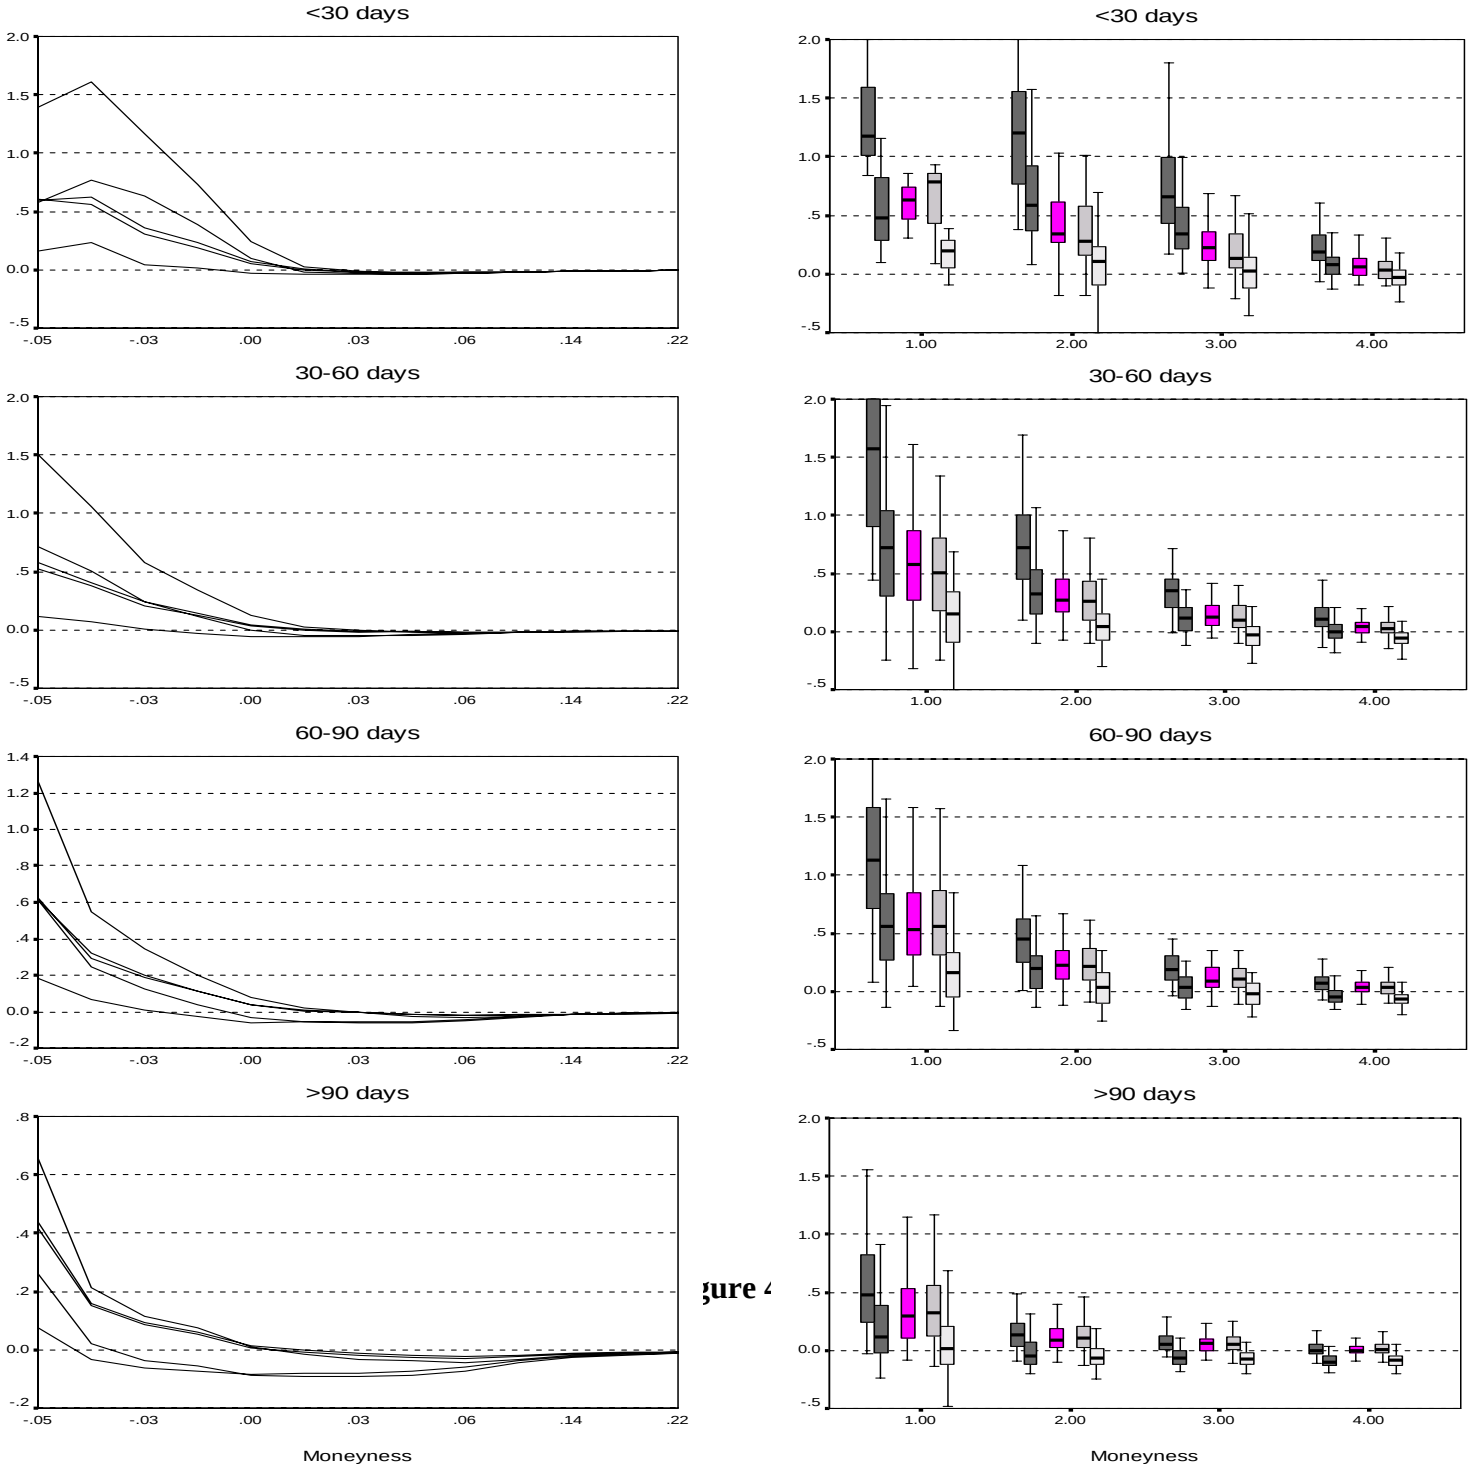

gure 4

Figure 2
Average Pricing Error by Time to Expiration

Each figure plots the average percentage error in prices across seven different maturity buckets, for the different moneyness categories. The thin solid line is the MERTON model, the dashed line is the GMERTON model, the more frequent dashed line in the NGARCH-Normal, the long dash-short dashed line is the RNGARCH-Jump, and the dark solid line is the NGARCH-Jump model.

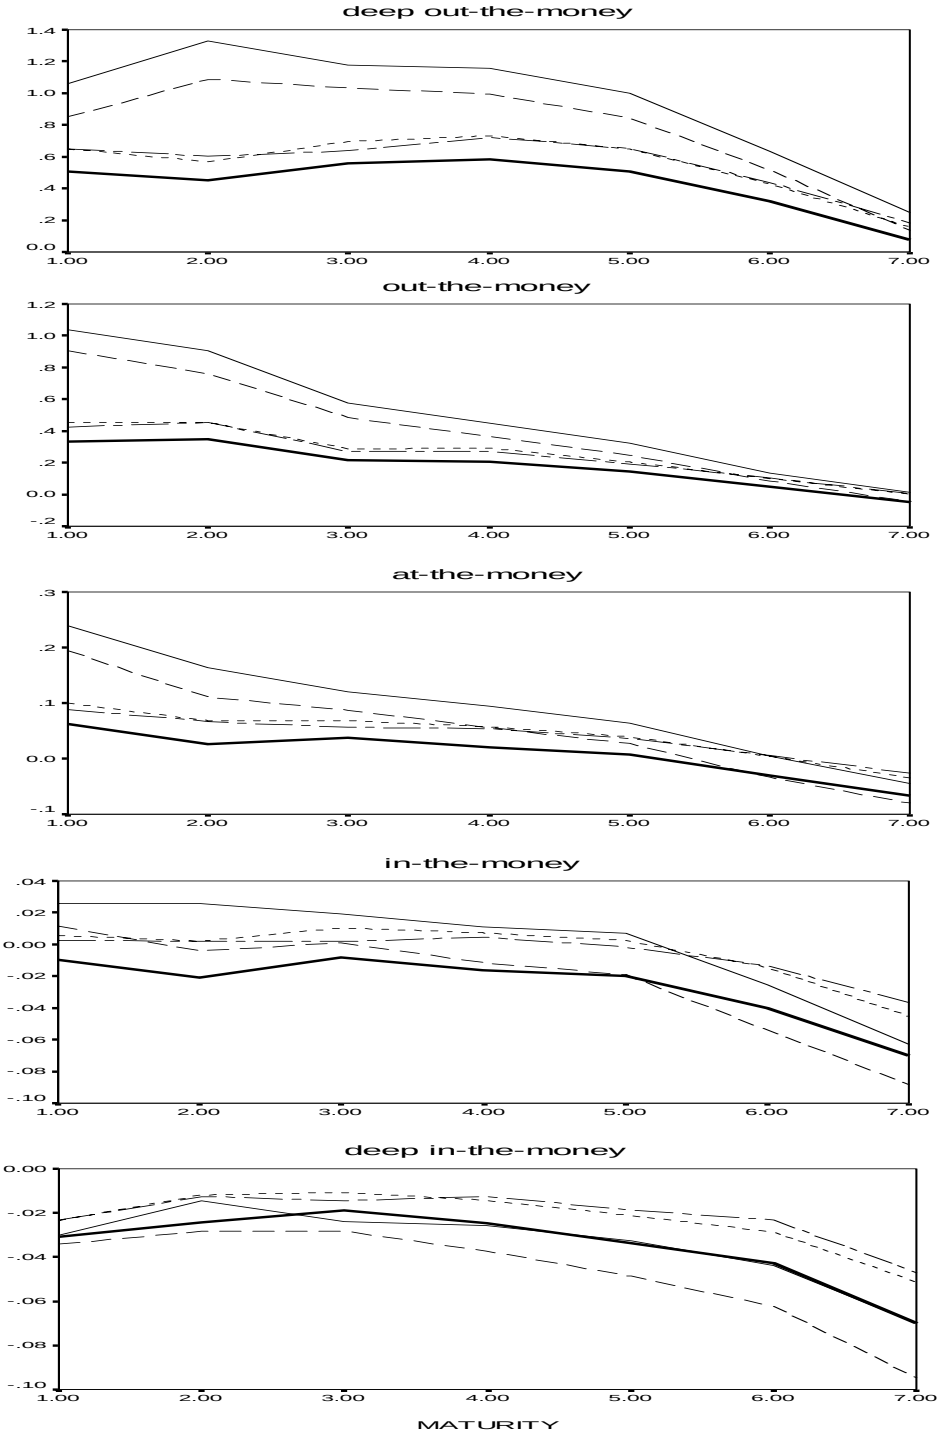

Figure 3
Box and Whiskers Plots of Net Profits of Delta Hedging over 15 days
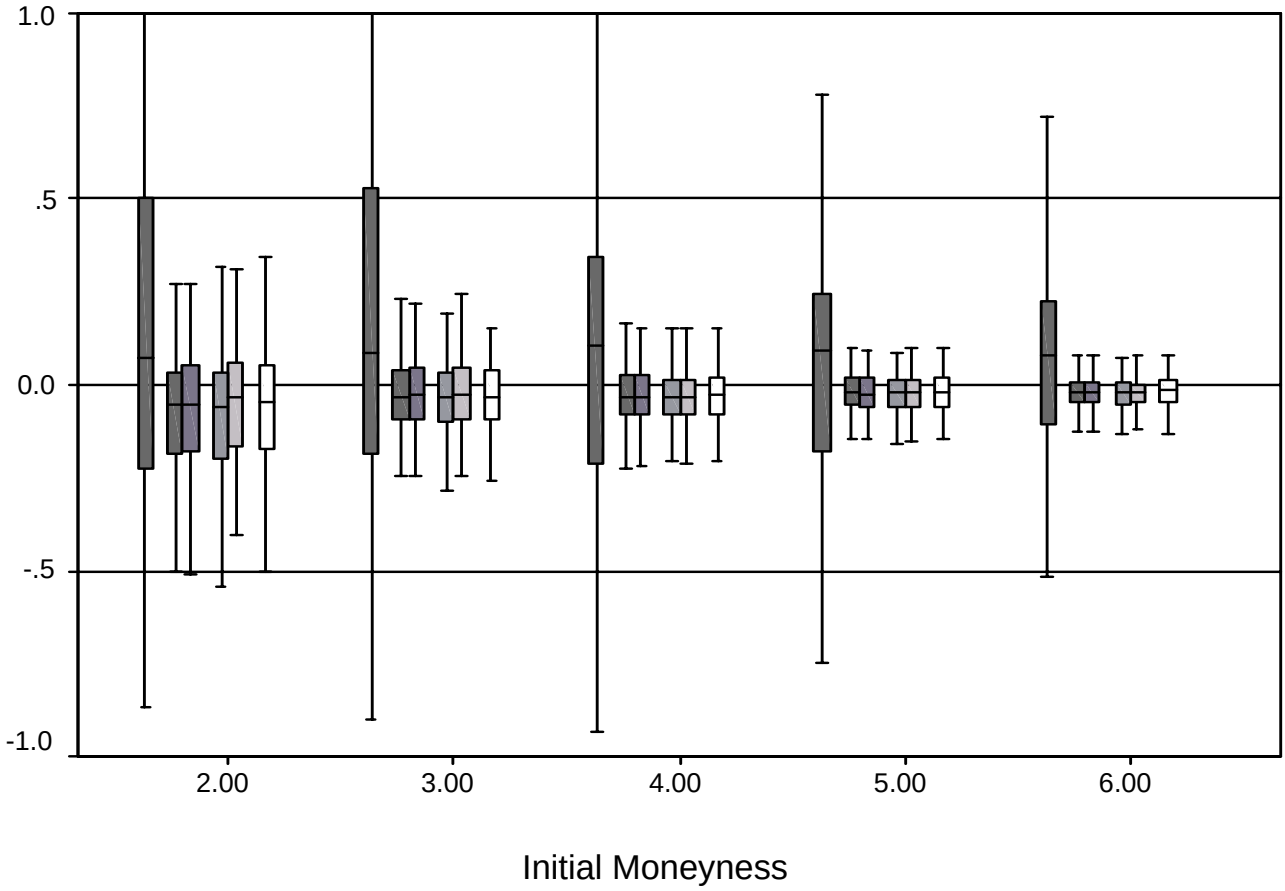
The figure compares the hedging errors of five different models to the unhedged error for all contracts in the out-of-sample hedging period as defined in the text. The plots are considered by initial moneyness. The left most box and whiskers plot is for the unhedged position in a one dollar investment in the option held for 15 days. The next five plots correspond to the NGARCH model, the RNGARCH-Jump model, the G-Merton model, the NGARCH-Jump model and the Black-Scholes model. The estimates for the first four model are all based on historical data. In contrast the hedge using the Black-Scholes model is based on the daily concurrent quoted implied volatility of each contract.