working paper

06 19

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

by Jin-Chuan Duan, Peter Ritchken, and Zhiqiang Sun

FEDERAL RESERVE BANK OF CLEVELAND

Working Paper 06-19 December 2006

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

by Jin-Chuan Duan, Peter Ritchken, and Zhiqiang Sun

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 models, where conditional returns are constrained to being normal, as well as mixed jump processes as used in Merton. The diffusion limits of our model have been shown to include jump diffusion models, stochastic volatility models and models with both jumps and diffusive elements in both returns and volatilities. 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 and improves option pricing performance. In addition, we provide the first ever hedging effectiveness tests of GARCH option models.

Keywords: GARCH option models, stochastic volatility models with jumps, pricing and hedging options

Jin-Chuan Duan is at the Rotman School of Management at the University of Toronto, and he can be reached at jcduan@rotman.utoronto.ca. Peter Ritchken is at the Weatherhead School of Management at Case Western Reserve University and is a research associate of the Federal Reserve Bank of Cleveland. He can be reached at peter.ritchken@case.edu. Zhiqiang Sun is at National City Corporation, and he can be reached at zhiqiang.sun@nationalcity.com.

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. This GARCH-Jump option pricing model is thus a generalization of the typical GARCH option pricing models with normal innovations, a pricing approach started by Duan (1995). We empirically test our model, and show that it not only fits the return data better than traditional GARCH models with normal innovations. Moreover, our model is better in removing more of the biases in option prices relative to the models with conditional normality.

Our new GARCH models are also interesting in that they serve as discrete time approximations for an array of continuous time jump diffusion models. Duan, Ritchken and Sun (2005) have derived a variety of continuous-time limiting models based on our GARCH-Jump processes. When the GARCH process is curtailed, but jumps allowed, the 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, the limiting model can be made to converge to continuous time stochastic volatility models, including Heston (1993), Hull and White (1987) and Scott (1987), among others. Finally, when jumps are permitted in our model, the limiting models contain jumps and diffusive elements in both returns and volatilities, along the lines of Eraker, Johannes and Polson (2003) and Duffie, Singleton and Pan (1999).

Just as the binomial model serves as a discrete time approximation for many underlying diffusion processes, our model serves as a useful approximation for underlying diffusive processes that permit jumps in returns and/or volatilities. Further, just as the appropriately redefined binomial model provides a useful mechanism for pricing American style options under a geometric Wiener process, our appropriately redefined risk neutralized discrete GARCH models provide a mechanism for pricing options under processes that also include random jumps in returns and/or volatilities. In light of the linkages between our GARCH-Jump models and their continuous time limits, our empirical results can be immediately linked to the huge literature on empirical performance of continuous time models.

Strictly speaking, the binomial model is a pure jump model where prices in each time increment can jump to one of two values. Similarly, GARCH models can be viewed as jump models, where in each time increment, the variance jumps to a new value based on the innovation that has just occurred. Here we emphasis "*Jump*" in our GARCH models to reflect the fact that the conditional return distribution in each time increment are compound Poisson random variables. This feature serves two purposes. First, it allows conditional distributions to have skewness and kurtosis that can readily be matched to the data. Second, as the time increments diminish, this feature leads to convergence of the return series towards processes that have diffusion and *jump* elements in either or both returns and volatilities.

This paper contributes to the literature in three ways. First, we establish the discrete time theory which allows us to price options when the underlying asset's innovations may be far from normal and when volatility is stochastic. This is important because most theoretical GARCH option models rely on normal innovations.[1] Second, we conduct an empirical analysis of these nested models that highlights the importance of incorporating jumps in returns and volatilities so as to better capture kurtosis and skewness in the time series dynamics and to better describe option prices and the volatility smile. Third, we provide a set of hedging tests among our GARCH option models. To our knowledge these hedging tests are the first to be performed using GARCH pricing models.

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.[2] While jumps in the return process can explain large daily shocks, these return shocks are highly 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. Eraker, Johannes and Polson (2003) examined 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. In contrast, Eraker (2004) 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 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.

The GARCH model has been extensively used in studying return time series. In recent years,

[1] We know of two exceptions. Duan (1999) developed a GARCH option model allowing for conditional skewness and kurtosis via a normal transformation technique. Christoffersen, Heston, and Jacobs (2004) developed a GARCH option pricing model using inverse Gaussian innovations.

[2] 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.

there has been an increasing use of the GARCH option pricing model to empirically examine its pricing performance. Heynen, Kemna and Vorst (1994), Duan (1996), Hardle and Hafner (2000), Heston and Nandi (2000), Hsieh and Ritchken (2000), Duan and Zhang (2001), Lehar, Scheicher and Schittenkopf (2002), Lehnert (2003) and Stentoft (2003) are some examples. More recently, Christoffersen and Jacobs (2004) examined a set of GARCH option models using the more general GARCH specifications given in Ding, Granger and Engle (1993) and Hentschel (1995). They concluded that while analysis of the return time series alone is in favor of more complex models, the option data suggest that the more parsimonious models with simple volatility clustering and leverage effects tend to have better performance. The GARCH option pricing models considered in Christoffersen and Jacobs (2004) all have conditionally normal innovations. Our study using the GARCH-Jump option pricing model thus adds to the empirical GARCH option pricing literature.

Our empirical analysis focuses on a nested set of models that contain interesting special cases. At one extreme, we consider models where in the limit 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 models containing no jumps but allow volatility to be time varying. Finally, we consider models where jump and diffusive risk is priced and whose continuous time limits 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. So, to the extent possible, we do not use option data to estimate parameter values, but rather view option prices as providing a set of information for which we can assess the ability of the pricing models. When we do use option data, in conjunction with time series information, to estimate parameters, we make sure that our tests are truly out of sample. Specifically, once all the parameters are estimated, we observe the future path of the asset, and based on the path, we 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. To our knowledge this is the first paper that has examined hedging effectiveness using GARCH models of any form. Our results for the $S\&P$ 500 demonstrate that incorporating jumps in volatility adds significantly to explaining the time series properties of the index, adds significantly to explaining patterns in option prices, and is capable of being well used to establish hedge ratios for dynamic rebalancing.

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 establish our nested models which represent interesting special cases. 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 nested GARCH with jumps models and present empirical evidence from time series of the $S\&P500$ index. In section 4 we 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 we 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 $\mathbf{F} = (\mathcal{F}_t)_{t\in\{0,1,\cdots,T\}}$ 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.[3] The random variables $X_t^{(j)}$ are independent for $j = 0, 1, 2, \cdots$ and $t = 1, 2, \cdots, T$.

[3]Maheu and McCurdy (2004) developed a GARCH-Jump model that allows for time variation and clustering in the jump intensity. They focused on describing the time series dynamics of individual stocks, rather than on the pricing of options.

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

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

Furthermore, for $t = 1, 2, \cdots, T$:

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

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 draw from the sum of $k+1$ normal distributions, while the logarithmic return of the asset also consists of a draw 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. For example, it could depend on all previous scale variables and standardized residuals. That is:

$$h_t = F(h_{t-i}, \bar{J}_{t-i}; i = 1, 2, \cdots) \tag{6}$$

We assume that the single period continuously compounded interest rate is constant, say, r.[4] Thus, the following restrictions must hold:

$$\begin{aligned} E^{\mathsf{P}}\left[\left.\frac{m_t}{m_{t-1}}\right| \mathcal{F}_{t-1}\right] &= e^{-r} && (7) \\ E^{\mathsf{P}}\left[\left.\frac{m_t}{m_{t-1}}\frac{S_t}{S_{t-1}}\right| \mathcal{F}_{t-1}\right] &= 1 && (8) \end{aligned}$$

[4]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.

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

$$\begin{aligned}
\alpha_t &= r - \frac{h_t}{2} - \sqrt{h_t} b\rho + \lambda\kappa\left(1 - K_t\right) && (10)\\
h_t &= F(h_{t-i}, \bar{J}_{t-i}; i = 1, 2, \cdots) && (11)\\
\kappa &= \exp\left(b\mu + \frac{1}{2} b^2\gamma^2\right) && (12)\\
K_t &= \exp\left(\sqrt{h_t}(\bar{\mu} + b\rho\gamma\bar{\gamma}) + \frac{1}{2} h_t \bar{\gamma}^2\right). && (13)
\end{aligned}$$

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

$$d\mathsf{Q} = e^{rT} \frac{m_T}{m_0} d\mathsf{P}. \tag{14}$$

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.

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

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

$$\frac{S_t}{S_{t-1}} = e^{\tilde{\alpha}_t + \sqrt{h_t}\tilde{J}_t} \tag{15}$$

where

$$\tilde{\alpha}_t = r - \frac{h_t}{2} + \tilde{\lambda}\,(1 - K_t) \tag{16}$$

$$h_t = F(h_{t-i}, \tilde{J}_{t-i} + b\rho; i = 1, 2, \cdots) \tag{17}$$

$$\tilde{J}_t = \tilde{X}_t^{(0)} + \sum_{j=1}^{\tilde{N}_t} \tilde{X}_t^{(j)} \tag{18}$$

$$\tilde{X}_t^{(0)} \sim N(0,1) \;\; \textit{for } t = 1, 2, \cdots, T$$

$$\tilde{X}_t^{(j)} \sim N(\bar{\mu} + b\rho\gamma\bar{\gamma}, \bar{\gamma}^2) \textit{ for } t = 1, 2, \cdots, T \textit{ and } j = 1, 2, \cdots$$

$$\tilde{X}_t^{(j)} \textit{ are independent for } t = 1, 2, \cdots, T \textit{ and } j = 0, 1, 2, \cdots$$

$\tilde{N}_t$ has a Poisson distribution with parameter $\tilde{\lambda} \equiv \lambda\kappa$ and K_t has been defined in Proposition 1.

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 $b\rho\gamma = 0$.[5] In other words, one should not in general expect the local risk-neutral valuation principle to apply. The expected value, $E^Q(\tilde{J}_t|\mathcal{F}_{t-1})$ and variance, $Var^Q(\tilde{J}_t|\mathcal{F}_{t-1})$ of $\tilde{J}_i$ are:

$$E^{\mathsf{Q}}(\tilde{J}_t) = \tilde{\lambda}\bar{\mu} + b\rho\gamma\bar{\gamma} \tag{19}$$

$$Var^{\mathsf{Q}}(\tilde{J}_t) = 1 + \tilde{\lambda}\tilde{\gamma}^2, \tag{20}$$

1.1 Updating Schemes for the Scaling Factor

For empirical work it is necessary to select specific structures for the scaling factor dynamic in equation (6). We consider two specific GARCH(1,1) dynamics.

[5]This result differs from the local risk-neutral valuation 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. 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).

1.1.1 The NGARCH Model

The first GARCH(1,1) model that we consider is of the form:

$$h_t = \beta_0 + \beta_1 h_{t-1} + \beta_2 h_{t-1} \left(\frac{\bar{J}_{t-1} - \lambda\bar{\mu}}{\sqrt{1+\lambda\hat{\gamma}^2}} - c \right)^2, \tag{21}$$

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. 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 $\beta_1 + \beta_2(1+c^2) < 1$. Both results are available in Duan (1997).

Notice that when $\lambda = 0$, the model reduces to the NGARCH-Normal process. In their empirical tests, Christoffersen and Jacobs (2004) found that this volatility dynamic performed the best among many GARCH option models with normal innovations. Their findings motivate this particular choice.

Using equations (19), (20) and (17), the updating scheme for the local scaling factor, h_t, specialized to equation (21), under measure Q, can be written as

$$h_t = \beta_0 + \beta_1 h_{t-1} + \beta_2^* h_{t-1} \left(\frac{\tilde{J}_{t-1} - (\tilde{\lambda}\bar{\mu} + b\rho\gamma\bar{\gamma})}{\sqrt{1+\tilde{\lambda}\tilde{\gamma}^2}} - c^* \right)^2 \tag{22}$$

where

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

Note that the fractional term inside the brackets has mean 0 and variance 1 under measure Q.

1.1.2 The TGARCH Model

The second GARCH specification that we consider is a TGARCH scheme. Let ϕ_t be an auxiliary state variable that fully determines h_t. We assume:

$$\begin{aligned}
\phi_t &= \beta_0 + \beta_1 \phi_{t-1} + \beta_2 \left| \frac{\bar{J}_{t-1} - \lambda\bar{\mu}}{\sqrt{1+\lambda\hat{\gamma}^2}} \right| + \beta_3 \max\left(-\frac{\bar{J}_{t-1} - \lambda\bar{\mu}}{\sqrt{1+\lambda\hat{\gamma}^2}}, 0 \right) \\
h_t &= \phi_t^2,
\end{aligned} \tag{23}$$

Again when $\lambda = 0$ this updating scheme reduces to the standard TGARCH scheme, that Hardle and Hafner (2000) found to be useful.

For this case the volatility process under measure Q, has the following dynamics:

$$\begin{aligned}
\phi_t &= \beta_0 + \beta_1 \phi_{t-1} + \beta_2^* \left| \frac{\tilde{J}_{t-1} - (\tilde{\lambda}\bar{\mu} + b\rho\gamma\bar{\gamma})}{\sqrt{1+\tilde{\lambda}\tilde{\gamma}^2}} + q \right| \\
&\quad + \beta_3^* \max\left(- \frac{\tilde{J}_{t-1} - (\tilde{\lambda}\bar{\mu} + b\rho\gamma\bar{\gamma})}{\sqrt{1+\tilde{\lambda}\tilde{\gamma}^2}} - q, 0 \right) \qquad (24) \\
h_t &= \phi_t^2,
\end{aligned}$$

where

$$\begin{aligned}
\beta_j^* &= \beta_j \sqrt{\frac{1+\tilde{\lambda}\tilde{\gamma}^2}{1+\lambda\hat{\gamma}^2}}, \qquad j = 2,3 \\
q &= \frac{b\rho(1+\gamma\bar{\gamma}) + \bar{\mu}\lambda(\kappa - 1)}{\sqrt{1+\tilde{\lambda}\tilde{\gamma}^2}}
\end{aligned}$$

In summary, when the local scaling factor h_t follows a NGARCH or TGARCH process, as in equation (21) or equation (23), then under measure Q, the updating schemes translates into a similar NGARCH and TGARCH processes. Proposition 2 allows us to easily come to specific pricing systems corresponding to different volatility dynamics.

1.2 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:

$$\begin{aligned}
\eta_t &= \lambda\kappa(1-K_t) - \lambda(1 - e^{\bar{\mu}\sqrt{h_t} + \frac{\bar{\gamma}^2 h_t}{2}}) - \sqrt{h_t}b\rho \qquad (25) \\
&\approx [\lambda\bar{\mu}(1-\kappa) - b\rho(1+\lambda\kappa\gamma\bar{\gamma})]\sqrt{h_t} + \lambda\bar{\gamma}^2(1-\kappa)\frac{h_t}{2}, \qquad (26)
\end{aligned}$$

where the approximation is justified if h_t is small.[6]

To gain some intuitions on the pricing model, first consider the case when $\kappa = 1$ and $\gamma = 0$. In this case, the risk premium, η_t reduces to $-b\rho\sqrt{h_t}$. This amounts to saying that the jump risk is fully diversifiable, which corresponds to the assumption made in Merton (1976). With $\kappa \neq 1$

[6] In our empirical studies we obtain h_t in the order of 10^{-6}.

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. Naik and Lee (1990) extended Merton's model to the case where jump risk is not diversifiable. In our model this is accomplished by releasing κ from 1 and/or γ from 0.

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:

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

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 affects both the interest rate and asset price.

1.3 The Nested Models

First, consider the case where $\kappa = 1$ and $\gamma = 0$. In this case the risk premium reduces to $\eta_t = -b\rho\sqrt{h_t}$. That is, the risk premium does not depend on jumps. With $\beta_1 = \beta_2 = 0$ in equation (21) or $\beta_1 = \beta_2 = \beta_3 = 0$ in equation (23) the scaling factor remains constant. Since jump risk is diversifiable, the local scaling factor is constant, and innovations, conditional on the number of jumps are normal, we refer to this model as the discrete-time Merton model, or MERTON, for short.

Second, consider the same model, but release κ and γ from 1 and 0. This implies that jump risk is priced. We call this model the generalized Merton model, or G-MERTON, for short.

The third set of models we consider are models with no jumps, i.e., $\lambda = 0$, but with our scaling factor being stochastic. In this case, innovations are normal random variables, and the risk premium is given by $\eta_t = -b\rho\sqrt{h_t}$. If the volatility dynamic is given by equation (21), the system becomes the NGARCH-Normal model. If volatility evolves according to equation (23), we have the TGARCH-Normal model. According to Duan (1997), these two models, in the limit, give rise to an extended version of the Hull and White (1987) and Heston (1993) stochastic volatility models, respectively.

The fourth set of models that we consider are models where $\kappa = 1$ and $\gamma = 0$ again, but the scaling factors are permitted to be stochastic and jumps are permitted. In these models, jump risk is diversifiable, volatility is stochastic and innovations are not normal. The two models are referred to as the Restricted NGARCH model and the Restricted TGARCH model.

The final set of models are the most general models where jump risk is priced, scaling factors are stochastic, jumps are present and innovations are not normal. These two models are referred to as the NGARCH-Jump and TGARCH-Jump models or as the full models. Duan, Ritchken and Sun (2005) have investigated the limiting behavior of these models as the time increment between consecutive updates is narrowed. They show that these models can be made to converge to continuous time models with diffusive elements and jumps in both returns and volatilities. For example, the TGARCH-Jump model can be viewed as a proxy for the following continuous-time process:

$$d\text{ln}S_t = f_{t-}dt + \sqrt{h_{t-}}dW_t + \left(\bar{\gamma}Z_{\pi_t} + \bar{\mu}\right)\sqrt{h_{t-}}d\pi_t \tag{27}$$

$$\begin{aligned} dh_t &= \left[\frac{\beta_3^2}{4(1+\lambda\hat{\gamma}^2)} + \left(\beta_2 + \frac{\beta_3}{2}\right)^2 \frac{\pi - 2}{\pi(1+\lambda\hat{\gamma}^2)} + 2\left(\beta_1 - 1\right)h_{t-}\right]dt \\ &\quad - \frac{\beta_3}{\sqrt{1+\lambda\hat{\gamma}^2}}\sqrt{h_{t-}}dW_t + (2\beta_2 + \beta_3)\sqrt{\frac{\pi-2}{\pi(1+\lambda\hat{\gamma}^2)}}\sqrt{h_{t-}}dB_t \\ &\quad + \frac{1}{1+\lambda\hat{\gamma}^2}\left[\beta_2\left|\bar{\gamma}Z_{\pi_t} + \bar{\mu}\right| + \beta_3 max(-\bar{\gamma}Z_{\pi_t} - \bar{\mu}, 0)\right]^2 d\pi_t. \end{aligned} \tag{28}$$

where

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

for all $0 \le t \le T$.

The above model is a mean-reverting square root process with jumps for h_t. By turning off jumps, the limiting model nests the square root stochastic volatility model given in Scott (1987) and Heston (1993). Without switching off jumps, the volatility dynamic in equation (28) is more general than that in Bakshi, Cao and Chen (1997), Bates (2000) or Pan (2002), for it allows for volatility jumps as well.

We therefore consider a total of 8 models, summarized in Table 1.

Table 1 Here

We will explore in the next section which of the models nested in our family can simultaneously explain both the time series of the $S\&P$ 500 index values 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 GARCH-Jump models using time series data on the $S\&P$ 500 index and dividends. Our main goal is to estimate models using return time series alone and then to evaluate the ability of these models to price and hedge options. We are particularly interested in evaluating the full NGARCH (TGARCH) model as well as their nested special cases.

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.[7] Since we used 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

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

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 Return 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

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

where the function α_t is given in Proposition 1, and the function $v(\cdot)$ is given by (21) or (23). The initial value of the local scaling factor is determined by

$$h_1 = V/(1 + \lambda\hat{\gamma}^2) \tag{29}$$

where V is the sample variance of the asset return and as defined earlier, $\hat{\gamma}^2 = \bar{\mu}^2 + \bar{\gamma}^2$.[8] Our model parameter set is

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

[8]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.

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 $\sigma_i^2(t) = h_t(1 + i\bar{\gamma}^2)$.[9] The log-likelihood function of the sample is:

$$L(\theta; y_1, ..., y_T) = \sum_{t=2}^{T} \ln \left[l(y_t|h_t, y_{t-1}) \right]. \tag{30}$$

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:

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

Hence,

$$\alpha_t \approx r - \frac{1}{2} h_t (1 + \lambda\kappa(\bar{\gamma}^2 + (\bar{\mu} + b\rho\gamma\bar{\gamma})^2)) - \sqrt{h_t}(b\rho + \lambda\kappa\bar{\mu} + \lambda\kappa b\rho\gamma\bar{\gamma}) \tag{31}$$

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 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 $\theta^* = \{\beta_0, \beta_1, \beta_2, c, \delta, \bar{\mu}, \bar{\gamma}, \lambda\}$, for the NGARCH models and $\theta^* = \{\beta_0, \beta_1, \beta_2, \beta_3, \delta, \bar{\mu}, \bar{\gamma}, \lambda\}$, for the TGARCH models.

Notice that when $\kappa = 1$ and $\gamma = 0$ we obtain the Restricted-Jump models. These models, as well as the MERTON model (with $\beta_1 = \beta_2 = 0$) can be fully estimated using the maximum likelihood method on the return time series alone. In addition, the two GARCH-Normal models

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

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 Restricted-Jump models to obtain estimates for the G-MERTON and the two full GARCH-Jump models.

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 two weeks, 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 (2003), 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 the return time series data as much as possible, and then assess the models not only on their return 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

$$\Delta_t = e^{-r(T-t)} E^{\mathsf{Q}} \left[\frac{S_T}{S_t} 1_{[\frac{S_T}{S_t} > \frac{X}{S_t}]} \right] \tag{32}$$

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 (32) 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 2 shows the parameter estimates based on the time series data over the first 1000 trading days for the 8 models. In particular, for all parameters that can be estimated from the time series alone, we report the point estimates and their standard deviations.

Table 2 Here

First, consider the 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, for the NGARCH models, 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 full NGARCH-Jump model and restricted NGARCH-Jump model. For the TGARCH models all the β values are significant.

Table 2 also reports the additional estimates for the G-MERTON and the full NGARCH-Jump and TGARCH-Jump models, when option data was used to identify κ, $b\rho$ and γ. The effects on option prices to changes in γ were found to be very minor, and hence the results reported here are obtained by fixing $\gamma = 1$. For the normal models, since $\lambda = 0$, $b\rho = \delta$, and these values are reported.

The option information allows us to extract information on jump risk premia. In particular according to equation (26), the contribution of the diffusion risk premium for the NGARCH-Jump model, $-b\rho\sqrt{h_t} = 0.01246\sqrt{h_t}$, while the jump risk premium, $\lambda(\bar{\mu}(1-\kappa) - b\rho\kappa\gamma\bar{\gamma})\sqrt{h_t} = 0.05986\sqrt{h_t}$. So the introduction of "jumps", that allow for kurtosis and skewness, accounts for about 82.7% of the total risk premium. For the TGARCH-Jump model, the contribution of the diffusive term to the risk premium is $-b\rho\sqrt{h_t} = 0.01623\sqrt{h_t}$, while the remaining components contribution is $0.07880\sqrt{h_t}$. So the introduction of "jumps", accounts for about 82.9% of the total risk premium.

Table 3 reports the skewness and kurtosis of the conditional daily return residual normalized by the square root of the local scaling factor, i.e., $y_t/\sqrt{h_t}$ for the various models. We know that the NGARCH and TGARCH-Normal models have conditionally normal distributions so the kurtosis of residuals should be 3. But Table 3 shows that the actual kurtosis is larger than 3.[10]

Table 3 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. In our model 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, for example, has a predicted kurtosis (skewness) of normalized daily return residuals equal to 4.119, (0.0276) close to the observed kurtosis of 4.107 (0.0264).

[10] 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+1}), 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 bins set up as follows: $1 = (< -0.05)$, $2 = (-0.05, -0.04)$, $3 = (-0.04, -0.03)$, $4 = (-0.03, -0.02)$, $5 = (-0.02, -0.01)$, $6 = (-0.01, 0.01)$, $7 = (0.01, 0.02)$, and so on up to $11 = (> 0.05)$. For ease of presentation sometimes we combined moneyness categories, and focused on out-the-money contracts since this is where the models produce significantly different results. 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. Our out-of-sample call option set consists of $17,891$ (3,710) contracts, when computed daily (weekly).

Table 4 shows the results of pairwise comparisons of the out-of-sample residuals for each option contract in the sample. In particular, the table reports the proportion of occasions that the model on the row outperformed the model in the column over the class of contracts indicated in the three panels.

Table 4 Here

The purpose of this table is to compare the performance of the 8 models and hopefully reduce the set to a fewer candidate set of models. As an example, from the first panel of deep out-the-money contracts we see that the G-Merton model outperforms the Merton Model in 92.3% of the contracts.[11]

From the tables we see that the models that incorporate historical option price information clearly outperform models that only used historical time series information on the stock. In particular, the G-Merton model outperforms the Merton Model and the Full NGARCH-Jump (Full TGARCH-Jump) model outperforms their restricted versions. It appears that incorporating option information in the estimation phase improves the out-of-sample performance.

The difference between the NGARCH and TGARCH models are different with an overall slight edge going to the NGARCH models. For example, for out-the money contracts in our three groups, (near, middle, and deep out-the-money contracts) the NGARCH-Normal model wins 47, 66 and 68% of the time, while the Restricted NGARCH-Jump model beats the Restricted

[11]The Table does not report the standard errors of these estimates, but in all cases sample sizes were fairly large (over 300) and standard errors were always less than 2%.

TGARCH model in 24, 35 and 44% of the time. Finally, the unrestricted NGARCH-Jump model outperforms the unrestricted TGARCH-Jump model in 55, 56, and 56% of the times. The results indicate that there may not be a major difference between the NGARCH and TGARCH models.

To explore this more fully, Figure 1 provides box and whisker plots for the percentage errors of the three NGARCH models side by side with the corresponding TGARCH models for each moneynesss category and separated by different maturities.

Figure 1 Here

As can be seen there are very little differences between the NGARCH and TGARCH- Normal models, between the restricted NGARCH and TGARCH models and between the two unrestricted models. The differences between the three classes of models, namely, normal, restricted and unrestricted, are much greater than the differences between NGARCH and TGARCH models. In light of these results, in what follows we only focus on the NGARCH models. This reduces the set of models from 8 to 5 and makes the presentation of the following results more manageable.

Figure 2 summarizes the pricing performance of the 5 remaining 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 2 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 3 shows the average percentage errors in prices for each model across expiration dates. Since the results depend on moneyness, four 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 3 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.[12] 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 more carefully the performance of the restricted NGARCH model with the unrestricted model. If the differences are small, then it suggests that the majority of option pricing information can

[12]For a good review of the alternative approaches, see Bates (2003).

indeed be identified from the time series of the underlying alone. If using historical option price information adds information, then the unrestricted models should perform better. Table 5 shows the mean absolute percentage error for the restricted and unrestricted NGARCH with Jumps models over several moneyness categories and for our maturity buckets.

Table 5 Here

The Table clearly reveals the superiority of the unrestricted model in computing better out-the-money prices. This holds true for all maturity options. Interestingly, however, for the in-the-money contracts, there is a small bias in favor of the restricted model. Recall, however, that in estimating the full models, the only additional information beyond the time series of index prices that was used were the short dated at the money call option prices. Since this time series provided significant improvements in predictions, clearly using additional historical option prices could improve the results further. The main conclusion from this analysis, however, is that for estimation purposes, incorporating the time series of option data does indeed improve the estimates for future option pricing.

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.

Table 6 shows the average 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 standard errors are also provided. The average absolute pricing error over all $17,891$ contracts was 6%. The median percentage error was 3.5%.

Table 6 Here

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.[13] In spite of differing objectives, our error terms appear to be of similar magnitudes to their reported values.

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 6 shows the means 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 4 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, the full NGARCH-Jump model and the Black-Scholes model. 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 4 Here

From the box and whiskers plot of the unhedged residuals, we see that over the sample period, the S&P 500 index 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.[14] To our knowledge, the results reported here are the first results that attempt to measure the effectiveness of hedges established using GARCH based models.

[13] 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.

[14] 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.

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 as a calibrating device.[15]

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 (2003), who explained this finding by postulating a negative risk premium for volatility risk.

Eraker (2004) 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 that relies on normal innovations to incorporate non-normal innovations. These GARCH-Jump models extend the literature in a very important way. Specifically, they contain, as special limiting cases, models of the underlying that contain jumps in returns and/or in volatilities. This is in contrast to the

[15] 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.

typical GARCH models based on normal innovations. Since these latter models only contain diffusive stochastic volatility models as limiting cases, it is not surprising that they are not capable of removing well known option pricing biases. We provide the theory of GARCH option pricing that permits contracts to be priced in the presence of skewed and leptokurtic innovations, and demonstrate that these advances are empirically significant.

Specifically, 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.

For pricing of options our simplest nested model, the Merton model, performed the worst. Capturing time varying volatility, and including priced jump risk lead to better results. However, all models based on normal innovations were dominated by models that allowed non-normal innovations. Unlike the findings of Christoffersen and Jacobs (2004), we demonstrate that complex models of the underlying that go beyond capturing simple volatility clustering and leverage effects, can add significantly to explaining the volatility smile. Further, we showed that our GARCH-Jump models are capable of pricing options well without requiring frequent recalibration. Indeed, our models were capable of good pricing up to a year after the parameters were estimated. Finally, the hedging effectiveness of the GARCH-Jump models was examined. The models were able to hedge European options very effectively.

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 jumps were frequent, with more than one "jump" a day. This implies that to capture fat tailed return distributions random mixing of normal innovations is necessary. 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. Although our empirical results related to a local volatility updating equation of NGARCH and TGARCH forms, our evidence suggests that the difference between the two structures is not that important for option pricing. What is important is that models should contain jumps and conditional local returns should not be normal.

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, *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. (2003) Empirical Option Pricing: A Retrospection, *Journal of Econometrics*, 116, 387-404.

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

Christoffersen, P., S. Heston and K. Jacobs (2004) Option Valuation with Conditional Skewness, forthcoming *Journal of Econometrics.*

Christoffersen, P. and K. Jacobs (2004) Which GARCH Model for Option Valuation? *Management Science*, 50, 1204-1221.

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

Ding, Z., C. Granger, R. Engle (1993) A Long Memory Property of Stock Market Returns and a New Model, *Journal of Empirical Finance,* 83-106.

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

Duan, J. (1996) Cracking the Smile, *Risk*, 9, 55-59.

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

Duan, J. (1999) Conditionally Fat Tailed Distributions and the Volatility Smile in Options, working paper, University of Toronto.

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 (2005) Approximating GARCH-Jump Models, Jump-Diffusion Processes, and Option Pricing, *Mathematical Finance,* forthcoming.

Duan, J. and H. Zhang (2001) Pricing Hang Seng Index Options – A GARCH Approach, *Journal of Banking and Finance*, 25, 1989-2014.

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

Eraker, B. (2004) Do Equity Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices, *Journal of Finance,* 59, p. 1367-1403

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.

Hentschel, L. (1995) All in the Family: Nesting Symmetric and Asymmetric GARCH Models, *Journal of Financial Economics*, 39, 71-104.

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

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, *Review of Financial Studies*, 13, 585-625.

Hsieh K. and P. Ritchken (2000) An Empirical Comparison of GARCH Option Pricing Models, working paper, Case Western Reserve University.

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

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

Jagannathan, R., A. Kaplin, and S. Sun (2003) An Evaluation of Multi-Factor CIR Models Using LIBOR, Swap Rates, and Cap and Swaption Prices, *Journal of Econometrics,* 116, 113-146.

Lehar, A., M. Scheicher and C. Schittenkopf (2002) GARCH vs. Stochastic Volatility Option Pricing and Risk Management, *Journal of Banking and Finance*, 26, 323-345.

Lehnert, T. (2003) Explaining Smiles: GARCH Option Pricing with Conditional Leptokurtosis and Skewness, *Journal of Derivatives*, 27-39.

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

Maheu, J. and T. McCurdy (2004) News Arrivals, Jump Dynamics and Volatility Components for Individual Stock Returns, *Journal of Finance*, 59, 755-793.

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

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

Naik, V. and M. Lee (1990) General Equilibrium Pricing of Options on the Market Portfolio with Discontinuous Returns, *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.

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

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

Stentoft, L. (2003) Pricing American Options when the Underlying Asset Exhibits Time-Varying Return Volatility, working paper, University of Aarhus.

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:

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

where the last equality follows from the fact that:

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

Continuing this process we obtain

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

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

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

Continuing this process, we obtain:

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

So, Q is a local risk neutral probability measure.

Proof of Proposition 2

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:

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

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

$$E^{\mathsf{Q}}[e^{cW_t}|\mathcal{F}_{t-1}] = \exp\left(c\alpha_t + r + a + \frac{1}{2}(c^2 h_t + b^2 + 2c\sqrt{h_t}b\rho) - \lambda(1 - \kappa K_t(c))\right) \tag{33}$$

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 (33), we obtain

$$E^{\mathsf{Q}}[e^{cW_t}|\mathcal{F}_{t-1}] = \exp\left(c\alpha_t + \frac{1}{2}(c^2 h_t + 2c\sqrt{h_t}b\rho) - \lambda\kappa(1 - K_t(c))\right) \tag{34}$$

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:

$$E^{\mathsf{Q}}[e^{cW_t}|\mathcal{F}_{t-1}] = \exp\left[c\tilde{\alpha}_t + \frac{1}{2}c^2 h_t - \tilde{\lambda}\,(1 - K_t(c))\right] \tag{35}$$

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 (35). Thus, under measure 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. Thus, $h_t = F(h_{t-i}, \tilde{J}_{t-i} + b\rho; i = 1, 2, \cdots)$. 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:
Taxonomy of Models

Model	Restrictions	Parameters estimated using
Jump Models:		
(1) Merton	$\beta_1 = \beta_2 = \beta_3 = 0,\ \kappa = 1, \gamma = 0$	Return time series
(2) G-Merton	$\beta_1 = \beta_2 = \beta_3 = 0$	Return time series & Options
Normal Models:		
(3) NGARCH-Normal	$\lambda = 0$	Return time series
(4) TGARCH-Normal	$\lambda = 0$	Return time series
Restricted Models:		
(5) RNGARCH	$\kappa = 1, \gamma = 0$	Return time series
(6) RTGARCH	$\kappa = 1, \gamma = 0$	Return time series
Full Models:		
(7) NGARCH-Jump		Return time series & Options
(8) TGARCH-Jump		Return time series & Options

Table 2:
Estimates for the Eight Models

The table shows the point estimates and standard deviations for the parameters of all eight models. The Merton Model, the NGARCH-Normal and the TGARCH-Normal Model can be estimated from the time series alone. For the two normal models bρ= δ. For the two restricted models we estimate δ, and restrict κ and γ to be 1 and 0 respectively. This uniquely identifies bρ. For the three unrestricted models (G-Merton, Full NGARCH-Jump and Full TGARCH-Jump) we release the two restrictions, but require the δ value to be consistent with the time series value. We incorporate the time series of the closest to 30 day at-the-money contracts, and then identify κ, bρ and γ using least squares as explained in the text. Since the calibrated γ values were not significantly different from 1, the optimizations reported are based under the constraint that γ=1.

	Jump			NGARCH			TGARCH	
Parameter	**Merton**	**G-Merton**	**Normal**	**Restricted**	**Full**	**Normal**	**Restricted**	**Full**
β_0	6.41E-06	6.41E-06	1.83E-06	1.65E-07	1.65E-07	-1.10E-04	-3.40E-05	-3.40E-05
	(2.96E-07)		(8.52E-07)	6.63E-09		(4.81E-05)	(3.03E-05)	
β_1	-		0.84795	0.84431	0.84431	0.95765	0.96597	0.96597
			(0.0040)	(0.0062)		(0.0224)	(4.75E-03)	
β_2	-		0.07962	0.07560	0.07560	2.56E-04	5.75E-05	5.75E-05
			(0.0035)	(0.0041)		(8.144E-05)	(4.21E-05)	
β_3	-		-	-	-	5.09E-04	1.53E-04	1.53E-04
						(1.37E-04)	(3.91E-05)	
c	-		0.66425	0.77139	0.77139	-	-	-
			(0.0412)	(0.0008)				
λ	1.4365	1.4365	-	2.20226	2.20226	-	2.1304	2.1304
	(0.0471)			(0.0004)			(0.0243)	
$\bar{\gamma}$	2.0705	2.0705	-	2.09608	2.09608	-	2.158	2.158
	(0.0451)			(0.0014)			(0.0084)	
$\bar{\mu}$	0.12941	0.12941	-	0.0332	0.0332	-	0.054841	0.054841
	(0.0202)			(0.0161)			(2.758E-04)	
δ	0.081681	0.081681	-0.02249	8.48E-04	8.48E-04	-0.029296	0.0218	0.0218
	(0.0293)		(0.0085)	(0.0064)		(0.033034)	(1.58E-04)	
bρ	-	-0.02572	-0.0225	-0.0723	-0.01246	-0.0293	-0.0950	-0.0162
κ	-	0.9818	-	1	0.8513	-	1	0.9008
γ	-	1	-	0	1	-	0	1
ML Value	3605.9		3616.6	3635.1		3612.0	3631.4	

Table 3:
Skewness and Kurtosis of J_t

The table presents the actual and theoretical conditional skewness and kurtosis values of the normalized innovation for all the models estimated from the time series of S&P 500 Index values alone. The data used was the daily return data on the S&P 500 index over the four year matching period for which option data was available, starting from January 1991.

		Pure Jump (Merton & G-Merton)	**NGARCH- Models** Normal	(Restricted & Full)	**TGARCH- Models** Normal	(Restricted & Full)
Skewness	**Actual**	0.2047	0.015	0.0264	0.15	0.0437
	Theoretical	0.1244	0	0.0276	0	0.0458
Kurtosis	**Actual**	4.9103	4.0724	4.1071	4.0888	4.0689
	Theoretical	4.5473	3	4.119	3	4.162

Table 4:
Pairwise Comparisons of Absolute Errors

The Table shows the proportion of contracts where the model indicated by the row had smaller absolute errors than the model indicated by the column. For example, for deep out-the-money options, the Merton Model outperformed the G-Merton Model for only 7.7% of the contracts. The number of contracts used in each table exceeded 500 in each moneyness category, and the standard errors were all less than 2%. The analysis is performed for all contracts in the middle of each week (Wednesday) over the fifty week periods after the parameters were estimated

Deep-Out-The-Money Options (Moneyness <-0.04)

	Merton	G-Merton	NGARCH	TGARCH	RNGARCH	RTGARCH	FNGARCH	FTGARCH
Merton	-	7.7%	12.3%	11.3%	12.1%	13.5%	27.6%	27.0%
G-Merton	92.3%	-	14.5%	13.9%	13.5%	12.9%	29.4%	28.8%
NGARCH	87.7%	85.5%	-	47.4%	32.3%	21.2%	36.9%	36.7%
TGARCH	88.7%	86.1%	52.6%	-	50.0%	22.2%	38.7%	38.5%
RNGARCH	87.9%	86.5%	67.7%	50.0%	-	24.0%	38.5%	37.9%
RTGARCH	86.5%	86.5%	78.8%	77.8%	76.0%	-	46.4%	45.8%
FNGARCH	72.4%	70.6%	63.1%	61.3%	61.5%	53.6%	-	55.2%
FTGARCH	73.0%	71.2%	63.3%	61.5%	62.1%	54.2%	44.8%	-

Medium Out-The-Money Options (Moneyness <-0.02)

	Merton	G-Merton	NGARCH	TGARCH	RNGARCH	RTGARCH	FNGARCH	FTGARCH
Merton	-	16.0%	17.0%	19.1%	16.6%	21.3%	29.5%	29.1%
G-Merton	84.0%	-	20.9%	22.3%	18.9%	25.2%	30.9%	29.9%
NGARCH	83.0%	79.1%	-	65.8%	25.4%	32.6%	38.5%	38.3%
TGARCH	80.9%	77.7%	34.2%	-	30.5%	27.9%	36.5%	38.1%
RNGARCH	83.4%	81.1%	74.6%	69.5%	-	35.2%	38.7%	39.5%
RTGARCH	78.7%	78.7%	67.4%	72.1%	64.8%	-	46.1%	46.5%
FNGARCH	70.5%	69.1%	61.5%	63.5%	61.3%	53.9%	-	56.1%
FTGARCH	70.9%	70.1%	61.7%	61.9%	60.5%	53.5%	43.9%	-

Near Out-The-Money Options (Moneyness < 0)

	Merton	G-Merton	NGARCH	TGARCH	RNGARCH	RTGARCH	FNGARCH	FTGARCH
Merton	-	20.8%	16.9%	20.5%	16.0%	23.5%	34.5%	35.2%
G-Merton	79.2%	-	19.5%	22.1%	19.5%	25.7%	35.5%	37.5%
NGARCH	83.1%	80.5%	-	67.8%	30.3%	40.4%	45.9%	46.6%
TGARCH	79.5%	77.9%	32.2%	-	30.3%	30.6%	42.7%	43.6%
RNGARCH	84.0%	80.5%	69.7%	69.7%	-	43.6%	47.6%	48.5%
RTGARCH	76.5%	76.5%	59.6%	69.4%	56.4%	-	53.7%	52.1%
FNGARCH	65.5%	64.5%	54.1%	57.3%	52.4%	46.3%	-	56.4%
FTGARCH	64.8%	62.5%	53.4%	56.4%	51.5%	47.9%	43.6%	-

Table 5:
Average Absolute Percentage Errors Out-of-Sample

The parameter values for the Restricted NGARCH Jump model are estimated using the time series of asset prices up to week 200 alone, while the unrestricted NGARCH-Jump model also uses information on the at-the-money option prices. In the 50 weeks, after the parameter values are estimated, the theoretical option prices are updated solely based on the path followed by the stock index. The table shows the mean (and standard error) of the absolute percentage errors over all contracts for each maturity-moneyness bucket for 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.

Moneyness	T<30days		30<T<60		60<T<90		T>90days		Total	
	RNGARCH	FNGARCH	RNGARCH	FNGARCH	RNGARCH	FNGARCH	RNGARCH	FNGARCH	RNGARCH	FNGARCH
M<-0.05	-	-	0.26	0.16	0.30	0.15	0.18	0.15	0.20	0.15
	-	-	(0.036)	(0.043)	(0.029)	(0.013)	(0.009)	(0.005)	(0.009)	(0.005)
-0.05<M<-0.04	0.42	0.12	0.36	0.17	0.21	0.12	0.14	0.14	0.19	0.14
	(0.138)	(0.047)	(0.050)	(0.027)	(0.025)	(0.012)	(0.013)	(0.008)	(0.014)	(0.007)
-0.04<M<-0.03	0.37	0.16	0.24	0.13	0.17	0.12	0.11	0.13	0.17	0.13
	(0.086)	(0.053)	(0.030)	(0.013)	(0.017)	(0.011)	(0.009)	(0.008)	(0.010)	(0.006)
-0.03<M<-0.02	0.37	0.19	0.20	0.11	0.13	0.11	0.09	0.12	0.15	0.12
	(0.054)	(0.035)	(0.019)	(0.008)	(0.011)	(0.008)	(0.007)	(0.007)	(0.009)	(0.005)
-0.02<M<-0.01	0.23	0.14	0.15	0.10	0.09	0.09	0.07	0.11	0.12	0.11
	(0.029)	(0.015)	(0.012)	(0.007)	(0.007)	(0.006)	(0.005)	(0.007)	(0.006)	(0.004)
-0.01<M<0.01	0.12	0.08	0.08	0.07	0.06	0.08	0.06	0.10	0.08	0.08
	(0.009)	(0.006)	(0.005)	(0.004)	(0.003)	(0.004)	(0.003)	(0.004)	(0.003)	(0.002)
M>0.01	0.02	0.03	0.02	0.03	0.04	0.04	0.04	0.05	0.03	0.04
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)
Total	0.06	0.05	0.06	0.05	0.07	0.06	0.07	0.08	0.07	0.06
	(0.004)	(0.002)	(0.003)	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)

Table 6:
Out-of-Sample Performance of the Full NGARCH-Jump model

The parameter values for the NGARCH-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 NGARCH-Jump option prices are updated solely based on the path followed by the stock index. The top panel shows the mean (and standard error) 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.

	Moneyness							
	M<-0.05	**-0.05<M<-0.04**	**-0.04<M<-0.03**	**-0.03<M<-0.02**	**-0.02<M<-0.01**	**-0.01<M<0.01**	**M>0.01**	**Total**
In Sample	0.15 (0.006)	0.15 (0.009)	0.14 (0.008)	0.14 (0.008)	0.13 (0.006)	0.10 (0.003)	0.05 (0.001)	0.08 (0.001)
1-10 weeks	0.17 (0.020)	0.10 (0.018)	0.10 (0.018)	0.09 (0.012)	0.08 (0.008)	0.05 (0.005)	0.04 (0.001)	0.05 (0.002)
10-20 weeks	0.12 (0.019)	0.09 (0.022)	0.11 (0.015)	0.08 (0.009)	0.07 (0.010)	0.05 (0.005)	0.03 (0.001)	0.04 (0.002)
20-30 weeks	0.12 (0.014)	0.12 (0.019)	0.10 (0.014)	0.10 (0.012)	0.08 (0.010)	0.06 (0.004)	0.03 (0.001)	0.05 (0.002)
30-40 weeks	0.20 (0.020)	0.13 (0.020)	0.12 (0.015)	0.12 (0.011)	0.11 (0.009)	0.10 (0.004)	0.03 (0.001)	0.06 (0.002)
40-50 weeks	0.14 (0.013)	0.17 (0.043)	0.13 (0.038)	0.14 (0.040)	0.12 (0.019)	0.09 (0.006)	0.03 (0.001)	0.04 (0.002)
Total	0.15 (0.005)	0.14 (0.007)	0.13 (0.006)	0.12 (0.005)	0.11 (0.004)	0.08 (0.002)	0.04 (0.000)	0.06 (0.001)

Figure 1:
Box Plots of NGARCH and TGARCH Option Pricing Errors

The figures show box plots of percentage errors for six models, plotted against moneyness for the four maturity buckets. In all cases, the ordering of the lines in each moneyness category is NGARCH and TGARCH-Normal models, followed by the Restricted NGARCH and TGARCH-Jump models, followed by the unrestricted NGARCH-Jump and TGARCH-Jump models . The moneyness categories are 1=(<-0.05), 2 = (-0.05,-0.04), 3=(-0.04,-0.03), 4=(-0.03,-0.02), 5= (-0.02,-0.01) and 6 = (-0.01,0.01). All residuals are out-of sample residuals. The parameters are estimated up to day 750, and the residuals are computed based on updating the path of the index over the subsequent 50 weeks.

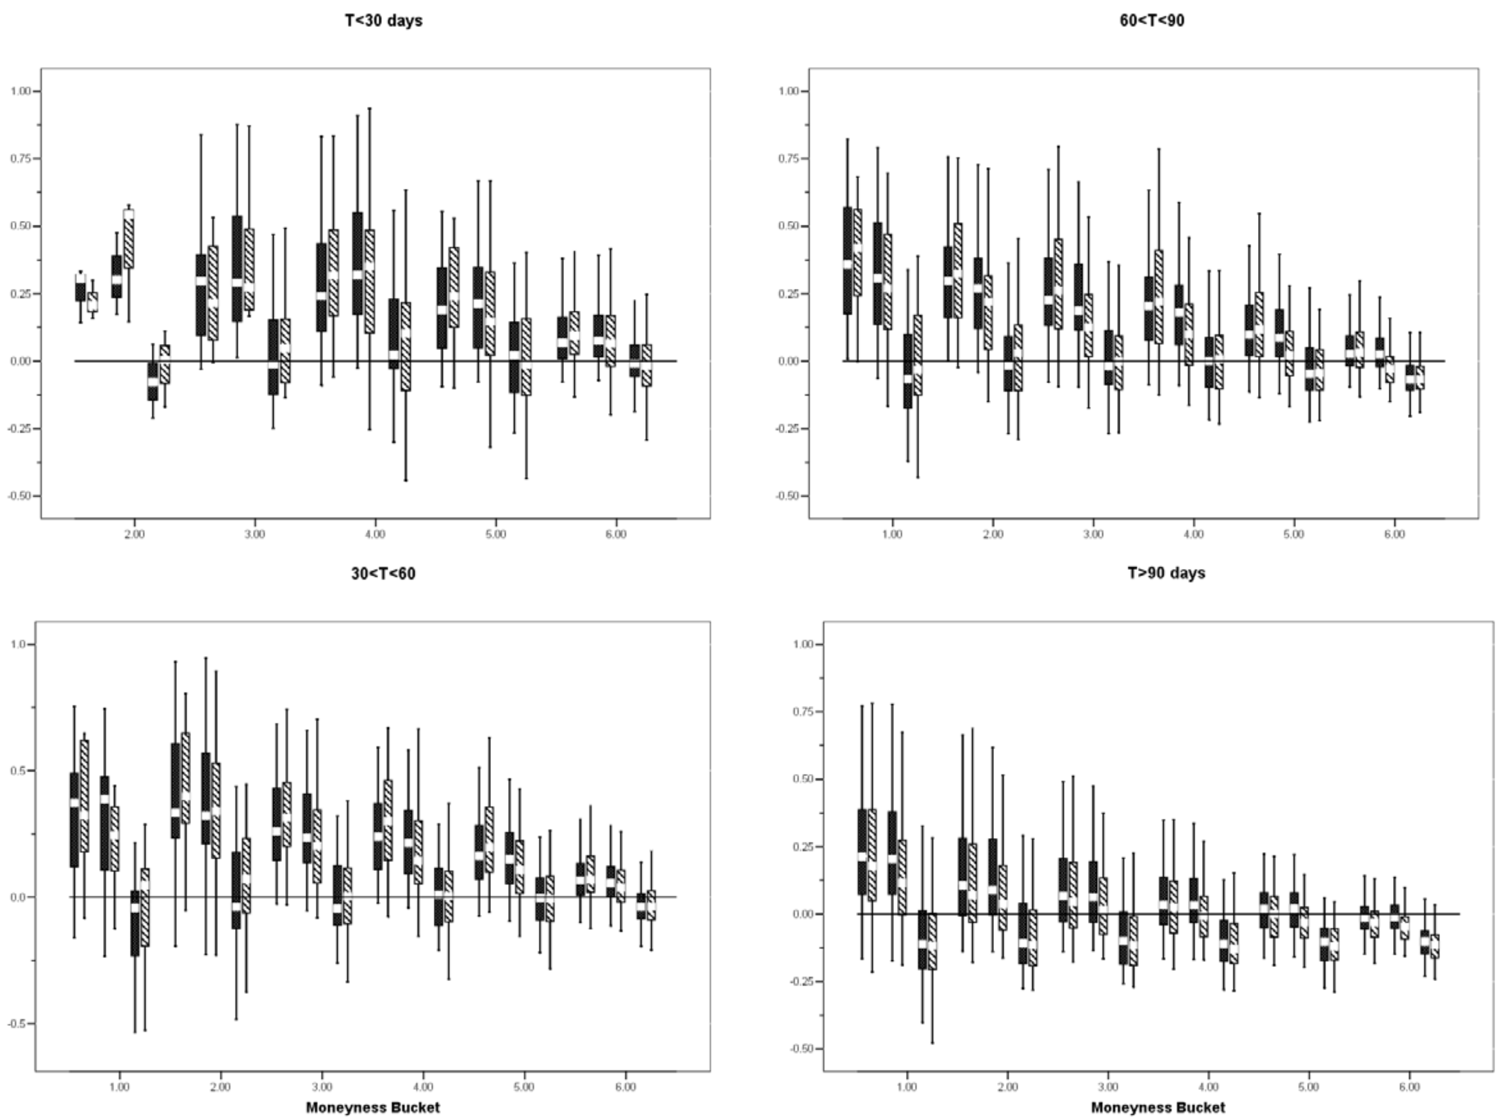

Figure 2:
Box Plots of Option Pricing Errors

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.

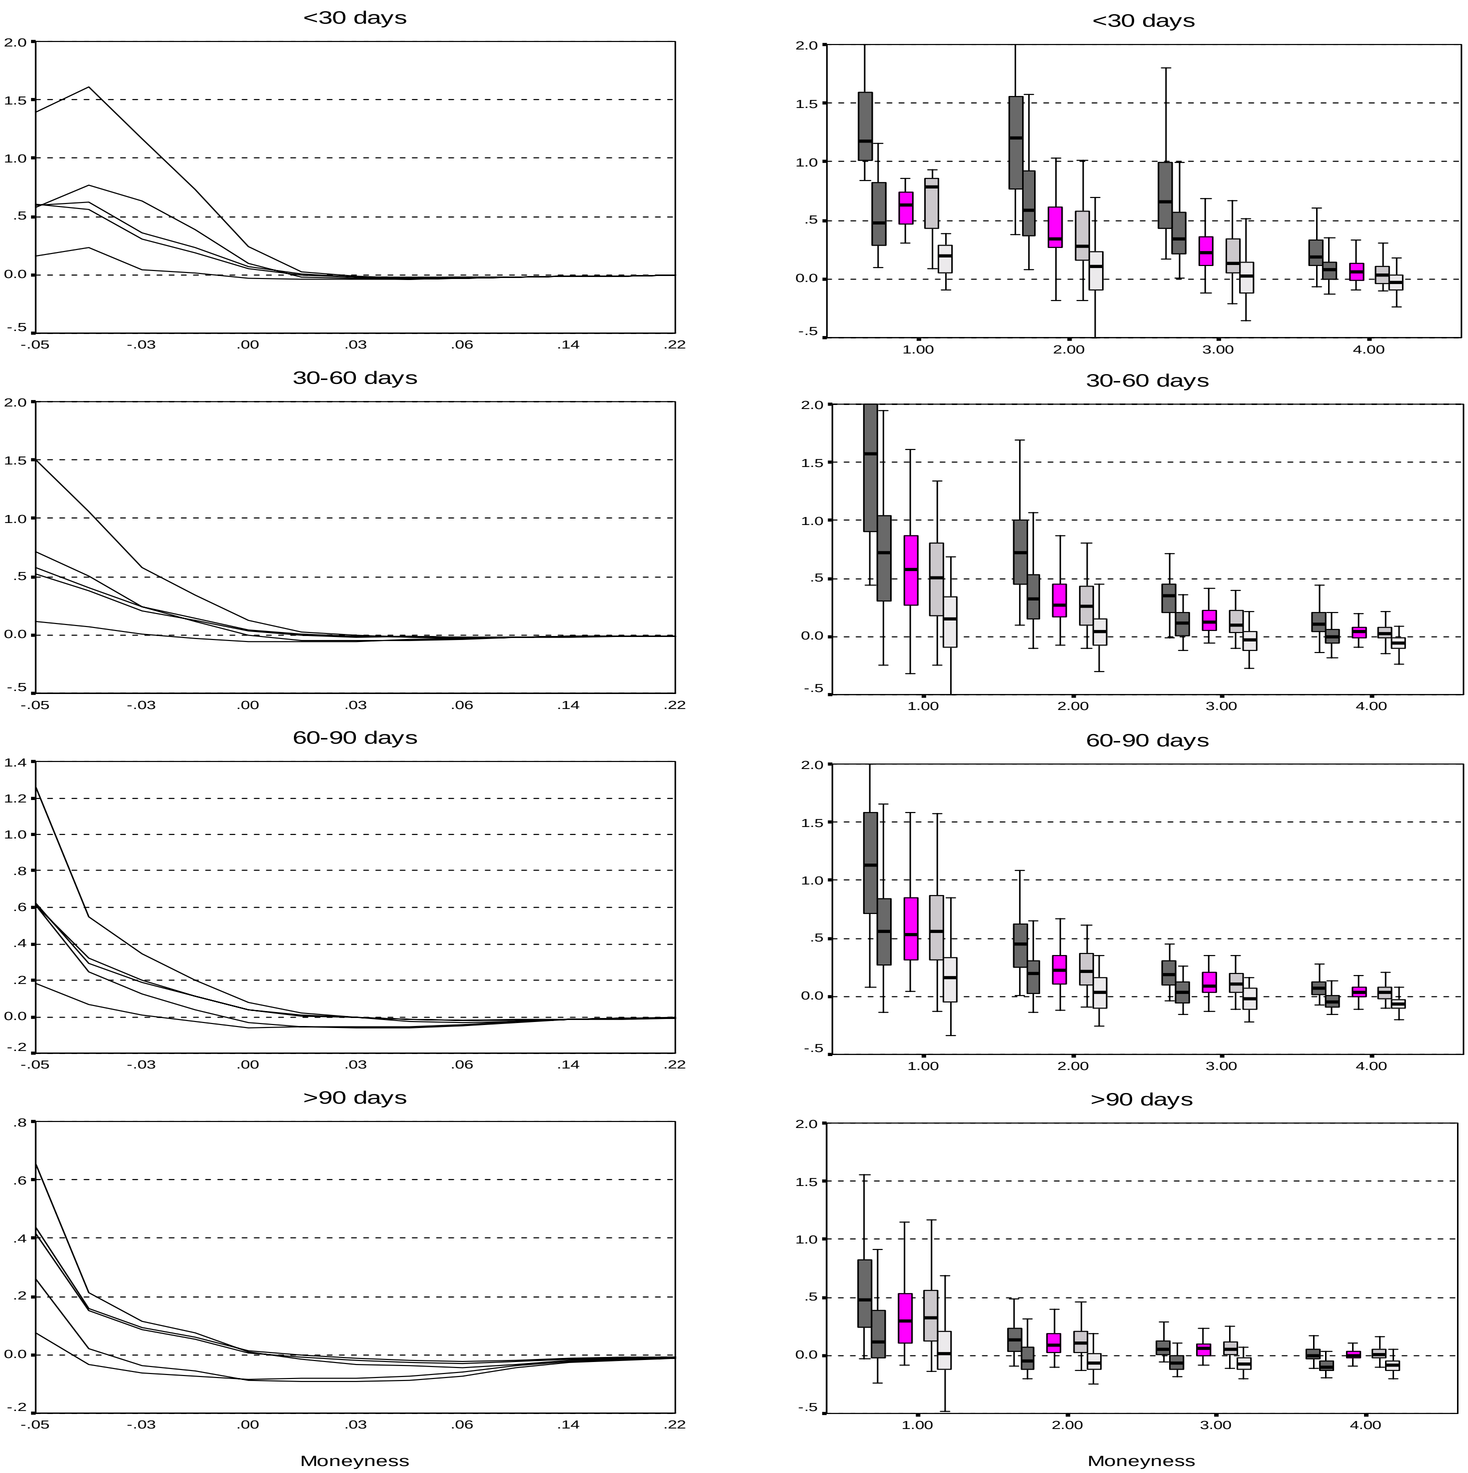

Figure 3:
Average Pricing Errors

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 G-MERTON 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 full NGARCH-Jump model.

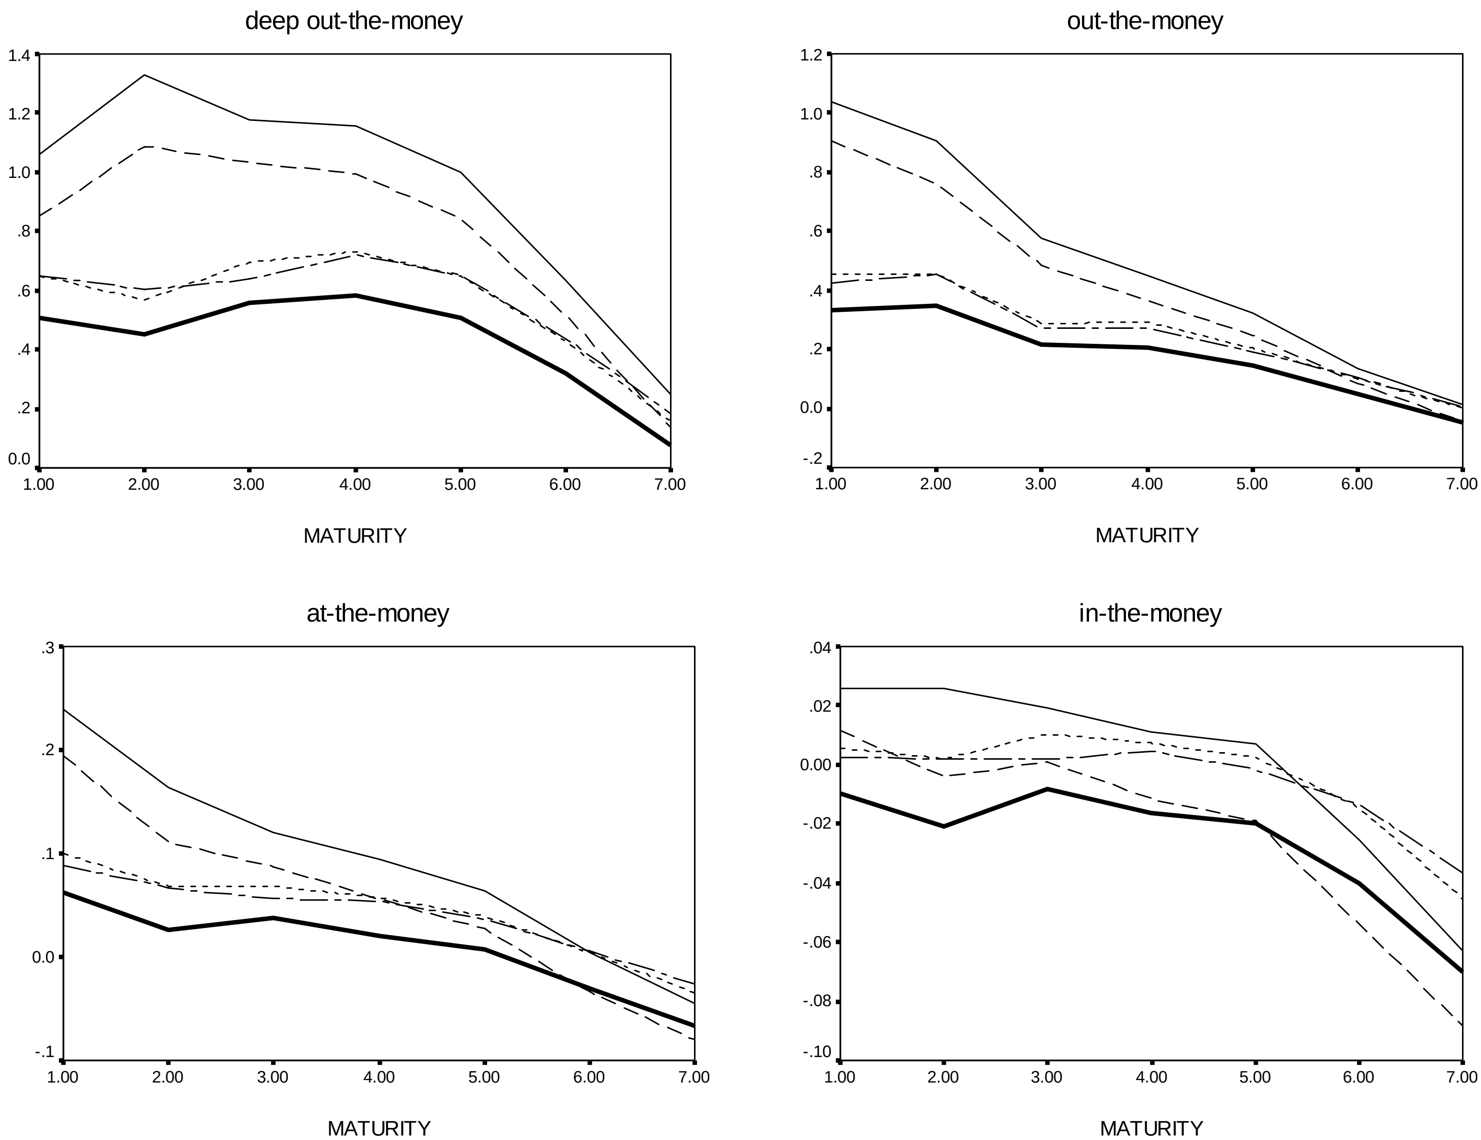

Figure 4:
Box Plots of Net Profits of Delta Hedging over 15 days

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 leftmost 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 ad-hoc 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.