Jump Starting GARCH:

Pricing Options with Jumps in Returns and Volatilities

J. Duan [*] P. Ritchken [†] Z. Sun [‡]

November 30, 2007

[*]Risk Management Institute and Department of Finance, National University of Singapore and Rotman School of Management, University of Toronto. Email: bizdjc@nus.edu.sg.

[†]Weatherhead School of Management, Case Western Reserve University, Cleveland, Ohio, 44106. Email: peter.ritchken@case.edu.

[‡]Fifth Third Asset Management,Cleveland, Ohio, 44114. Email: zhiqiang.sun@53.com.

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

ABSTRACT

This paper considers the pricing of options when there are jumps in the pricing kernel and correlated jumps in asset returns and volatilities. Our model nests Duan's GARCH option models where conditional returns are constrained to being normal, as well as extends Merton's jump-diffusion model by allowing return volatility to exhibit GARCH-like behavior. Empirical analysis on the S&P 500 index returns reveals that the incorporation of jumps in returns and volatilities improves significantly the performance of the GARCH model in capturing the observed time series of the S&P 500 index returns. Moreover, the corresponding GARCH option pricing model with the jump component delivers a better performance on pricing the S&P 500 index options.

(*GARCH, options, stochastic volatility, jumps*)

In this paper we introduce a new family of GARCH models driven by compound Poisson innovations and derive the corresponding option pricing theory. Because the compound Poisson innovations are analogous to the increments in a continuous-time compound Poisson process, we refer to this new class as the GARCH-Jump model. 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 priced in a way to reflect changing volatility and jumps in both returns and volatilities. This GARCH-Jump option pricing model is a natural generalization of the typical GARCH option pricing models with normal innovations, a pricing approach originated in Duan (1995). We empirically test the model, and show that it fits the return data better than the traditional GARCH model with normal innovations and outperforms the inverse Gaussian GARCH model recently proposed by Christoffersen, Heston and Jacobs (2006) (CHJ). Moreover, our model is better in removing more of the biases in option prices.

Just as the binomial model serves as a discrete-time approximation for many underlying diffusion processes, the class of GARCH-Jump models can serve as discrete-time approximations for an array of continuous-time jump diffusion models. As shown in Duan, Ritchken and Sun (2006), a variety of continuous-time limiting models can in fact be derived using our GARCH-Jump processes; for example, (1) when the GARCH feature is disabled but jumps allowed, the limiting model nests the jump-diffusion model of Merton (1976), (2) when jumps are suppressed, 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, and (3) when jumps are permitted, 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).

Furthermore, just as the appropriately defined binomial model provides a useful mechanism for pricing American style options under the geometric Brownian motion assumption, our appropriately defined risk neutralized discrete-time GARCH models provide a mechanism for pricing options when returns and/or volatilities experience random jumps. The option theoretical results developed in this paper has in fact been utilized by Duan, Ritchken and Sun (2006) to derive the limiting option pricing models.

This paper contributes to the literature in three aspects. First, we propose a new class of GARCH models based on compound Poisson innovations and establish the discrete-time option pricing theory which allows us to price derivatives when the underlying asset's innovations may be far from normal and when volatility is stochastic. This is important because our approach offers a unique GARCH option model with non-normal innovations that can be naturally linked to stochastic volatility model with jumps.[1] Second, we conduct an empirical analysis to demon-

[1]We know of three alternative ways of introducing non-normal innovations into the GARCH option pricing

strate the importance of incorporating jumps in returns and volatilities so as to better capture kurtosis and skewness in the time series return dynamics. Third, we employ a more comprehensive approach to estimating parameters and comparing the performance of different option models, that incorporates both the time series of asset returns as well as the cross section of option prices.

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) argue for volatility-jump models because implied volatilities move too abruptly for a diffusion.[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 recommended 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 showed 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, 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), Duan and Zhang (2001), Lehar, Scheicher and Schittenkopf (2002), Lehnert (2003), Stentoft (2005) and Hsieh and Ritchken (2005), are some examples. Christoffersen and Jacobs (2004) examined a set of GARCH option models using the more general

model. Duan (1999,2002) developed two versions of the GARCH option model allowing for conditional skewness and kurtosis via a normal transformation technique and the entropy principle, respectively. Christoffersen, Heston, and Jacobs (2006) 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) provided empirical tests of alternative option models, none of which contain jumps in volatility. Naik (1993) considered 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) provided empirical support for some stochastic volatility models with jumps in returns and volatility. For alternative models see Alexander (2004), Brigo and Mercurio (2002), Brigo, Mercurio and Rapisarda (2004), Carr and Wu (2003), and Madan, Carr and Chang (1998).

GARCH specification 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 that model is also considered where jump risk is priced. At the other extreme we consider models that contain no jumps but allows volatility to be time varying. Finally, we consider models where jump and diffusive risks are priced and whose continuous-time limits contain jumps in both returns and volatilities.

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 and option pricing issues in the discrete-time GARCH-Jump framework. In section 3 we examine our nested GARCH-Jump models and present empirical evidence from time series of the $S\&P$ 500 index. In section 4 we employ an estimation method that combines time series data with cross sectional data on option prices. We investigate the in-sample fits of our nested models and compare them with the inverse Gaussian GARCH option model by Christoffersen, Heston and Jacobs (2006) that also employs non-normal innovations. Finally, we investigate how the option pricing models perform when the analysis is conducted using data in the out-of-sample period of up to 5 years after the model parameters have been estimated. 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 $(\mathcal{F}_t; t \in \{0, 1, \cdots, T\})$ 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:

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

where a_t is the conditional mean growth rate of the pricing kernel and b is a scaling parameter to adjust the volatility of the pricing kernel; and 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

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

and N_t is distributed as a Poisson random variable with parameter λ_t, which may be stochastic but is known at time $t-1$ (i.e., $\mathcal{F}_{t-1}$-measurable).[3] The random variables $X_t^{(j)}$ are independent for $j = 0, 1, 2, \cdots$ and $t = 1, 2, \cdots, T$. Hence:

$$\begin{aligned} E^{\mathsf{P}}\left[J_t | \mathcal{F}_{t-1}\right] &= \lambda_t \mu \\ Var^{\mathsf{P}}\left[J_t | \mathcal{F}_{t-1}\right] &= 1 + \lambda_t(\mu^2 + \gamma^2). \end{aligned}$$

Let r_t denote the single period (from $t-1$ to t) continuously compounded risk-free interest rate. Equilibrium implies that

$$E^{\mathsf{P}}\left[\frac{m_t}{m_{t-1}}\middle| \mathcal{F}_{t-1}\right] = e^{-r_t} \tag{4}$$

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

$$E^{\mathsf{P}}\left[\frac{m_t}{m_{t-1}} | \mathcal{F}_{t-1}\right] = e^{a_t + b^2/2 + \lambda_t(\kappa - 1)}, \tag{5}$$

where

$$\kappa = e^{\mu + b^2\gamma^2/2}.$$

Combining equations (12) and (5), we have:

$$r_t = -\left(a_t + b^2/2 + \lambda_t(\kappa - 1)\right). \tag{6}$$

[3] Maheu and McCurdy (2004), building on a model by Bates and Craine (1999), offered one interesting specification for time-varying λ_t.

Notice that the pricing kernel has local variance,

$$Var^{\mathsf{P}}\left[\ln\left(\frac{m_t}{m_{t-1}}\right)|\mathcal{F}_{t-1}\right] = b^2\left[1+\lambda_t(\mu^2+\gamma^2)\right], \tag{7}$$

and is time-varying as long as the jump intensity is time-varying. For the special case when $\kappa = 1$, or equivalently when $\mu = -b\gamma^2/2$, the effects of the jump in the pricing kernel play no role on the interest rate. For all other values, the jump process explicitly affects both the interest rate and asset price.

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

$$\frac{S_t}{S_{t-1}} = e^{\alpha_t+\sqrt{h_t}\bar{J}_t} \tag{8}$$

where $\bar{J}_t$ is a standard normal random variable plus a Poisson random sum of normal random variables; α_t is part of the conditional mean return to be determined later in Proposition 1; and h_t is a local scaling variable with its precise definition given later. In particular:

$$\bar{J}_t = \bar{X}_t^{(0)} + \sum_{j=1}^{N_t}\bar{X}_t^{(j)} \tag{9}$$

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

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

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.

Since:

$$\begin{aligned}E^{\mathsf{P}}\left[\bar{J}_t|\mathcal{F}_{t-1}\right] &= \lambda_t\bar{\mu}\\ Var^{\mathsf{P}}\left[\bar{J}_t|\mathcal{F}_{t-1}\right] &= 1+\lambda_t(\bar{\mu}^2+\bar{\gamma}^2),\end{aligned}$$

the local variance of the logarithmic returns for date t, viewed from date $t-1$ is

$$Var^{\mathsf{P}}\left[\ln\left(\frac{S_t}{S_{t-1}}\right)|\mathcal{F}_{t-1}\right] = h_t(1+\lambda_t\hat{\gamma}^2), \tag{10}$$

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 factor. In general, the local scaling factor h_t can be any predictable process. For example, it could depend on all previous scaling factors and shocks. That is:

$$h_t = F(h_{t-i}, \bar{J}_{t-i}; i = 1, 2, \cdots) \tag{11}$$

Equilibrium implies that

$$E^{\mathsf{P}}\left[\frac{m_t}{m_{t-1}}|\mathcal{F}_{t-1}\right] = e^{-r_t} \tag{12}$$

$$E^{\mathsf{P}}\left[\frac{m_t}{m_{t-1}}\frac{S_t}{S_{t-1}}|\mathcal{F}_{t-1}\right] = 1 \tag{13}$$

These 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:*

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

where

$$\alpha_t = r_t - \frac{h_t}{2} - \sqrt{h_t}b\rho + \lambda_t\kappa\left(1 - K_t\right) \tag{15}$$

$$h_t = F(h_{t-i}, \bar{J}_{t-i}; i = 1, 2, \cdots) \tag{16}$$

$$K_t = \exp\left(\sqrt{h_t}(\bar{\mu} + b\rho\gamma\bar{\gamma}) + \frac{1}{2}h_t\bar{\gamma}^2\right). \tag{17}$$

Proof: See Appendix

1.1 Pricing Derivatives

It is both customary and arguably more desirable to price derivative claims using 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

$$d\mathsf{Q} = \exp\left(\sum_{t=1}^{T} r_t\right)\frac{m_T}{m_0}d\mathsf{P}. \tag{18}$$

Lemma 1

(i) Q *is a probability measure.*

(ii) For any $\mathcal{F}_t$ *measurable contingent claim* Z_t*, its time-(t-1) price is*

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

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:*

$$\frac{S_t}{S_{t-1}} = e^{\tilde{\alpha}_t + \sqrt{h_t}\tilde{J}_t} \tag{19}$$

where

$$\tilde{\alpha}_t = r_t - \frac{h_t}{2} + \tilde{\lambda}_t\left(1 - K_t\right) \tag{20}$$

$$h_t = F(h_{t-i}, \tilde{J}_{t-i} + b\rho; i = 1, 2, \cdots) \tag{21}$$

$$\tilde{J}_t = \tilde{X}_t^{(0)} + \sum_{j=1}^{\tilde{N}_t} \tilde{X}_t^{(j)} \tag{22}$$

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

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

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

$\tilde{N}_t$ *has a Poisson distribution with parameter* $\tilde{\lambda}_t \equiv \lambda_t \kappa$ *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 not equal to the local variance

under the original P measure unless $\kappa = 1$ and $b\rho\gamma = 0$. To see this, note that 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:

$$E^{\mathsf{Q}}(\tilde{J}_t|\mathcal{F}_{t-1}) = \tilde{\lambda}_t(\bar{\mu} + b\rho\gamma\bar{\gamma}) \tag{23}$$

$$Var^{\mathsf{Q}}(\tilde{J}_t|\mathcal{F}_{t-1}) = 1 + \tilde{\lambda}_t\tilde{\gamma}^2, \tag{24}$$

where

$$\tilde{\gamma}^2 = (\bar{\mu} + b\rho\gamma\bar{\gamma})^2 + \bar{\gamma}^2.$$

Hence the local variance of the innovation under measure Q is $h_t(1 + \tilde{\lambda}_t\tilde{\gamma}^2)$, which differs from the local variance of the innovation under measure P.[4] In other words, one should not in general expect the local risk-neutral valuation principle to apply.

Our theoretical results in Propositions 1 and 2 have already been utilized by Duan, Ritchken and Sun (2006) to construct approximating GARCH-jump models with the purpose of deriving limiting forms by shrinking the time interval. They show that the GARCH-jump model with different GARCH specifications can be made to converge to continuous-time models with diffusive elements and jumps in both returns and volatilities. The limiting model can be more general than that in Bakshi, Cao and Chen (1997), Bates (2000) or Pan (2002), for it allows for volatility jumps as well. By turning off jumps, the limiting model nests the square root stochastic volatility model given in Scott (1987) and Heston (1993).

For empirical work it is necessary to select specific structures for the scaling factor's dynamic in equation (11). We adopt the NGARCH(1,1) dynamic for our empirical analysis. Notice that when $\lambda_t = 0$, 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.

Our NGARCH(1,1) model with a compound Poisson innovation is of the form:

$$h_t = \beta_0 + \beta_1 h_{t-1} + \beta_2 h_{t-1}\left(\frac{\bar{J}_{t-1} - \lambda_{t-1}\bar{\mu}}{\sqrt{1 + \lambda_{t-1}\hat{\gamma}^2}} - c\right)^2, \tag{25}$$

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).

[4] 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 that of Duan (1995).

Using equations (23), (24) and (21), the updating scheme for the local scaling factor, h_t, specialized to equation (25), under measure Q, can be written as

$$h_t = \beta_0 + \beta_1 h_{t-1} + \beta^*_{2,t-1} h_{t-1} \left(\frac{\tilde{J}_{t-1} - \tilde{\lambda}_{t-1}(\bar{\mu} + b\rho\gamma\bar{\gamma})}{\sqrt{1 + \tilde{\lambda}_{t-1}\tilde{\gamma}^2}} - c^*_{t-1} \right)^2 \tag{26}$$

where

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

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

In summary, when the local scaling factor h_t follows a NGARCH process, then under measure Q, the updating scheme translates into a similar NGARCH process. A similar result applies to different GARCH specifications because Proposition 2 allows us to easily come to specific pricing system corresponding to different volatility dynamics. The specific result above allows us to examine whether extending the NGARCH-Normal model to the NGARCH-Jump model will further reduce the option pricing biases significantly.

1.2 The Nested Models

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

$$E^{\mathsf{P}} \left(\frac{S_t}{S_{t-1}} | \mathcal{F}_{t-1} \right) = e^{(r_t + \eta_t)}$$

where the risk premium η_t is given by:

$$\eta_t = \lambda_t \kappa (1 - K_t) - \lambda_t (1 - e^{\bar{\mu}\sqrt{h_t} + \frac{\bar{\gamma}^2 h_t}{2}}) - \sqrt{h_t} b\rho \tag{27}$$

Recall, from equation (17) that:

$$\begin{aligned} 1 - K_t &= 1 - \exp(\sqrt{h_t}(\bar{\mu} + b\rho\gamma\bar{\gamma}) + \frac{1}{2} h_t \bar{\gamma}^2) \\ &\approx -\sqrt{h_t}(\bar{\mu} + b\rho\gamma\bar{\gamma}) - \frac{1}{2} h_t ((\bar{\mu} + b\rho\gamma\bar{\gamma})^2 + \bar{\gamma}^2). \end{aligned} \tag{28}$$

where the approximation is justified if h_t is small.[5] Substituting this approximation into the risk premium equation, we obtain:

$$\eta_t \approx [\lambda_t \bar{\mu}(1 - \kappa) - b\rho(1 + \lambda_t \kappa \gamma\bar{\gamma})] \sqrt{h_t} + \lambda_t \bar{\gamma}^2 (1 - \kappa) \frac{h_t}{2}. \tag{29}$$

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

This form of the risk premia allows us to gain some insight into the pricing model.

(i) The Merton Model

First consider the case when $\kappa = 1$ and $\gamma = 0$. In this case, $\mu = 0$ and there are no jumps in the pricing kernel. The risk premium, η_t reduces to $-b\rho\sqrt{h_t}$. That is, the risk premium does not depend on jumps. With $\beta_1 = \beta_2 = 0$ in equation (25) 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 (1976) model, or Merton, for short.

(ii) The Generalized Merton Model

Second, consider the same model, but release κ and γ from 1 and 0. 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 $\kappa \neq 1$ and $\gamma = 0$ in the pricing kernel, the sensitivity of the risk premium to $\bar{\gamma}$ is very small. That is, the randomness about the jump size adds minimally to the risk premium.

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

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

Here, the uncertainty of the jump size, as measured by $\bar{\gamma}$, adds to the risk premium as does the intensity. This implies that jump risk is priced. With $\beta_1 = \beta_2 = 0$ in equation (25) the scaling factor remains constant. We call this model the generalized Merton model.

(iii) The NGARCH-Normal Model

The third model we consider has no jumps, i.e., $\lambda_t = 0$, but with our scaling factor being stochastic. In this case, innovations are normal random variables, and the risk premium is given by $\eta_t = -b\rho\sqrt{h_t}$. The system is referred to as the NGARCH-Normal model, or the NGARCH model for short.

(iv) The Restricted NGARCH-Jump Model

The fourth model keeps $\kappa = 1$ and $\gamma = 0$ again, but the scaling factor is permitted to be stochastic and jumps in prices are allowed. In this model, jump risk is diversifiable, volatility is stochastic and innovations are non-normal. The model is referred to as the Restricted NGARCH-Jump model, or restricted JGARCH for short.

(v) The NGARCH-Jump Model

The final model is the most general model where jump risk is priced, scaling factor is stochastic, jumps are present and innovations are not normal. The model is referred to as the NGARCH-Jump model, or JGARCH for short.

We will explore in the next section which of the models nested in our family can 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.

1.3 The Inverse-Gaussian GARCH Option Model

Christoffersen, Heston and Jacobs (2006) (hereafter CHJ) proposed a discrete-time GARCH option pricing model that has non-normal innovations. We will compare the performance of our GARCH-Jump model with their model, which we refer to as the IG-GARCH model. Under the data generating measure, their model has the form

$$\frac{S_t}{S_{t-1}} = e^{r+vh_t+\eta\epsilon_t} \tag{30}$$

$$h_t = \beta_0 + \beta_1 h_{t-1} + \beta_2 \epsilon_{t-1} + \beta_3 \frac{h_{t-1}^2}{\epsilon_{t-1}} \tag{31}$$

where ϵ_t has an inverse-Gaussian distribution with degrees of freedom $\delta_t = h_t/\eta^2$. That is

$$f_{IG}(\epsilon_t;\delta_t) = \frac{\delta_t}{\sqrt{2\pi\epsilon_t^3}} e^{-(\sqrt{\epsilon_t}-\delta_t/\sqrt{\epsilon_t})^2/2} \times 1_{\{\epsilon_t>0\}} \tag{32}$$

If $\eta < 0$, then the conditional distribution has negative skewness. Note that the physical system has six parameters, i.e., $\{\beta_0, \beta_1, \beta_2, \beta_3, v, \eta\}$.

CHJ assumed a pricing kernel under which options can be priced using a risk-neutral dynamic that continues to be an IG-GARCH process but with different parameters. In particular,

$$\frac{S_t}{S_{t-1}} = e^{r+v^*h_t^*+\eta^*\epsilon_t^*} \tag{33}$$

$$h_t^* = \beta_0^* + \beta_1 h_{t-1}^* + \beta_2^* \epsilon_{t-1}^* + \beta_3^* \frac{h_{t-1}^{*2}}{\epsilon_{t-1}^*} \tag{34}$$

where ϵ_t^* has an inverse-Gaussian distribution with parameter, $\delta_t^* = h_t^*/\eta^{*2}$, and $v^* = v(\eta^*/\eta)^{-3/2}$, $\epsilon_{t+1}^* = \epsilon_{t+1}(\eta^*/\eta)^{-1}$, $\beta_0^* = \beta_0(\eta^*/\eta)^{-3/2}$, $\beta_2^* = \beta_2(\eta^*/\eta)^{3/2}$, $\beta_3^* = \beta_3(\eta^*/\eta)^{-5/2}$. In addition, since the true drift parameter is eliminated from the risk neutral pricing condition, the following constraint must be imposed:

$$v^* = \frac{1}{\eta^{*2}}(\sqrt{1-2\eta^*} - 1).$$

As compared to the physical system, the resulting risk neutral system has five free parameters, i.e., $\{\beta_0^*, \beta_1^*, \beta_2^*, \beta_3^*, \eta^*\}$. The overall system (physical and risk-neutral) has seven parameters, however, because the risk-neutral system has in effect introduced one extra parameter η^*. This is evident from the above relationships linking the five risk-neutral parameters to the six physical parameters. In short, in the empirical study of the IG-GARCH model using the combined data of returns and option prices, the relevant parameter set is $\{\beta_0, \beta_1, \beta_2, \beta_3, v, \eta, \eta^*\}$.

2 Data

We examine the performance of the above models using time series data on the S&P 500 index and dividends, as well as option price information on the S&P 500 index.

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 options, spans the time period from January 1996 to June 2005. The data comes from Ivy OptionMetrics database which is a comprehensive database covering US index and equity options markets. From this database we extract daily highest closing bid and lowest closing ask prices across all exchanges for each option contract on the S&P 500 index. In addition we record the time to expiration, the strike price, and the closing price of the S&P 500 index. We only consider option contracts that have a time to expiration greater than 10 days and less than 120 days. We also exclude those option contracts that are so far away from the money that there are liquidity concerns. In particular we exclude an option when the midpoint of the bid-ask price is below intrinsic value, when the vega of the option is below 0.5, and when the implied volatility calculation fails to converge. We also record the interest rate information on each day, and the dividend yield that OptionMetrics uses to compute the Black-Scholes implied volatilities.

The interest rates that are available from OptionMetrics correspond to the continuously compounded zero-rates derived from LIBOR rates and settlement prices of CME Eurodollar futures. For any given option, the appropriate interest rate corresponds to the zero rate that has maturity equal to the option's expiration, and is obtained by interpolating between the two closest zero rates on the zero curve.

In order to price the options we need to adjust the index level according to the dividends paid out over the time to expiration. Harvey and Whaley (1992), and Bakshi, Cao and Chen (1997), used 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 was 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 when the option quotes were obtained. There are other methods for establishing the dividend adjusted index level. The first 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 with the same maturity. The second 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. OptionMetrics uses a regression approach that exploits put-call parity conditions repeatedly over a number of contracts and over a ten-day period to extract an implied dividend yield that they apply to all option contracts. We use their extracted

dividend yield in all our computations. For a discussion of different approaches see Jackwerth and Rubinstein (1996).

We extract option prices on a weekly basis, each Wednesday, over the 9.5-year period (January 1996 to June 2005). For the underlying S&P 500 index we obtain the time series of daily returns going back to January 1970. We have 43,377 option contracts that pass all the filters, with the number of contracts each year ranging from a low of 2,306 for the six month period in 2005 to a high of 5,069 contracts in 2002.

We split up our data set into two, an "in-sample" period extending over the first 5 years, and an "out-of-sample" period covering the remaining 5 years.

We define moneyness as the index price relative to the strike price. We construct histograms of moneyness and days to expiration. The top graph of Figure 1 shows the distribution of moneyness over the in-sample period, and the bottom figure shows the distribution of expiration dates. The mode of the distribution is at the short maturities. Over 90% of maturities are within 90 days. We construct buckets of maturities, defining bin 1 as those contracts that expire within 30 days; bin 2, those that expire in the 31-60 day range; bin 3, those that expire in the 61-90 day range and bin 4 are those contracts that have longer expiration dates.

Figure 1 Here

Table 1 shows the number of option contracts by moneyness and by maturity and the average implied volatilities for each of the moneyness-maturity bins. We see from the table that the largest variation in volatilities across strike prices occurs for the shorter-term option contracts.

Table 1 Here

The top graph in Figure 2 plots the average monthly volatility smile by moneyness over the ten-year time period for all contracts less than 30 days. The exact shape of the volatility smile fluctuates over time. From this figure we see that deep in-the-money calls generally have the largest volatilities. The bottom graph in Figure 2 shows the time series behavior of at-the-money implied volatilities over the ten-year period. As can be seen volatilities have fluctuated from 10% to over 30%.

Figure 2 Here

3 Estimation of Models Using Time Series of Returns and/or Option Prices

Our first set of experiments are concerned with using the time series return data on the S&P 500 index alone to compare the performance of some of the models nested in the GARCH-Jump family, and to contrast these models with the IG-GARCH model.

We apply the maximum-likelihood method on the time series of index returns to estimate the model parameters. Our model parameter set is $\theta = \{\beta_0, \beta_1, \beta_2, c, b\rho, \kappa, \gamma, \bar{\mu}, \bar{\gamma}, \lambda\}$. Define $R_t \equiv \ln(S_t/S_{t-1})$. The conditional probability density function for R_t is:

$$l\left(R_t|h_t\right) = \sum_{i=0}^{\infty} \frac{\lambda^i}{i!} e^{-\lambda} f\left(R_t - \alpha_t; \mu_i(t), \sigma_i^2(t)\right) \tag{35}$$

where α_t is given in Proposition 1; the NGARCH(1,1) local scaling factor, h_t, is by equation (25); and $f\left(\cdot; \mu_i(t), \sigma_i^2(t)\right)$ is the normal density function with mean $\mu_i(t) = i\bar{\mu}\sqrt{h_t}$ and variance $\sigma_i^2(t) = h_t(1 + i\bar{\gamma}^2)$.[6] The initial value of the local scaling factor is set according to

$$h_1 = V/(1 + \lambda\hat{\gamma}^2) \tag{36}$$

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

The log-likelihood function for the sample of asset prices is:

$$L(\theta; S_1, S_2, \cdots, S_T) = \sum_{t=2}^{T} \ln\left[l\left(R_t|h_t\right)\right]. \tag{37}$$

The maximum likelihood estimator for θ is the solution of maximizing the above log-likelihood function. Given the asset price time series, $\{S_t\}_{1 \le t \le T}$, we can compute the log-likelihood function recursively, and solve this optimization problem numerically.

In principle, the entire set of parameters can be identified by only using a return time series. In practice, however, two of them are hard to pin down empirically. To understand this assertion, notice that conditional on α_t the log-likelihood function is fully determined by $\bar{\mu}$, $\bar{\gamma}^2$, λ, and the parameters driving the variance updates – β_0, β_1, β_2, and c. To fully characterize the log-likelihood function, we of course need to know α_t, which in turn requires knowledge of three extra parameters, namely μ, γ and $b\rho$. Recall from equation (28) that: $1 - K_t \approx -\sqrt{h_t}(\bar{\mu} + b\rho\gamma\bar{\gamma}) - \frac{1}{2}h_t((\bar{\mu} + b\rho\gamma\bar{\gamma})^2 + \bar{\gamma}^2)$. Substituting this expression into equation (15) leads

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

[7] Note that $Var^{\mathrm{P}}(\bar{J}_t) = (1 + \lambda\hat{\gamma}^2)$. Thus, we have in essence removed the extra volatility arising from the jump component in setting the initial value of the scaling factor process.

to:

$$\begin{aligned}\alpha_t &\approx r - \frac{1}{2}h_t(1+\lambda\kappa(\bar{\gamma}^2+(\bar{\mu}+b\rho\gamma\bar{\gamma})^2)) - \sqrt{h_t}(b\rho+\lambda\kappa\bar{\mu}+\lambda\kappa b\rho\gamma\bar{\gamma}) \\ &\approx r - \sqrt{h_t}(b\rho+\lambda\kappa\bar{\mu}+\lambda\kappa b\rho\gamma\bar{\gamma}) \qquad (38)\end{aligned}$$

where the second approximate equality comes from the fact that h_t is much smaller than $\sqrt{h_t}$ because $\sqrt{h_t}$ takes on small values already. 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. Since only the sum $(b\rho+\lambda\kappa\bar{\mu}+\lambda\kappa b\rho\gamma\bar{\gamma})$ matters in the log-likelihood function, two of the three parameters – $b\rho$, κ and γ – are indeterminate. In the estimation, we thus introduce a composite parameter $\delta = b\rho+\lambda\kappa\bar{\mu}+\lambda\kappa b\rho\gamma\bar{\gamma}$, and actually estimate the parameters in the set, $\theta^* = \{\beta_0, \beta_1, \beta_2, c, \delta, \bar{\mu}, \bar{\gamma}, \lambda\}$.

Notice that when $\kappa = 1$ and $\gamma = 0$ we obtain the restricted NGARCH-Jump model. This model 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 NGARCH-Normal model can be readily estimated from the time series alone.

For the IG-GARCH model, the parameter set is $\theta^* = \{\beta_0, \beta_1, \beta_2, \beta_3, v, \eta\}$, and maximum likelihood estimation is employed with the conditional density function in equation (32). Recall that $R_t = \ln(S_t/S_{t-1})$. Let $\epsilon_t = \frac{R_t-\mu_t}{\eta}$ where $\mu_t = r + vh_t$. Then if ϵ_t is distributed as inverse Gaussian, the density of R_t is

$$\begin{aligned}f(R_t;\delta_t) &= f_{IG}\left(\frac{R_t-\mu_t}{\eta};\delta_t\right)\left|\frac{\partial\epsilon_t}{\partial R_t}\right| \\ &= \frac{\delta_t}{\sqrt{2\pi(R_t-\mu_t)/\eta)^3}}e^{-\frac{1}{2}\left(\sqrt{(R_t-\mu_t)/\eta}-\frac{\delta_t}{\sqrt{(R_t-\mu_t)/\eta}}\right)^2}\left|\frac{1}{\eta}\right| \\ &= \frac{h_t}{\eta^2\sqrt{2\pi(R_t-\mu_t)/\eta)^3}}e^{-\frac{1}{2}\left(\sqrt{(R_t-\mu_t)/\eta}-\frac{h_t}{\eta^2\sqrt{(R_t-\mu_t)/\eta}}\right)^2}\left|\frac{1}{\eta}\right|.\end{aligned}$$

3.1 Empirical results using the return time series

Table 2 shows the parameter estimates based on the time series of daily S&P 500 index returns starting from 1970 through 2005 and repeated over different subsamples.

Table 2 Here

For the Merton model, the NGARCH model, the restricted NGARCH-Jump model and the IG-GARCH model, the parameter estimates and their standard deviations are reported. They

are followed by the maximum log-likelihood value and the Akaike Information Criterion (AIC) which measures model performance by accounting for goodness of fit and parsimony. A smaller AIC indicates a better performance.

For the two models (Merton and NGARCH-Normal) nested in the restricted NGARCH-Jump model, one can use the likelihood ratio test to evaluate whether there is any incremental value in adding complexity to the Merton model. Under the null hypothesis that the more complex model does not add significantly, two times the difference in the log-likelihood values should distribute as a chi-square distribution with the degrees of freedom equal to the difference in the numbers of parameters of the two models.

The likelihood ratio test based on Table 2 indicates that, at the 1% level, the restricted NGARCH-Jump model is a significant improvement over the Merton model. This suggests that return data exhibits the GARCH feature. Further, the likelihood ratio test reveals that the effect of adding jumps to the NGARCH model leads to significant improvements. Specifically, Table 2 reveals that λ is significantly different from 0, indicating that the incorporation of jumps is significant and the conditional distribution exhibits skewed and heavy-tailed behavior. For the restricted NGARCH-Jump model, the non-linear term c, capturing the so called leverage effect, continues to be significant after adding jumps to the NGARCH model. In addition, the results are similar for the two sub-samples. The results based on the AIC criterion also lead to the same conclusion. Taken together, the restricted NGARCH-Jump model is clearly the dominant model in this nested class.

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 over 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 particularly 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 results in Table 2 reveals that the performance of the IG-GARCH model according to the AIC criterion is markedly worse than either the NGARCH-Normal or restricted NGARCH-Jump model in the whole sample as well as in two sub-samples. The fact that it is even worse than the NGARCH-Normal model is particularly worth noting. Although the IG-GARCH model allows for non-normal conditional innovations, it relies on a GARCH specification that is less compatible with the data than is the NGARCH model. As a result, the NGARCH model restricted to normal conditional innovations can still outperform the IG-GARCH model in terms of the log-likelihood value even before accounting for the compensation factor due to fewer model parameters.

Figure 3 shows the quantile-quantile (Q-Q) plots for the three special cases of the NGARCH-Jump model and for the IG-GARCH model. For this analysis we begin by first transforming returns using the appropriate conditional distribution function corresponding to each model, say $F_{t-1}(R_t;\hat{\theta})$. The resultant sequence of random variables should be approximately independent, identically distributed uniform random variables. Then, applying the inverse standard normal distribution function, these variables are transformed into standard normal random variables. The overall transformation is $\varepsilon_t = \Phi^{-1}(F_{t-1}(R_t;\hat{\theta}))$. If a model is adequate, then the sequence of ε_t should form an i.i.d. sample of standard normal random variables.

Figure 3 Here

The top figures shows the Q-Q plots of the theoretical normalized residuals against the actual normalized residuals for the Merton model and the IG-GARCH model, while the lower panel shows the plots from the NGARCH-Normal (NGARCH) and the restricted NGARCH-Jump (JGARCH) models. The straight lines connect the 25th and 75th percentiles. The scales on these plots are all identical.

The figures show that there are small differences between the IG-GARCH and the NGARCH-Normal models. Both do not fit the tails of the distribution well. The Merton model seems to fit better, although there are still deviations in both tails.[8] The restricted GARCH-Jump model seems to eliminate the bias in the upper tail, but some smaller bias remains in the lower tail.

Panel A of Table 3 provides further descriptive statistics on the normalized residuals for each of the four models. Note that the skewness and kurtosis of the normalized residuals corresponding to the restricted NGARCH-Jump model are closest to the prediction of standard normality, and the number of outliers is the smallest. Formal Kolmogorov-Smirnov tests for normality, reject the null hypothesis for all the models at the 1% level of significance, with the exception of the restricted NGARCH-Jump model. The Merton model, although being rejected, seems to perform better than the other two models. We can thus conclude that the compound Poisson distribution shared by the Merton and GARCH-Jump models is more compatible with return data than the Inverse-Gaussian distribution adopted by CHJ.

Table 3 Here

Panel B of Table 3 reports the Ljung-Box statistics for testing the significance of the autocorrelation among the residuals and squared residuals for all four models. The tests are conducted for the full sample period and for two subsample periods after the stock market crash of 1987.

[8]This result is not particularly surprising because the Q-Q plot mainly focuses on the distributional characteristics rather than on the time series dependency.

The autocorrelation of residuals produced by all models is significant at the 5% level. This significance can be traced back to the very high autocorrelation in the raw returns that existed in the early data period from 1970 to 1987, likely due to non-synchronous trading among the S&P 500 component stocks in this earlier period. In the two subperiods after 1987, all models produced residuals that were not significantly autocorrelated.

As can be seen from the table, the autocorrelation among the squared residuals produced by the Merton model is significant for the entire sample as well as for each of the two subsamples. The two NGARCH models were effective in removing autocorrelation in the squared residuals. In contrast, the IG-GARCH model still yields significantly autocorrelated squared residuals for the whole sample, indicating that the IG-GARCH volatility specification is less compatible with the data.

3.2 Empirical results using time series of returns and option prices

There are several ways in which option pricing models can be evaluated.[9] First, one can estimate all model parameters with the underlying asset price series. The estimated parameters can then be used in the option pricing system (under measure Q) to price all options. This implementation can be likened to estimating historical volatilities from the return data and then evaluating the Black-Scholes formula based on these estimates. This approach is obviously quite demanding on an option pricing model because option data have not been used in the parameter estimation. For the NGARCH-Jump model, as discussed earlier, there is also an econometric complexity associated with only using the return data. Specifically, for this model, some parameters (i.e., $b\rho$, κ and γ) are difficult to identify from the return time series alone. This approach is thus not practically advisable.

The second way of implementation uses option data in estimation. The underlying return series is used to update the local volatility but it does not directly enter into the likelihood function in parameter estimation. This approach is in a way like running a nonlinear regression by treating the underlying asset price as an exogenous variable. In the Black-Scholes framework, this is analogous to obtaining just one implied volatility by pooling the information from many options together. For the GARCH option pricing model, one can obtain the time series of volatilities, corresponding to a set of parameters and conditional on stock price moves, and then periodically price some panels of option contracts. The optimization criterion typically focuses on minimizing sums of squared option pricing errors. This approach clearly overlooks the information embedded in the underlying return time series because such series also provides useful information concerning the basic premise of an option pricing model; that is, the assumed dynamic for the underlying asset price. The implementation of the GARCH option pricing model

[9] For an excellent review of empirical option pricing see Bates (2003).

adopted by Heston and Nandi (2000), Christoffersen and Jacobs (2004), Hsieh and Ritchken (2005) all fall in this category.

The third and better way of implementing an option pricing model is to combine the option and underlying asset prices over the sample period to conduct a joint estimation of the model. Thus, both option and underlying asset prices directly enter into the joint likelihood function describing the overall system. For ease of exposition, we will refer to this method as the *panel estimation* approach which has been adopted in our empirical analysis.

The panel estimation approach can be implemented in many different ways. The part concerning the underlying return time series is rather straightforward because it follows directly from our GARCH-Jump specification. For option prices, one does have a choice, namely using the prices directly, or transforming them to the Black-Scholes implied volatilities. In essence, the choice amounts to opting for a particular specification for the cross-sectional error structure. Since option prices naturally vary according to their strike prices and maturities, implied volatilities offer a cross-sectional data standardization that is intuitively appealing and conveniently tied-to common industry practice of quoting implied volatilities in referring to option contracts. In short, our specific panel estimation approach take a vector of observations at a time point whose first entry is always the underlying asset price and remaining entries are the implied volatilities of options included in our sample. Note that the dimension of this data vector depends on the availability of options at a particular time point, as well as our decision as to which options to include. We derive the joint likelihood function for such kind of data vectors over the sample period and use the joint likelihood function to conduct the maximum likelihood parameter estimation.

Denote the data vector at time by $D_t(m_t)$; that is, $D_t(m_t) = (S_t, IV_{t,1}, \cdots, IV_{t,m_t})'$ where m_t is the number of options at time t being included and $IV_{t,j}$ is the Black-Scholes implied volatility of the j-th option at time t. The joint log-likelihood function thus becomes

$$\begin{aligned} & L(\theta, \omega; D_1(m_1), \cdots, D_T(m_T)) \\ = & \sum_{t=2}^{T} \ln\left[l\left(R_t|h_t\right)\right] - \left(\frac{\ln(2\pi)}{2} + \ln(\omega)\right) \sum_{t=1}^{T} m_t - \frac{1}{2\omega^2} \sum_{t=1}^{T} \sum_{j=1}^{m_t} \left(IV_{t,j} - \hat{IV}_{t,j}(\theta)\right)^2 \end{aligned}$$

where the expression for $l\left(R_t|h_t\right)$ is available in equation (35); and $\hat{IV}_{t,j}(\theta)$ is the Black-Scholes implied volatility computed from the GARCH option price of the j-th option at time t evaluated at parameter θ. The first term on the right-hand side is the log-likelihood associated with the underlying stock price time series. The summation begins from $t = 2$ because the first stock price is only used to anchor the conditional distribution for the second stock price. The second and third terms on the right-hand side is the log-likelihood function of the option data conditional on the underlying stock price series. We have assumed that the errors in terms of implied volatilities share the same magnitude and are independent across options and over time. The

error distribution is assumed to be normal with a standard deviation of ω. Needless to say, this restriction can be relaxed, such as allowing for time and/or cross-sectional dependency, at the expense of introducing more parameters.

For the NGARCH-Jump model, the combined data of asset and option prices can be used to identify the full parameter set governing this option pricing model. In this case, $\theta = \{\beta_0, \beta_1, \beta_2, c, b\rho, \kappa, \gamma, \bar{\mu}, \bar{\gamma}, \lambda\}$. The whole set of parameters to be estimated is thus composed of θ and ω. In the case of the IG-GARCH model, the full parameter set governing this option pricing model is $\theta = \{\beta_0, \beta_1, \beta_2, \beta_3, v, \eta, \eta^*\}$. Again the whole set of parameter to be estimated consists of θ and ω.

We use the time series of daily returns up to the first Wednesday in January 1996 to update our local scaling factor. Then, given the dividend adjusted index at that date, and using the term structure of riskless rates at that date, we numerically compute the theoretical option prices of selected call and put contracts for a given model. Since there are no simple analytical expressions for the NGARCH option pricing models, their prices are generated by Monte Carlo simulation. For a fair comparison, we also use the same stream of random numbers to generate the prices for the IG-GARCH model. We use 5,000 sample paths to generate Monte Carlo option prices. To improve the quality of Monte-Carlo prices, we apply the analytical pricing formula on the last day of the contract because the analytical solution does exist for one-day option contracts. We then compute the implied volatilities using the Black Scholes model, and the resulting values are used in the log-likelihood function. We then use the daily returns to update the local volatility to the next time point on which the next set of option prices is available. The process is repeated and the individual log-likelihood is added to the overall log-likelihood function. We continue this process for all days in the estimation period.

We split our option data into two – the "in-sample" period starting in 1996 and extending through 2000 and the "out-of-sample" period with the remaining data. Our initial parameter estimates are taken from the return time series analysis that uses all information up to 1996. The likelihood function then consists of the time series component together with information from options taken every fourth week up to the end of 2000. As a result we have used 60 panels of cross sectional option vectors, together with the 5 years of daily return time series information in the parameter estimation. The option contracts used in the optimization include all contracts with maturities between 10 and 90 days and moneyness within 5% of the index value.

Once the optimized parameter values are obtained, we use them and the daily time series of the S&P 500 index to compute the full time series of the local scaling factor over both the in- and out-of-sample periods. Then, for each Wednesday, given the index level and the updated scaling factor, together with interest rate and dividend information, we can compute the theoretical implied volatilities and compare to their actual counterparts. We do this for each week in the in-sample period and for each week in the 5 years corresponding to the "out-of-sample" period.

Further, we compute prices for the longer maturity contracts (maturities that exceed 90 days) and moneyness factors that were never used in the optimization. The residuals obtained from different models then form the basis for our comparisons of the models. An option model is viewed positively if the in-sample fits are more precise and less biased, and if, conditional on future up-to-date index values, the "out-of-sample" price predictions are also more precise and less biased.

Before presenting the results, it should be noted that our objective is not to fit the volatility smile precisely at any one date. Were this curve fitting exercise the goal, we would merely choose parameter sets each week that minimize the sum of squared errors (in implied volatilities) for different models. By using 5 years of data to choose one set of parameters, our fits of implied volatilities at any week in the in-sample period will naturally be less precise than models obtained by more frequent recalibrations. Further, in our "out-of-sample" period, we would expect the performance of all models to deteriorate over time, especially when parameter estimates are not updated, say after 5 years. However, our goal here is only to compare the relative performance of these models under identical conditions. Any model that is successful in these tests can then be further scrutinized by evaluating its performance with more frequent parameter updates.

Table 4 reports the parameter estimates for the generalized Merton, the NGARCH-Normal, the IG-GARCH, and the NGARCH-Jump models, together with their log-likelihood values in the in-sample setting. Note that in the case of the generalized Merton model, jump risk is allowed to be priced. In terms of NGARCH-Jump model, we no longer need to apply the restricted form because all parameters can be identified with the added information from option prices.

Table 4 Here

The likelihood ratio tests reveal that at the 1% level of significance, the NGARCH-Jump model improves significantly over the NGARCH-Normal model, and is also superior to the generalized Merton model. These results echo the earlier findings based on the return time series only. In short, both the GARCH effect and jumps play critical roles in capturing data features.

For comparing the non-nested models, we again resort to the AIC. The results in Table 4 show a clear dominance of the NGARCH-Jump model over the IG-GARCH model. The NGARCH-Normal also dominates the IG-GARCH model. Strikingly perhaps, the NGARCH-Normal model can dominate the IG-GARCH model based on the log-likelihood values even before factoring in its smaller set of parameters (6 vs. 8). The results from the combined data of returns and option prices indicate that the restricted form of the volatility dynamic essential to the IG-GARCH model is at odds with the data. In summary, simply incorporating skewed and/or heavy-tailed innovations into a GARCH model need not work well on return and/or option data. A suitable

volatility specification may be more important than the form of the conditional distribution for return innovations.

The two parameters, κ and γ, can be identified under the NGARCH-Jump model with the combined data. Their individual parameter estimates, $\kappa = 1.186$ and $\gamma = 0.41392$, differ from the theoretical values of 1 and 0 predicted when the jump risk is not priced, but not statistically significant. For the IG-GARCH model, the coefficient multiplying the inverse Gaussian innovation hardly changes from the physical measure to the risk-neutral measure (from $\eta = 5.99 \times 10^{-4}$ to $\eta^* = 5.92 \times 10^{-4}$).

Once the parameters are estimated, theoretical prices for all option contracts, including those with strike prices more than 5% away from the money, and with maturities exceeding 90 days are computed on each Wednesday, and their Black-Scholes implied volatilities established. The residuals are then established as the difference between theoretical and observed implied volatilities. Figure 4 shows a box and whiskers plot of these residuals by moneyness for each of the four models. Thirteen moneyness buckets were established with mean moneyness in each bucket ranging from 0.85 to 1.07, with the at-the-money contracts being in group 9. The figure clearly reveals that the generalized Merton model is incapable of removing significant biases along the moneyness dimension. At-the-money contracts are priced with little bias, but deep out-the-money contracts are overpriced and deep in-the-money contracts are underpriced.

Figure 4 Here

The plots of the NGARCH-Normal and IG-GARCH models are fairly similar to that of the generalized Merton model, in that the patterns of skewness are the same although of a much smaller magnitude. The NGARCH-Jump model has the least bias of all the models and clearly produces small biases for deep in-the-money contracts.

Figure 5 shows the box and whiskers plots of the residuals by maturity buckets for each of the four models also using all contracts in the in-sample period. For each maturity bucket, there are four box and whiskers plots representing from left to right the generalized Merton model, the NGARCH-Normal model, the IG-GARCH model and the NGARCH-Jump model, respectively.

Figure 5 Here

This figure shows large biases and large interquartile ranges for the generalized Merton model. For short-maturity contracts the NGARCH-Jump model displays some bias. For the other maturities the bias is slight and the interquartile ranges are smaller than the other models. Except for the generalized Merton model, the models seem to work roughly the same viewed from the maturity dimension.

Table 5 shows the proportion of times the NGARCH-Jump model produces a smaller absolute error than the other models. These comparisons are conducted over both moneyness and maturity buckets. For ease of displaying the results the moneyness categories are reduced to 5, with group 3 being at-the-money. Nonparametric tests of the null hypothesis that the proportion of wins is 50% against the alternative that the proportion is larger than 50% are conducted. The bold-faced numbers indicate those moneyness-maturity combinations where the reported proportion is significantly different, at the 1% level, from 50%. We also report the results of pairwise comparisons between the NGARCH-Normal and IG-GARCH models and between these models and the generalized Merton model. The left column reports the results for the in-sample period and the right side reports the results for the out-of-sample period.

Table 5 Here

The results show how poorly the Merton model performs against all other models. It also shows that the overall performance the NGARCH-Normal model is comparable to that of the IG-GARCH model in the in-sample test but it dominates in the out-of-sample analysis. Overall the NGARCH-Jump model outperforms all others in both "in-sample" and "out-of-sample" tests.

Our last set of analyses in this section deals with "out-of-sample" performance using the last five years of data. Figure 6 compares the time series of at-the-money volatilities averaged over a month with their theoretical counterparts for each of the four models. For the first 60 months these are "in-sample"' values. After that, the implied volatilities are "out-of-sample". The Merton model has a constant volatility, so the implied volatilities do not fluctuate. The figure shows the realization of volatilities for the three GARCH models. It is evident from these plots that the NGARCH-Jump model has the best performance whereas the Merton model performs the worst.

Figure 6

Our "in-sample" estimation was conducted using 60 cross-sections of option prices monthly over a period of 5 years. It is unlikely that, without more frequent recalibration, a model will provide reasonable fits to all options over the entire 5 years. In Figure 7 we investigate the bias produced by each model over each year in pricing at-the-money options. For each of the first 5 years, the box plots of the residuals for each model are "in-sample" residuals. For the years after 2000, the residuals are "out-of-sample". A model is viewed positively if the box plots for a given model over each of the first five years are unbiased and if the deterioration of the fit in the "out-of-sample" period is slow.

Figure 7 Here

The box plots in each year are ordered with NGARCH-Jump followed by NGARCH-Normal, IG-GARCH, and Merton. As can be seen there is considerable time variation over the years which the models do not capture. Of all the plots, the NGARCH-Jump model produces box plots with the least bias. However, all the models reveal that error terms are highly persistent. In the first three years of the "out-of -sample" period the fit of the NGARCH-Jump model is especially good compared to the other models. In the last two years the performance of all models deteriorates, with significant overpricing of all contracts.

To investigate whether the NGARCH-Jump model performs relatively better in the out-of-sample period for each month we compute the mean squared errors (MSE). We then compute the ratio of the MSE of each model relative to the MSE of the NGARCH-Jump model. In Figure 8 we present the time series of these ratios for the NGARCH-Normal model and for the IG-GARCH model. Notice that for both of these time series the ratios are for the most part significantly above 1. These results indicate that the NGARCH-Jump model performs well out-of-sample relative to the NGARCH-Normal and IG-GARCH models.

Figure 8

The lower panel of Figure 8 shows the time series of the monthly ratio of the MSE for the IG-GARCH model relative to the MSE for the NGARCH-Normal model. The graph reveals that in most months the ratio is above 1 indicating the superiority of the NGARCH-Normal model.

The final analysis investigates the out-of-sample bias across moneyness. For each of our 13 moneyness categories, we compute the MSE of residuals for each model for each month. The ratio of the MSEs of the NGARCH-Normal and IG-GARCH models relative to the NGARCH-Jump model are computed and box plots of these ratios are presented in Figure 9. The median ratios are above 1 for all moneyness contracts for the IG-GARCH model. The NGARCH-Normal model is more competitive with the NGARCH-Jump model, but only for contracts close to bin 9 which represents the at-the-money category.

Figure 9

The errors reported in this study 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 in conjunction with panel data on options to estimate one set of parameters, while they refit their parameters based on out-the-money

contracts. Further, we do not reestimate our model dynamically. Indeed, at the extreme, we obtain option prices 5 years after parameters have been estimated. 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. In spite of differing objectives, our error terms appear to be of similar magnitudes to their reported values.

4 Conclusion

In this paper we extend Duan's (1995) GARCH option pricing model that relies on normal innovations to allow for a particular type of non-normal innovations that can be related to jumps. This class of GARCH-Jump models extends 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 sharp contrast to the typical GARCH models based on normal innovations. Since these latter models in their limiting forms only exhibit a diffusive stochastic volatility behavior, it is not surprising that they are incapable of removing some well-known option pricing biases. We provide in this paper a theory of GARCH option pricing that permits contracts to be priced in the presence of skewed and leptokurtic innovations, and demonstrate that this generalization is empirically significant.

Specifically, using data on the S&P 500 index and the set of European options written on this index, we have provided empirical tests of the ability of GARCH-Jump models to price options. We show that introducing jumps that allow for heavy tails and higher kurtosis adds significantly to explaining the time series behavior of the S&P 500 index and its options.

For pricing of options our simplest nested model, the Merton model, performed the worst. Either allowing for time-varying volatility or including priced jump risk leads to better performance. 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 over time. Further, we show that the GARCH-Jump model is capable of pricing options well without requiring frequent model recalibration. Indeed, our models were capable of good pricing even after one or two years had passed after the parameters have been estimated.

In general, distinguishing between stochastic volatility and jumps is difficult. Our empirical results show that jumps were frequent, with more than one "jump" a day. This implies that to capture heavy-tailed return distributions random mixing of normal innovations as in the

GARCH-Jump model can be a productive approach. Although our empirical results are based on a local volatility updating equation of the NGARCH form, the theory holds for other GARCH specifications. We have, for example, examined a threshold GARCH model with jumps as an alternative to the NGARCH model with jumps. While not reported here, our preliminary results indicate that the difference between the two volatility structures is not particular important for European option pricing. What is more important is that both models should allow for jumps so that local returns conditionally are not constrained to be normally distributed.

References

Alexander, C. (2004) Normal Mixture Diffusion with uncertain volatility: Modeling short and long-term smile effects, *Journal of Banking and Finance*, 28, 2957-2980.

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., 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.

Bates, D. and R. Craine (1999) Valuing the Futures market Clearinghouse's Default Exposure during the 1987 Crash, *Journal of Money, Credit, and Banking*, 31, 248-272.

Brigo, Damiano., Fabio Mercurio, and F. Rapisarda (2004) Smile at Uncertainty, *Risk*, 17, 97-101.

Brigo, D., and F. Mercurio (2002) Lognormal-Mixture Dynamics and Calibration to Market Volatility Smiles, *International Journal of Theoretical and Applied Finance,* 5, 427-446.

Carr, P., and L. Wu (2003) The Finite Moment Log Stable Process and Option Pricing, *Journal of Finance,* 58, 753-777.

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 (2006) Option Valuation with Conditional Skewness, *Journal of Econometrics*, 131, 253-284.

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

Ding, Z., C. Granger, R. Engle (1993) A Long Memory Property of Stock Market Returns and a New Model, *Journal of Empirical Finance*, 1, 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. (2002) Nonparametric Option Pricing by Transformation, 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 (2006) Approximating GARCH-Jump Models, Jump-Diffusion Processes, and Option Pricing, *Mathematical Finance*, 16, 21-52.

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.

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 (2005) An Empirical Comparison of GARCH Models, *Review of Derivatives Research*, 8, 129-150.

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

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

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.

Madan, D., P. Carr, and E. Chang (1998) The Variance Gamma Process and Option Pricing, *European Finance Review,* 2, 79-105.

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

McCulloch, H. (2003) The Risk-Neutral Measure and Option Pricing under Log-Stable Uncertainty, *working paper,* Ohio State University.

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. (2005) Pricing American Options when the Underlying Asset Follows GARCH Processes, *Journal of Empirical Finance*, 12, 576-611.

Appendix

Proof of Proposition 1

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

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

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

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

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

Now consider the pricing equation for the asset. We have, from equation (13),

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

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:

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

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

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

and the result follows.

Proof of Lemma 1

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

(i) Q is a probability measure since:

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

where the last equality follows from the fact that:

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

Continuing this process we obtain

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

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

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

Continuing this process, we obtain:

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

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:

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

We know that

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

Using these results, we obtain

$$
E^{\mathsf{Q}}[e^{cW_t}|\mathcal{F}_{t-1}] = \exp\left(c\alpha_t + r_t + a + \frac{1}{2}(c^2 h_t + b^2 + 2c\sqrt{h_t}b\rho) - \lambda_t(1 - \kappa K_t(c))\right) \tag{39}
$$

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

Now, let $c = 0$. Then,

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

or, equivalently,

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

Substituting this expression into equation (39), we obtain

$$
E^{\mathsf{Q}}[e^{cW_t}|\mathcal{F}_{t-1}] = \exp\left(c\alpha_t + \frac{1}{2}(c^2 h_t + 2c\sqrt{h_t}b\rho) - \lambda_t\kappa(1 - K_t(c))\right) \tag{40}
$$

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

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

from which:

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

Hence:

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

Let

$$
\begin{aligned}
\tilde{\alpha}_t &= r_t - \frac{1}{2}h_t + \tilde{\lambda}_t\left(1 - K_t(1)\right) \\
\tilde{\lambda}_t &= \lambda_t\kappa
\end{aligned}
$$

We can write:

$$
E^{\mathsf{Q}}[e^{cW_t}|\mathcal{F}_{t-1}] = \exp\left[c\tilde{\alpha}_t + \frac{1}{2}c^2 h_t - \tilde{\lambda}_t\left(1 - K_t(c)\right)\right] \tag{41}
$$

Now consider the following system:

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

where

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

It is straightforward to verify that the moment generating function of $\tilde{W}_t$ is the same as that in equation (41). Thus, under measure $\mathtt{Q}$, W_t is distributionally equivalent to $\tilde{W}_t$.

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

Figure 1: Moneyness and Expiry Distributions

The top figure shows the distribution of moneyness of all contracts, where moneyness is defined as the index price divided by the strike price. The bottom figure shows the histogram of days to expiration for all our contracts, over the time period from January 1996 to June 2005.

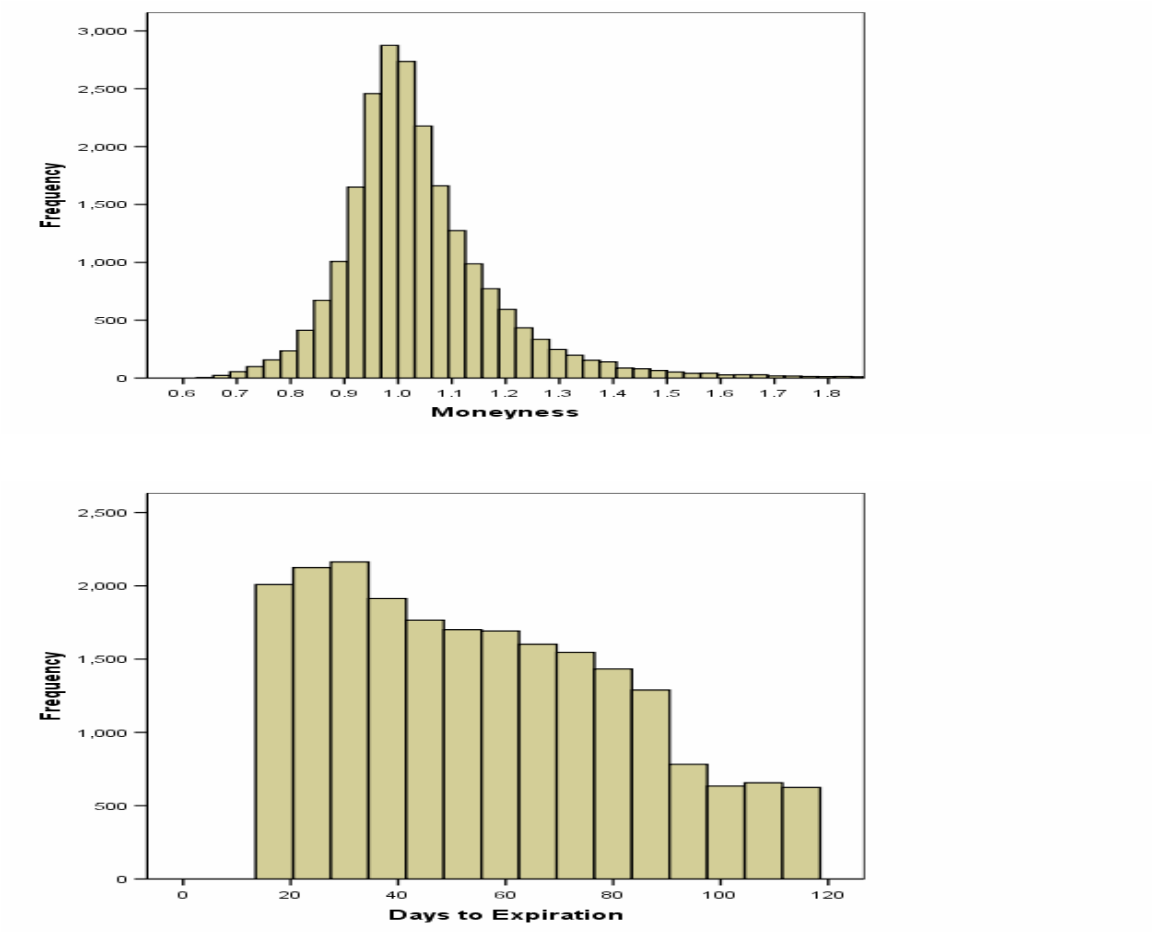

Figure 2: Implied Volatility

The top figure shows the time series of the volatility smile. In particular, the average actual implied volatility over the moneyness groups is computed for each month in the data set, and these values are plotted over time. The bottom graph plots the average implied volatility for each moneyness group over the entire data period. The final graph plots the time series of the implied volatilities of the at-the-money contracts for each month over the time period from January 1996 (month 1) to June 2005.

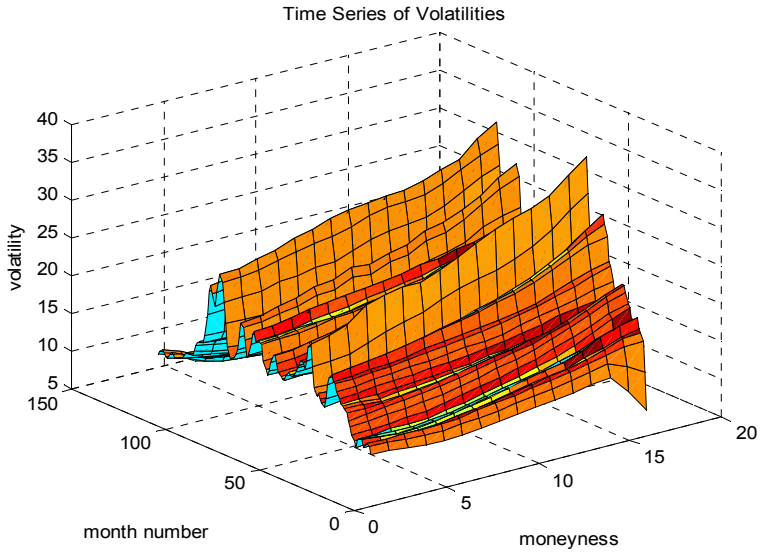

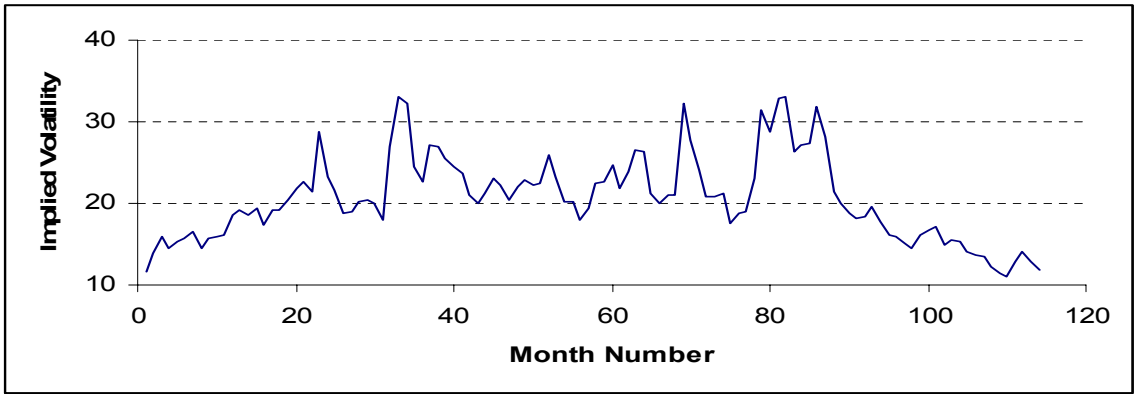

Figure 3
Q-Q Plots for Different Models

Each daily return residual is first transformed into a uniform random variable using the appropriate conditional distribution function. The uniform random variables are then converted to standard normal random variables using the inverse normal transformation. The plots below are Q-Q plots that show deviations from normality.

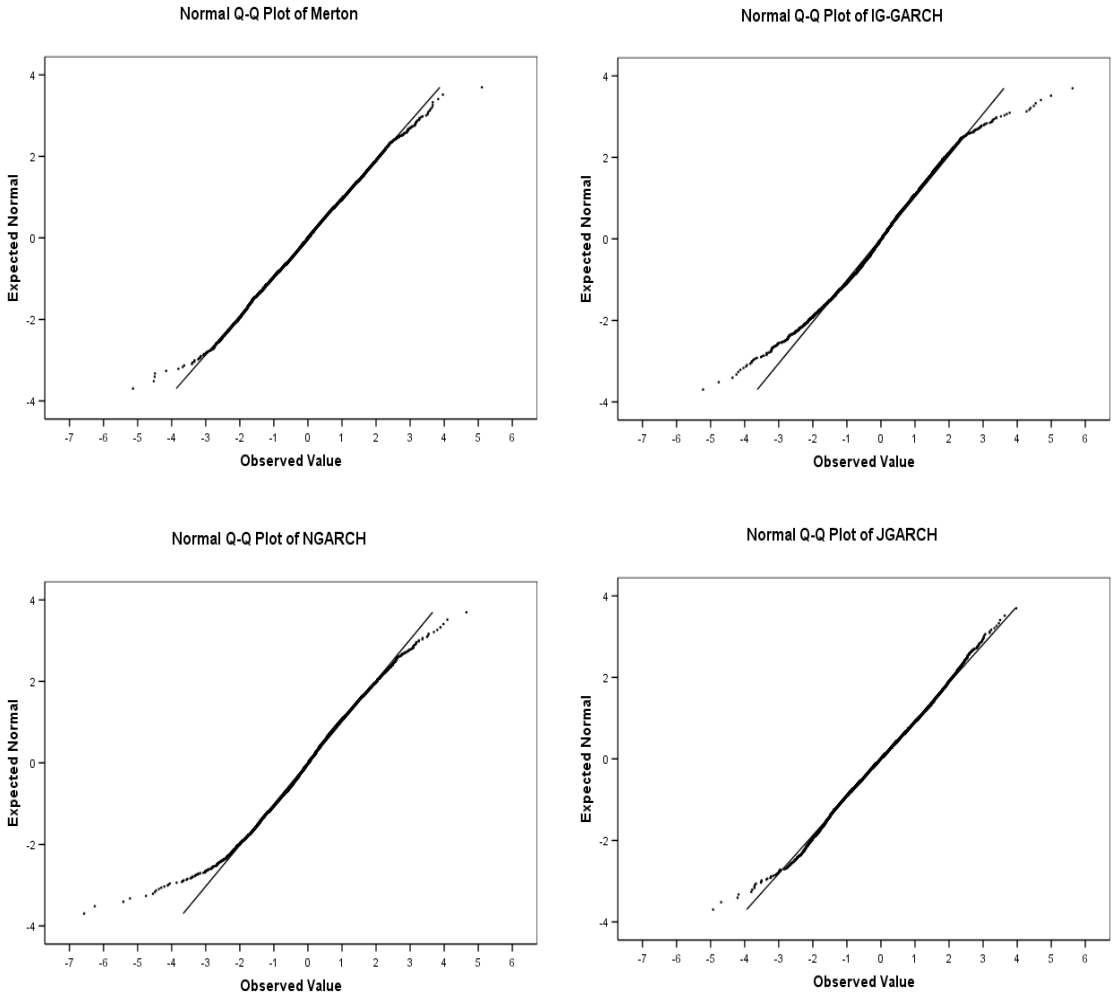

Figure 4: In-Sample Residuals by Moneyness.

The plots show pricing errors by moneyness for all models. In-sample pricing errors are computed as the difference between theoretical and actual implied volatilities, so negative values imply underpricing by the model. The plots show the residuals for the NGARCH-Jump (JGARCH), NGARCH-Normal (NGARCH), IG-GARCH and Merton models respectively. The residuals are in-sample residuals covering the first half of the data period, namely the five years from January 1996. The moneyness bins were chosen so that the number of contracts in each bin was about equal. Bin 1 consists of deep out-the-money options with moneyness less than 0.86. The successive upper bounds for the bins are 0.90, 0.93, 0.95, 0.96, 0.98, 0.99, 1.00, 1.01, 1.03, 1.04, 1.06 and 1.20.

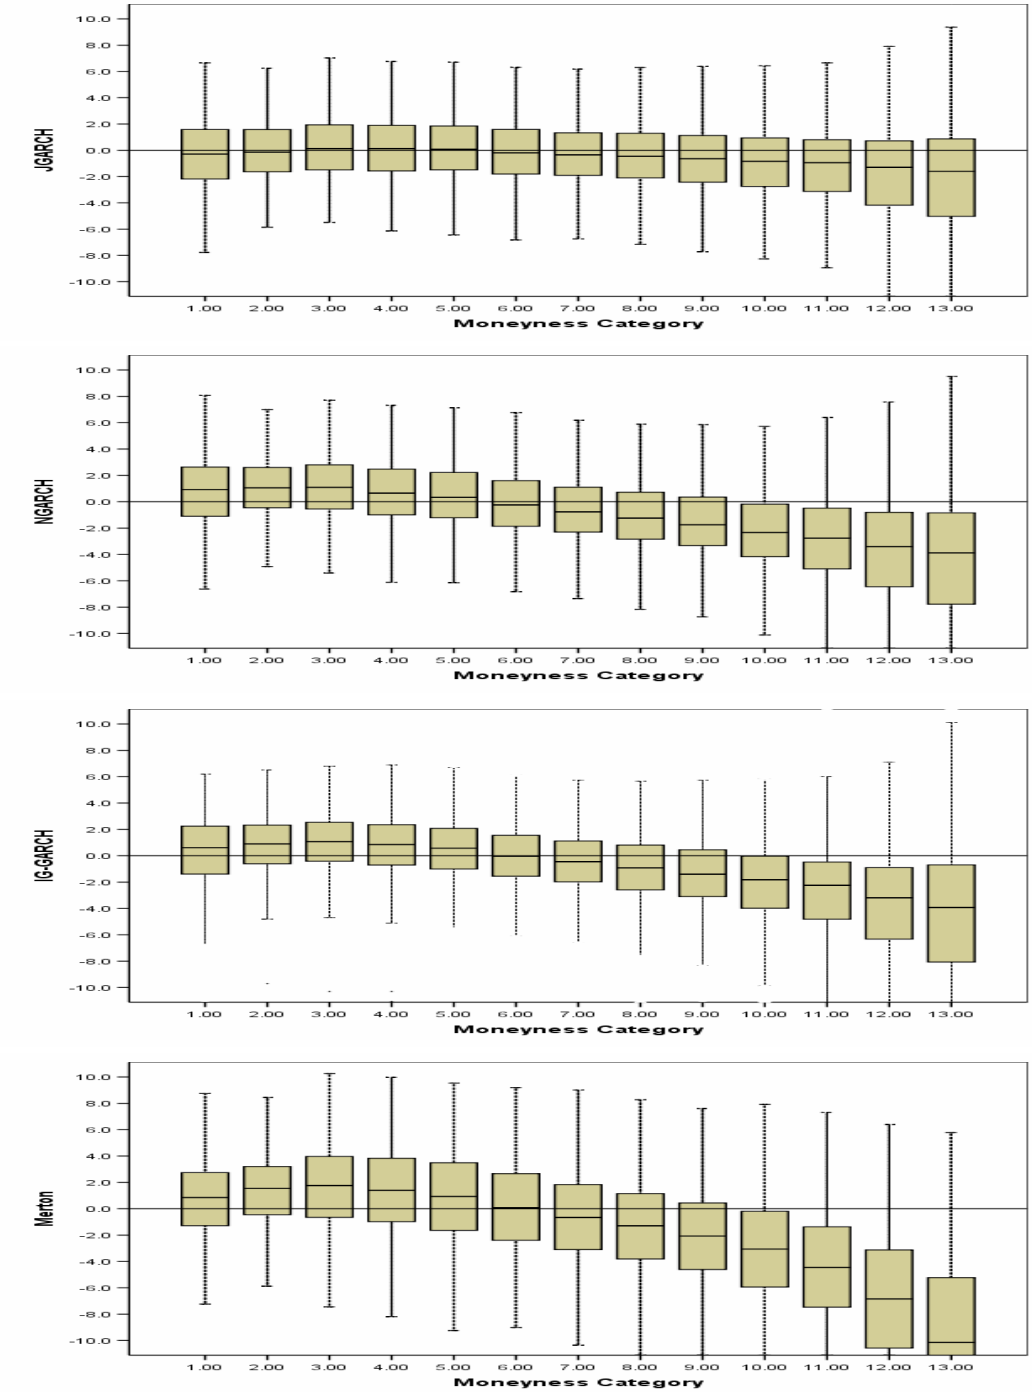

Figure 5: Maturity Bias by Four Models

The figure shows box-whisker plots of the pricing errors produced by maturity for four models. In-sample errors are computed as the difference between theoretical and actual implied volatilities, so negative values imply underpricing by the model. For each maturity bucket there are four box plots ordered from left to right by the Merton model, the IG-GARCH model, the NGARCH-Normal model and the NGARCH-Jump model. The residuals are in-sample residuals covering the first half of the data period, namely the five years from January 1996. The contracts with maturities exceeding 90 days were not used in estimation, and thus can be viewed as out-of-sample residuals.

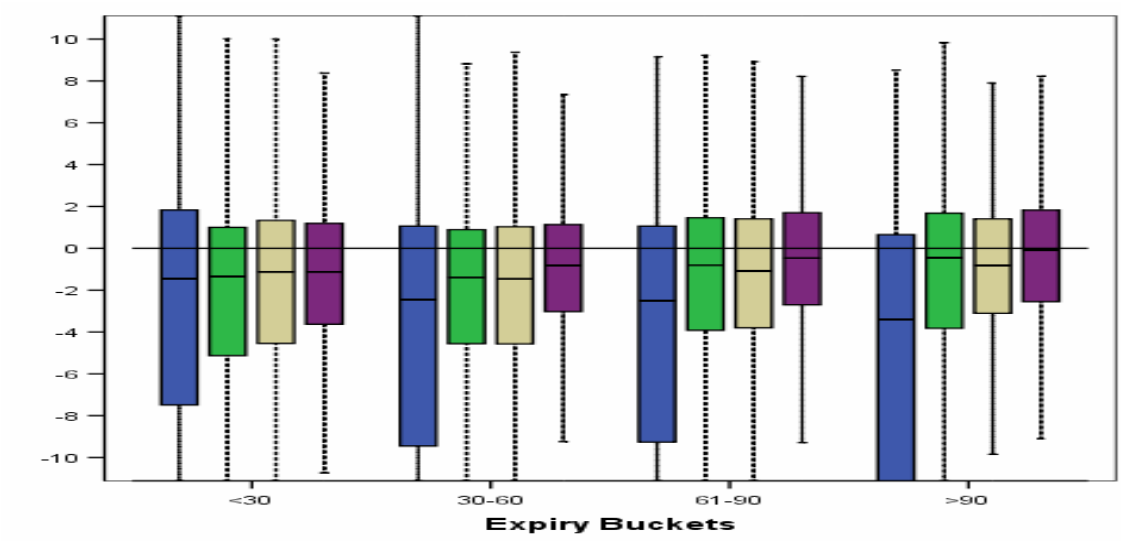

Figure 6: Time Series of Implied Volatilities for Each Model

For each month we obtain the theoretical prices for all contracts within one half percent of the index value. We update the volatility process daily based on the realized index price. Given the theoretical option prices we compute the Black-Scholes implied volatility for each contract. For each model we compute the average Black-Scholes implied volatility for the month. The time series of these averages are compared with the actual average Black-Scholes volatilities for the month. The first figure is the Merton model, where the volatility is constant; the second figure is for the NGARCH-Normal model; the next figure is the IG-GARCH model and the bottom figure is the NGARCH-Jump model.

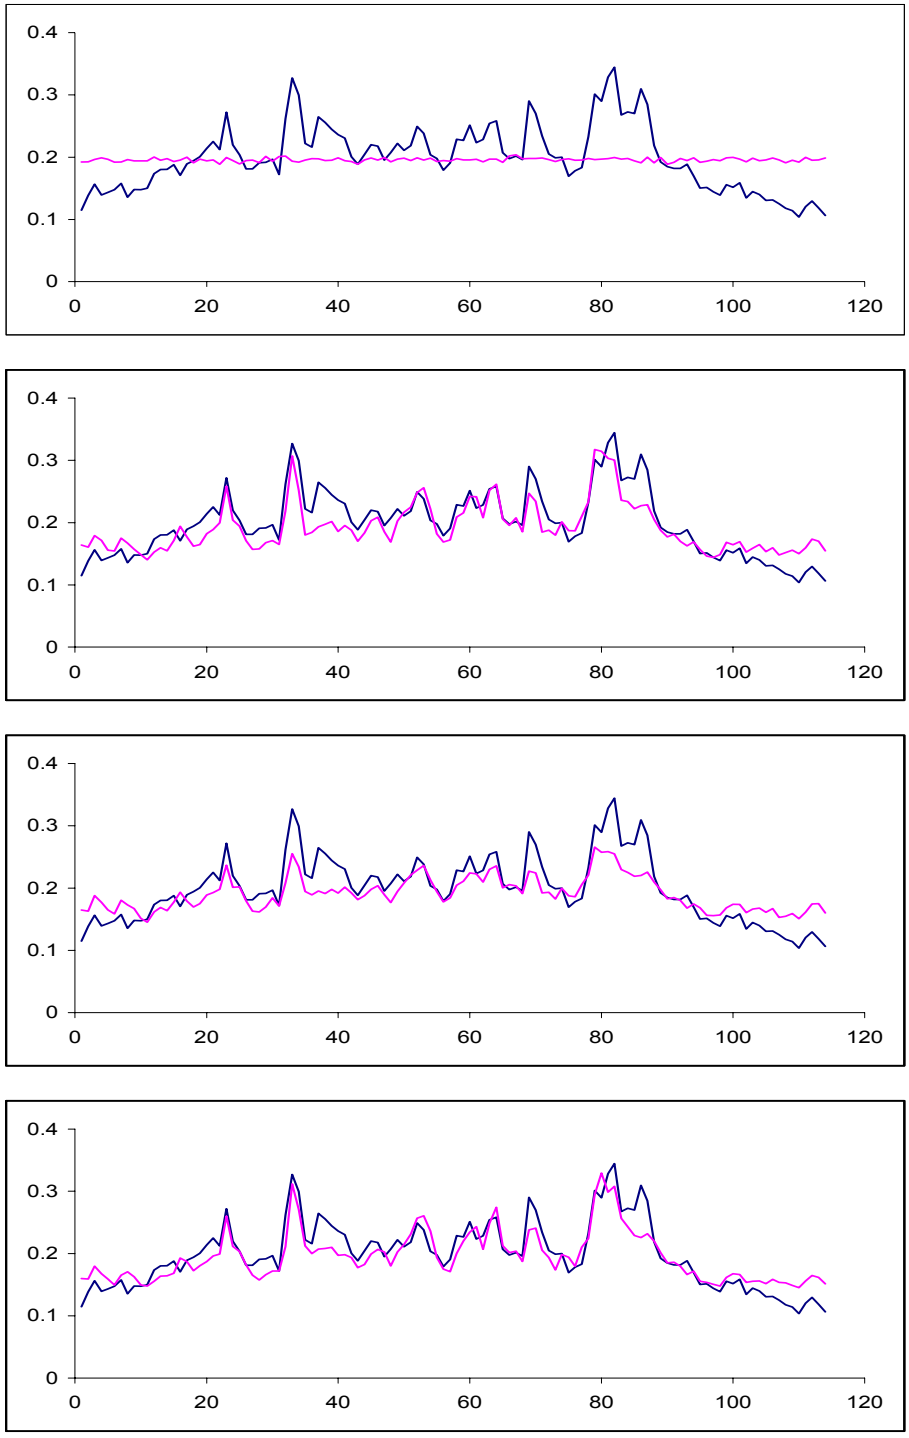

Figure 7: Residuals by Year and by Model

The figure shows the box plots of the residuals for at-the-money contracts for each model and for each year in the in-sample period for the first 5 years through year 2000, and in the out-of-sample period for the next 5 years. The box plots in each year are from left to right are the NGARCH-Jump, NGARCH-Normal, IG-GARCH, and Merton models, respectively. The parameters are estimated using option contracts with expiry dates less than 90 days over the period 1996 through 2000. The residuals used to compute the box plots are based off all contracts within one half percent of the strike price. The out-of-sample prices obtained after 2000 are based on parameters estimated using the data up through 2000.

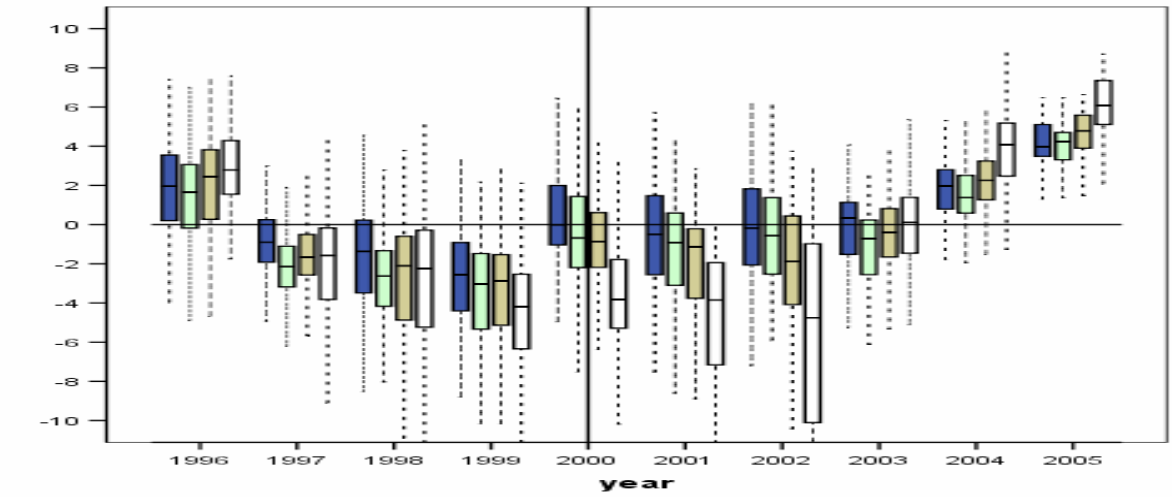

Figure 8: Relative Out-of-Sample Performance of Models over Time

The top figure shows the time series of monthly ratios of MSEs of the NGARCH-Normal (dashed line) and IG-GARCH (solid line) relative to the MSE of the NGARCH-Jump model. Panel B shows the time series of the monthly ratios of the MSEs of the IG-GARCH model relative to the MSEs of the NGARCH-Normal model. Ratios bigger than 1 indicate the superiority of the numeraire model. Month number starts from the first out-of-sample month in January 2001.

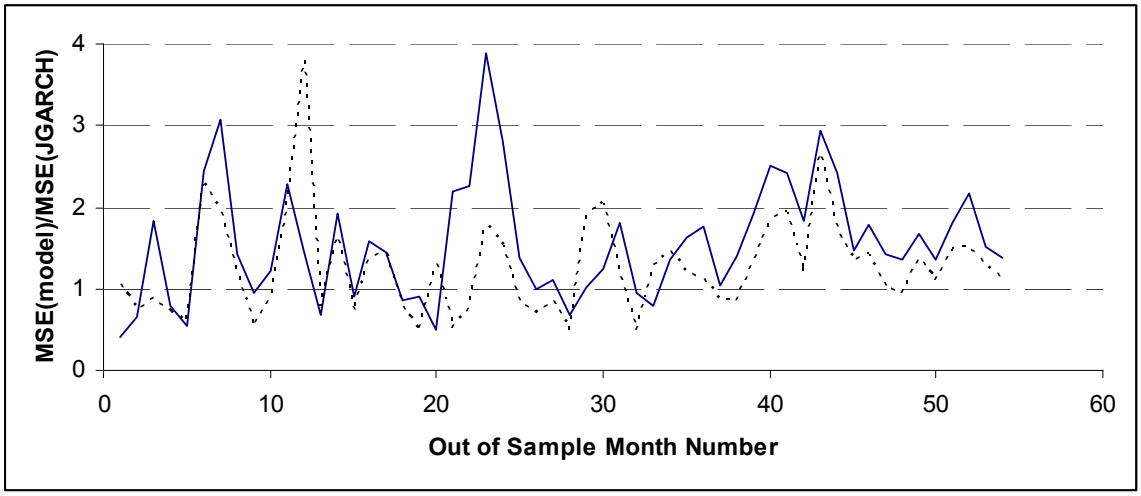

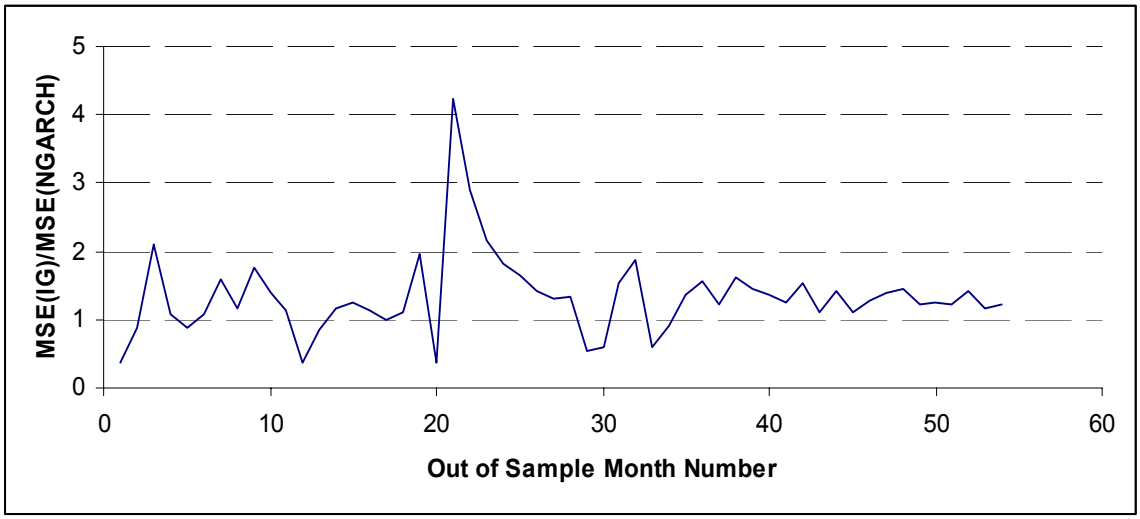

Figure 9: Relative Out-of-Sample Performance of Models by Moneyness

The top figure displays a box plot of monthly ratios of MSE of NGARCH-Normal and IG-GARCH residuals relative to MSE of NGARCH-Jump residuals in each moneyness bin for all out-of-sample months. Ratios bigger than 1 indicate the superiority of the JGARCH model. The bottom figure displays a box plot of monthly ratios of the MSE for the IG-GARCH model relative to the MSE of the NGARCH-Normal model for all contracts in each moneyness bin, for all out-of-sample months. Ratios bigger than 1 indicate the superiority of the NGARCH-Normal model.

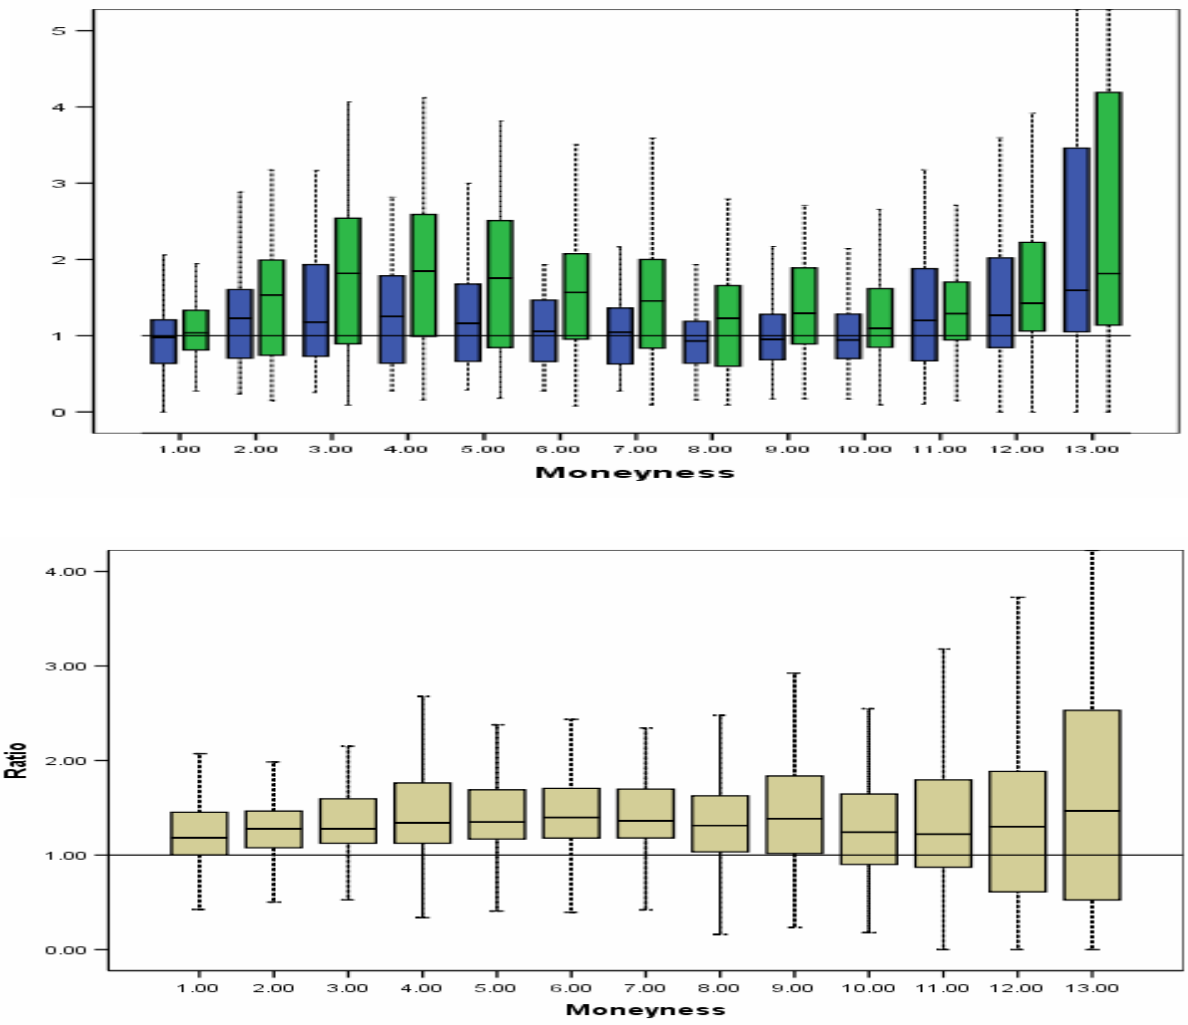

Table 1: Distribution of Contracts by Moneyness and Maturity

The top table shows the number of contracts in each moneyness-maturity bin in the in-sample from 1996 through 2000. The bottom table shows the average implied volatilities.

Moneyness	Days to Expiration				
	<31	31-60	61-90	>90	Total
<0.86	191	438	326	264	1219
0.86 - 0.90	234	480	345	158	1217
0.90 - 0.93	272	491	331	125	1219
0.93 - 0.95	294	498	315	113	1220
0.95 - 0.96	285	511	313	109	1218
0.96 - 0.98	285	511	308	114	1218
0.98 - 0.99	273	526	311	108	1218
0.99 - 1.00	262	526	320	110	1218
1.00 - 1.01	274	521	324	104	1223
1.01 - 1.03	273	538	307	99	1217
1.03 - 1.04	261	532	308	117	1218
1.04 - 1.06	239	541	319	120	1219
1.06 - 1.09	225	537	326	132	1220
1.09 - 1.11	191	550	328	147	1216
1.11 - 1.15	166	550	339	166	1221
1.15 - 1.20	143	521	359	195	1218
1.20 - 1.30	141	507	356	214	1218
>1.30	125	458	335	301	1219
Total	4134	9236	5870	2696	21936

Moneyness	Days to Expiration				
	<31	31-60	61-90	>90	Total
<0.86	33.52	24.27	20.29	18.78	23.47
0.86 - 0.90	22.39	18.04	17.56	17.43	18.66
0.90 - 0.93	18.23	17.28	17.85	18.15	17.73
0.93 - 0.95	17.03	16.92	17.72	18.35	17.29
0.95 - 0.96	16.88	17.44	18.18	18.62	17.61
0.96 - 0.98	17.21	17.90	18.65	18.89	18.02
0.98 - 0.99	17.85	18.71	19.24	19.97	18.77
0.99 - 1.00	18.52	19.55	19.86	20.19	19.47
1.00 - 1.01	19.59	20.04	20.29	21.08	20.09
1.01 - 1.03	20.25	20.87	21.05	20.65	20.76
1.03 - 1.04	21.29	21.69	21.79	21.74	21.63
1.04 - 1.06	22.64	22.56	22.73	22.14	22.58
1.06 - 1.09	25.26	23.69	23.54	23.29	23.90
1.09 - 1.11	28.49	25.31	24.69	24.03	25.49
1.11 - 1.15	33.52	27.85	25.93	24.94	27.69
1.15 - 1.20	39.55	31.24	28.32	26.29	30.56
1.20 - 1.30	51.11	36.92	31.56	29.49	35.69
>1.30	67.06	53.83	42.90	38.14	48.31
Total	24.39	23.97	23.04	23.63	23.76

Table 2: Parameter Estimates from the Time Series of Daily Returns

The tables report the parameter estimates for the models based on time series over the time period starting in 1970 and over two sub-samples. The Merton and NGARCH-Jump models are their restricted versions.

1. Sample: 1/2/1970 - 12/30/2005

Parameter	Merton		NGARCH		JGARCH		IG-GARCH	
	Estimate	Std. Dev.	Estimate	Std. Dev.	Estimate	Std. Dev.	Estimate	Std. Dev.
β_0	2.73E-05	6.06E-07	1.13E-06	7.81E-08	1.71E-07	1.99E-08	-1.44E-06	1.661E-08
β_1			9.08E-01	2.11E-03	0.91946	0.0033	-19.13300	5.819E-13
β_2			5.62E-02	1.21E-03	0.04843	0.0012	4.08E-06	1.277E-10
β_3							2.47E+07	2.305E-19
c			6.60E-01	1.20E-02	0.67556	0.0124		
bρ	-0.08347	6.71E-03	-0.03960	0.00327				
λ	0.61160	1.60E-02			2.26560	0.0047		
$\bar{\gamma}$	1.9728	1.02E-02			1.35490	0.0073		
$\bar{\mu}$	-0.12294	7.86E-03			-0.0733	0.0107		
δ					-2.56E-01	0.0141		
η							-6.20E-04	1.070E-07
ν							1614.5	4.410E-14
ML Value	29745.3		30297.4		30457.6		30238.3	
AIC Test	-59480.6		-60584.8		-60899.2		-60464.6	

2. Sample: 1/2/1980 - 12/30/2005

Parameter	Merton		NGARCH		JGARCH		IG-GARCH	
	Estimate	Std. Dev.	Estimate	Std. Dev.	Estimate	Std. Dev.	Estimate	Std. Dev.
β_0	3.48E-05	2.99E-07	1.62E-06	2.58E-08	1.34E-07	3.42E-09	3.32E-08	2.232E-09
β_1			8.89E-01	7.59E-04	0.90912	0.0011	-19.20300	1.630E-13
β_2			6.39E-02	7.77E-04	0.05407	0.0009	4.10E-06	3.438E-11
β_3							2.47E+07	1.311E-19
c			7.22E-01	6.42E-03	0.71245	0.0039		
bρ	-0.06888	2.28E-03	-0.04088	0.00139				
λ	0.50025	7.56E-03			2.27930	0.0037		
$\bar{\gamma}$	1.9412	5.07E-03			1.7435	0.0024		
$\bar{\mu}$	-0.091127	4.74E-03			-0.35940	0.0026		
δ					-1.07E-01	0.0021		
η							-6.14E-04	2.581E-08
ν							1632.4	5.645E-15
ML Value	21258.9		21595.2		21755.3		21560.8	
AIC Test	-42507.8		-43180.4		-43494.6		-43109.6	

3. Sample: 1/2/1990 - 12/30/2005

Parameter	Merton		NGARCH		JGARCH		IG-GARCH	
	Estimate	Std. Dev.	Estimate	Std. Dev.	Estimate	Std. Dev.	Estimate	Std. Dev.
β_0	2.74E-05	2.12E-07	1.15E-06	1.26E-08	1.82E-07	2.47E-09	1.66E-08	2.155E-09
β_1			8.86E-01	5.07E-04	0.88888	0.0007	-19.17000	9.200E-14
β_2			5.64E-02	4.19E-04	0.05509	0.0004	4.07E-06	3.498E-11
β_3							2.47E+07	7.312E-20
c			9.41E-01	6.63E-03	0.92969	0.0017		
bρ	-0.05992	3.05E-03	-0.03169	0.00129				
λ	0.70446	6.77E-03			2.39300	0.0030		
$\bar{\gamma}$	1.9441	2.56E-03			1.3460	0.0008		
$\bar{\mu}$	-0.10091	5.30E-03			-0.13336	0.0032		
δ					-3.94E-01	0.0008		
η							-6.08E-04	2.258E-08
ν							1647.4	6.145E-15
ML Value	13060.4		13407.6		13464.7		13369.9	
AIC Test	-26110.8		-26805.2		-26913.4		-26727.8	

Table 3: Descriptive Statistics on Residuals for Time Series Models

The top table reports descriptive statistics on the normalized residuals for the four different time series models. If the models are appropriately specified, the residuals should be standard normal random variables. The bottom table presents the Ljung-Box tests on the residuals and squared residuals for lags of up to 5 business days. The tests of hypothesis on no autocorrelation were conducted at the 5% level, and the significant values are indicated with stars. In addition to the full sample, tests are done in the two sub-periods after the stock market crash in October 1987.

Descriptive Statistics of Normalized Residuals

Residual	Merton	NGARCH	JGARCH	IG-GARCH
mean	0.003	0.000	0.001	-0.009
std	1.048	0.991	1.067	0.980
max	5.106	4.652	3.972	5.623
min	-5.135	-6.570	-4.927	-5.222
skew	0.019	-0.152	-0.005	-0.172
excess kurtosis	0.248	1.247	-0.111	1.198
Test of Normality				
Kolmogorov-Smirnov Statistic	0.012	0.025	0.010	0.033
p value	0.005	0.000	0.027	0.000

Ljung-Box tests for autocorrelation on residuals and squared residuals

1970-2005

		JGARCH		IG-GARCH		NGARCH		Merton	
	Lag	Auto-Correlation	Box-Ljung Statistic	Auto-Correlation	Box-Ljung Statistic	Auto-Correlation	Box-Ljung Statistic	Auto-Correlation	Box-Ljung Statistic
Residuals	1	0.099	89.91*	0.093	78.66*	0.099	89.44*	0.072	47.34*
	2	-0.002	89.94*	-0.004	78.81*	-0.001	89.46*	-0.025	53.04*
	3	-0.011	91.01*	-0.011	79.86*	-0.011	90.52*	-0.021	57.07*
	4	-0.012	92.30*	-0.011	80.95*	-0.012	91.84*	-0.016	59.26*
	5	-0.010	93.16*	-0.008	81.51*	-0.010	92.70*	-0.003	59.37*
Squared Residuals	1	0.019	3.294	0.051	23.23*	0.014	1.775	0.110	109.93*
	2	0.004	3.414	0.040	37.54*	0.000	1.776	0.160	343.32*
	3	0.002	3.439	0.029	45.03*	-0.001	1.779	0.096	426.66*
	4	-0.003	3.514	0.010	45.91*	-0.004	1.914	0.034	437.31*
	5	0.005	3.709	0.044	63.28*	0.003	2.000	0.158	663.32*

1988-1995

		JGARCH		IG-GARCH		NGARCH		Merton	
	Lag	Auto-Correlation	Box-Ljung Statistic	Auto-Correlation	Box-Ljung Statistic	Auto-Correlation	Box-Ljung Statistic	Auto-Correlation	Box-Ljung Statistic
Residuals	1	0.036	2.612	0.029	1.721	0.036	2.569	0.022	0.953
	2	-0.001	2.613	-0.002	1.728	-0.002	2.574	-0.014	1.340
	3	-0.039	5.760	-0.041	5.196	-0.039	5.638	-0.047	5.881
	4	-0.029	7.520	-0.021	6.132	-0.031	7.584	-0.009	6.031
	5	-0.017	8.106	-0.016	6.637	-0.018	8.235	-0.015	6.497
Squared Residuals	1	0.004	0.027	0.014	0.403	0.019	0.716	0.068	9.40*
	2	-0.009	0.186	-0.008	0.526	-0.002	0.727	0.028	10.96*
	3	-0.014	0.558	-0.010	0.713	-0.008	0.858	0.029	12.67*
	4	-0.004	0.592	-0.001	0.714	0.005	0.913	0.032	14.80*
	5	-0.007	0.690	0.008	0.833	-0.005	0.964	0.068	24.19*

1996-2005

		JGARCH		IG-GARCH		NGARCH		Merton	
	Lag	Auto-Correlation	Box-Ljung Statistic	Auto-Correlation	Box-Ljung Statistic	Auto-Correlation	Box-Ljung Statistic	Auto-Correlation	Box-Ljung Statistic
Residuals	1	0.011	0.320	0.008	0.171	0.012	0.344	-0.014	0.460
	2	-0.012	0.663	-0.013	0.568	-0.011	0.663	-0.027	2.359
	3	-0.026	2.425	-0.026	2.213	-0.026	2.387	-0.028	4.308
	4	-0.014	2.952	-0.007	2.355	-0.014	2.871	0.006	4.385
	5	-0.034	5.872	-0.035	5.370	-0.034	5.753	-0.041	8.643
Squared Residuals	1	-0.025	1.546	0.004	0.043	-0.027	1.889	0.186	86.89*
	2	0.026	3.217	0.044	4.867	0.024	3.382	0.187	175.50*
	3	-0.007	3.352	0.008	5.032	-0.009	3.571	0.176	253.38*
	4	0.009	3.567	0.012	5.421	0.009	3.784	0.120	289.99*
	5	0.015	4.170	0.040	9.451	0.015	4.378	0.181	372.58*

Table 4: Estimates Using Time Series and Options

The parameter estimates based on the time series analysis were used as initial starting points for the panel estimation optimization problem, where option prices were incorporated into the analysis. The options used were restricted to contracts within 5% of the strike and maturities less than 90 days. Option prices were taken monthly, and incorporated with daily return data as discussed in the text. The data used spanned the period from January 1996 through to December 2000. For each model we report the point estimates, the standard deviations, the log-likelihood function (ML-value) and the Akaike Information Criterion (AIC)

1/1996-12/2000

Parameter	Merton		NGARCH		JGARCH		IG-GARCH	
	Estimate	Std. Dev.	Estimate	Std. Dev.	Estimate	Std. Dev.	Estimate	Std. Dev.
β_0	1.38E-04	2.28E-06	2.34E-06	1.06E-08	1.59E-07	1.81E-08	9.33E-09	1.592E-06
β_1			0.92994	6.23E-04	0.95110	0.0004	-19.47	6.118E-01
β_2			0.03600	3.35E-04	0.03004	0.0051	3.89E-06	4.717E-08
β_3							2.68E+07	1.739E+06
c			0.73599	0.085	-0.07674	0.0958		
$b\rho$	-0.07033	3.57E-01	-0.06254	0.08321	-0.40798	0.1808		
λ	0.91867	1.12E+00			0.38638	0.0722		
$\bar{\gamma}$	5.28E-07	4.91E-05			4.32200	0.2603		
$\bar{\mu}$	0.3295	3.59E-02			-2.0920	0.9460		
δ								
η							-5.99E-04	1.068E-05
η^*							-5.92E-04	1.095E-05
ν							1649.4	2.937E+01
γ	9.3426E-07	2.69E-11			0.41392	0.6866		
K0	0.89314	1.09E+00			1.18600	0.2239		
ω	0.041002	6.07E-07	0.02920	0.000054338	0.02469	5.063E-05	0.02819	9.155E-05
ML Value	3821.10		4533.90		4852.10		4465.05	
AIC Test	-7634.20		-9063.80		-9690.20		-8922.10	

Table 5: Pairwise Comparisons of Models

We compute the residuals for each contract using different models. For each pair of models, we record which model produces the smaller absolute residual. We report the proportion of times the model indicated first, produces smaller absolute errors than the model indicated second. For example, for deep in the money contracts with maturities exceeding 90 days, the NGARCH-Jump (JGARCH) model beat the NGARCH-Normal (NGARCH) model in the in-sample period 59% of the time. Values significantly different from 50% are bold-faced.

In-Sample Comparisons

Expiry	deep out	out	at	in	deep in	Total
JGARCH vs NGARCH						
<30	**0.47**	**0.57**	0.53	**0.64**	**0.73**	**0.58**
30-60	**0.55**	0.53	**0.59**	**0.68**	**0.75**	**0.64**
61-90	**0.59**	0.54	**0.59**	**0.62**	**0.65**	**0.61**
>90	**0.61**	0.48	0.47	0.50	**0.62**	**0.59**
Total	**0.55**	**0.54**	**0.57**	**0.64**	**0.70**	**0.61**
JGARCH vs IG-GARCH						
<30	**0.46**	0.49	0.53	**0.66**	**0.74**	**0.58**
30-60	0.51	0.52	**0.61**	**0.69**	**0.75**	**0.64**
61-90	**0.58**	**0.56**	**0.59**	**0.63**	**0.71**	**0.63**
>90	**0.61**	**0.61**	0.54	**0.68**	**0.75**	**0.67**
Total	**0.53**	**0.53**	**0.58**	**0.67**	**0.74**	**0.63**
JGARCH vs Merton						
<30	0.54	**0.61**	**0.65**	**0.68**	**0.86**	**0.67**
30-60	**0.60**	**0.67**	**0.70**	**0.72**	**0.89**	**0.75**
61-90	**0.62**	**0.67**	**0.70**	**0.74**	**0.89**	**0.76**
>90	**0.66**	**0.79**	**0.74**	**0.68**	**0.92**	**0.80**
Total	**0.60**	**0.67**	**0.69**	**0.71**	**0.89**	**0.74**
NGARCH vs IG-GARCH						
<30	**0.45**	**0.45**	0.47	0.55	**0.49**	**0.47**
30-60	0.49	0.49	**0.43**	**0.44**	**0.48**	**0.47**
61-90	**0.53**	0.53	0.47	**0.45**	0.50	0.50
>90	**0.55**	0.58	0.51	0.54	0.54	**0.55**
Total	0.50	0.50	**0.46**	**0.48**	**0.49**	**0.49**
NGARCH vs Merton						
<30	**0.56**	**0.65**	**0.65**	**0.57**	**0.64**	**0.61**
30-60	**0.61**	**0.66**	**0.62**	**0.62**	**0.73**	**0.66**
61-90	**0.61**	**0.65**	**0.63**	**0.67**	**0.80**	**0.70**
>90	**0.63**	**0.80**	**0.81**	**0.69**	**0.85**	**0.77**
Total	**0.60**	**0.67**	**0.65**	**0.63**	**0.75**	**0.68**
IG-GARCH vs Merton						
<30	**0.59**	**0.72**	**0.71**	**0.70**	**0.64**	**0.65**
30-60	**0.64**	**0.74**	**0.68**	**0.64**	**0.72**	**0.69**
61-90	**0.65**	**0.71**	**0.69**	**0.69**	**0.80**	**0.72**
>90	**0.63**	**0.71**	**0.67**	**0.65**	**0.82**	**0.73**
Total	**0.63**	**0.72**	**0.69**	**0.67**	**0.75**	**0.70**

Out-Sample Comparisons

Expiry	deep out	out	at	in	deep in	Total
JGARCH vs NGARCH						
<30	**0.47**	**0.66**	0.51	**0.44**	**0.60**	**0.53**
30-60	**0.54**	**0.69**	**0.54**	**0.44**	**0.60**	**0.57**
61-90	**0.58**	**0.63**	0.53	0.48	**0.60**	**0.58**
>90	**0.64**	**0.59**	0.51	0.44	**0.57**	**0.60**
Total	**0.56**	**0.67**	**0.53**	**0.44**	**0.59**	**0.57**
JGARCH vs IG-GARCH						
<30	**0.58**	**0.71**	**0.63**	0.49	**0.66**	**0.62**
30-60	**0.61**	**0.76**	**0.70**	**0.57**	**0.67**	**0.65**
61-90	**0.63**	**0.72**	**0.70**	**0.63**	**0.66**	**0.66**
>90	**0.71**	**0.80**	**0.74**	**0.71**	**0.64**	**0.69**
Total	**0.63**	**0.74**	**0.69**	**0.57**	**0.66**	**0.65**
JGARCH vs Merton						
<30	**0.67**	**0.87**	**0.82**	**0.81**	**0.85**	**0.78**
30-60	**0.71**	**0.86**	**0.82**	**0.82**	**0.87**	**0.80**
61-90	**0.75**	**0.81**	**0.81**	**0.83**	**0.88**	**0.81**
>90	**0.73**	**0.83**	**0.81**	**0.70**	**0.84**	**0.78**
Total	**0.72**	**0.85**	**0.82**	**0.81**	**0.87**	**0.80**
NGARCH vs IG-GARCH						
<30	**0.61**	**0.63**	**0.66**	**0.57**	**0.46**	**0.57**
30-60	**0.63**	**0.78**	**0.71**	**0.65**	**0.54**	**0.63**
61-90	**0.65**	**0.75**	**0.72**	**0.66**	**0.54**	**0.63**
>90	**0.65**	**0.80**	**0.76**	**0.78**	**0.58**	**0.65**
Total	**0.64**	**0.74**	**0.70**	**0.64**	0.53	**0.62**
NGARCH vs Merton						
<30	**0.70**	**0.86**	**0.83**	**0.80**	**0.67**	**0.74**
30-60	**0.70**	**0.85**	**0.83**	**0.81**	**0.75**	**0.76**
61-90	**0.73**	**0.81**	**0.81**	**0.82**	**0.81**	**0.78**
>90	**0.65**	**0.85**	**0.84**	**0.76**	**0.83**	**0.74**
Total	**0.70**	**0.84**	**0.82**	**0.80**	**0.76**	**0.76**
IG-GARCH vs Merton						
<30	**0.65**	**0.89**	**0.85**	**0.82**	**0.72**	**0.75**
30-60	**0.67**	**0.85**	**0.85**	**0.80**	**0.77**	**0.75**
61-90	**0.74**	**0.79**	**0.75**	**0.76**	**0.80**	**0.76**
>90	**0.65**	**0.78**	**0.71**	**0.56**	**0.75**	**0.69**
Total	**0.68**	**0.84**	**0.82**	**0.78**	**0.77**	**0.75**