CREATES
Center for Research in Econometric Analysis of Time Series

CREATES Research Paper 2008-41

American Option Pricing using GARCH models and the Normal Inverse Gaussian distribution

Lars Stentoft

School of Economics and Management
University of Aarhus
Building 1322, DK-8000 Aarhus C
Denmark

Aarhus School of Business
University of Aarhus
Handelshøjskolen
Aarhus Universitet

UNIVERSITY OF COPENHAGEN

American Option Pricing using GARCH models and the Normal Inverse Gaussian distribution*

Lars Stentoft†

HEC Montréal and CREATES

August 19, 2008

Abstract

In this paper we propose a feasible way to price American options in a model with time varying volatility and conditional skewness and leptokurtosis using GARCH processes and the Normal Inverse Gaussian distribution. We show how the risk neutral dynamics can be obtained in this model, we interpret the effect of the riskneutralization, and we derive approximation procedures which allow for a computationally efficient implementation of the model. When the model is estimated on financial returns data the results indicate that compared to the Gaussian case the extension is important. A study of the model properties shows that there are important option pricing differences compared to the Gaussian case as well as to the symmetric special case. A large scale empirical examination shows that our model outperforms the Gaussian case for pricing options on three large US stocks as well as a major index. In particular, improvements are found when considering the smile in implied standard deviations.

JEL Classification: C22, C53, G13

Keywords: GARCH models, Normal Inverse Gaussian distribution, American Options, Least Squares Monte Carlo method.

*The author thanks the Associate Editor and two anonymous referees from the *Journal of Financial Econometrics* for valuable comments and suggestions. The author also thanks participants at the 3rd Nordic Econometric Meeting and the Montréal Actuarial and Financial Mathematics Day, participants and discussants at the 2006 SCSE and NFA meetings, as well as seminar participants at HEC Montréal and Brock University. Financial support from the Danish Social Science Research Council (Grant 95015328), from CREATES, funded by the Danish National Research Foundation, and from Institut de Finance Mathématique de Montréal (IFM2) is gratefully appreciated.

†Address correspondance to Lars Stentoft, 3000 chemin de la Cote-Sainte-Catherine, Montreal (Quebec) Canada H3T 2A7, or e-mail: lars.stentoft@hec.ca.

1 Introduction

A recent application of the generalized autoregressive conditional heteroskedasticity, or GARCH, models of Engle (1982) and Bollerslev (1986) has been in the context of derivative pricing. The obvious reason for this is that for this particular type of assets correct estimation and modelling of the volatility of the underlying asset return process is of paramount importance. Since the GARCH models offer a very flexible framework for this it constitutes an obvious extension to the constant volatility framework of Black & Scholes (1973) and Merton (1973). In terms of option pricing the added flexibility comes at a cost since with time varying volatility the market is no longer complete. However, in Duan (1995) a GARCH option pricing model is derived under the assumption of conditionally Gaussian distributed innovations and under some familiar assumptions on investor preferences. The paper also shows that the implied volatility backed out from option prices generated with a Gaussian GARCH model shows patterns similar to those found empirically, and the Gaussian GARCH option pricing model could thus potentially explain the mispricings found empirically when e.g. constant volatility models are used.

When the Gaussian GARCH models are compared to e.g. the constant volatility model smaller pricing errors are obtained empirically. In particular, this is found for European style options on the Standard & Poor's 500 Index in Bollerslev & Mikkelsen (1996), Bollerslev & Mikkelsen (1999), Heston & Nandi (2000), Christoffersen & Jacobs (2004), and Hsieh & Ritchken (2005). In addition to this the GARCH models are found to perform well for European style options on the German DAX index in Härdle & Hafner (2000), on the Hang Seng Index in Duan & Zhang (2001), and on the FTSE 100 Index in Lehar, Scheicher & Schittenkopf (2002). More recently the GARCH option pricing model has been used to price American style options using an application of the Least Squares Monte-Carlo method of Longstaff & Schwartz (2001) to take account of the early exercise feature in Stentoft (2005). Again the finding is that GARCH models provide smaller pricing errors than the constant volatility alternative for options on a US index and three major individual stocks. However, this last paper also shows that systematic mispricing still occurs when Gaussian GARCH models are used empirically for option pricing. In particular, it is found that short term deep out of the money options remains underpriced by the model even when time varying volatility is considered.

One obvious extension is to consider models which are driven by non Gaussian innovations. In the vast empirical finance literature such models are well known within the GARCH framework where alternative assumptions on the conditional distribution have been suggested and extensively analyzed. The best known examples are probably the use of the Student's t-distribution in Bollerslev (1987) and of the Generalized Error distribution in Nelson (1991). Both of these distributions can potentially account for the excess kurtosis often found in the standardized residuals from the GARCH models, and a general finding in these studies is that financial return distributions indeed do have fatter tails than the Gaussian distribution. In addition to these leptokurtic distributions various skewed distributions, like e.g. the Skewed Student's t-distribution, have also been used (see e.g. Lambert & Laurent (2001) which also lists a number of different contributions in this area). While the empirical results are less abundant it does appear that skewness is potentially important for financial asset return processes.

The theoretical foundation for option pricing in this more general framework was provided in Duan (1999) which extends the Gaussian GARCH option pricing model to situations with conditional leptokurtic distributions. In particular, the choice of distribution in that paper is the Generalized Error Distribution and the paper examines the effect on the model predicted option prices compared to what is obtained in the Gaussian case. The results of the paper indicate that conditional leptokurtosis plays an important role in determining the option prices. These findings are confirmed in Hafner & Herwartz (2001) which uses a Student's t-distribution as the conditional distribution in the GARCH framework and which obtains qualitatively the same results. However, empirical applications of such models have been very limited most likely because of the added complexity of the framework.

In this paper we propose a feasible way to price American options in a framework with time varying volatility and conditional skewness and leptokurtosis based on the framework of Duan (1999), we analyze the properties of the model, and we provide a large scale empirical application to individual stock options. In particular, the first objective of this paper is to extend the application of the GARCH option pricing model from Duan (1999) to the situation where the conditional distribution is allowed to be skewed in addition to being leptokurtic. To achieve this we choose to work with the Normal Inverse Gaussian distribution which can accommodate both of these features. This distribution was introduced in the finance literature recently and used together with GARCH models in e.g. Barndorff-Nielsen (1997), Andersson (2001), and Jensen & Lunde (2001). The model is referred to as the NIG GARCH model

The second objective of the paper is to discuss in detail how the NIG GARCH option pricing model can be implemented numerically in a computationally efficient way since the model involves some additional complexities. The method we proposed extends and simplifies the methods suggested in Duan (1999). In particular, our proposed method avoids the use of Monte Carlo simulation to evaluate particular functions and it exploits the use of simple linear approximations. This ensures that the method is simple to implement and because of the approximations also relatively fast. We also reiterate how a simulation approach can be used to price options with early exercise features using the approach suggested in Longstaff & Schwartz (2001) which was used in a Gaussian GARCH context in Stentoft (2005). Since this style of options occurs in many situation and since both the most liquid index options and all individual options traded on the Chicago Board Options Exchange are American style options this is of great importance. However, the added flexibility of the NIG GARCH model would also have potentially important effects on the valuation of European options.

Based on the NIG GARCH option pricing framework and the procedure for the numerical implementation, the third objective of the paper is to analyze the properties of the proposed model. In particular, since the parameters of the distribution can be linked to the conditional skewness and leptokurtosis in an intuitive way it becomes possible to gauge the importance of either of these distributional features for option pricing purposes. This can be done simply by comparing the model prices to those obtained under a Gaussian alternative. The analysis indicates that by allowing the model to be governed by a skewed and leptokurtic conditional distribution added flexibility is provided compared to that of the Gaussian alternative.

Finally, since it appears that the proposed model can potentially explain the shortcomings of alternative models the paper applies the model empirically. To be specific, the pricing performance for options written on three major US stocks and the Standard & Poor's 100 Index is analyzed. The results obtained show that improvements are found compared to the Gaussian special case. In particular, we find large improvements in the pricing performance for out of the money options in general when compared to a Gaussian alternative. Moreover, the NIG GARCH models provide some reconciliation of the smile in the implied standard deviation from actual option prices. This is particularly so for the individual stock options.

It should be noted that while we use historical data on the underlying asset to obtain the parameters to be used in option valuation these parameters could potentially be backed out from historical option data also. In particular, this procedure which consists of calibrating the option pricing model to existing data is used in Heston & Nandi (2000), Christoffersen & Jacobs (2004), and Hsieh & Ritchken (2005) among others for the Gaussian case. In Christoffersen, Heston & Jacobs (2006) and Lehnert (2003) GARCH type models with skewed and leptokurtic innovations are used although neither of these uses the exact framework of Duan (1999). However, in all these cases the options under consideration are European style and the models considered are limited to those which allow for approximately closed form pricing formulas. Such formulas are unfortunately not available for the American style options and for this reason the calibration procedure would be much more computationally complex. However, using the efficient algorithms developed here it is certainly possible to price these type of options also and we expect that similar improvements can be obtained.

Finally, we note that in addition to the proposed discrete time GARCH models other models with time varying volatility have been proposed. In particular, within the context of derivatives pricing the continuous time Stochastic Volatility (SV) models are important and models with non Gaussian distributions have been derived recently (see e.g. Barndorff-Nielsen & Shephard (2001), Carr & Wu (2004), Eberlein (2001), and Huang & Wu (2004)). One of the main advantages the continuous time SV framework offers over the discrete time GARCH framework is that in some cases analytical formulas exist for European style derivatives. However, again this is unfortunately not the case for the American style derivatives considered here. Furthermore, while the GARCH framework only depends on the observable return process the application of any SV model would, in addition, require the unobserved volatility as a state variable. This complicates not only the estimation procedure but also the actual option pricing procedure. For these reasons the present paper focuses on the discrete time GARCH framework instead of the continuous time SV framework.

The rest of the paper is organized as follows: Section 2 presents the discrete time NIG GARCH model which allows for conditional skewness and excess kurtosis in addition to time varying volatility. In this section the appropriate model to be used for option pricing is also derived. In Section 3 we discuss how the model may be implemented numerically using linear approximations and we provide maximum likelihood estimation results. Section 4 explains how options may be priced using simulation techniques. This section also contains a study of the potential gains obtained with the NIG GARCH model compared to the Gaussian GARCH benchmark in terms of option pricing. Section 5 presents the results of our empirical application and documents the improvements obtained with the proposed model. Finally, Section 6 offers concluding remarks and some directions for future research.

2 The skewed and leptokurtic generalized GARCH framework

In this paper a discrete time economy is considered. We denoted by S_t the asset prices, by d_t dividends, and by $R_t = \ln\left(\frac{S_t}{S_{t-1}} + d_t\right)$ the continuously compounded return process. We assume that the return process can be modelled using what we will refer to as a generalized GARCH framework. In particular, we specify the dynamics for R_t as

$$R_t = m_t\left(\cdot;\theta_m\right) + \sqrt{h_t}\varepsilon_t \tag{1}$$

$$h_t = g\left(h_s, \varepsilon_s; -\infty < s \leq t-1; \theta_h\right) \tag{2}$$

$$\varepsilon_t | \mathcal{F}_{t-1} \sim D\left(0, 1; \theta_D\right), \tag{3}$$

where $\mathcal{F}_{t-1}$ is the information set containing all information up to and including time $t-1$.

In equation (1) we use $m_t\left(\cdot;\theta_m\right)$ to denote the conditional mean, which is allowed to be governed by a set of parameters θ_m provided that the process is measurable with respect to the information set $\mathcal{F}_{t-1}$. Likewise, in (2) the parameter set θ_h governs the variance process. This process is allowed to depend on lagged values of the innovations to the return process, lagged values of the volatility itself, as well as transformations hereof. Finally, in (3) we use $D\left(0,1;\theta_D\right)$ to denote a zero mean and unit variance distribution function which is also allowed to depend on a set of parameters θ_D. For notational convenience, in the following we let θ denote the set of all parameters in θ_m, θ_h, and θ_D.

The above framework is very general and nests many well known GARCH specifications. First of all, we observe that the specification in (1) allows for quite general forms of the conditional mean. In particular, for the GARCH in Mean specification suggested in e.g. Duan (1995) the parameter set is given by $\theta_m = \left\{r, \lambda, -\frac{1}{2}\right\}$ and the functional form can be specified as

$$m_t\left(\cdot;\theta_m\right) = r + \lambda\sqrt{h_t} - \frac{1}{2}h_t, \tag{4}$$

where r is the risk free interest rate and $-\frac{1}{2}h_t$ is the well known correction for continuously compounded return in a Gaussian world. A particular nice feature of this specification is that the parameter λ may be interpreted as the unit risk premium.

Secondly, in order to relate the flexible variance specification in (2) to some existing models we may start by noting that for the well known GARCH specification θ_h consists of the set of parameters $\{\omega, \beta, \alpha\}$ and (2) would have the following functional form:

$$h_t = \omega + \beta h_{t-1} + \alpha h_{t-1} \varepsilon_{t-1}^2. \tag{5}$$

Obviously, using more lags can be considered as simple extensions. Alternatively, we may wish to consider specifications which can potentially accommodate asymmetric responses to negative and positive return innovations. Such models are generally said to allow for a leverage effect, which refers to the tendency for changes in stock prices to be negatively correlated with volatility.

One such extension to the GARCH model considered in (5) is the non-linear asymmetric GARCH model, or NGARCH, of Engle & Ng (1993). The particular specification of the variance process for this model is given by

$$h_t = \omega + \beta h_{t-1} + \alpha h_{t-1} \left(\varepsilon_{t-1} + \gamma\right)^2, \tag{6}$$

and for this specification we have $\theta_h = \{\omega, \alpha, \beta, \gamma\}$. In the NGARCH model the leverage effect is modelled through the parameter γ, and if $\gamma < 0$ leverage effects are said to be found. It is clear that this model nests the ordinary GARCH specification, which obtains when $\gamma = 0$, and the model thus allows us to compare the contribution of the leverage effect directly by comparison to the GARCH specification.

2.1 The NIG GARCH model

While the generalized GARCH framework can be used with a host of conditional distributions like the Generalized Error Distribution or Student's t-distribution, in the present paper we chose to use the Normal Inverse Gaussian, or NIG for short, distribution. Following Barndorff-Nielsen (1997) and Jensen & Lunde (2001) the $NIG(\mu, \delta, a, b)$ distribution can be defined in terms of the so called location and scale invariant parameters as

$$f_{NIG}(x; \mu, \delta, a, b) = \frac{a}{\pi\delta} \exp\left(\sqrt{a^2 - b^2} + b\frac{(x-\mu)}{\delta}\right) q\left(\frac{x-\mu}{\delta}\right)^{-1} K_1\left(aq\left(\frac{x-\mu}{\delta}\right)\right), \tag{7}$$

where $q(z) = \sqrt{1+z^2}$ and $K_1(\cdot)$ is the modified Bessel function of third order and index 1. For the distribution to be well defined we obviously need to ensure that $0 \leq |b| \leq a$ and $0 < \delta$.

In (7), μ is a location parameter and δ is a scale parameter, and if we define $\rho = b/a$ the mean and variance of a $NIG(\mu, \delta, a, b)$ distributed variable are given as

$$E(X) = \mu + \frac{\rho\delta}{\sqrt{1-\rho^2}}, \text{ and} \tag{8}$$

$$Var(X) = \frac{\delta^2}{a\left(1-\rho^2\right)^{3/2}}. \tag{9}$$

Thus, since the distribution is a location-scale family, that is $X \sim NIG(\mu, \delta, a, b) \Leftrightarrow \frac{X-\mu}{\delta} \sim NIG(0, 1, a, b)$, a zero mean and unit variance NIG distributed variable can be obtained by restricting μ and δ to

$$\mu = \frac{-\rho\delta}{\sqrt{1-\rho^2}}, \text{ and} \tag{10}$$

$$\delta = \sqrt{a\left(1-\rho^2\right)^{3/2}}. \tag{11}$$

In the following we will refer to this standardized distribution as the $NIG(a, b)$ distribution and this distribution will be used as the zero mean unit variance distribution in (3).

From Jensen & Lunde (2001) we also have expressions for the skewness and kurtosis of the standardized $NIG(a, b)$ distribution which are given by

$$Skew(X) = \frac{3\rho}{\sqrt{a}\sqrt[4]{1-\rho^2}}, \text{ and} \tag{12}$$

$$Kurt(X) = 3\left(1 + \frac{4\rho^2+1}{a\sqrt{1-\rho^2}}\right). \tag{13}$$

Thus, it is is clear that neither skewness nor kurtosis is affected by the location and scale parameters which explain why the a and b parameters are referred to as being location and scale invariant. From (12) and (13) it is also seen that we may interpret a and b as shape parameters with a determining the leptokurtosis and b the asymmetry. In particular, the distribution is symmetric when $b = 0$. Furthermore, as much leptokurtosis as desired may be accommodated within the standardized $NIG(a, b)$ distribution as a tends to zero. On the other hand, in the limit as a tends to infinity the Gaussian distribution is obtained.

It should be noted that the NIG distribution has been used in a GARCH framework previously in e.g. Forsberg & Bollerslev (2002) and Jensen & Lunde (2001). However, these applications can all be thought of as special cases of the generalized framework above. In particular, with $b = 0$ the restrictions in (10) and (11) simplify to yield $\mu = 0$ and $\delta = \sqrt{a}$. If it is further assumed that the mean specification in (1) is identically equal to zero and the volatility specification in (2) is given by $h_t = \omega + \beta h_{t-1} + \alpha R_{t-1}^2$ we essentially obtain the model used for the ECU/USD exchange rate in Forsberg & Bollerslev (2002). Likewise, the model used in Jensen & Lunde (2001) is obtained by restricting the mean equation in (1) to be given by $R_t = \mu + \lambda\sqrt{h_t} + \sqrt{h_t}\varepsilon_t$ which is obtained in the above framework when $m_t(\cdot; \theta_m) = \mu + \lambda\sqrt{h_t}$.

2.2 Option pricing with the NIG GARCH model

In Duan (1999) the Local Risk Neutral Valuation Relationship, or LRNVR for short, of Duan (1995) is generalized to situations where the innovations to the asset return process are potentially non Gaussian. In particular, the Generalized LRNVR, or GLRNVR for short, applies if this distribution is skewed and leptokurtic. In our setting this corresponds to the general situation when using the $NIG(a, b)$ distribution in (3) with b different from zero and a finite. However, the relationship obviously also holds when we consider the limiting case of a Gaussian distribution.

Sufficient conditions for the GLRNVR to hold are stated in Proposition 3 in Duan (1999) and involve restrictions on the utility function of the representative agent. Under these assumptions Proposition 4 in the paper states that the risk neutralized dynamics of the system in (1) – (3) are given as

$$R_t = m_t(\cdot;\theta_m) + \sqrt{h_t}\varepsilon_t \tag{14}$$

$$h_t = g(h_s, \varepsilon_s; -\infty < s \leq t-1, \theta_h) \tag{15}$$

$$\varepsilon_t = F_{NIG}^{-1}[\Phi(Z_t - \lambda_t)], \tag{16}$$

where Z_t, conditional on $\mathcal{F}_{t-1}$, is a standard Gaussian variable under the risk neutral measure $\mathcal{Q}$, and where λ_t is the solution to

$$E^{\mathcal{Q}}\left[\exp\left(m_t(\cdot;\theta_m) + \sqrt{h_t}F_{NIG}^{-1}[\Phi(Z_t - \lambda_t)]\right)\Big|\mathcal{F}_{t-1}\right] = \exp(r_t). \tag{17}$$

In the above equation r_t denotes the one period risk free interest rate at time t, and although this rate has to be deterministic it may in fact be time varying. Note that the same mean is used in the risk neutral process as in (1) and instead the risk neutralization is obtained through a change in the innovation term as specified in (16).

In the pricing system described by (14) – (17) above F_{NIG}^{-1} denotes the inverse cumulative distribution function associated with the standardized $NIG(a,b)$ distribution, whereas Φ denotes the standard Gaussian cumulative distribution function. Thus, in the limit when a tends to infinity such that the underlying distribution corresponds to the Gaussian it follows that $F_{NIG}^{-1}[\Phi(z)] = z$, for any z, and in this situation the innovations in the risk neutral world remain Gaussian although with non zero mean. However, when the underlying distribution departs from the Gaussian the transformation in $F_{NIG}^{-1}[\Phi(z)]$ yields innovations under the risk neutral measure with the appropriate properties to be used when pricing e.g. options. In particular, note that when $\lambda_t = 0$ the innovations in (16) corresponds to random draws from the $NIG(a,b)$ distribution.

2.2.1 The restriction to conditional normality

In the limiting special case of conditional Gaussianity it may be observed immediately that the restriction in (17) simplifies to

$$\exp\left(m_t(\cdot;\theta_m) - \lambda_t\sqrt{h_t} + \frac{1}{2}h_t\right) = \exp(r_t),$$

since $m_t(\cdot;\theta_m)$ is measurable with respect to the time $t-1$ information set by assumption and since Z_t is a standard Gaussian variable conditional on the information set $\mathcal{F}_{t-1}$ under the risk neutral probability measure $\mathcal{Q}$. Thus, upon rearranging it follows that

$$\lambda_t = \frac{m_t(\cdot;\theta_m) + \frac{1}{2}h_t - r_t}{\sqrt{h_t}}, \tag{18}$$

is the solution required in (17). It is thus immediately seen that λ_t is a simple function of the mean specification.

2.2.2 Interpreting the price of risk in the general case

In the general case we cannot expect to find such a simple expression for λ_t as the one given in (18). However, we may start out by noting that since $m_t(\cdot;\theta_m)$ is measurable with respect to the time $t-1$ information set by assumption we may rewrite (17) as

$$m_t(\cdot;\theta_m) + \ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t} F_{NIG}^{-1}\left[\Phi\left(Z_t - \lambda_t\right)\right]\right)\Big|\, \mathcal{F}_{t-1}\right] = r_t. \tag{19}$$

Next, for simplicity define $\Lambda(Z_t - \lambda_t) \equiv \ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t} F_{NIG}^{-1}\left[\Phi\left(Z_t - \lambda_t\right)\right]\right)\big|\, \mathcal{F}_{t-1}\right]$. The question then becomes one of approximating this function. Here we proceed using a first order Taylor series expansion around Z_t to obtain the following expression

$$\begin{aligned}
\Lambda(Z_t - \lambda_t) &\simeq \Lambda(Z_t) + \Lambda'(Z_t)(Z_t - \lambda_t - Z_t) \\
&= \Lambda(Z_t) - \lambda_t \Lambda'(Z_t).
\end{aligned}$$

We now analyze each of these terms.

First of all, by definition the first term is given by $\Lambda(Z_t) = \ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t} F_{NIG}^{-1}\left[\Phi\left(Z_t\right)\right]\right)\big|\, \mathcal{F}_{t-1}\right]$. However, since Z_t, conditional on $\mathcal{F}_{t-1}$, is a standard Gaussian variable under the risk neutral measure $\mathcal{Q}$, the transformation $X_t = F_{NIG}^{-1}\left[\Phi\left(Z_t\right)\right]$ yields an $NIG(a,b)$ distributed variable. Hence, it may be realized that $\Lambda(Z_t)$ in fact corresponds to the cumulant generating function $K_{NIG}(t) = \ln E^{\mathcal{Q}}\left[\exp(tX_t)\right|\mathcal{F}_{t-1}]$ of an $NIG(a,b)$ distributed variable evaluated at $t = \sqrt{h_t}$. For the NIG distribution, as it is the case with the Gaussian distribution, closed form expressions exist for this function which can be evaluated explicitly. In particular, for the Gaussian case this function is given by $K_N(t) = \frac{1}{2}t^2$.

Secondly, we examine the term involving the first derivative of $\Lambda(Z_t)$. Using the chain rule and assuming sufficient regularity so that we can change the order of differentiation and integration we obtain

$$\begin{aligned}
\Lambda'(Z_t) &= \frac{1}{E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t}X_t\right)|\mathcal{F}_{t-1}\right]} E^{\mathcal{Q}}\left[\sqrt{h_t}\exp\left(\sqrt{h_t}X_t\right)\frac{\partial F_{NIG}^{-1}\left[\Phi\left(Z_t\right)\right]}{\partial Z_t}|\mathcal{F}_{t-1}\right] \\
&= \sqrt{h_t}E^{\mathcal{Q}}\left[\frac{\partial F_{NIG}^{-1}\left[\Phi\left(Z_t\right)\right]}{\partial \Phi\left(Z_t\right)}\frac{\partial \Phi\left(Z_t\right)}{\partial Z_t}|\mathcal{F}_{t-1}\right] = \sqrt{h_t}E^{\mathcal{Q}}\left[\frac{f\left(\Phi\left(Z_t\right)\right)}{f_{NIG}\left(F_{NIG}^{-1}\left[\Phi\left(Z_t\right)\right]\right)}|\mathcal{F}_{t-1}\right].
\end{aligned}$$

Since Z_t, conditional on $\mathcal{F}_{t-1}$, is a standard Gaussian variable under the risk neutral measure $\mathcal{Q}$, $\Phi(Z_t)$ corresponds to a particular percentile of this distribution. Hence, the term in square brackets essentially corresponds to the ratio between two densities evaluated at the same percentiles. Although a closed form expression may not exist for the expected value of this ratio, for given values of a and b in the NIG distribution this is easily approximated.

Finally, putting it all together we obtain the following approximate specification of λ_t in the NIG GARCH framework

$$\lambda_t \simeq \frac{m_t(\cdot;\theta_m) + K_{NIG}\left(\sqrt{h_t}\right) - r_t}{\sqrt{h_t}} E^{\mathcal{Q}}\left[\frac{f\left(\Phi\left(Z_t\right)\right)}{f_{NIG}\left(F_{NIG}^{-1}\left[\Phi\left(Z_t\right)\right]\right)}|\mathcal{F}_{t-1}\right]^{-1}. \tag{20}$$

Thus, in general the same structure applies to λ_t as in the Gaussian special case shown in (18) and it can be given the same overall interpretation. However, because of the departure from a Gaussian distribution an extra term is present. While this expectation approaches unity as a tends to infinity, for the $NIG\,(a,b)$ distribution used here it is less than one and hence λ_t is scaled upward. This may be interpreted as a reward required by investors for holding assets with non Gaussian payoffs.

2.3 A feasible specification of the NIG GARCH option pricing model

In principle, the system above is completely self-contained. However, when it comes to actually implementing it problems may occur as a result of the requirement that λ_t be the solution to (17) since this may be extremely difficult to solve. In particular, a simple functional relationship like the one obtained in (18) for λ_t, given r_t, $m_t\,(\cdot;\theta_m)$, and h_t, in the Gaussian special case will not be available as (20) indicates.

In the present paper we turn the problem slightly around and use the property that even in this situation there is a one to one correspondence between potential λ_t's and potential mean specifications $m_t\,(\cdot;\theta_m)$ even though this is no longer simple. In particular, we may rewritten (19) as

$$m_t\,(\cdot;\theta_m) = r_t - \ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t}F_{NIG}^{-1}\left[\Phi\left(Z_t-\lambda_t\right)\right]\right)\middle|\,\mathcal{F}_{t-1}\right]. \tag{21}$$

It follows that instead of solving (17) for λ_t given r_t, $m_t\,(\cdot;\theta_m)$, and h_t, we may specify λ_t directly and "imply" the corresponding mean specification, $m_t\,(\cdot;\theta_m)$. It is immediately clear that this suggestion potentially complicates the estimation procedure since (21) would have to be evaluated numerically in search of the parameter estimates when estimating the system in (1) – (3). In fact, in the numerical part of Duan (1999) it is argued that this procedure becomes extremely demanding in terms of computing time.

In the present paper we provide results showing that the procedure of implying the solution to (21) is in fact feasible and the estimation procedure does not necessarily become computationally excessively demanding since this expectation can be calculated easily following the procedures outlined below in Section 3.1. While it is possible to specify λ_t in a quite flexible manner generally a choice has to be made about the specific structure of λ_t. In the present paper we will assume that $\lambda_t = \lambda$ for any t, that is constant through time. Furthermore, we will assume that the interest rate is constant, that is $r_t = r$ for any t also. These two assumptions clearly simplify the implementation of the model significantly and serve as obvious first benchmarks.

2.3.1 The restriction to conditional normality

While the restriction to a constant λ and r not only serves as a relevant first approximation in the general framework it also implies that we obtain the specification suggested in Duan (1995) in the Gaussian case. In particular, from (18) we see that when λ is specified an explicit form for the mean may be derived under the assumption of Gaussianity. To be specific, what is obtained is the relationship in (4) which may be substituted directly into the system in (1) – (3). In this case the system to be estimated simply becomes

$$R_t = r + \lambda\sqrt{h_t} - \frac{1}{2}h_t + \sqrt{h_t}\varepsilon_t \tag{22}$$

$$h_t = g(h_s, \varepsilon_s; -\infty < s \leq t-1, \theta_h) \tag{23}$$

$$\varepsilon_t | \mathcal{F}_{t-1} \sim N(0,1) \tag{24}$$

and from this we see that in this situation λ may be easily interpreted as the unit risk premium.

Moreover, by substituting (4) into (14) – (16) it follows that the system to be used for pricing purposes becomes

$$R_t = r + \lambda\sqrt{h_t} - \frac{1}{2}h_t + \sqrt{h_t}\varepsilon_t \tag{25}$$

$$h_t = g(h_s, \varepsilon_s; -\infty < s \leq t-1, \theta_h) \tag{26}$$

$$\varepsilon_t = (Z_t - \lambda), \tag{27}$$

where Z_t, conditional on $\mathcal{F}_{t-1}$, is a standard Gaussian variable under the risk neutral measure $\mathcal{Q}$. When the volatility dynamics are assumed to be of the simple GARCH form this limiting specification in fact corresponds exactly to the specification used in Duan (1995). Thus, this model will serve as an obvious benchmark for this paper.

2.3.2 Risk neutral innovations in the general case

As shown above the riskneutral innovations remain Gaussian, although with a non-zero mean, in the standard model. However, this is not generally the case. In particular, the riskneutral innovations in the NIG GARCH model are given by equation (16), and unless $\lambda = 0$ this will not yield $NIG(a,b)$ distributed innovations. The implications of λ may again be examined using a Taylor series expansion. In particular, using a first order approximation of $F_{NIG}^{-1}[\Phi(Z_t - \lambda)]$ around Z_t we obtain

$$\begin{aligned} F_{NIG}^{-1}[\Phi(Z_t - \lambda)] &\simeq F_{NIG}^{-1}[\Phi(Z_t)] + \frac{\partial F_{NIG}^{-1}[\Phi(Z_t)]}{\partial Z_t}(Z_t - \lambda - Z_t) \\ &= F_{NIG}^{-1}[\Phi(Z_t)] - \lambda \frac{f(\Phi(Z_t))}{f_{NIG}\left(F_{NIG}^{-1}[\Phi(Z_t)]\right)}, \end{aligned} \tag{28}$$

where Z_t, conditional on $\mathcal{F}_{t-1}$, is a standard Gaussian variable under the risk neutral measure $\mathcal{Q}$.

In the approximation above we immediately recognize the first term as the $NIG(a,b)$ distributed variable $X_t = F_{NIG}^{-1}[\Phi(Z_t)]$ from Section 2.2.2. Likewise, the last term essentially corresponds to the ratio between two densities evaluated at the same percentile. Up to an approximation error the risk neutral innovation is a combination of these two factors.

Figure 1 plots the logarithm of this ratio for two specific versions of the $NIG(a,b)$ distribution against the quantiles. From this it is seen that for Z_t's in the extremes the ratio generally exceeds one whereas for Z_t's close to the mean the ratio is less than one. On the other hand, in the Gaussian case the ratio is trivially equal to one. Thus, in a risk neutral world where investors require compensation for non Gaussian returns the compensation required in the return process when the distribution is leptokurtic is relatively more important when returns are extreme.

3 Implementation and estimation of the NIG GARCH option pricing model

One of the important properties of the GLRNVR framework of Duan (1999) is that a close link is provided between the observed asset return process and the process which has to be used for valuation of the corresponding options. To be specific, note that by substituting (21) into the system in (1) − (3) it is possible to obtain the following specification for the return process to be used for estimation:

$$R_t = r - \ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t}F_{NIG}^{-1}\left[\Phi\left(\varepsilon_t-\lambda\right)\right]\right)\Big|\,\mathcal{F}_{t-1}\right] + \sqrt{h_t}\varepsilon_t \tag{29}$$

$$h_t = g\left(h_s, \varepsilon_s; -\infty < s \leq t-1, \theta_h\right) \tag{30}$$

$$\varepsilon_t|\,\mathcal{F}_{t-1} \sim NIG\left(a,b\right). \tag{31}$$

Comparing this system to the one used for pricing in (14)−(16) it is immediately clear that it is in fact possible to estimate all the necessary parameters from the historical returns. Thus, one of the major strengths of the proposed NIG GARCH framework is that cumbersome calibration procedures involving matching model option prices to observed prices to derive the model parameters may be avoided.

However, before we can actually implement the NIG GARCH option pricing model we need to obtain procedures for evaluating the transformation of the random variables Z_t through $F_{NIG}^{-1}\left[\Phi\left(Z_t-\lambda\right)\right]$ as well as for evaluating the logarithm of the expectation of the scaled exponential value of this, that is $\ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t}F_{NIG}^{-1}\left[\Phi\left(Z_t-\lambda\right)\right]\right)\big|\,\mathcal{F}_{t-1}\right]$. Note though that such procedures would be needed even if we were to use a calibration based method. In the following sections we detail first how each of these expressions can be calculated and approximated and we provide estimation results for the suggested models.

3.1 Approximation procedures

While numerical methods for calculating the necessary functionals in the symmetric situation are provided in Duan (1999) the paper claims that especially for estimation purposes these become extremely demanding in terms of computational time. In this section we extend the methods of Duan (1999) to a situation with conditional skewness and in each case we explore the transformations for the standardized $NIG\left(a,b\right)$ distribution used in this paper. More importantly, however, we improve on the methods and we suggest methods for how to approximate the two functions efficiently. When it comes to actually implementing the algorithm these approximations are particularly important because they increase the computational speed.

3.1.1 Computing and approximating $F_{NIG}^{-1}[\Phi(Z_t - \lambda)]$

In Duan (1999) a numerical scheme for calculating $F_D^{-1}[x]$ is proposed, where D represent the chosen distribution. This method involves approximating the Cumulative Distribution Function, or CDF, $F_D[x_i]$ at a sequence of points $i = 1, 2, ...$ using the Probability Density Function, or PDF, f_D. The approximation is piecewise linear in the x's and the inverse value can therefore be found using the two values of x_i's surrounding any value of x. For any finite number of points the calculated value is an approximation to the true value of $F_D^{-1}[x]$ although it should converge to the true value as the number of points used in the approximation tends to infinity.

In the following we suggest to use a further approximation step which speeds up the calculations. This step is based on the observation that $F_{NIG}^{-1}[\Phi(Z_t - \lambda)]$ is a smooth function of Z_t and that it may therefore be well approximated using a low order polynomial in Z_t. The method we propose consists of the following three steps:

1. Calculate an approximation to the CDF of the chosen distribution using an extension of the procedure outlined in Duan (1999). In particular, since the distribution under consideration is skewed a choice has to made as to how the algorithm should be initialized. In the present paper this is done by selecting a large negative value and building the approximation from this. If the value chosen is large enough the CDF should be approximately zero. In order to improve on the approximation, however, a numerical integration of the PDF to this point is performed.

2. Evaluate the expression $F_{NIG}^{-1}[x]$ at a number of points equidistantly distributed between zero and one in probability according to the piecewise linear approximation obtained above. This corresponds to evaluating the function $F_{NIG}^{-1}[\Phi(y)]$ at the corresponding quantiles, y, of the Gaussian distribution. Hence it may be realized that it also corresponds to evaluating the expression $F_{NIG}^{-1}[\Phi(Z - \lambda)]$ at $Z = y + \lambda$.

3. Calculate an approximation to $F_{NIG}^{-1}[\Phi(Z - \lambda)]$ which may be realized to be a simple smooth function of Z. In particular, this can be done by regressing the values of $F_{NIG}^{-1}[\Phi(Z - \lambda)]$ from the step above on a set of transformations of the values of Z.

In Figure 2 the approximations to $F_{NIG}^{-1}[\Phi(Z - \lambda)]$ from step two above are plotted against the values of Z. In the procedure of Duan (1999) we use $5,000$ points each spaced with 0.005 symmetrically around zero thus covering the interval from -12.5 to 12.5 to approximate the CDF. The sample of Z equals the $9,999$ quantile values from the Gaussian distribution and $\lambda = 0.05$. The top panel of the figure shows the one to one relation between the Z's and $F_{NIG}^{-1}[\Phi(Z - \lambda)]$ in the Gaussian case. However, the figure also shows that the relationship remains smooth and well behaved even when leptokurtic or skewed distributions are used as is indicated by the middle and bottom panel which shows the relationship between the Z's and $F_{NIG}^{-1}[\Phi(Z - \lambda)]$ with the symmetric $NIG(2, 0)$ and the $NIG(2, 0.2)$ distributions respectively.

With these findings in mind we suggest that the function $F_{NIG}^{-1}[\Phi(Z-\lambda)]$ can be approximated well by a low order polynomial in Z. In our experience good approximations can be obtained with a fifth order polynomial approximation in the Z's, and this particular choice of approximation is used in the following. In particular, when performing this simple regression the ensuing R^2 is 100% indicating that all the variation is explained. The resulting predicted values are superimposed as a line in the two bottom plots.

When compared to the suggestion in Duan (1999) we note that the use of this approximation should speed up the computational work significantly because it involves less computational complexity. In particular, the approximation will have to be calculated only once for given values of the parameters of the chosen distribution. For the rest of the procedure simple linear operations are used. Thus, if a large number of evaluations is needed the computational complexity is essentially that of evaluating a fifth degree polynomial. This is, however, not the case if the approximation is not used since each evaluation involves a search among the values of the approximated CDF constructed in step 2. Since the computational complexity of searching among n elements is approximately proportional to $\log(n)$ this method will be dominated by that of using the approximation unless n is very low and hence the approximation to the CDF very coarse.

3.1.2 Computing and approximating $\ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t}F_{NIG}^{-1}[\Phi(Z_t-\lambda)]\right)\middle|\mathcal{F}_{t-1}\right]$

In Duan (1999) it is suggested that one can calculate the k'th element of the conditional expectation $E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t^{(k)}}F_{NIG}^{-1}[\Phi(Z_t-\lambda)]\right)\middle|\mathcal{F}_{t-1}\right]$ using a "vector" Monte Carlo simulation with N independent standard Gaussian random variables, $\left\{Z_t^i\right\}_{i=1}^N$, for which the inverse is calculated. Thus, the k'th element should be approximated by $\frac{1}{N}\sum_{i=1}^N \exp\left(\sqrt{h_t^{(k)}}F_{NIG}^{-1}\left[\Phi\left(Z_t^i-\lambda\right)\right]\right)$. However, since this procedure will have to be repeated for each value of h_t it will be extremely time consuming, and this even so if the above approximation to $F_{NIG}^{-1}[\Phi(Z_t-\lambda)]$ is used.

In the following we suggest an approximation which speeds up the calculations. The procedure exploits the fact that h_t is measurable with respect to the current information set and therefore the logarithm of the expectation, $\ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t}F_{NIG}^{-1}[\Phi(Z_t-\lambda_t)]\right)\middle|\mathcal{F}_{t-1}\right]$, may be considered a function of this variable. In particular, the expression turns out to be a relatively simple function of $\sqrt{h_t}$ which can be easily approximated using a low order polynomial in $\sqrt{h_t}$. The method we propose consists of the following three steps:

1. Select a number of appropriate values for h. This can be done either by selecting random numbers from an appropriate distribution or by simply choosing a range of potential variances. In particular, if we believe that variances are IG distributed a random draw from this distribution may be used or a number of quantiles may be used. Secondly, select appropriate numbers for Z. While these may also be selected as a random draw from the Gaussian distribution a more sensible choice is to use the quantiles from the approximation procedure in the previous section.

2. Calculate the product of each h and the vector of $F_{NIG}^{-1}[\Phi(Z-\lambda)]$ based on the approximation scheme

above and exponentiate the values. The logarithm of the mean of this is an estimate of the value of $\ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h}F_{NIG}^{-1}\left[\Phi\left(Z-\lambda\right)\right]\right)\middle|\mathcal{F}_{t-1}\right]$ for each h. In other words, we estimate the logarithm of the expectation by the logarithm of the mean of the exponential values of the product of the particular h and the approximation $F_{NIG}^{-1}\left[\Phi\left(Z-\lambda\right)\right]$.

3. Calculate an approximation to $\ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h}F_{NIG}^{-1}\left[\Phi\left(Z-\lambda\right)\right]\right)\middle|\mathcal{F}_{t-1}\right]$ which may be realized to be a relatively simple smooth function of $\sqrt{h}$. In particular, this can be done by regressing the values of $\ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h}F_{NIG}^{-1}\left[\Phi\left(Z-\lambda\right)\right]\right)\middle|\mathcal{F}_{t-1}\right]$ from the step above on a set of transformations of the values of $\sqrt{h}$.

In Figure 3 the approximations to the function $\ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h}F_{NIG}^{-1}\left[\Phi\left(Z-\lambda\right)\right]\right)\middle|\mathcal{F}_{t-1}\right]$ from step two above are plotted against values of $\sqrt{h}$. The sample of h equals the 999 quantile values from the Inverse Gaussian distribution with $a=2$ and $F_{NIG}^{-1}\left[\Phi\left(Z-\lambda\right)\right]$ is calculated as specified in the previous section. The top panel of the figure shows the approximation calculated in step two above for the Gaussian case together with the true relationship which is superimposed as a solid line. The panel shows that the approximation is very close to the exact expectation which is given as $\frac{1}{2}h-\lambda\sqrt{h}$. However, the figure also shows that the relationship remains smooth and well behaved even when leptokurtic or skewed distributions are used as is indicated by the middle and bottom panels which show the relationship between the $\sqrt{h}$'s and $\ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h}F_{NIG}^{-1}\left[\Phi\left(Z-\lambda\right)\right]\right)\middle|\mathcal{F}_{t-1}\right]$ with the symmetric $NIG\left(2,0\right)$ and the $NIG\left(2,0.2\right)$ distributions, respectively.

With these findings in mind we suggest that the function $\ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h}F_{NIG}^{-1}\left[\Phi\left(Z-\lambda\right)\right]\right)\middle|\mathcal{F}_{t-1}\right]$ can be approximated well by a low order polynomial in $\sqrt{h}$. In our experience good approximations can be obtained with a third order polynomial approximation in the $\sqrt{h}$'s, and this particular choice of approximation is used in the following. In particular, when performing this simple regression the ensuing R^2 is 100% indicating that all the variation is explained. The resulting predicted values are superimposed as a line in the two bottom plots.

Again we note that compared to the suggestion in Duan (1999) the use of the above approximation should speed up the computational work significantly. In particular, as it was the case with the inverse transformations the approximation will need to be calculated only once and for the rest of the procedure simple linear operations are used. This is, however, not the case if the approximation is not used since each evaluation would involve repeating step 2 above. Although these calculations could be based on the same values of the inverse approximations in $F_{NIG}^{-1}\left[\Phi\left(Z-\lambda\right)\right]$, they would still involve the use of the computationally expensive logarithmic and exponential functions a large number of times, which will be computationally inefficient.

3.2 Estimation Procedure

With the approximations outlined above it becomes possible to estimate the system in (29) – (31) using a Maximum Likelihood based approach. In this section we provide estimation results for three major US stocks, specifically General Motors (GM), International Business Machines (IBM), and Merck & Company Inc (MRK), as well as for the Standard & Poor's 100 Index in Tables 1 through 4. The data used for the individual stocks consists of time series of continuously compounded returns from July 1962 through 1995 which have been corrected for dividend payments. The source of the data is the 1997 Stock File from CRSP (see CRSP (1998)). The data used for the Standard & Poor's 100 Index spans the period from January 1973 through 1995 and we use the total return index including dividend payments obtained from Datastream. As the short rate, r, we take the 1 month LIBOR rate on December 29, 1995, at which time is was 5.4% annualized and continuously compounded. The data used corresponds to what was used in Stentoft (2005) with the sample period extended back from 1976 as long as possible given the available data.

In terms of estimation results, the Gaussian case serves as an appropriate benchmark for comparison since we know the exact form of the transformations. The results presented show that it is in fact possible to obtain parameter estimates which are statistically equal to those from the exact procedure in a matter of minutes. The estimation procedure is obviously more time consuming for the general case with $NIG(a, b)$ innovations. However, it remains feasible and results can be obtained in a matter of minutes on a standard laptop computer. To speed up the estimation process we apply the concept of variance targeting originally proposed in Engle & Mezrich (1996) which ensures that the implied unconditional level of the variance corresponds to the historical volatility actually observed. The procedure can be implemented in our framework simply by setting $\omega = Var(R_t) * \left(1 - \alpha\left(1 + \gamma^2\right) - \beta\right)$, with $\gamma = 0$ when the volatility is of the GARCH type and $\alpha = \beta = \gamma = 0$ in the constant volatility case.

3.2.1 Estimation results for the Gaussian GARCH model

Columns two, three, and four of the tables provide estimation results for the Gaussian special case using the estimation procedure derived in the present paper. In terms of the estimates relating to the variance specification we note that all the parameters are significant and that the asymmetry parameter in the NGARCH specification has the expected sign. We also note that allowing for GARCH volatility specifications is important and does lower the test statistics for ARCH and serial dependency in the standardized returns as well as the squared standardized returns.

However, the bottom part of the tables containing the usual test statistics also shows that a major problem with the Gaussian GARCH model is that the assumption of normality can be rejected for all series irrespective of the variance specification. This is evident from the J-B test statistics which are all significant at any reasonable level. More detailed tests, although not reported here, indicate that it is particularly in terms of excess kurtosis that departures are found although significant skewness is also found to be present in the standardized residuals for all the series but MRK. The problem with non Gaussianity is also evident from the left hand panels of Figure 4 which plots the quantiles of the standardized residuals from the Gaussian NGARCH model against the theoretical values. These plots clearly indicate that the standardized residuals are leptokurtic and may even be skewed.

3.2.2 Estimation results for the NIG GARCH model

Columns five and six, respectively seven and eight, of the four tables provide the corresponding estimation results for the symmetric as well as the general NIG GARCH models. We start out by examining the parameters related to the $NIG\,(a, b)$ distribution for which the tables clearly show that departures from the Gaussian distribution are important. In particular, the point estimate for the a parameter which governs the degree of leptokurtosis implies a conditional distribution which is quite leptokurtic for all the series. Furthermore, for most of the series the b parameter, which governs the skewness of the conditional distribution, is estimated significantly different from zero indicating that a skewed distribution is warranted by the data. The results in the tables also allow us to conclude that the symmetric although leptokurtic version of the NIG model which has $b = 0$ in general is too simple a model for the series considered here with the exception of the index series.

The importance of the distributional flexibility may also be analyzed using the QQ plots in Figure 4 plotting the quantiles of the standardized residuals from the various models using the NGARCH variance specification. These plots can be readily compared in terms of fit across distribution. When this is done it is seen that the fit is in general improved upon with the NIG GARCH models compared to the Gaussian version. In particular, the improvement is pronounced in the tails of the conditional distributions for the individual stock returns.

We next examine the estimates of the parameters related to the volatility specifications. With respect to these the tables show that the estimated α's and β's are quite stable across the underlying distribution. However, while the leverage parameter γ in the NGARCH model remains important the point estimate obtained once a fat tailed distribution is used is somewhat lower. This holds especially for the individual stock return series. The exception, however, to this is the index series for which it is found that allowing for fat tailed and skewed distributions does not lower the γ estimate of the leverage effect. Combining this observation with the fact that for this series the skewness parameter b in the general NIG GARCH model was estimated insignificantly different from zero indicates that these two effects are different.

Finally, when comparing the different NIG GARCH specifications to the equivalent Gaussian specifications the tables show large increases in the likelihood values when allowing for skewness and leptokurtosis indicating the importance of the extra modelling flexibility. However, a better comparison may be performed by comparing the Schwarz Information Criterion which is reported in the last line of each panel in each table. This criterion takes into consideration the additional number of parameters and the preferred model is the one with the minimum value. The panels show that for all the return series the model achieving the lowest value is in fact the NIG NGARCH model. We note that when only Gaussian models, respectively symmetric $NIG\,(a,0)$ models, are considered the NGARCH specification is also the preferred one for all series. On the other hand, when either of the volatility specifications is compared across distributions the general $NIG\,(a,b)$ version is always the preferred one.

Thus, in all respects our estimation results indicate a consistent ranking of the volatility specifications across the different distributions and of the distributions across different volatility specifications in favor of the general $NIG\,(a,b)$ distribution and the NGARCH volatility specification. To be specific, we obtain parameter estimates for the $NIG\,(a,b)$ distribution which are significantly different from the Gaussian special case, the fit for the standardized residuals is in general improved, and when comparing the Schwarz Information Criterion the NIG GARCH models are preferred to the Gaussian alternatives. Since these conclusions hold for all the returns series we may conclude that the skewed and leptokurtic NIG GARCH model is an important extension to the Gaussian GARCH model. This is so particularly for the individual stock return series considered in the present paper whereas a leptokurtic only symmetric $NIG\,(a,0)$ GARCH model may suffice for the index return series.

3.2.3 Estimates of the unit risk premium from returns data

The four tables also contain estimates of the λ parameter, which we in the Gaussian special case could interpret as the unit risk premium. It is well known that this premium is difficult to estimate from returns data alone and this is confirmed for the data considered here. In fact λ is estimated significantly different from zero only for MRK across all variance specifications and distributions. For the other three series the parameter is generally insignificantly estimated particularly in the most general specification with NGARCH variance and an $NIG\,(a,b)$ distribution. Because of the problems with estimating λ in the following we will assume that this is fixed at either $\lambda = 0$ or $\lambda = 0.025$ where the latter value corresponds to the average estimate for the NIG NGARCH model.

4 Pricing procedure for and properties of the NIG GARCH model

By substituting (21) into (14) – (16) it is possible to derive the following specification for the risk neutralized return process to be used for pricing purposes:

$$R_t = r - \ln E^{\mathcal{Q}} \left[\exp\left(\sqrt{h_t} F_{NIG}^{-1} \left[\Phi\left(Z_t - \lambda \right) \right] \right) \Big| \, \mathcal{F}_{t-1} \right] + \sqrt{h_t}\varepsilon_t, \tag{32}$$

$$h_t = g\left(h_s, \varepsilon_s; -\infty < s \leq t-1, \theta_h \right), \tag{33}$$

$$\varepsilon_t = F_{NIG}^{-1} \left[\Phi\left(Z_t - \lambda \right) \right]. \tag{34}$$

With the approximation schemes outlined above and with the appropriate parameter values this system yields the dynamics to be used to price options in an NIG GARCH framework. However, one potential problem with the framework is that no closed form solutions for the option prices exist even in the case where European options are considered. The reason is that in this general framework it is difficult if not impossible to obtain an analytical expression for the distribution of the future value of the underlying asset. Therefore, alternative numerical procedures are needed for option pricing applications.

In the following we describe a numerical procedure which can be readily used with the NIG GARCH option pricing model. The proposed method uses a simulation framework along the lines of that used in Stentoft (2005) which we generalize to the case of non Gaussian distributions. The framework is then used to study the properties of the NIG GARCH model, and the results are compared to the Gaussian case.

4.1 Pricing Procedure

While it is possible to apply different approaches to obtain option prices given the structure of the system in (32) − (34) the obvious choice in the present setting is to use a simulation based approach. In particular, random future paths can be easily simulated under the appropriate risk neutral dynamics according to this system. Once we have obtained a sample of paths the Least Squares Monte Carlo method, or LSM method, of Longstaff & Schwartz (2001) can be used to obtain the American option prices. Thus, the numerical procedure to be used involves a simulation step and a pricing step.

4.1.1 Simulation step

In the first step a large number of future paths for the underlying asset and for the conditional variances is simulated. In particular, this is done by successively applying the system in (32) − (34) until the maturity date of the option. Such a simulation can be easily performed for a large number of paths at each step with minimum added computational complexity.

Although the simulation involves transforming the Gaussian innovations, the Z's, at every step along all paths this is in fact feasible when the approximations from Section 3.1 are used. In particular, because the approximations need only be calculated once at the beginning of the simulation the computational complexity remains approximately linear in the number of paths and in the number of steps in the simulation.

4.1.2 Pricing step

Given the full sample of random paths the pricing step is initiated at the maturity date of the option. At this time it is possible to decide along each path if the option should be exercised since the future value trivially equals zero. Hence, the pathwise payoffs may be easily determined. Working backwards through time, however, this is not so. In particular, assuming exact knowledge of the future value of the underlying asset along each path and thereby perfect foresight would lead to biased estimates of the option prices. The

LSM method offers a solution to this which uses all the cross sectional information at a given time step in the simulation to estimate the expected value of holding the option for one more period conditional on the state. Since the estimates used are based on all the paths no assumption of pathwise perfect foresight is made and this mitigates the potential high bias.

To be specific, at any point in time where early exercise is to be considered a cross sectional regression is performed where the future values along each path are regressed on transformations of the current stock prices and volatility levels. The fitted values from this regression are then used as estimates of the pathwise conditional expected value of holding the option for one more period. Based on this a decision of whether to exercise or not can be made along each path by comparing the conditional expected value of continuing to hold the option to the value of immediate exercise. If immediate exercise yields superior payoff this is the optimal choice along this path. Once this decision has been recorded along each path we can move back through the simulation to the previous early exercise point and perform a new cross sectional regression with the appropriate values based on the previous choices.

The validity of using the LSM simulation approach in a GARCH framework was established in Stentoft (2005) in which the method is compared to the Markov Chain method of Duan & Simonato (2001) and to the lattice based algorithm of Ritchken & Trevor (1999). That paper shows that appropriate American option prices can be obtained by using a simple second order polynomial in the stock price and in the level of the volatility together with the cross product and a constant term in the cross sectional regression in the LSM method. Since the extension from the Gaussian GARCH model to the NIG GARCH model involves no extra complexity in terms of the dynamics to be used for pricing, this simulation based approach should remain equally valid in our setting.

4.2 Pricing Properties

To illustrate the pricing properties we perform a study using the same basic setup as was used in Stentoft (2005). In particular, we use a set of artificial options and the parameter values corresponding approximately to the average of the estimates in Tables 1 to 4. Specifically we set $\beta = 0.94$ and $\alpha = 0.05$ in the GARCH models and $\beta = 0.94$, $\alpha = 0.04$, and $\gamma = -0.5$ in the NGARCH model. Furthermore, we set $\lambda = 0.025$ which is approximately the average of the estimates for the NIG NGARCH specifications. Finally, in order for the American feature to be interesting for the call options dividends need to be included. Thus, we include dividend payments in the simulation at a rate of 3% annually paid out continuously. For simplicity we set the interest rate r equal to 3% also which yields a riskneutral process with zero drift. In the pricing step we consider early exercise on a daily basis and in order to approximate the conditional expectation function we use the suggested specification from Stentoft (2005) for the cross sectional regression.

Since the primary goal of this study is to examine the effect on option prices when the underlying distribution is skewed and leptokurtic we report the pricing results compared to that of the Gaussian GARCH models which constitute a natural benchmark. Thus, in the following all comparisons will be made between this model and a symmetric NIG distribution where the parameter b is restricted to zero and a general NIG distribution with b different form zero and finite a. Thus, while the latter distribution allows for both skewness and leptokurtosis we also consider a leptokurtic but symmetric distribution in order to assess the importance of either of these departures from the Gaussian distribution. With respect to the parameter values we chose these in order to correspond to the empirical averages and set $a = 2$ and $b = 0.2$.

4.2.1 Pricing results for GARCH models

The actual pricing results for the Gaussian model correspond to what has been previously reported in e.g. Stentoft (2005). Thus, in this section we will report the call pricing errors of this model relative to that of a symmetric NIG distribution as well as the general fat tailed and skewed NIG distribution with non zero b and finite a. Plots of this pricing error for the two volatility specifications for two different maturities are shown as a function of moneyness in Figure 5. In all the panels the results for the Gaussian model relative to the symmetric $NIG\,(a, 0)$ model are indicated by the dark gray line with squares whereas the results for the Gaussian model relative to the skewed and leptokurtic $NIG\,(a, b)$ model are indicated by the light gray line with triangles. Before we proceed it should be noted that in all respects similar results are obtained with put options.

Considering first the results for the GARCH model relative to the symmetric but leptokurtic $NIG\,(a, 0)$ distribution the two top panels of the figure show a very interesting pattern across moneyness. In particular, the panels indicate that, relative to a leptokurtic distribution, the Gaussian GARCH models will underprice out of the money options as well as in the money options and overprice at the money options. This pattern is particularly pronounced for the short term options indicated in the top left panel. This result is qualitatively similar to e.g. Heston (1993) where the pricing performance of a stochastic volatility model is compared to that of the constant volatility Black-Scholes-Merton model. Intuitively, the leptokurtosis from the conditional distribution plays a separate and additional role to that of having a time varying volatility in terms of generating two fat tails.

Next, we consider the effect of allowing for asymmetries in the distribution using the general $NIG\,(a, b)$ distribution. Compared to the situation with a symmetric distribution the top panels of the figure show that the underpricing of the Gaussian GARCH model is now even worse for the out of money options. On the other hand, there is now little or no underpricing of the in the money options. Intuitively out of the money call options benefit significantly from the fat right tail generated by the positive conditional skewness of the $NIG\,(a, b)$ while in the money options are affected less. Again this type of results is similar to e.g. Heston (1993) which finds the same effect when allowing for positive correlation between the return process and the variance process in a stochastic volatility framework.

Finally, comparing the results of the GARCH specification to that obtained with the NGARCH model, which allows for a leverage effect, the bottom panel shows that relative to the Gaussian model the pricing performance persists. In particular, the overall pattern is only marginally changed when the volatility specification is changed in order to include asymmetries. Thus, it seems that this feature is consistently related to the properties of the underlying distribution and not necessarily to the properties of the variance process. We may thus conclude that the NIG GARCH model provides an important extension to the Gaussian GARCH model when it comes to option pricing. In particular, allowing for skewness and leptokurtosis yields a more flexible model than the Gaussian benchmark.

4.2.2 Implied standard deviations for GARCH models

An often used alternative to the above pricing error measures is that of considering differences in implied standard deviations, or ISDs, for various option prices. In particular, when analyzing the pricing performance of e.g. the constant volatility Black-Scholes-Merton models this if often done by documenting so-called smiles or smirks in the ISD. Such patterns cannot be generated by a constant volatility model and hence it must be wrong. In this section we compare the standard deviations implied by the option prices from the GARCH and NGARCH models implemented with the three distributions considered.

Figure 6 plots the ISDs for these models against moneyness for the two different maturities. The left hand panels of the figure show the differences in ISDs for the short term options with approximately one month to maturity whereas the right hand panels show the results for options with around six months maturity which is equivalent to 126 trading days. To back out the volatility a standard binomial model was used in order to take account of the early exercise feature in the options.

First, comparing the plots in Figure 6 across maturity it is seen that irrespective which of the two volatility specifications is used the effects on ISDs vanishes with maturity. In particular, while it is possible to observe large differences between the ISD across moneyness for the short term options this is much less so for the long term options. In fact, for the GARCH model any relationship between ISD and moneyness is negligible at the six month horizon.

Next, when comparing across models in Figure 6 it is indicated that although the ISDs are close for the at the money options there are differences for the in the money and out of the money options. To be specific, the panels show that the ISDs are generally larger for in the money options in a NGARCH model than in a GARCH model. For the out of the money options the results are the opposite.

Finally, comparing between distributions the figure clearly indicates that differences in the ISDs are obtained. In particular, this is so for the out of the money options where the level of the standard deviation implied by the option prices from an NIG GARCH model is larger than those from a symmetric but leptokurtic model. These on the other hand are larger than what is obtained with the Gaussian model. Again this is a pattern which seems independent of the particular specification of the variance process, and hence confirms that allowing for fat tailed and skewed distributions provides a potentially important extension to the Gaussian model due to the added flexibility.

5 Empirical performance of option pricing with the NIG GARCH model

In Stentoft (2005) the Gaussian GARCH model was used to price options on a sample of individual stocks and a major stock index. The paper showed that these models reduce the pricing errors compared to the constant volatility Black-Scholes-Merton model. However, the paper also showed that some pricing errors remain. As we have noted above the study of the properties of the NIG GARCH model shows that this provides a more flexible specification. Thus, in this section we use the same sample of options as in Stentoft (2005) to test the empirical performance of the NIG GARCH option pricing model.

5.1 Data and option pricing procedure

The sample used for our empirical application consists of weekly observation on options traded between 1991 and 1995 on the three individual stocks considered in Section 3.2, i.e. General Motors (GM), International Business Machines (IBM), and Merck & Company Inc. (MRK), as well as options on the Standard & Poor's 100 Index (OEX). These are American style options which are traded on the Chicago Board Options Exchange (CBOE) and were made available through the Berkeley Options Data Base (BODB). At any particular day the sample contains an end of day option quote observed immediately before 3PM for all existing contracts, that is for each combination of strike price and maturity.

The data used for estimation of the NIG GARCH models as well as the estimation procedure corresponds to that used for estimation previously in Section 3.2. However, at any given day where options are priced only the available historical return data is used. This means that when options on a given Wednesday are to be priced the sample period used for estimation is July 1962, to the Tuesday immediately before this Wednesday for the individual stocks. Due to restrictions on data availability the sample used for estimation for the index commences in January 1973 only. Thus, as we move forward in time this sample will increase. The estimated parameters are then used in the pricing procedure on this particular day.

In the simulations we assume that the future dividend payments are known in advance, both in terms of the ex-dividend day and the size of the dividends. We also make the assumption that they spill over fully on the simulated prices of the underlying asset. Thus, if day t is an ex-dividend day, on this day the dividend payment is deducted from the simulated prices. Finally, as the risk free interest rate we take the LIBOR rate on the last day used for estimation. Thus, although the same constant interest rate is used both in the estimation and in the simulation at any given day, in fact the interest rate does vary from one Wednesday to the next Wednesday. We note that both the assumption on dividends and on interest rates could be changed as long as they remain known at the time of pricing.

With respect to all other aspects of the actual implementation of the LSM method we follow the outline from Section 4.1.2. The only exception is with respect to the possible early exercises for call options. The reason for treating this differently is that we know theoretically that for assets which pay out dividends discretely early exercise can only be optimal immediately before the asset goes ex-dividend. Hence, for this type of options early exercise does not need to be considered at other points in time. On the other hand, for the put options all that is known is that this type of options are more likely to be exercised immediately after the asset goes ex-dividend. Thus, for the put options we maintain the assumption that early exercise is in fact possible once a day.

5.2 Overall pricing results for the NIG GARCH models

The actual pricing results for the NIG GARCH are presented in Table 5 assuming a value of $\lambda = 0.025$ for all series. This table reports results for the four assets individually and each panel provides results for puts and calls separately. To gauge the pricing performance we use two of the relevant metrics from the literature on both prices and ISDs. In particular, letting P_k and $\tilde{P}_k$ denote the k'th observed price respectively the k'th model price, the table reports the bias, $BIAS \equiv K^{-1} \sum_{k=1}^{K} \left(P_k - \tilde{P}_k \right)$, and the root mean squared error, $RMSE \equiv \sqrt{K^{-1} \sum_{k=1}^{K} \left(P_k - \tilde{P}_k \right)^2}$. When using the metrics on the ISDs these are simply used instead of the prices in the two formulas.

Comparing first the results for the individual stock options in Panels A through C the table shows a very similar pattern. In particular, in the majority of the cases the NIG NGARCH model provides the smallest errors. The exception to this is for call options on IBM where the GARCH model performs slightly better in terms of dollar BIAS and RMSE and in terms of ISD BIAS. For the ISD, though, the performance of the NGARCH model is virtually the same. For MRK the call dollar BIAS is also marginally smaller with a GARCH model than with a NGARCH model.

Next, when considering the actual size of the errors the panels also show that the improvements obtained with the NIG NGARCH model are potentially large compared to the alternative volatility specifications. This is especially clear when comparing to the CV specification, which is a special case of the flexible framework used here. However, even when comparing the two models with time varying volatility, i.e. the GARCH and NGARCH specifications, large improvements are found for several metrics. The improvements are generally largest for the put options where the dollar BIAS is 20% smaller for the NGARCH models than for the GARCH models for both GM and IBM.

Finally, comparing the results for the index options in Panel D it is seen that although the NGARCH model is the best performing model for all but one of the metrics the errors are substantially larger for these options. Moreover, the dollar BIAS for the two types of options may have opposite signs even when using the same type of model. To be specific, the panel shows that whereas the GARCH and NGARCH models underprice put options they overprice call options in terms of dollar errors, although by a somewhat smaller amount. The ISDs seem to allow for a somewhat easier comparison, and generally indicate a slightly better performance for the models allowing for a time varying volatility than for models with constant volatility.

5.2.1 Relative pricing performance of the NIG GARCH model

In Table 6 we provide results on the performance of the GARCH models using the two other distributions considered in this paper, i.e. the symmetric NIG and the Gaussian distributions. In the table only the results for the ISDs are reported. The left hand part of the table reports the results obtained with the symmetric $NIG\left(a,0\right)$ distribution relative to those for the general $NIG\left(a,b\right)$ from Table 6. The right hand part compares the Gaussian model to the $NIG\left(a,b\right)$ in the same way. Thus, when looking at the table, cells with values larger than one indicates metrics and option types for which the general NIG GARCH model performs the best.

Comparing first the two NIG models the table shows that the two models perform almost equally well on average. However, there does appear to be some difference between the two types of options with the skewed model performing better when it comes to pricing call options and outperforming the symmetric NIG model in the majority of the cases.

Next, when considering the relative performance of the Gaussian model the table indicates that this model is in general inferior to the general NIG model. In particular, this is the case for all the call options as well as for the put options on GM except when the volatility is assumed constant. The differences are also quite big in many cases. As an example, consider the BIAS for the GARCH model for GM which is 20% larger with a Gaussian model than with the NIG model.

5.3 Fitting the smile in option ISDs

The average pricing performance in terms of dollar errors or even in terms of ISDs is only one possible metric for comparison. However, when it comes to option pricing it is perhaps of more interest to examine how the models fit across moneyness. In particular, option prices are often quoted in terms of implied volatilities, and often such volatility quotes vary with moneyness. The ultimate test of an option pricing model may well be to fit this pattern which is known as the volatility smile. In Section 4 it was shown that the shape of the ISD for the NIG GARCH models did differ from that of the Gaussian GARCH model, and this was particularly so for out of the money options. Thus, in the rest of this section we examine how well the models fit the smile found empirically in option ISDs.

5.3.1 The existence of a smile in empirical ISDs

Before we analyze the NIG GARCH models' ability to fit the smile in empirical ISDs it seems natural to start by documenting the existence of such a smile for the options considered here. This is done in Figures 7 and 8 for put and call options for the four assets, respectively. Instead of plotting the actual ISD BIAS in terms of moneyness we only report the average within specific categories of moneyness. In the figures D-ITM stands for deep in the money and corresponds to options which are more than 7.5% in the money, and ITM stands for in the money options and corresponds to options which are between 7.5% and 2% in the money. For out of the money, or OTM, and deep out of the money, or D-OTM, the same applies and thus the category of ATM, or at the money options are essentially those with a moneyness between 2% in the money and 2% out of the money. For both put and call options moneyness is simply calculated as the ratio of the underlying value to the strike price.

In both figures the solid line shows the BIAS in the ISDs for a constant volatility Black-Scholes-Merton type model compared to the ISDs of the actual data. Since the CV model by construction yields a constant ISD across moneyness any shape observed in the BIAS stems from a shape in ISD of the actual data. From looking at the figures we may thus conclude that there is a clear smiling pattern in individual stock options. The smile is present for puts as well as call options, although it appears to be slightly asymmetric for the call options.

For the index options the figures indicate a very asymmetric smile, or smirk, for both put and call options. This smirk has the highest ISDs for the D-OTM category for the put options, whereas for the call options the smirk is highest for the D-ITM options. For the put options this finding corresponds well to the case where D-OTM options are used heavily as insurance against market crashes and therefore trade at somewhat elevated prices. Because of the relationship between D-OTM puts and D-ITM calls this would also in part explain the high ISD for these contracts.

5.3.2 The conditional distribution and the smile

As the previous analysis has shown overall the NGARCH model is the best perform model within the NIG GARCH framework. For this reason we start out by analyzing these models potential for accommodating the smile or smirk found in our option data. This is also done in Figures 7 and 8 which show the ISD BIAS for the NGARCH models with the three different proposed distributions.

The first thing to note from the figures is that the smile in the BIAS is less pronounced for the models with time varying volatility than for the CV model when considering the individual stock options. This is particularly so for GM where the BIAS is decreased from between 5% and 10% with the CV model to between 1% and 3% for the put options and 0.5% and 4% for the call options. For MRK the decrease is less pronounced with the results for IBM in between.

The second thing to note is then when comparing the performance in terms of BIAS for the different distributions the NIG NGARCH model performs quite well. To be specific, whereas the figures show that for the D-OTM puts and D-ITM calls the differences are only minor for the D-ITM puts and D-OTM calls, the NIG NGARCH model outperforms the Gaussian version. This is particularly the case for call options on GM and IBM where D-OTM BIAS is reduced from 3.2% to 2.1% and from 4.6% to 3.6%, respectively. More importantly, though, is the fact that due to this reduction the smile effect is significantly less pronounced for this model than for the Gaussian counterpart.

Finally, we note that when considering the index options the figures do seem to indicate that the NGARCH models are capable of diminishing the smirk in the ISDs. However, while the NGARCH models do well for ITM and ATM puts and ATM and OTM calls they do less well than even the CV model for other categories of moneyness. Also the plots indicate that it is in fact the Gaussian NGARCH model which performs the best for this particular series.

5.3.3 Alternative sources of asymmetries

The NIG GARCH model used above with $\lambda = 0.025$ is the most versatile model allowing for not only skewness in the conditional distribution and in the variance process but also for an effect through the risk neutralized innovations. Thus, it may be of interest to see if any of these effects in particular contributes to the models' ability to explain the smile or smirk in real option ISDs. In this section we analyze each of these factors in turn. Note that since the best performing model for the index was the Gaussian NGARCH model this will be the benchmark for that particular series. The results of the comparison are given in Figures 9 and 10 which compares the BIAS for the most elaborate model to the simplifications. Thus, for the individual stock options the figures allow us to compare the performance of the NIG NGARCH model to the NIG GARCH model and the NIG NGARCH model with λ set equal to zero.

We first consider the individual stock options and the effect of the leverage effect by comparing the NIG NGARCH model to the GARCH alternative. The results overall indicate that the GARCH model outperforms the NGARCH model for D-ITM puts and D-OTM calls whereas it underperforms the NGARCH model for D-OTM puts and D-ITM calls. Although at first this result may seem somewhat difficult to interpret looking at the graphs we may actually conclude that in terms of mitigating the smile effect the NGARCH specification performs better since the smile in this BIAS is less pronounced.

Next, to consider the effect of the λ parameter which can approximately be interpreted as the unit risk premium the conclusion is in favour of a premium in the model. In particular, although the differences are minor in all cases the BIAS for the model with $\lambda = 0.025$ lies below that of the restricted model with $\lambda = 0$. This is in fact to be expected since a positive price of risk increases the unconditional level of the volatility of the risk neutral process which in turn increases option prices overall. However, from the figures it is not possible to determine if a positive risk premium particularly improves for one category of options or for puts versus call options.

Finally, we consider the index options for which we note that the same overall conclusions are obtained as with the individual options. In particular, the effect of setting λ equal to zero seems of minor importance. Furthermore, with respect to the GARCH application the performance relative to the NGARCH specification is such that when it comes to explaining the smile the NGARCH model does the best job.

6 Conclusion

In this paper we propose a feasible way to price American options in a framework with time varying volatility and conditional skewness and leptokurtosis. To be specific, we use a Generalized Autoregressive Conditional Heteroskedasticity framework where the conditional distribution is the Normal Inverse Gaussian. We show how the Generalized Local Risk Neutral Valuation Relationship of Duan (1999) can be used in this framework to derive the riskneutral dynamics.

Because of the departure from a Gaussian framework the application of the NIG GARCH model is computationally complex. In particular, in order to implement the model it is necessary to derive methods for computing functions of normal variates and conditional first moments of these. In the present paper computationally easy and efficient approximations for doing this are suggested which rely only on approximations of smooth functions with low order polynomials. We provide estimation results which support the use of the Normal Inverse Gaussian distribution as an important extension of the Gaussian model for modelling asset returns.

Next we show that the ability of the NIG GARCH model to accommodate conditional skewness and leptokurtosis can explain some of the systematic pricing errors found in recent empirical work on option pricing with time varying volatility. In particular, the NIG GARCH model has the ability to generate higher prices than the Gaussian GARCH model for out of the money put and call options particularly of short maturity. Since such findings have been documented empirically the results support the use of the Normal Inverse Gaussian distribution as an important extension of the Gaussian model for option pricing purposes.

Finally, we perform an empirical evaluation of the NIG GARCH models' potential when it comes to pricing American options on a number of individual US stocks as well as a major US index. The results provide evidence favouring the NIG GARCH model over the Gaussian alternative for the individual stock options. In particular, improvements are found when comparing the implied standard deviations of the model prices to those from the empirical data and the NIG GARCH model significantly reduces the smile effect often found when applying option pricing models to this type of data.

In terms of future research we note that this paper uses a very flexible estimation and pricing framework which can accommodate many potential extensions of and alternatives to the NIG GARCH models considered here. Examining such extensions within the framework should definitely be the focus of future research and we are currently conducting a large scale comparison of alternative conditional distributions.

References

Andersson, J. (2001), 'On the Normal Inverse Gaussian Stochastic Volatility Model', *Journal of Business and Economic Statistics* **19**(1), 44–54.

Barndorff-Nielsen, O. E. (1997), 'Normal Inverse Gaussian Distributions and Stochastic Volatility Modelling', *Scandinavian Journal of Statistics* **24**, 1–13.

Barndorff-Nielsen, O. E. & Shephard, N. (2001), 'Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics', *Journal of the Royal Statistical Society, Series B* **63**, 167–241.

Black, F. & Scholes, M. (1973), 'The Pricing of Options and Corporate Liabilities', *Journal of Political Economy* **81**, 637–654.

Bollerslev, T. (1986), 'Generalized Autoregressive Conditional Heteroskedasticity', *Journal of Econometrics* **31**, 307–327.

Bollerslev, T. (1987), 'A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return', *Review of Economics and Statistics* pp. 542–547.

Bollerslev, T. & Mikkelsen, H. O. (1996), 'Modelling and Pricing Long Memory in Stock Market Volatility', *Journal of Econometrics* **73**, 151–184.

Bollerslev, T. & Mikkelsen, H. O. (1999), 'Long-Term Equity Anticipation Securities and Stock Market Volatility Dynamics', *Journal of Econometrics* **92**, 75–99.

Carr, P. & Wu, L. (2004), 'Time-changed Lévy processes and option pricing', *Journal of Financial Economics* **71**, 113–141.

Christoffersen, P., Heston, S. & Jacobs, K. (2006), 'Option valuation with Conditional skewness', *Journal of Econometrics* **131**, 253–284.

Christoffersen, P. & Jacobs, K. (2004), 'Which GARCH Model for Option Valuation?', *Management Science* **50**(9), 1204–1221.

CRSP (1998), *CRSP Stock File Guide*, Center for Research in Security Prices, Graduate School of Business, University of Chicago.

Duan, J.-C. (1995), 'The GARCH Option Pricing Model', *Mathematical Finance* **5**(1), 13–32.

Duan, J.-C. (1999), 'Conditionally Fat-Tailed Distributions and the Volatility Smile in Options', *Working paper, Hong Kong University of Science and Technology* .

Duan, J.-C. & Simonato, J.-G. (2001), 'American Option Pricing under GARCH by a Markov Chain Approximation', *Journal of Economic Dynamics and Control* **25**, 1689–1718.

Duan, J.-C. & Zhang, H. (2001), 'Pricing Hang Seng Index options around the Asian financial crisis - A GARCH approach', *Journal of Banking and Finance* **25**, 1989–2014.

Eberlein, E. (2001), Application of generalized hyperbolic Lévy motions to finance, *in* O. E. Barndorff-Nielsen, T. Mikosch & S. Resnick, eds, 'Lévy Processes: Theory and Apllications', Birkhauser Verlag, Boston, pp. 319–336.

Engle, R. F. (1982), 'Autoregressive Conditional Heteroscedasticity With Estimates of The Variance of United Kingdom Inflation', *Econometrica* **50**(4), 987–1007.

Engle, R. F. & Ng, V. K. (1993), 'Measuring and Testing the Impact of News on Volatility', *Journal of Finance* **48**(5), 1749–1778.

Engle, R. & Mezrich, J. (1996), 'GARCH for Groups', *Risk* **9**(8), 36–40.

Forsberg, L. & Bollerslev, T. (2002), 'Bridging the Gap Between the Distribution of Realized (ECU) Volatility and ARCH Modeling (of the EURO): The GARCH Normal Inverse Gaussian Model', *Journal of Applied Econometrics* **17**, 535–548.

Hafner, C. M. & Herwartz, H. (2001), 'Options pricing under linear autoregressive dynamics, heteroskedasticity, and conditional leptokurtosis', *Journal of Empirical Finance* **8**, 1–34.

Härdle, W. & Hafner, C. M. (2000), 'Discrete time option pricing with flexible volatility estimation', *Finance and Stochastics* **4**, 189–207.

Heston, S. L. (1993), 'A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options', *Review of Financial Studies* **6**(2), 327–343.

Heston, S. L. & Nandi, S. (2000), 'A Closed-Form GARCH Option Valuation Model', *Review of Financial Studies* **13**(3), 585–625.

Hsieh, K. C. & Ritchken, P. (2005), 'An Empirical Comparison of GARCH Option Pricing Models', *Review of Derivatives Research* **8**, 129–150.

Huang, J.-Z. & Wu, L. (2004), 'Specification Analysis of Option Pricing Models Based on Time-Changed Lévy Processes', *Journal of Finance* **LIX**(3), 1405–1439.

Jensen, M. B. & Lunde, A. (2001), 'The NIG-S&ARCH model: a fat-tailed, stochastic, and autoregressive conditional heteroskedastic volatility model', *Econometrics Journal* **4**, 319–342.

Lambert, P. & Laurent, S. (2001), 'Modelling Financial Time Series Using GARCH-Type Models with a Skewed Student Distribution for the Innovations', *Working paper, Université de Liège* .

Lehar, A., Scheicher, M. & Schittenkopf, C. (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* **10**(3), 27–39.

Longstaff, F. A. & Schwartz, E. S. (2001), 'Valuing American Options by Simulation: A Simple Least-Squares Approach', *Review of Financial Studies* **14**, 113–147.

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

Nelson, D. B. (1991), 'Conditional Heteroskedasticity In Asset Returns: A New Approach', *Econometrica* **59**(2), 347–370.

Ritchken, P. & Trevor, R. (1999), 'Pricing Options under Generalized GARCH and Stochastic Volatility Processes', *Journal of Finance* **59**(1), 377–402.

Stentoft, L. (2005), 'Pricing American Options When the Underlying Asset follows GARCH Processes', *Journal of Empirical Finance* **12**(4), 576–611.

Figures and Tables

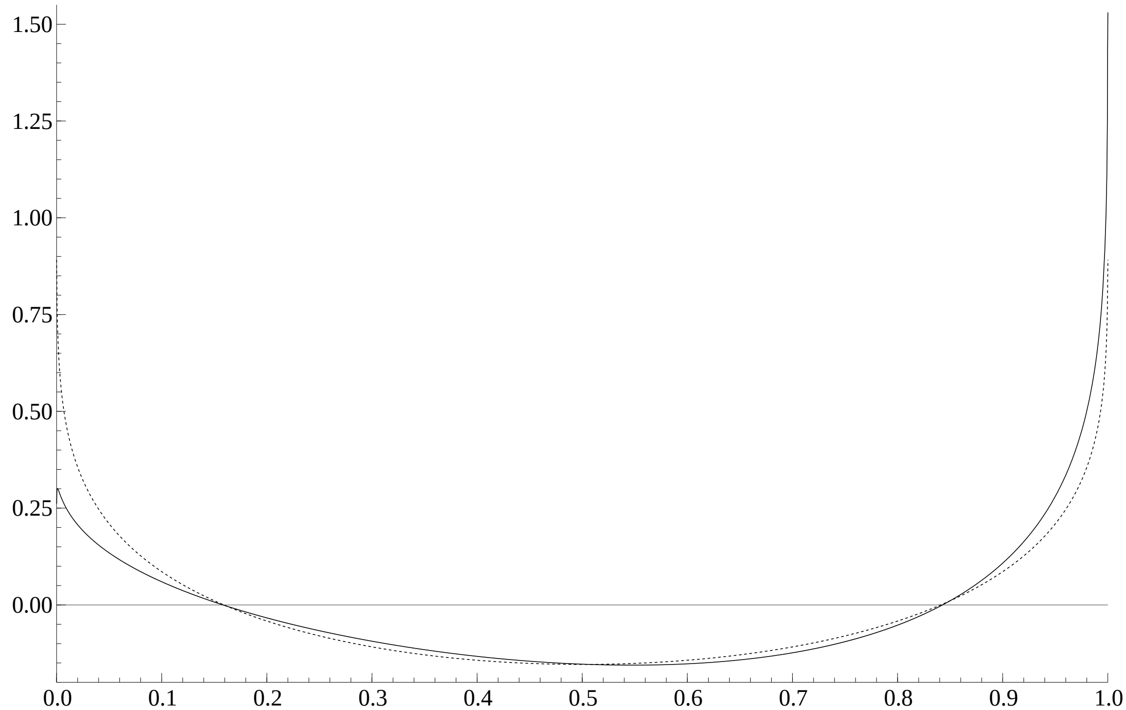

Figure 1: Plots of the logarithm of the ratio between the $NIG\,(2,0)$ density and the Gaussian density (the dotted line) and the $NIG\,(2,0.2)$ density and Gaussian density (solid line) against the quantile values.

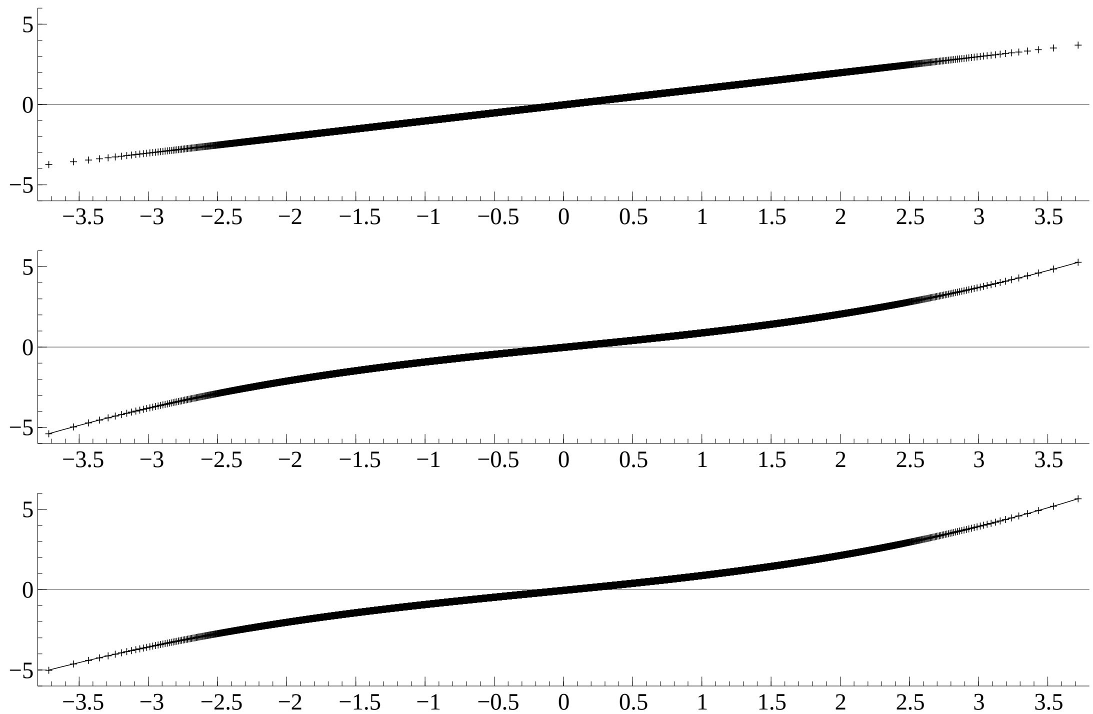

Figure 2: Plots of the approximated values of $F_{NIG}^{-1}\left[\Phi\left(Z-\lambda\right)\right]$ with $\lambda = 0.025$ against Z, where Z is equal to the $9,999$ quantile values from the Gaussian distribution. The top panel shows the results for the Gaussian distribution, the middle panel shows the results for the symmetric $NIG\left(2,0\right)$, and the bottom panel shows the results for the general $NIG\left(2,0.2\right)$ distribution. Superimposed on the two bottom plots as a solid line is the estimated relationship from the suggested fifth degree polynomial approximation.

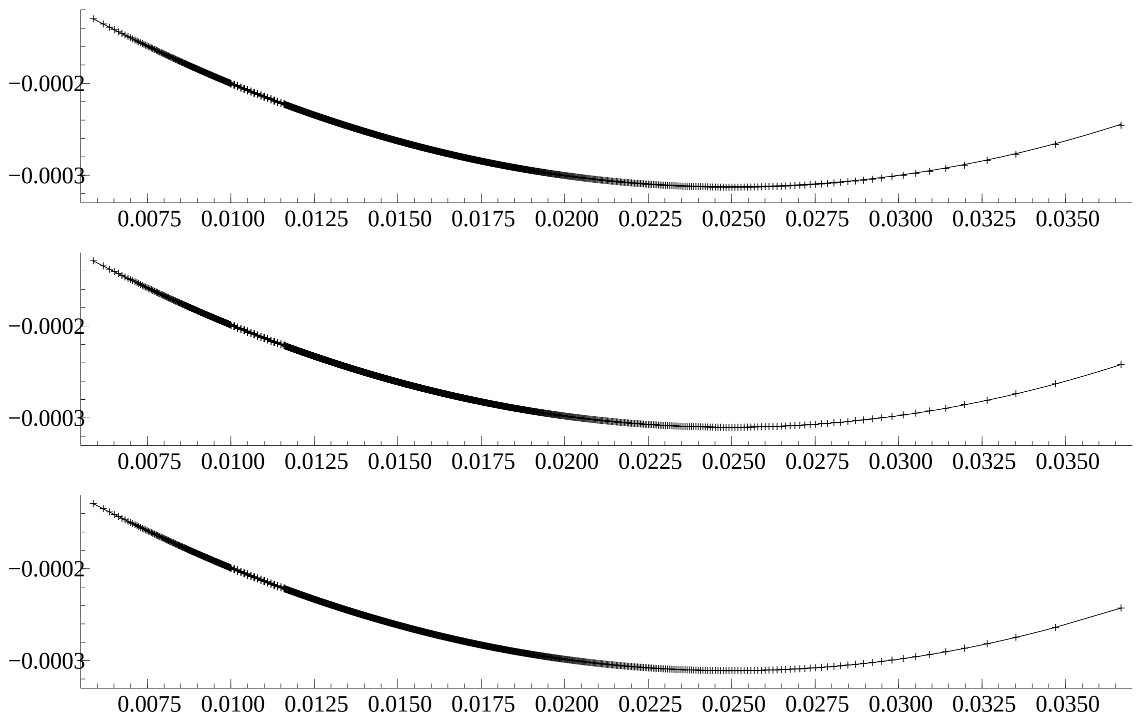

Figure 3: Plots of the approximated values of the function $\ln E^{\mathcal{Q}}\left[\exp\left(\sqrt{h_t}F_{NIG}^{-1}\left[\Phi\left(Z_t-\lambda\right)\right]\right)\middle|\,\mathcal{F}_{t-1}\right]$ with $\lambda = 0.025$ against $\sqrt{h_t}$, where h_t is equal to the 999 quantile values of the Inverse Gaussian distribution with $a = 2$ and where $F_{NIG}^{-1}\left[\Phi\left(Z_t-\lambda\right)\right]$ is calculated according to the suggested approximation. The top panel shows the results for the Gaussian distribution, the middle panel shows the results for the symmetric $NIG\left(2,0\right)$ distribution, and the bottom panel shows the results for the general $NIG\left(2,0.2\right)$ distribution. Superimposed as a solid line on the Gaussian plot is the true relation given by $-\lambda\sqrt{h_t}+\frac{1}{2}h_t$. In the two bottom plots the solid line indicates the estimated relationship from the suggested third degree polynomial approximation.

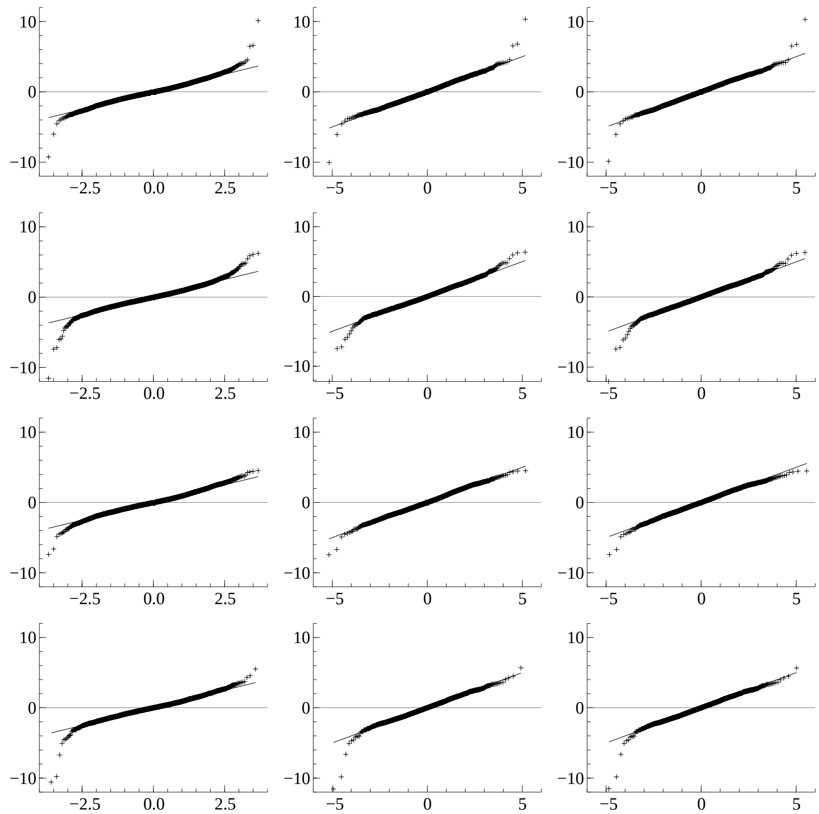

Figure 4: QQ plots of the residuals from the NGARCH models with various distributions. From top to bottom the results are for GM, IBM, MRK, and OEX. From left to right the results are for the Gaussian, the symmetric $NIG\,(a, 0)$, and the general $NIG\,(a, b)$ distribution where a and b correspond to the maximum likelihood estimates from Tables 1 through 4.

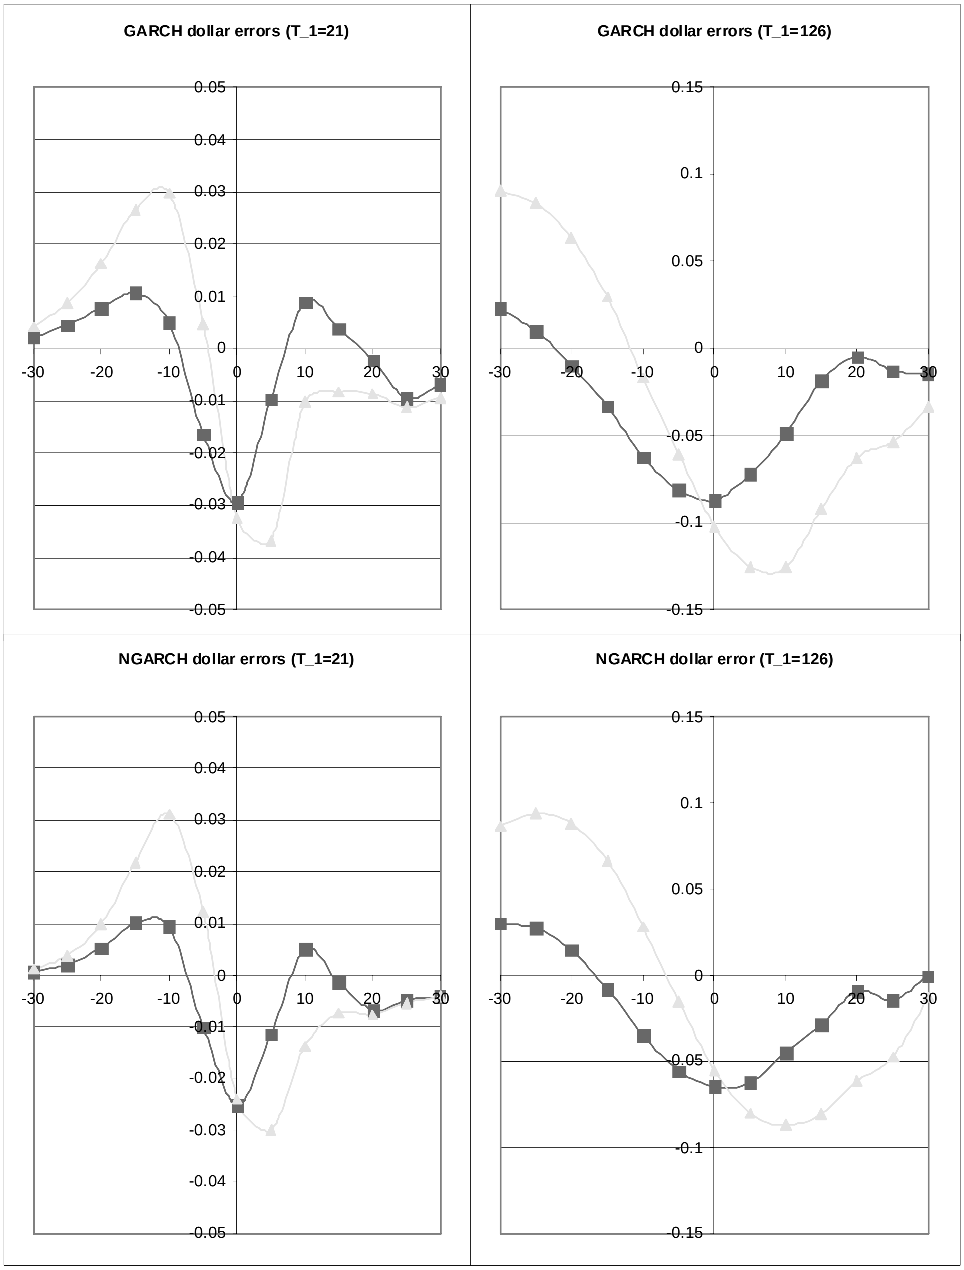

Figure 5: Plots of the dollar pricing errors, calculated as the $NIG\,(2,0)$, respectively $NIG\,(2,0.2)$ price minus the Gaussian price, against moneyness, calculated as the price of the underlying minus the strikeprice. Top plots are for GARCH models whereas bottom plots are for NGARCH models. Left hand plots are for options with one month to expiration whereas right hand plots are for options with six months to expiration. The dark gray line marked with squares compares the Gaussian model to the symmetric $NIG\,(2,0)$ model whereas the lighter gray line marked with triangles compares it to the general $NIG\,(2,0.2)$ model.

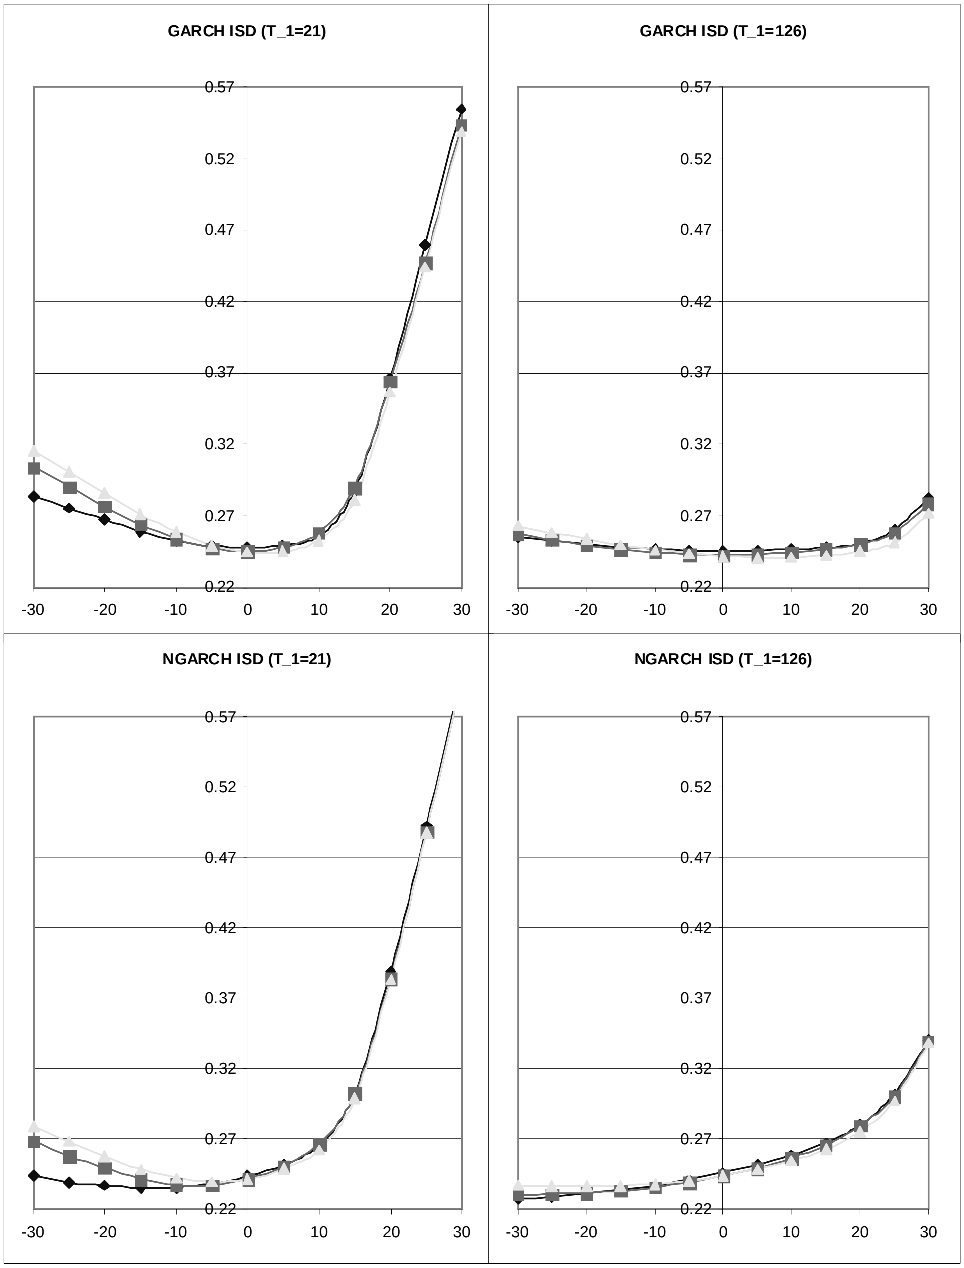

Figure 6: Plots of the implied standard errors, ISDs, against moneyness, calculated as the price of the underlying minus the strikeprice. Top plots are for GARCH models whereas bottom plots are for NGARCH models. Left hand plots are for options with one month to expiration whereas right hand plots are for options with six months to expiration. The black line marked with diamonds is for the Gaussian models, the dark gray line marked with squares is for the symmetric $NIG(2,0)$ model, and the lighter gray line marked with triangles is for the general $NIG(2,0.2)$ model.

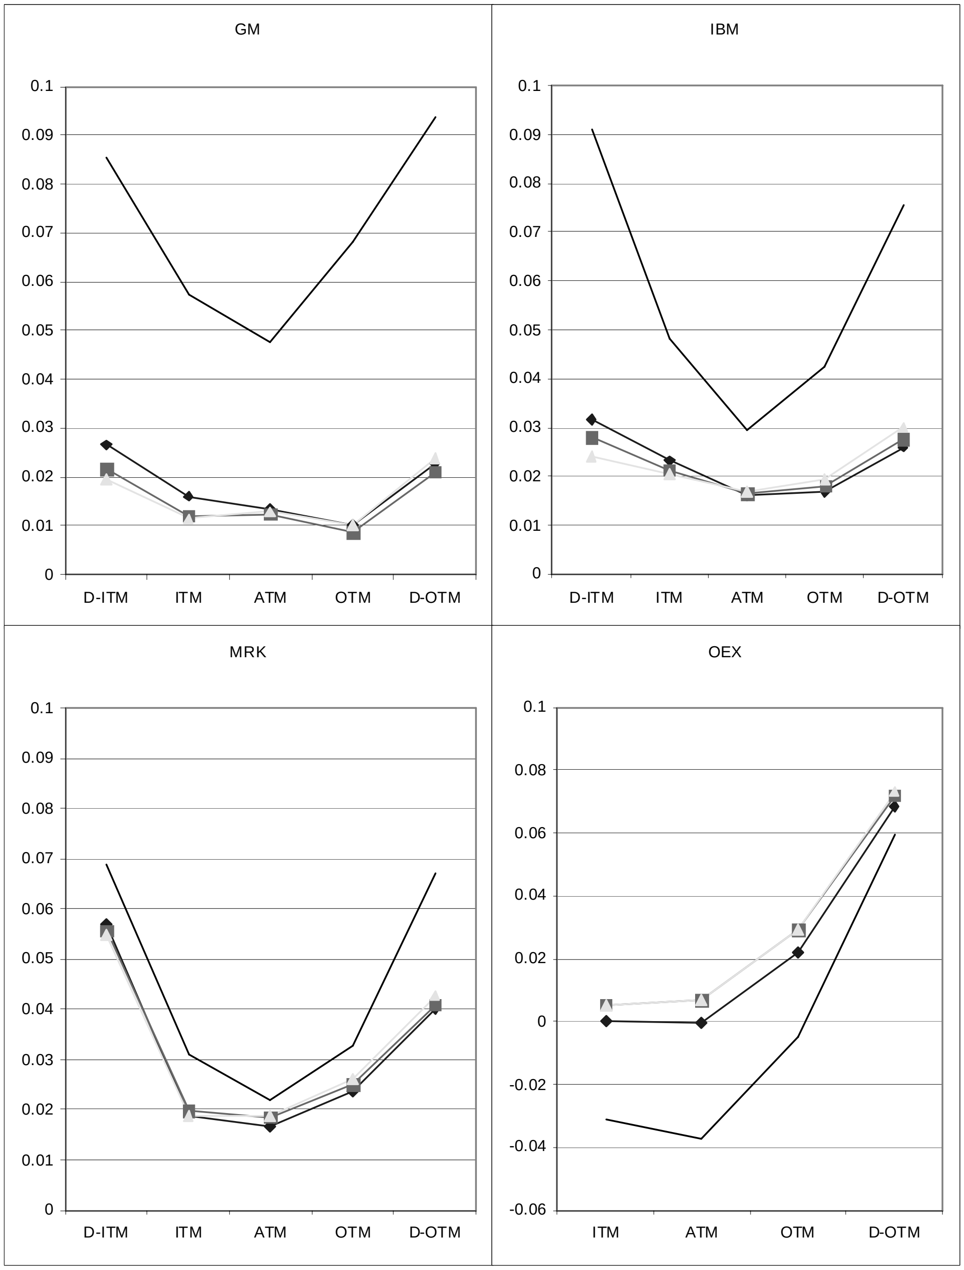

Figure 7: Plots of the bias in implied standard deviations, ISDs, from NGARCH models calculated as the actual ISD minus the model ISD, against different categories of moneyness for put options. The black line marked with diamonds is for the Gaussian model, the dark gray line marked with squares is for the symmetric $NIG(a, 0)$ model, and the lighter gray line marked with triangles is for the general $NIG(a, b)$ model. The black line with no markings is for the Gaussian constant volatility model which corresponds to the Black-Scholes-Merton model.

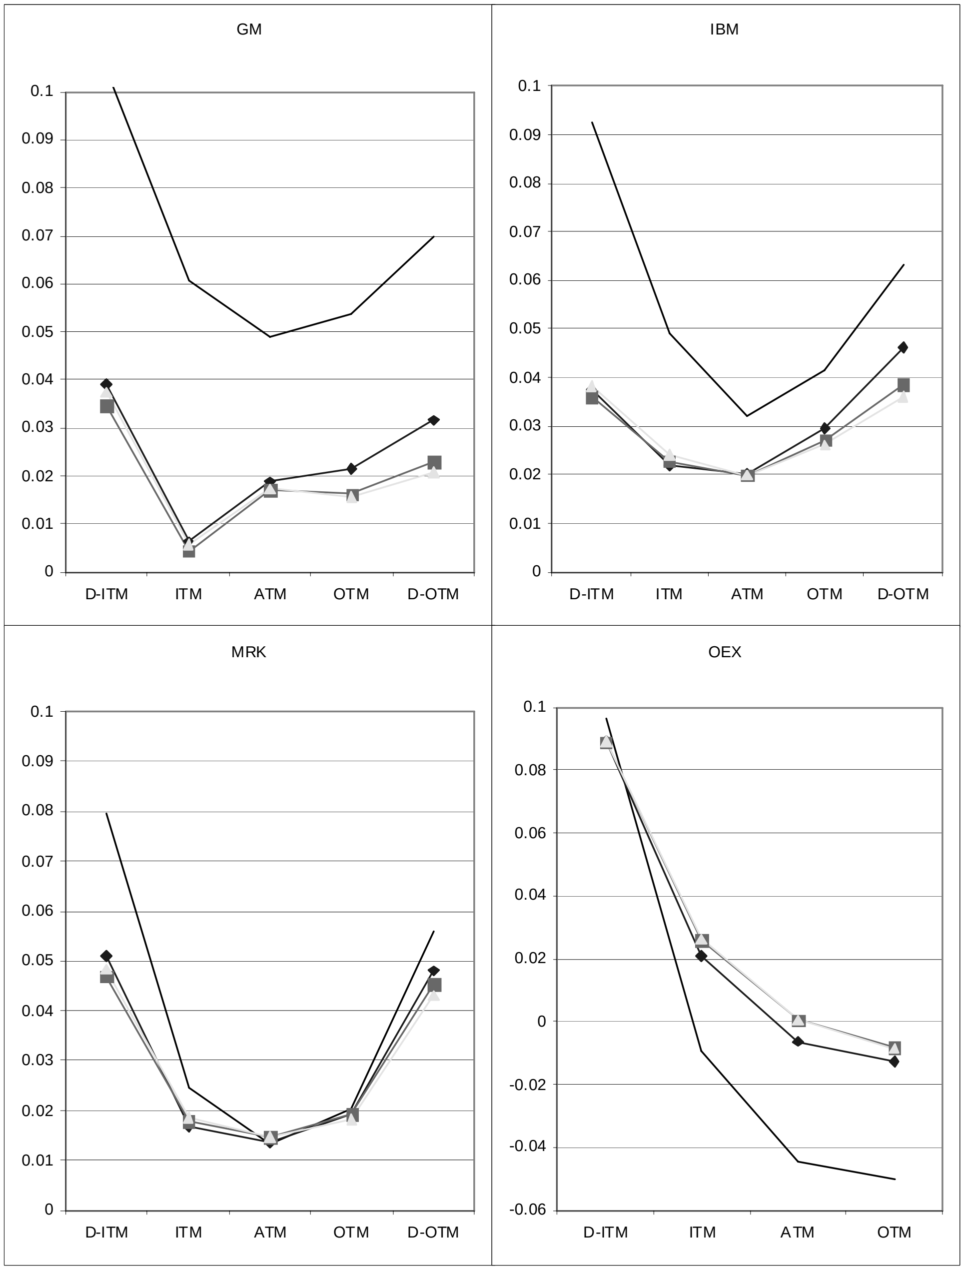

Figure 8: Plots of the bias in implied standard deviations, ISDs, from NGARCH models calculated as the actual ISD minus the model ISD, against different categories of moneyness for call options. The black line marked with diamonds is for the Gaussian model, the dark gray line marked with squares is for the symmetric $NIG(a, 0)$ model, and the lighter gray line marked with triangles is for the general $NIG(a, b)$ model. The black line with no markings is for the Gaussian constant volatility model which corresponds to the Black-Scholes-Merton model.

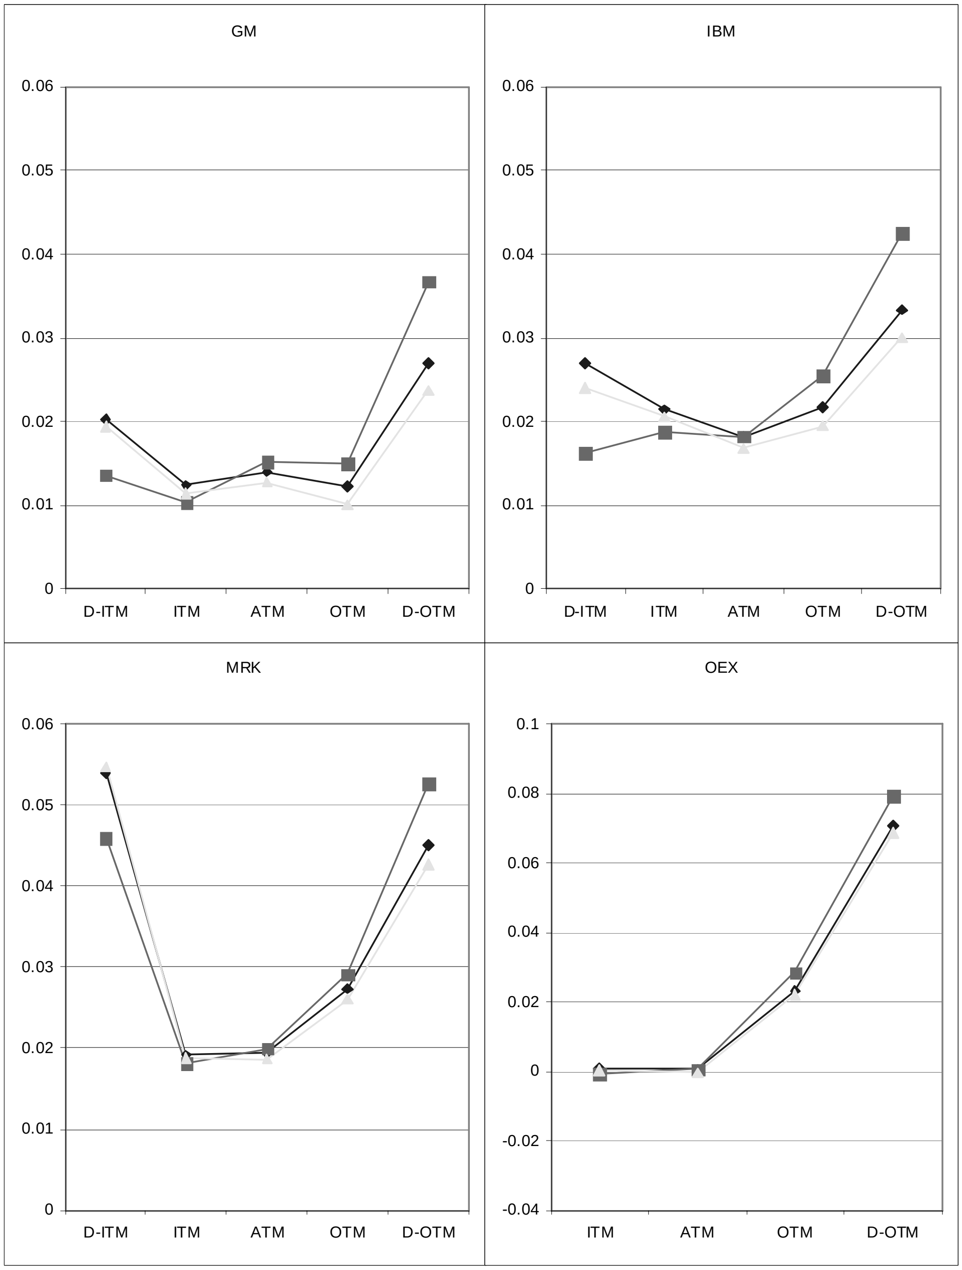

Figure 9: Plots of the bias in implied standard deviations, ISDs, calculated as the actual ISD minus the model ISD, against different categories of moneyness for put options. The light gray line marked with triangles is for the preferred NGARCH model, the dark gray line marked with squares is for the GARCH special case, and the black line marked with diamonds is for the special case of the NGARCH model with λ set equal to zero. For GM, IBM, and MRK the results are shown for the general $NIG(a, b)$ distribution whereas for OEX the results are for the Gaussian distribution.

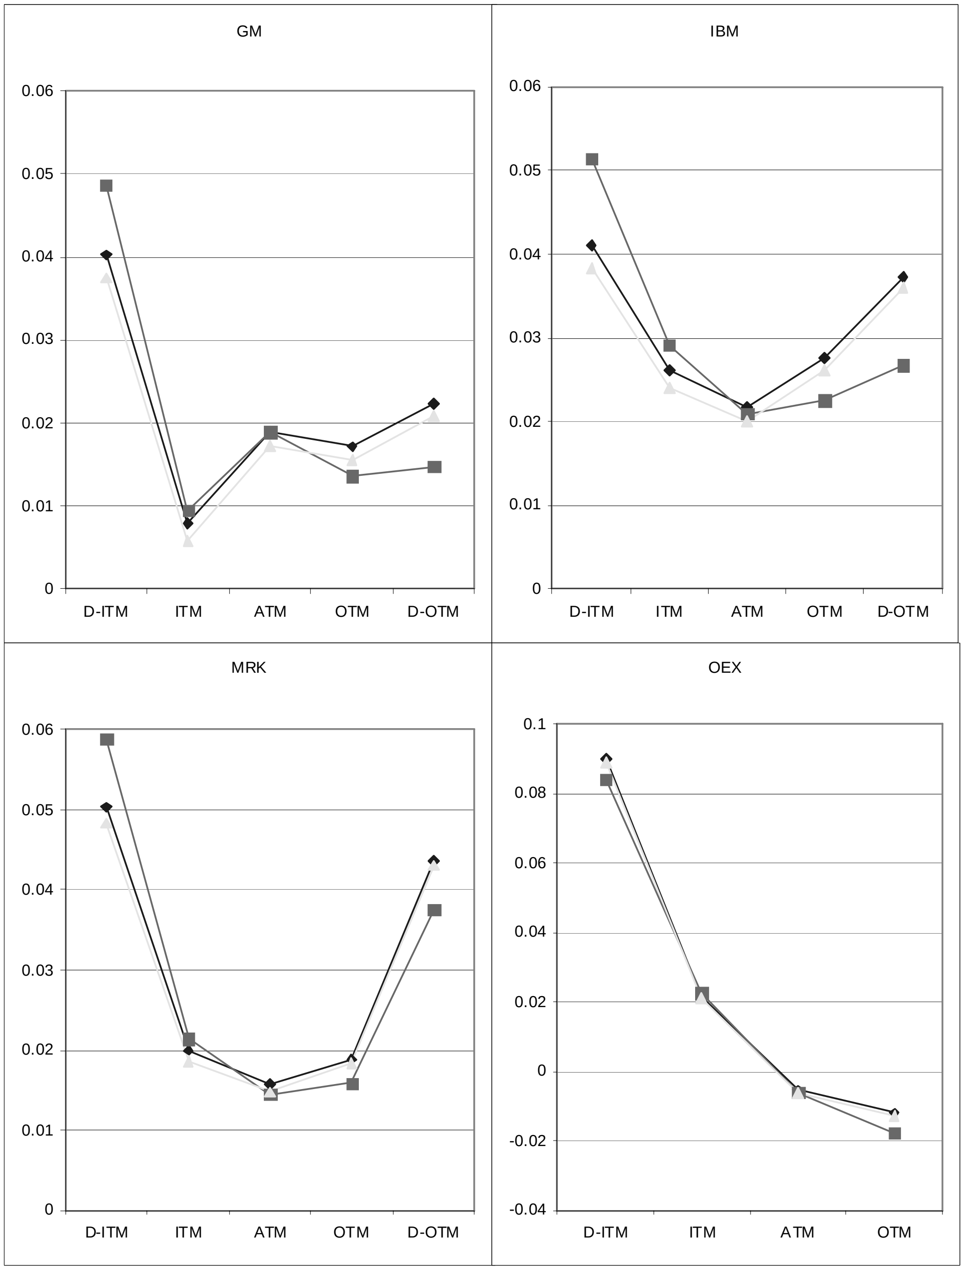

Figure 10: Plots of the bias in implied standard deviations, ISDs, calculated as the actual ISD minus the model ISD, against different categories of moneyness for call options. The light gray line marked with triangles is for the preferred NGARCH model, the dark gray line marked with squares is for the GARCH special case, and the black line marked with diamonds is for the special case of the NGARCH model with λ set equal to zero. For GM, IBM, and MRK the results are shown for the general $NIG(a, b)$ distribution whereas for OEX the results are for the Gaussian distribution.

Table 1: Estimation results for General Motors (GM)

	Gaussian Models			Symmetric NIG Models		NIG Models	
Model	CV	GARCH	NGARCH	GARCH	NGARCH	GARCH	NGARCH
Loglik	-15482.1	-14840.6	-14804.6	-14626.5	-14610.1	-14619.2	-14603.2
	Estim.	Estim.	Estim.	Estim.	Estim.	Estim.	Estim.
λ	0.0152	0.0302	0.0171	0.0070	0.0015	0.0210	0.0131
	(0.011)	(0.011)	(0.012)	(0.010)	(0.010)	(0.011)	(0.010)
β		0.9322	0.9425	0.9506	0.9550	0.9488	0.9530
		(0.021)	(0.013)	(0.011)	(0.009)	(0.011)	(0.010)
α		0.0554	0.0380	0.0421	0.0323	0.0434	0.0341
		(0.016)	(0.011)	(0.008)	(0.009)	(0.008)	(0.009)
γ			-0.5684		-0.4871		-0.4721
			(0.153)		(0.137)		(0.129)
a				2.0085	2.1435	2.0155	2.1239
				(0.216)	(0.235)	(0.212)	(0.231)
b						0.1744	0.1783
						(0.051)	(0.053)
	Stat.	Stat.	Stat.	Stat.	Stat.	Stat.	Stat.
J-B	34649	4867.27	3809.18				
	[0.000]	[0.000]	[0.000]				
Q(20)	65.51	60.67	57.86	60.02	57.77	60.04	57.85
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
$Q^2(20)$	1070.06	30.33	36.16	46.16	47.60	43.56	44.22
	[0.000]	[0.034]	[0.007]	[0.000]	[0.000]	[0.001]	[0.001]
ARCH5	175.823	4.035	4.771	7.402	7.386	6.884	6.664
	[0.000]	[0.001]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
SIC	3.673	3.521	3.512	3.470	3.466	3.468	3.465
							(*)

Notes: This table reports Quasi Maximum Likelihood Estimates (QMLE) using daily returns from July 1962 through 1995 assuming a risk-free interest rate of 5.4% corresponding to the value on December 29, 1995. Robust standard errors are reported in parentheses below the estimates. Loglik denotes the log likelihood value and J-B is the value of the Jarque-Bera normality test for the standardized residuals. $Q(20)$ is the Ljung-Box portmanteau test for up to 20'th order serial correlation in the standardized residuals, whereas $Q^2(20)$ is for up to 20'th order serial correlation in the squared standardized residuals. Finally, ARCH5 denotes the ARCH test from Engle (1982). In square brackets below all test statistics p-values are reported. The last row reports the Schwarz Information Criteria, with an asterix denoting the minimum value.

Table 2: Estimation results for International Business Machines (IBM)

	Gaussian Models			Symmetric NIG Models		NIG Models	
Model	CV	GARCH	NGARCH	GARCH	NGARCH	GARCH	NGARCH
Loglik	-14963.5	-14413.7	-14362.4	-14140.2	-14121.4	-14133.3	-14115.1
	Estim.	Estim.	Estim.	Estim.	Estim.	Estim.	Estim.
λ	0.0136	0.0373	0.0216	0.0109	0.0021	0.0244	0.0137
	(0.011)	(0.013)	(0.013)	(0.010)	(0.010)	(0.011)	(0.011)
β		0.9333	0.9272	0.9423	0.9378	0.9408	0.9360
		(0.022)	(0.018)	(0.010)	(0.009)	(0.010)	(0.009)
α		0.0537	0.0463	0.0464	0.0426	0.0473	0.0441
		(0.016)	(0.011)	(0.007)	(0.006)	(0.007)	(0.006)
γ			-0.5643		-0.4511		-0.4407
			(0.098)		(0.083)		(0.083)
a				2.0303	2.1553	2.0301	2.1341
				(0.212)	(0.226)	(0.214)	(0.227)
b						0.1743	0.1759
						(0.062)	(0.062)
	Stat.	Stat.	Stat.	Stat.	Stat.	Stat.	Stat.
J-B	102111	11374.52	6561.72				
	[0.000]	[0.000]	[0.000]				
Q(20)	42.11	28.14	28.00	28.20	27.96	28.14	27.97
	[0.003]	[0.106]	[0.109]	[0.105]	[0.110]	[0.106]	[0.110]
$Q^2(20)$	390.50	15.68	17.06	16.09	17.02	15.98	16.99
	[0.000]	[0.615]	[0.519]	[0.586]	[0.521]	[0.594]	[0.524]
ARCH5	57.319	1.135	1.354	1.485	1.487	1.444	1.423
	[0.000]	[0.339]	[0.238]	[0.191]	[0.191]	[0.205]	[0.212]
SIC	3.550	3.419	3.407	3.355	3.350	3.353	3.349
							(*)

Notes: See the notes to Table 1.

Table 3: Estimation results for Merck & Company Inc. (MRK)

	Gaussian Models			Symmetric NIG Models		NIG Models	
Model	CV	GARCH	NGARCH	GARCH	NGARCH	GARCH	NGARCH
Loglik	-15087.7	-14759.5	-14742.9	-14604.9	-14596.6	-14594.6	-14586.7
	Estim.	Estim.	Estim.	Estim.	Estim.	Estim.	Estim.
λ	0.0372	0.0490	0.0370	0.0285	0.0231	0.0450	0.0387
	(0.011)	(0.011)	(0.011)	(0.010)	(0.010)	(0.011)	(0.011)
β		0.9204	0.9278	0.9180	0.9257	0.9188	0.9255
		(0.022)	(0.014)	(0.018)	(0.012)	(0.018)	(0.012)
α		0.0516	0.0420	0.0528	0.0447	0.0525	0.0455
		(0.012)	(0.008)	(0.010)	(0.008)	(0.010)	(0.008)
γ			-0.4272		-0.3524		-0.3397
			(0.117)		(0.105)		(0.104)
a				2.0436	2.1117	2.0439	2.0922
				(0.205)	(0.211)	(0.203)	(0.209)
b						0.2086	0.2100
						(0.055)	(0.055)
	Stat.	Stat.	Stat.	Stat.	Stat.	Stat.	Stat.
J-B	3480	983.17	817.01				
	[0.000]	[0.000]	[0.000]				
Q(20)	108.55	94.50	94.00	94.58	94.17	94.56	94.19
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
$Q^2(20)$	1163.18	26.57	30.00	26.30	29.25	26.27	29.12
	[0.000]	[0.088]	[0.037]	[0.093]	[0.045]	[0.094]	[0.047]
ARCH5	79.434	1.931	2.489	1.835	2.203	1.843	2.130
	[0.000]	[0.086]	[0.029]	[0.102]	[0.051]	[0.101]	[0.059]
SIC	3.579	3.502	3.498	3.465	3.463	3.462	3.461
							(*)

Notes: See the notes to Table 1.

Table 4: Estimation results for Standard & Poor's 100 Index (OEX)

	Gaussian Models			Symmetric NIG Models		NIG Models	
Model	CV	GARCH	NGARCH	GARCH	NGARCH	GARCH	NGARCH
Loglik	-8430.4	-7746.5	-7722.6	-7578.6	-7565.6	-7578.1	-7565.3
	Estim.	Estim.	Estim.	Estim.	Estim.	Estim.	Estim.
λ	0.0259	0.0451	0.0269	0.0325	0.0223	0.0371	0.0256
	(0.013)	(0.014)	(0.014)	(0.012)	(0.012)	(0.013)	(0.012)
β		0.9354	0.9267	0.9546	0.9451	0.9544	0.9449
		(0.023)	(0.025)	(0.006)	(0.009)	(0.007)	(0.009)
α		0.0538	0.0492	0.0396	0.0390	0.0397	0.0393
		(0.018)	(0.015)	(0.005)	(0.006)	(0.005)	(0.006)
γ			-0.5001		-0.4825		-0.4801
			(0.091)		(0.095)		(0.095)
a				2.0740	2.1595	2.0655	2.1451
				(0.297)	(0.312)	(0.295)	(0.308)
b						0.0597	0.0517
						(0.067)	(0.066)
	Stat.	Stat.	Stat.	Stat.	Stat.	Stat.	Stat.
J-B	518971	7443.90	5780.44				
	[0.000]	[0.000]	[0.000]				
Q(20)	60.92	45.62	41.44	45.23	41.68	45.20	41.67
	[0.000]	[0.001]	[0.003]	[0.001]	[0.003]	[0.001]	[0.003]
$Q^2(20)$	514.40	16.16	12.29	25.13	17.49	25.05	17.36
	[0.000]	[0.581]	[0.832]	[0.121]	[0.490]	[0.124]	[0.499]
ARCH5	76.409	1.381	0.781	3.120	1.761	3.100	1.733
	[0.000]	[0.228]	[0.563]	[0.008]	[0.117]	[0.008]	[0.123]
SIC	2.888	2.654	2.646	2.597	2.592	2.597	2.592
							(*)

Notes: See the notes to Table 1. For the Standard & Poor's 100 Index data is available only from January 1973 through 1995.

Table 5: Overall performance of the NIG GARCH option pricing model

Panel A: GM

	Dollar errors				ISD errors			
	Put	(611)	Call	(1193)	Put	(611)	Call	(1193)
Model	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	0.325	0.415	0.331	0.418	0.069	0.085	0.064	0.079
GARCH	0.105	0.255	0.121	0.267	0.018	0.047	0.020	0.049
NGARCH	0.084	0.238	0.117	0.256	0.015	0.043	0.019	0.043

Panel B: IBM

	Dollar errors				ISD errors			
	Put	(1795)	Call	(2901)	Put	(1795)	Call	(2901)
Model	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	0.384	0.571	0.417	0.631	0.052	0.077	0.051	0.076
GARCH	0.229	0.523	0.260	0.567	0.026	0.054	0.028	0.054
NGARCH	0.182	0.514	0.275	0.572	0.022	0.052	0.029	0.051

Panel C: MRK

	Dollar errors				ISD errors			
	Put	(544)	Call	(1288)	Put	(544)	Call	(1288)
Model	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	0.199	0.345	0.170	0.370	0.038	0.062	0.031	0.063
GARCH	0.184	0.313	0.163	0.335	0.030	0.049	0.026	0.053
NGARCH	0.171	0.290	0.169	0.322	0.028	0.047	0.025	0.046

Panel D: OEX

	Dollar errors				ISD errors			
	Put	(4804)	Call	(3481)	Put	(4804)	Call	(3481)
Model	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	-0.878	1.697	-1.917	2.447	-0.010	0.049	-0.042	0.052
GARCH	0.628	1.003	-0.186	0.836	0.029	0.051	-0.002	0.024
NGARCH	0.572	0.941	-0.106	0.789	0.026	0.048	0.000	0.024

Notes: This table reports results on the overall pricing performance of the NIG GARCH option pricing model using the Bias and RMSE metrics in Section 5 for dollar price differences as well as for differences in implied standard deviations.

Table 6: Relative performance of the symmetirc NIG and Gaussian GARCH option pricing model

Panel A: GM

	Symmetric NIG relative to NIG				Gaussian relative to NIG			
	Put	(611)	Call	(1193)	Put	(611)	Call	(1193)
Model	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	0.997	0.997	1.005	1.003	0.988	0.995	1.038	1.087
GARCH	0.947	0.985	0.968	0.938	1.213	1.043	1.280	1.071
NGARCH	0.954	0.991	0.990	0.990	1.104	1.014	1.230	1.102

Panel B: IBM

	Symmetric NIG relative to NIG				Gaussian relative to NIG			
	Put	(1795)	Call	(2901)	Put	(1795)	Call	(2901)
Model	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	0.998	1.005	1.010	1.009	0.988	0.998	1.037	1.057
GARCH	0.966	0.997	1.010	1.003	0.986	1.027	1.070	1.074
NGARCH	0.966	0.994	1.018	1.007	0.952	1.013	1.117	1.091

Panel C: MRK

	Symmetric NIG relative to NIG				Gaussian relative to NIG			
	Put	(544)	Call	(1288)	Put	(544)	Call	(1288)
Model	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	0.995	0.997	1.024	1.044	0.980	0.995	1.059	1.126
GARCH	0.996	1.024	0.993	0.983	0.929	0.980	1.010	1.128
NGARCH	0.985	0.992	1.007	1.007	0.941	0.981	1.026	1.083

Panel D: OEX

	Symmetric NIG relative to NIG				Gaussian relative to NIG			
	Put	(4804)	Call	(3481)	Put	(4804)	Call	(3481)
Model	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	1.044	0.985	0.997	0.998	0.947	1.063	1.005	1.008
GARCH	0.990	0.990	0.973	0.997	0.832	0.996	3.146	1.087
NGARCH	0.990	0.989	3.160	0.997	0.734	0.969	190.715	1.068

Notes: This table reports results on the pricing performance of the symmetric NIG GARCH option pricing model and the Gaussian GARCH option pricing model relative to that of the general NIG GARCH option pricing model. Comparison is performed using the Bias and RMSE metrics in Section 5 for the implied standard deviations.

Research Papers 2008

2008-26: Martin Møller Andreasen: Ensuring the Validity of the Micro Foundation in DSGE Models

2008-27: Tom Engsted and Thomas Q. Pedersen: Return predictability and intertemporal asset allocation: Evidence from a bias-adjusted VAR model

2008-28: Frank S. Nielsen: Local polynomial Whittle estimation covering non-stationary fractional processes

2008-29: Per Frederiksen, Frank S. Nielsen and Morten Ørregaard Nielsen: Local polynomial Whittle estimation of perturbed fractional processes

2008-30: Mika Meitz and Pentti Saikkonen: Parameter estimation in nonlinear AR-GARCH models

2008-31: Ingmar Nolte and Valeri Voev: Estimating High-Frequency Based (Co-) Variances: A Unified Approach

2008-32: Martin Møller Andreasen: How to Maximize the Likelihood Function for a DSGE Model

2008-33: Martin Møller Andreasen: Non-linear DSGE Models, The Central Difference Kalman Filter, and The Mean Shifted Particle Filter

2008-34: Mark Podolskij and Daniel Ziggel: New tests for jumps: a threshold-based approach

2008-35: Per Frederiksen and Morten Ørregaard Nielsen: Bias-reduced estimation of long memory stochastic volatility

2008-36: Morten Ørregaard Nielsen: A Powerful Test of the Autoregressive Unit Root Hypothesis Based on a Tuning Parameter Free Statistic

2008-37: Dennis Kristensen: Uniform Convergence Rates of Kernel Estimators with Heterogenous, Dependent Data

2008-38: Christian M. Dahl and Emma M. Iglesias: The limiting properties of the QMLE in a general class of asymmetric volatility models

2008-39: Roxana Chiriac and Valeri Voev: Modelling and Forecasting Multivariate Realized Volatility

2008-40: Stig Vinther Møller: Consumption growth and time-varying expected stock returns

2008-41: Lars Stentoft: American Option Pricing using GARCH models and the Normal Inverse Gaussian distribution