Mandelbrot and the Smile

Thorsten Lehnert [*]

Limburg Institute of Financial Economics (LIFE), Maastricht University,
P.O. Box 616, 6200 MD Maastricht, The Netherlands

First version: January 2005

Abstract

It is a well-documented empirical fact that index option prices systematically differ from Black-Scholes prices. However, previous research provides inconclusive results whether the observed volatility smile could be explained by a discrete-time dynamic model of stock returns with skewed, leptokurtic innovations. The improvements in pricing errors are particularly pronounced for out-of-the money put options, while the models partly underperform a Gaussian alternative for near-the-money options. Motivated by theses empirical evidence, I develop a new GARCH option-pricing model with a more flexible innovation structure. In an application of the model to DAX index options, I test the relative performance of the approach against a standard nested GARCH specification and the well-known practitioners Black-Scholes model. I show that the performance of the truncated Lévy GARCH option pricing model is superior to existing approaches.

* Correspondence to: Tel: +31-43-3883838; Fax: +31-43-3884875. E-mail: t.lehnert@berfin.unimaas.nl.

1. Introduction

There is extensive empirical evidence that observed market prices of traded options systematically differ from Black-Scholes prices. Out-of-the-money calls and puts are relatively overpriced compared to at-the-money options. A fact that is often represented by the well-known volatility smile 'implied out' from observed option prices. For index options, the smile is skewed towards out-of-the money puts. As a result, the implied risk-neutral density function is leptokurtic and heavily skewed to the left. The difference between the actual return distribution of the underlying and the risk-neutral distribution implied out from option prices can only be explained by extremely high levels of risk aversion. Market participants seem to overestimate the probability of extreme downward movements and are willing to purchase overpriced out-of-the money put options.

At the empirical level, several studies have shown that option valuation models with conditional heteroskedasticity and negative correlation between volatility and spot returns capture the particular pattern and significantly improve upon the performance of the Black-Scholes model. The discrete-time GARCH option pricing model has shown to be a flexible, empirically successful model (see among others Heynen, et al (1994), Duan (1996), Heston and Nandi (2000), and Ritchken and Trevor (1999)). Recently, an increasing number of numerical methods for this class of option pricing models become available (see Hanke (1997), Ritchken and Trevor (1999), Duan and Simonato (1998), Duan et al. (2001) and Heston and Nandi (2000)). Heston and Nandi (2000) developed a closed form solution of a GARCH option-pricing model. They show that the single lag version of their model contains Heston's (1993) model as a continuous-time limit, but the discrete-time counterpart is much easier to apply with available data.

However, despite the rather sophisticated modeling approach, a Gaussian model cannot adequately account for the particular pattern observed in option prices. Using the generalized GARCH option-pricing framework of Duan (1999), Lehnert (2003) showed that conditional leptokurtosis and skewness reinforces the effects of conditional heteroskedasticity and asymmetry in the volatility process. His GARCH option pricing model driven by skewed generalized error distributed innovations outperforms the closed-form GARCH option pricing model of Heston and Nandi in-sample as well as out-of-sample. The improvements in pricing errors are particularly pronounced for out-of-the money put and call options, while the model

partly underperforms the Gaussian model for near-the-money options. The results are partly in line with recent results obtained by Christoffersen et al. (2003). They developed and empirically test a GARCH option pricing model with conditional skewness. While they demonstrate the importance of conditional skewness and jumps for the pricing of out-of-the-money puts, their closed-form Inverse Gaussian GARCH option pricing model significantly underperforms a standard Gaussian model for several other types of options. In contrast to Lehnert (2003), Christoffersen et al. (2003) conclude that the overall pricing performance is inferior to the standard Gaussian model. Therefore, the empirical evidence does not necessarily suggest that modeling jumps in returns and volatility in addition to stochastic volatility is the appropriate approach for the purpose of option valuation.

In this paper, I argue that most is to be gained by modeling deviations from normality. In order to investigate the research question further, I empirically investigate whether the use of a more flexible innovations distribution is able to improve the performance for out-of-the-money options without worsening the performance for near-the-money options. A possible candidate might be the truncated Lévy distribution often studied in physics (see Lehnert and Wolff (2004)). Mandelbrot (1963) first proposed the idea that price changes are distributed according to a Lévy stable law. This model was frequently criticized, because the tails are now too fat from a financial modeling perspective and the infinite variance makes it impossible to apply the Central Limit Theorem. The problems with these kinds of distributions are the power law tails, which decay too slowly. This problem can be overcome by taking the Lévy distribution in the central part and introducing a cutoff in the far tails that is faster than the Lévy power law tails. The Lévy distribution with a cutoff and exponentially declining tails was introduced in the physics literature by Mantegna and Stanley (1994) and is known as a truncated Lévy distribution. The exponential decay in the tails ensures that all relevant moments are finite.

In an in-sample and out-of-sample analysis, I compare my model with two benchmarks: a standard nested Gaussian alternative specification and the ad-hoc Black-Scholes model of Dumas, Fleming and Whaley (1998) (DFW) and conclude which method is superior in describing the observed market smile in DAX index options. The paper is arranged as follows. Section 2 describes the GARCH option-pricing framework. Section 3 discusses the data and methodology. In Section 4, I provide the empirical results and finally Section 5 concludes.

2. Econometric Framework

2.1. Option Pricing under GARCH

In a Gaussian discrete-time economy the value of the index at time t, S_t, can be assumed to follow the following dynamics (see e.g. Duan (1995)):

$$r_t = \ln\left(\frac{S_t}{S_{t-1}}\right) + d_t = r_f + \lambda\sigma_t + \sigma_t e_t$$

$$e_t \mid \Omega_{t-1} \sim N(0,1) \text{ under probability measure } P \tag{1}$$

$$\ln(\sigma_t^2) = \omega_0 + \omega_1 \ln(\sigma_{t-1}^2) + \omega_2\left(\mid e_{t-1} \mid - \gamma\, e_{t-1}\right)$$

where d_t is the dividend yield of the index portfolio, r_f is the risk-free rate, λ is the price of risk; Ω_{t-1} is the information set in period t-1 and the combination of ω_2 and γ captures the leverage.

Duan (1995) shows that under the Local Risk Neutral Valuation Relationship (LRNVR) the conditional variance remains unchanged, but under the pricing measure Q the conditional expectation of r_t is equal to the risk free rate r_f:

$$E^Q\left[\exp(r_t)\mid\Omega_{t-1}\right] = \exp(r_f), \tag{2}$$

Therefore, the LRNVR transforms the physical return process to a risk-neutral dynamic. The risk-neutral Gaussian GARCH process reads:

$$r_t = r_f - \tfrac{1}{2}\sigma_t^2 + \sigma_t\varepsilon_t$$

$$\varepsilon_t \mid \Omega_{t-1} \sim N(0,1) \text{ under risk-neutralized probability measure } Q \tag{3}$$

$$\ln(\sigma_t^2) = \omega_0 + \omega_1 \ln(\sigma_{t-1}^2) + \omega_2\left(\mid \varepsilon_{t-1} - \lambda \mid - \gamma\left(\varepsilon_{t-1} - \lambda\right)\right)$$

where the term $-\tfrac{1}{2}\sigma_t^2$ gives additional control for the conditional mean. In Equation (3), ε_t is not necessarily normal, but to include the Black-Scholes model as a special case we typically assume that ε_t is a Gaussian random variable.

2.2. GARCH option pricing with conditional skewness and leptokurtosis

Lévy flights have been observed experimentally in physical systems and have been used very successfully to describe for instance the spectral random walk of a single molecule embedded in a solid. In all these cases an unavoidable cutoff in the tails of the distribution is always present. One possible cutoff is the exponential function, for which the characteristic function (CF) of the so-called truncated Lévy distribution (TLD) has been developed (Koponen (1995)). For financial data the cutoff region is early in the tails, which ensures the finiteness of all relevant moments. The standardized CF with location parameter equal to zero and scale parameter c equal to one reads (Nakao (2000)):

$$\psi_{TLD}(k,\alpha,\delta,\beta) = \frac{\delta^{\alpha} - \left(k^2 + \delta^2\right)^{\alpha/2}}{\cos(\pi\alpha/2)} \cos\left(\alpha\arctan\left(\frac{|k|}{\delta}\right)\right)\left[1 + i\,\text{sgn}(k)\beta\tan\left(\alpha\arctan\left(\frac{|k|}{\delta}\right)\right)\right] \tag{4}$$

where α is the characteristic exponent determining the shape of the distribution and especially the fatness of the tails ($0 < \alpha \le 2$, but $\alpha \ne 1$) and δ is the cutoff parameter, which determines the speed of the decay in the tails and as a result the cutoff region. The parameter β ($\beta \in [-1,1]$) determines the skewness when $\beta \ne 0$, the distribution is skewed to the right when $-1 < \beta < 0$ and skewed to the left when $0 < \beta < 1$. For $\delta \to +0$ the TLD reduces to the Lévy distribution and for $\delta \to +0$, $\beta=0$ and $\alpha=2$ the TLD reduces to the Normal distribution with scale parameter c. For comparison purposes, Figure 1 shows the density of a truncated Lévy distribution with reasonable parameter values for financial return data and the special case of a Gaussian density. Both densities are standardized, such that the scale parameter c equals one. Accurate numerical values for the density ψ_{TLD} can be calculated by Fourier-transforming the CF and evaluating the integral numerically. I use Romberg integration, which allows ex-ante specification of the tolerated error and in fact a calculation of the density as precise as necessary (see Lambert and Lindsey (1999)).

The analogue of the standard deviation σ in the family of Lévy distributions is the scale parameter c. If I replace the standard deviation σ by the scale parameter c, I allow the conditional

scale parameter c_t to be serially correlated and to vary over time. If e_t, conditional on Ω_{t-1} is a skewed truncated Lévy distributed random variable, then the risk-neutral GARCH process reads[1]:

$$r_t = r_f - \ln\left(E^Q\left(\exp\left(c_t\eta_t\right)|\Omega_{t-1}\right)\right) + c_t\eta_t$$

$$\varepsilon_t | \Omega_{t-1} \sim N_{Lévy}(\mu = 0, c = 1) \text{ under risk-neutralized probability measure Q.}$$

$$\eta_t = TLD^{-1}\left(\phi_{Lévy}\left(\varepsilon_t - \lambda\right); \mu = 0, c = 1, \alpha, \delta, \beta\right)$$

$$\ln(c_t) = \omega_0 + \omega_1\ln(c_{t-1}) + \omega_2\left(|\eta_{t-1}| - \gamma\,\eta_{t-1}\right),$$

(5)

where $\phi_{Lévy}[.]$ stands for a standardized normal cumulative distribution with zero mean and scale parameter c equal to 1, $TLD^{-1}[.]$ stands for the inverse skewed truncated Lévy cumulative distribution with standardized mean equal to 0, scale parameter c equal to 1, tail parameter α, skewness parameter β and δ controls for the exponential decay in the tails. The term $E^Q(.|\Omega)$ gives additional control for the conditional mean and can be evaluated numerically. The unconditional volatility level is equal to $\sqrt{\exp\left(\dfrac{\omega_0 + \omega_2\,E\left[\left(|\eta| - \gamma\,\eta\right)\right]}{1 - \omega_1}\right)}$ and can be evaluated numerically. The parameter ω_1 measures the persistence of the variance process.

A European call option with exercise price X and time to maturity T has at time t price equal to:

$$c_t = \exp(-rT)E_t^Q\left[\max\left(S_t - X, 0\right)|\Omega_{t-1}\right]$$

(6)

For this kind of derivative valuation models with a high degree of path dependency, computationally demanding Monte Carlo simulations are commonly used for valuing derivative securities. I use the recently proposed simulation adjustment method, the empirical martingale simulation (EMS) of Duan and Simonato (1998), which has been shown to substantially accelerate the convergence of Monte Carlo price estimates and to reduce the so called

---

[1] See Lehnert (2003) for details regarding GARCH option pricing with conditional leptokurtosis and skewness.

'simulation error'. Monte Carlo simulations are also frequently and successfully used in various risk management applications (e.g. see Bams, Lehnert and Wolff (2002)).

2.3. Alternative option valuation approach

I compare both GARCH option-pricing models with the so-called practitioners Black-Scholes model (see e.g. Dumas, Fleming and Whaley (1998) or Christoffersen and Jacobs (2002b)). I allow each option to have its own Black-Scholes implied volatility depending on the exercise price X and time to maturity T and use the following functional form for σ:

$$\sigma_{ij} = \pi_0 + \pi_1 M_i + \pi_2 M_i^2 + \pi_3 T_j + \pi_4 T_j^2 + \pi_5 M_i T_j, \tag{7}$$

where σ_{ij} denotes the implied volatility and M_i the 'moneyness', $\dfrac{X_i}{F_j}$, (F_j is the forward price) of an option for the i-th exercise price and j-th maturity. For every exercise price and maturity I can compute the implied volatility and derive option prices using the Black-Scholes model.

I calibrate the parameters of the various model by minimizing the root mean squared absolute pricing error between the market prices and the theoretical option prices:

$$RMSE = \sqrt{\frac{1}{N} \min \sum_{i=1}^{n} \sum_{j=1}^{m_i} \left(\hat{c}_{i,j} - c_{i,j} \right)^2} \ \text{ or } \ RMSE = \sqrt{\frac{1}{N} \min \sum_{i=1}^{n} \sum_{j=1}^{m_i} \left(\hat{p}_{i,j} - p_{i,j} \right)^2}, \tag{8}$$

where N is the total number of call options evaluated, the subscript i refers to the n different maturities and subscript j to the m_i different strike prices in a particular maturity series i. Depending on the moneyness of the option, put or call prices are used. Motivated by the actual trading volume, I use put prices for option with moneyness of less than 1 and call prices for options with moneyness of more than 1.

As starting values for the calibration, I make use of the time-series estimates from the augmented GARCH model using approximately three years (752 trading days) of historical returns. Additionally, I use the time-series parameter estimate of the price of risk parameter δ for the option calibration (see e.g. Christoffersen and Jacobs (2002a)).

3. Data and Methodology

I use daily closing DAX 30 index options and futures prices for a period from January 2000 until December 2000. The raw data set is directly obtained from the EUREX, European Futures and Options Exchange. The market for DAX index options and futures is the most active index options and futures market in Europe. Therefore, it is a good market for testing option-pricing models.

For index options the expiration months are the three nearest calendar months, the three following months within the cycle March, June, September and December, as well as the two following months of the cycle June, December. For index futures, the expiration months are the three nearest calendar months within the cycle March, June, September and December. The last trading day is the third Friday of the expiration month, if that is an exchange trading day; otherwise on the exchange-trading day immediately prior to that Friday.

I exclude options with less than one week and more than 25 weeks until maturity and options with a price of less than 2 Euro to avoid liquidity-related biases and because of less useful information on volatilities. Instead of using a static rule and exclude options with absolute moneyness |K/F-1| of more than 10% (see DFW), I exclude options with a daily turnover of less than 10000 Euro (see Lehnert (2003)). Among others DFW argue that options with absolute moneyness of more than 10% are not actively traded and therefore contain no information on volatilities. Therefore, an obvious solution is to filter the available option prices and include all options that are actively traded, inside or outside the 10% absolute moneyness interval. In particular, in volatile periods deep out-of-the money options are highly informative if they are actively traded. As a result, each day I use a minimum of 3 and a maximum of 4 different maturities for the calibration.

The DAX index calculation is based on the assumption that the cash dividend payments are reinvested. Therefore, when calculating option prices, theoretically I don't have to adjust the index level for the fact that the stock price drops on the ex-dividend date. But the cash dividend payments are taxed and the reinvestment does not fully compensate for the decrease in the stock price. Therefore, in the conversion from e.g. futures prices to the implied spot rate, one observe empirically a different implied dividend adjusted underlying for different maturities. For this reason, I always work with the adjusted underlying index level implied out from futures or option market prices.

In particular I'm using the following procedure for one particular day to price options on the following trading day:

First, I compute the implied interest rates and implied dividend adjusted index rates from the observed put and call option prices. I'm using a modified put-call parity regression proposed by Shimko (1993). The put-call parity for European options reads:

$$c_{i,j} - p_{i,j} = [S_t - PV(D_j)] - X_i e^{-r_j(T_j - t)} \tag{9}$$

where $c_{i,j}$ and $p_{i,j}$ are the observed call and put closing prices, respectively, with exercise prices X_i and maturity $(T_j\text{-}t)$, $PV(D_j)$ denotes the present value of dividends to be paid from time t until the maturity of the options contract at time T_j and r_j is the continuously compounded interest rate that matches the maturity of the option contract. Therefore, I can infer a value for the implied dividend adjusted index for different maturities, $S_t\text{-}PV(D_j)$, and the continuously compounded interest rate for different maturities, r_j. In order to ensure that the implied dividend adjusted index value is a non-increasing function of the maturity of the option, I occasionally adjust the standard put-call parity regression. Therefore, I control and ensure that the value for $S_t\text{-}PV(D_j)$ is decreasing with maturity, T_j. Since I'm using closing prices for the estimation, one alternative is to use implied index levels from DAX index futures prices assuming that both markets are closely integrated.

Second, I estimate the parameters of the particular models by minimizing the loss function (8). Given reasonable starting values, I price European call options with exercise price X_i and maturity T_j. Using well-known optimization methods (e.g. Newton-Raphson method), I obtain the parameter estimates that minimize the loss function. The goodness of fit measure for the optimization is the mean squared valuation error criterion.

Third, having estimated the parameters in-sample, I turn to out-of-sample valuation performance and evaluate how well each day's estimated models value the traded options at the end of the following day. I filter the available option prices according to our criteria for the in-sample calibration. The futures market is the most liquid market and the options and the futures market are closely integrated, therefore it can also be assumed that the futures price is more informative for option pricing than just using the value of the index. For every observed futures

closing price I can derive the implied underlying index level and evaluate the option. Given a futures price F_j with time to maturity T_j, spot futures parity is used to determine S_t-$PV(D_j)$ from

$$S_t - PV(D_j) = F_j \mathrm{e}^{-r_j T_j} \qquad (10)$$

where $PV(D_j)$ denotes the present value of dividends to be paid from time t until the maturity of the options contract at time T_j and r_j is the continuously compounded interest rate (the interpolated EURIBOR rate) that matches the maturity of the futures contract (or time to expiration of the option). If a given option price observation corresponds to an option that expires at the time of delivery of a futures contract, then the price of the futures contract can be used to determine the quantity S_t-$PV(D_j)$ directly.

The maturities of DAX index options do not always correspond to the delivery dates of the futures contracts. In particular for index options the two following months are always expiration months, but not necessarily a delivery month for the futures contract. When an option expires on a date other than the delivery date of the futures contract, then the quantity S_t-$PV(D_j)$ is computed from various futures contracts. Let F_1 be the futures price for a contract with the shortest maturity, T_1 and F_2 and F_3 are the futures prices for contracts with the second and third closest delivery months, T_2 and T_3, respectively. Then the expected future rate of dividend payment d can be computed via spot-futures parity by:

$$d = \frac{r_3 T_3 - r_2 T_2 - \log(F_3 / F_2)}{(T_3 - T_2)} \qquad (11)$$

Hence, the quantity S_t-$PV(D) = S_t \mathrm{e}^{-dT}$ associated with the option that expires at time T in the future can be computed by[2]

$$S_t \mathrm{e}^{-dT} = F_1 \mathrm{e}^{(-(r_1 - d)T_1 - dT)}. \qquad (12)$$

---

[2] See e.g. the appendix in Poteshman (2001) for details.

This method allows us to perfectly match the observed option price and the underlying dividend adjusted spot rate. Given the parameter estimates and the implied dividend adjusted underlying I can calculate option prices and compare them to the observed option prices of traded index options. For the out-of-sample part the same loss functions for call options are used. The prediction performance of the various models are evaluated and compared by using the root mean squared valuation error criterion. I compare the predicted option values with the observed prices for every traded option. I repeat the whole procedure over the out-of-sample period and conclude, which model minimizes the out-of-sample pricing error.

4. Empirical results

In this section, I present estimation and evaluation results using 254 days of option data from the year 2000. Each trading day, on average 85 option prices are used for the calibration and evaluation of the models, with a minimum of 62 and a maximum of 155. The numbers of option contracts in the in-sample and out-of-sample data set are reported across moneyness and maturity in Table 1.

[Table 1]

In total more than 21400 option contracts have been evaluated. The number contracts are nearly equally distributed over the five different moneyness categories, except the fewer number options with moneyness of less than 0.92. Most options are short term (<21 trading days until maturity) or long term (>63 days until maturity), but there is also a substantial number options traded in the medium term.

All models are calibrated using the Euro root mean squared error loss function (8). For the GARCH models, theoretical option prices are obtained using (5), and the local scale parameter c_{t+1} is estimated together with the other parameters. I additionally use a time-series parameter estimate for the option calibration: the price of risk parameter λ. The joint identification of λ and γ is possible but not reasonable, since both parameters control for asymmetry by modifying the news impact curve.

Table 2 shows the averages of daily parameter estimates of the three models analyzed, obtained by minimizing the squared Euro pricing error, using data for 253 trading days in the year 2000.

[Table 2]

For the models that are investigated, estimating the parameters using a single day of option prices is a convenient approach. The objective function is always well behaved and no numerical problems have been encountered. It is well-known fact that the parameters of the ad-hoc Black-Scholes model can be significantly estimated, but substantially vary over time, already on a day-to-day basis. Similar results are obtained for the particular option data set under investigation (results not reported). For the GARCH option pricing models, parameter estimates are obtained that are more stable and in line with the ones expected from a time-series calibration. The value of 0.98 for ω_1 suggests strong mean reversion of the volatility process. The significant estimates for ω_2 and γ suggest that the volatility process is asymmetric, meaning that returns and volatility are negatively correlated and resulting in negative skewness in the simulated multiperiod index returns. However, in the case of the truncated Lévy GARCH option pricing model, this effect is augmented by negative skewness in the innovations distribution. Additionally, also the other parameter estimates suggest that the data dictate a non-normal innovations distribution: the tail fatness parameter α is different from 2 and the 'cutoff' parameter δ is different from 0. A value of around 0.25 for δ implies that the exponential decay is introduced earlier in the tails, which rejects the extremely fat-tailed Lévy distribution as a possible alternative.

The in-sample relative pricing errors are in the range of 1–2 Euros; usually slightly smaller for the truncated Lévy GARCH option pricing model (on average around €1.5) and the ad-hoc Black-Scholes model (on average around €2) compared to the Gaussian GARCH option valuation model (on average around €2.2). Table 3 reports the results for the in-sample pricing errors of the various models.

[Table 3]

In general, the in-sample results are in line with the empirical findings of Christoffersen et al. (2003) and Lehnert (2003): The more flexible innovations structure improves the pricing performance of the GARCH option pricing model significantly. The results are consistent over different levels of moneyness and for various maturities. The good pricing performance of the GARCH model with conditional skewness and leptokurtosis cannot be confirmed for the Gaussian alternative: on average the model underperforms the ad-hoc Black-Scholes model. However, it is a well-known fact that the ad-hoc Black-Scholes model typically overfits the data in-sample, but when evaluated out-of-sample, it typically underperforms GARCH-type option pricing approaches (Heston and Nandi (2000)). Additionally Christoffersen et al. (2003) show that a GARCH model with conditional skewness and jumps overfits the data in-sample, but underperforms a Gaussian model when evaluated out-of-sample. Therefore, we cannot rely on the in-sample results, but an out-of-sample analysis has to be conducted.

The out-of-sample valuation errors are presented in Table 4. Results strongly suggest that the findings of Christoffersen et al. (2003) cannot be confirmed for a more recent data set of DAX options: the out-of-sample valuation errors of the truncated Lévy GARCH model are on average lower compared to the Gaussian alternative and the ad-hoc Black-Scholes model.

[Table 4]

The improvements in the pricing performance are in particularly pronounced for short-term out-of-the-money put options. Christoffersen et al. (2003) presented empirical evidence that the pricing performance of the GARCH model with skewness and jumps is inferior to the Gaussian alternative especially for longer-term options. Nevertheless, it is remarkable that also with the approach proposed in this study, the pricing performance worsens for longer-term options, but still remains superior compared the Gaussian alternative. One might conclude that the more flexible innovations structure of the truncated Lévy GARCH option pricing model does not seem to overfit the data in-sample, but results in significant valuation improvements for all types of options. Therefore, one might argue that most is to be gained from modeling deviations from normality, but less is to be gained from modeling jumps in returns and volatility in addition to stochastic volatility. In addition, the findings of Heston and Nandi (2000) can be confirmed: the

alternative Gaussian GARCH model outperforms the ad-hoc Black-Scholes model out-of-sample. As expected, the alternative approach of modeling the scale parameter c instead of the variance does not result in a different pricing performance of the Gaussian GARCH option valuation model.

However, one result is in particularly interesting: for all approaches the pricing performance dramatically worsens when the maturity of the option contract increases. This is in line with the results of Christoffersen et al. (2003) and shows that still not even the more sophisticated GARCH approach adequately captures the volatility dynamics underlying option prices.

5. Conclusions

This paper presents a new option valuation model that is based on a return dynamic that contains conditional skewness and leptokurtosis as well as conditional heteroskedasticity and a leverage effect. The truncated Lévy GARCH option pricing model nests a standard GARCH model, which contains Gaussian innovations, and the empirical comparison between our new model and the standard GARCH model investigates the importance of modeling conditional skewness and leptokurtosis. Our empirical results are strongly in favor of the new modeling approach: the truncated Lévy GARCH option pricing model achieves a better fit than standard models in-sample and out-of-sample. The improvements in the pricing performance are particularly pronounced for short-term deep out-of-the-money puts, but it also performs better than standard models for longer terms and for several other types of options.

References

Bams, D., T. Lehnert and C.C.P. Wolff (2002), "An Evaluation Framework for Alternative VaR Models", forthcoming *Journal of International Money and Finance*.

Christoffersen, P. and K. Jacobs (2002a), "Which Volatility Model for Option Valuation?", Working Paper, McGill University.

Christoffersen, P. and K. Jacobs (2002b): "The Importance of the Loss Function in Option Valuation", forthcoming *Journal of Financial Economics*.

Christoffersen, P., S. Heston and K. Jacobs (2003), "Option Valuation with Conditional Skewness", Working Paper, McGill University.

Duan, J.-C. (1995), "The GARCH Option Pricing Model", Mathematical Finance 5, 13-32.

Duan, J.-C. (1996), "Cracking the Smile", RISK 9, 55-59.

Duan, J.-C. and J.-G. Simonato (1998), "Empirical Martingale Simulation for Asset Prices", *Management Science* 44, 1218-1233.

Duan, J.-C. (1999), "Conditional Fat-Tailed Distributions and the Volatility Smile in Options", Rotman School of Management, University of Toronto, Working Paper.

Duan, J.-C. Gauthier G. and J.-G. Simonato (2001), "Asymptotic Distribution of the Empirical Martingale Simulation Option Price Estimator", to appear in *Management Science*.

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

Heston, S.L. (1993), "A closed-form solution for options with stochastic volatility with applications to bond and currency options", *The Review of Financial Studies*, 6, 327-343.

Heston, S.L. and S. Nandi (2000), "A Closed-Form GARCH Option Valuation Model", *The Review of Financial Studies* 3, 585-625.

Heynen, R., A. Kemma and T. Vorst (1994), "Analysis of the term structure of implied volatilities", *Journal of Financial and Quantitative Analysis*, 29, 31-56.

Koponen, (1995), "Analytic approach to the problem of convergence of truncated Lévy flights towards the Gaussian stochastic process", *Physical Review E*, Volume 52, Number 1, 1197-1199.

Lambert P. and J.K. Lindsey (1999), "Analysing financial returns by using regression models based on non-symmetric stable distributions", *Journal of Applied Statistics*, 48, 409-424.

Lehnert, T. (2003)", "Explaining Smiles: GARCH Option Pricing with Conditional Leptokurtosis and Skewness, *Journal of Derivatives*, 10, 3, 27-39.

Lehnert, T. and C.C.P. Wolff (2004), "Scale-Consistent Value-at-Risk", forthcoming *Finance Research Letters*.

Mandelbrot, B. (1963), "The Variation of Certain Speculative Prices," *Journal of Business*, 36, 394-419.

Mantegna, R. N. and H. E. Stanley (1994), "Stochastic process with ultraslow convergence to a Gaussian: the truncated Lévy flight", *Phys. Rev. Lett.*, 73, 2946-2949.

Nakao, H., (2000), "Multi-scaling properties of truncated Lévy flights", Manuscript, Graduate School of Mathematical Sciences, University of Tokyo.

Poteshman, A.M. (2001), "Underreaction, Overreaction, and Increasing Misreaction to Information in the Option Market", *Journal of Finance*, 56, 3, 851-876.

Ritchken, P. and R. Trevor (1999), "Pricing options under generalized GARCH and stochastic volatility processes", *Journal of Finance*, 54, 377-402.

Shimko, D. (1993), "Bounds of probability", *RISK*, 6, 33-37.

Figure 1: Truncated Lévy Density

The graph depicts a comparison of a truncated Lévy density and a normal density. Both densities are standardized, such that the scale parameter c is equal to one.
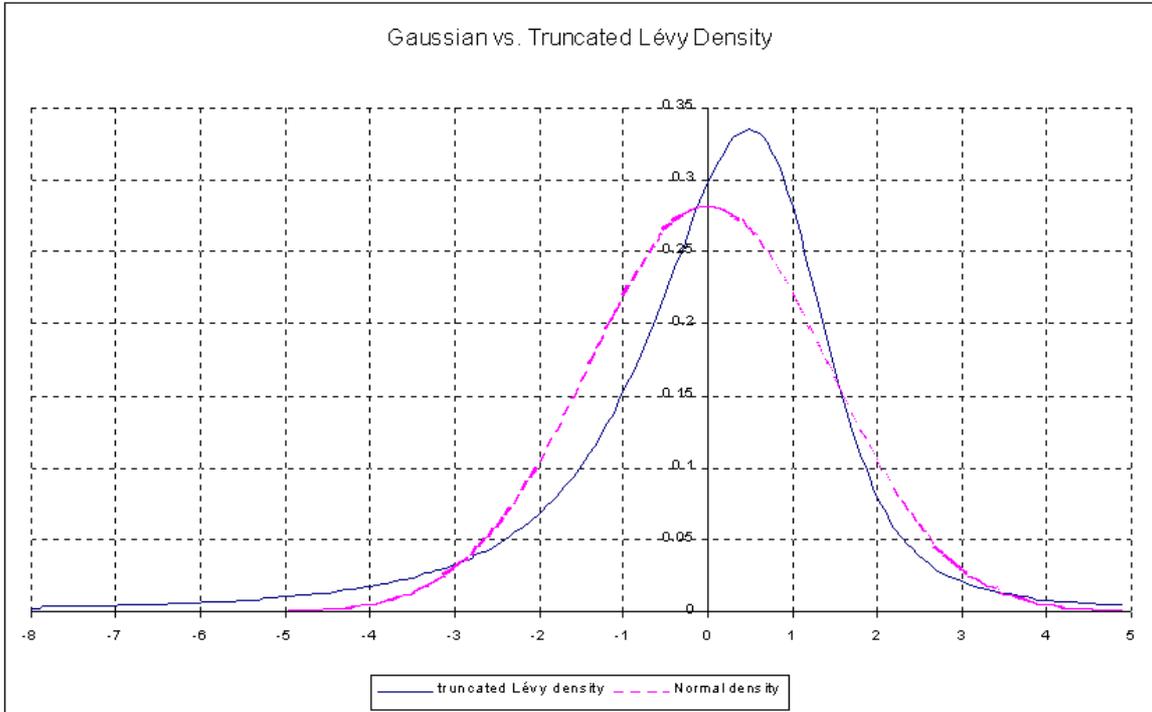

Gaussian vs. Truncated Lévy Density

— truncated Lévy density   — — — Normal density

Table 1: Number of Observations

| Moneyness | Days to Expiration | | | | | | | | | Total |
| | < 21 | | | [21,63] | | | > 63 | | | |
	Total	Daily Median	Daily Max	Total	Daily Median	Daily Max	Total	Daily Median	Daily Max	
< 0.92	909	3	16	634	3	6	1712	7	16	**3255**
[0.92,0.96)	2918	9	44	1017	5	7	728	3	6	**4663**
[0.96,1.00)	2034	7	26	1816	7	17	1074	4	15	**4924**
[1.00,1.04)	753	7	36	818	3	10	2791	10	42	**4362**
> 1.04	1444	6	12	1030	5	7	1728	7	15	**4202**
Total	**8058**	**5**	**44**	**5315**	**4**	**17**	**8033**	**6**	**42**	**21406**

Notes: This Table reports the number of observations for different moneyness and maturity during the period January 2000 until December 2000. 'Moneyness' is defined in the following way: a put option is said to be near- or at-the-money if $K/F \in [0.96,1.00]$ or $K/F \in [1.00,1.04)$, out of the money if $K/F \in [0.92,0.96)$, deep out-of-the-money if $K/F < 0.92$ and in the money if $K/F > 1.04$, where K is the strike price and F is the forward price. Similar terminology is defined for calls by replacing K/F by F/K. 'Days to Expiration' is the number of trading days until maturity. 'Total' is the total number of options priced during the period August 2000 until March 2001 of the particular maturity and/or exercise price. 'Median' and 'Max' are the median and the maximum number of options priced on one trading day of the particular moneyness and/or maturity.

Table 2: Parameter Estimates

	Ad-hoc Black-Scholes		Gaussian GARCH	Truncated Lévy GARCH
Parameter	Mean	Parameter	Mean	Mean
π_0	0.9027	ω_0	-0.2274	-0.1548
π_1	-1.0026	ω_1	0.9713	0.9801
π_2	0.3213	ω_2	0.0925	0.0602
π_3	-1.1039	γ	0.6821	0.9781
π_4	-0.1133	$c_{\tau+1}$	0.0123	0.0124
π_5	0.2394	α	2	1.7211
		δ	0	0.2532
		β	0	0.1987

Notes. The table presents the average parameters estimates of the daily estimations of the various models during the period January 2000 until December 2000.

Table 3: In-Sample Analysis: Average Valuation Errors across moneyness and maturity

	Ad-hoc Black-Scholes				EGARCH-Normal				EGARCH-truncated Lévy			
	Days to Expiration				Days to Expiration				Days to Expiration			
Moneyness	< 21	[21,63]	> 63	**Total**	< 21	[21,63]	> 63	**Total**	< 21	[21,63]	> 63	**Total**
< 0.92	3.04	1.02	1.69	**1.81**	2.23	1.34	2.41	**2.13**	1.81	0.99	1.57	**1.53**
[0.92,0.96)	1.68	1.96	0.79	**1.45**	1.89	2.22	1.06	**1.66**	1.64	1.19	0.78	**1.36**
[0.96,1.00)	1.76	1.78	2.02	**1.84**	1.96	1.80	1.93	**1.88**	1.27	1.78	1.71	**1.75**
[1.00,1.04)	2.98	0.73	2.45	**2.00**	2.35	0.94	1.62	**1.92**	2.01	0.76	1.44	**1.51**
> 1.04	1.29	1.94	0.96	**1.36**	1.51	2.43	1.06	**1.41**	1.12	1.68	0.94	**1.21**
Total	**2.06**	**1.45**	**1.57**	**1.92**	**2.09**	**1.77**	**1.81**	**1.98**	**1.63**	**1.34**	**1.32**	**1.46**

Notes: This Table reports the average absolute in-sample pricing errors of the alternative models for different moneyness and maturity during the period January 2000 until December 2000. 'Moneyness' is defined in the following way: an option is said to be near- or at-the-money if $K/F \in$ [0.96,1.00] or $K/F \in$ [1.00,1.04), out of the money if $K/F \in$ [0.92,0.96), deep out-of-the-money if $K/F < 0.92$ and (deep) in the money if $K/F >$ 1.04, where K is the strike price and F is the forward price. 'Days to Expiration' is the number of trading days until maturity.

Table 4: Out-of-Sample Analysis: Average Valuation Errors across moneyness and maturity

	Ad-hoc Black-Scholes				EGARCH-Normal				EGARCH-truncated Lévy			
	Days to Expiration				Days to Expiration				Days to Expiration			
Moneyness	< 21	[21,63]	> 63	Total	< 21	[21,63]	> 63	Total	< 21	[21,63]	> 63	Total
< 0.92	3.14	7.65	7.23	**6.27**	2.64	6.42	6.15	**5.76**	1.95	5.77	5.46	**4.83**
[0.92,0.96)	3.74	5.29	8.17	**5.74**	3.62	5.67	7.54	**5.87**	3.47	4.40	7.35	**5.58**
[0.96,1.00)	5.38	7.30	3.15	**5.43**	5.42	7.02	3.57	**5.49**	4.78	6.69	2.67	**4.98**
[1.00,1.04)	4.28	8.21	5.42	**6.07**	4.10	7.65	4.52	**5.79**	3.33	7.46	4.26	**5.31**
> 1.04	6.11	5.23	6.49	**5.99**	6.02	5.50	5.98	**5.75**	5.87	4.32	5.79	**5.46**
Total	**4.62**	**6.84**	**6.20**	**6.05**	**4.49**	**6.58**	**6.19**	**5.73**	**4.01**	**6.01**	**5.43**	**5.14**

Notes: This Table reports the average absolute out-of-sample pricing errors of the alternative models for different moneyness and maturity during the period January 2000 until December 2000. 'Moneyness' is defined in the following way: an option is said to be near- or at-the-money if $K/F \in [0.96,1.00]$ or $K/F \in [1.00,1.04)$, out of the money if $K/F \in [0.92,0.96)$, deep out-of-the-money if $K/F < 0.92$ and (deep) in the money if $K/F > 1.04$, where K is the strike price and F is the forward price. 'Days to Expiration' is the number of trading days until maturity.