NORGES TEKNISK-NATURVITENSKAPELIGE UNIVERSITET

Working paper no. N1-2010
Department of Mathematical Sciences, NTNU

Pricing of barrier options under a NIG market model using numerical path integration

by

Arvid Naess, Eivind Aukrust and Sjur Westgaard

PREPRINT
STATISTICS NO. 13/2010

NORWEGIAN UNIVERSITY OF SCIENCE AND TECHNOLOGY
TRONDHEIM, NORWAY

This preprint has URL
`http://www.math.ntnu.no/preprint/statistics/2010/S13-2010.pdf`
Arvid Naess has homepage: `http://www.math.ntnu.no/~arvidn`
E-mail: `arvidn@math.ntnu.no`
Address: Department of Mathematical Sciences, Norwegian University of Science and Technology, NO-7491 Trondheim, Norway.

Pricing of barrier options under a NIG market model using numerical path integration

Arvid Naess*, Eivind G. Aukrust† & Sjur Westgaard‡

Abstract

In recent years it has been demonstrated that the Normal Inverse Gaussian (NIG) process can be used to model stock returns. The characteristics of the NIG market model compares favourably with empirical findings in the financial markets. In the paper it is briefly described how vanilla options can be priced efficiently, and how the corresponding model option prices match with actual option surfaces. More importantly, it is demonstrated how discretely monitored barrier options can be calculated very fast and accurately under the NIG market model using the numerical path integration approach. Several numerical examples are presented to highlight accuracy and efficiency.

Keywords: Barrier options, discrete monitoring, NIG dynamics, numerical path integration.

*Department of Mathematical Sciences, Norwegian University of Science and Technology, NO-7491 Trondheim, Norway, e-mail: arvidn@math.ntnu.no

†Norges Bank Investment Management, Equity Trading, NO-0240 Oslo, Norway, email: ega@nbim.no

‡Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, NO-7491 Trondheim, Norway, email: sjur.westgaard@hist.no

1 Introduction

Due to the many shortfalls of the Black-Scholes model (Black and Scholes, 1973), financial modelers are constantly trying to find more realistic models for the dynamics of financial instruments. Stochastic volatility and jump-diffusion models are frequently referenced in the literature, see e.g. McDonald (2002); Rebonato (2004); Hull (2006); Haug (2007). An alternative approach is to use different kinds of Lévy processes as a modelling tool. Madan et al. (1998) introduced the three parameter Variance Gamma process, and Barndorff-Nielsen (1994) introduced the three parameter Normal Inverse Gaussian (NIG) process applied to derivatives pricing. The present paper will cover option pricing under the NIG market model. The NIG model allows for both skewness and kurtosis in the stock return distributions, and since it only has three parameters, the calibration process is comfortable.

The topic of this paper is pricing of barrier options, which is the most popular class of exotic options (Haug, 2007). Within the Black and Scholes framework closed form pricing formulas for barrier options exist (Merton, 1973; Haug, 2007). In this paper we present a new way of pricing path-dependent barrier options under the NIG market model. Since no closed form solutions are available, we do this by using a numerical method based on path integration. While most texts suggest using numerical integration for pricing vanilla options, Monte Carlo methods seem to be the method most often suggested for pricing path dependent options like the barrier options. As standard numerical integration is used for pricing vanilla options, numerical path integration can be used for pricing many path dependent options. While Monte Carlo methods are flexible and easy to implement, they tend to be computationally slow, even if variance reduction techniques can be used to speed up the calculations. The motivation for a faster numerical implementation is thus strong. The numerical path integration approach presented in this paper is very easy to understand and implement, and it can be used to calculate barrier option prices in about one second on a regular laptop.

This paper is organized as follows. In Section 2 we describe the NIG process, while Section 3 discusses the NIG market model. Section 4 describes the data for selected equity markets used for the calculations. In Section 5 is discussed the calibration of the NIG models to plain vanilla options, while Section 6 details the pricing of barrier options by path integration using the calibrated models from Section 5. The numerical results for the chosen markets are presented in Section 7. Some concluding remarks are given in Section 8.

2 The Normal Inverse Gaussian Process

The Normal Inverse Gaussian (NIG) distribution is a three parameter probability distribution that allows for both skewness and higher kurtosis than the Normal (Gaussian) distribution. These properties make it a potentially suitable tool for modelling of stock returns (Barndorff-Nielsen, 1994).

We will use the notation $X \sim \mathrm{NIG}(\alpha, \beta, \delta)$ to signify that the random variable X has a NIG distribution with the three parameters α, β and δ, where $\alpha > 0$, $-\alpha < \beta < \alpha$, and $\delta > 0$. The probability density function of a NIG variable is given by (Schoutens, 2003):

$$f_{NIG}(x;\alpha,\beta,\delta) = \frac{\alpha\delta}{\pi}\exp(\delta\sqrt{\alpha^2-\beta^2}+\beta x)\frac{K_1(\alpha\sqrt{\delta^2+x^2})}{\sqrt{\delta^2+x^2}}\,, \tag{1}$$

where $K_1(\cdot)$ denotes the modified Bessel function of the second kind, for details see e.g. Abramowitz and Stegun (1968). We define the NIG process with parameters α, β, δ as,

$$X^{NIG} = \left\{X_t^{NIG},\, t \geqslant 0\right\} \tag{2}$$

with $X_0^{NIG} = 0$, and with stationary and independent NIG distributed increments, with $X_{s+t}^{NIG} - X_s^{NIG} \sim \mathrm{NIG}(\alpha,\beta,\delta t)$.

The name of the NIG process is due to the fact that we can relate the NIG process to an Inverse Gaussian time-changed Brownian motion. The density function for the IG(a, b) distribution is given by $(x > 0)$

$$f_{IG}(x;a,b) = \left[\frac{a}{2\pi x^3}\right]^{1/2}\exp\left(\frac{-a(x-b)^2}{2b^2x}\right). \tag{3}$$

Let I_t be an IG process with parameters a=1 and $b = \delta\sqrt{\alpha^2-\beta^2}$ with α >0, $|\alpha| > \beta$ and $\delta > 0$, and W_t be a standard Brownian motion. In (Schoutens, 2003) it is stated that the stochastic process

$$X_t = \beta\delta^2 I_t + \delta W_{I_t} \tag{4}$$

is a NIG process with parameters α, β and δ. This representation of a NIG process makes it easy to simulate. An example of a NIG sample path is shown in Fig. 1.

The characteristics of the NIG process are summarized in Table 1. It is seen that it is not easy to directly recognize what effect each of the three parameters has on the four different statistical moments. The intuitive understanding of the parameters is that α determines the tail behavior, β the skewness and δ is the time scaling parameter.

	NIG$(\alpha,\ \beta,\ \delta)$
Mean:	$\delta\beta/\sqrt{\alpha^2-\beta^2}$
Variance:	$\alpha^2\delta(\alpha^2-\beta^2)^{-3/2}$
Skewness:	$3\beta\alpha^{-1}\delta^{-1/2}(\alpha^2-\beta^2)^{-1/4}$
Kurtosis:	$3(1+\frac{\alpha^2+4\beta^2}{\delta\alpha^2\sqrt{\alpha^2-\beta^2}})$

Table 1: NIG distribution characteristics

3 The NIG Market Model

We want to model the log returns of stocks as a NIG-distribution. The model for the stock is thus given by:

$$S_t = S_0\exp(X_t^{NIG})\,, \tag{5}$$

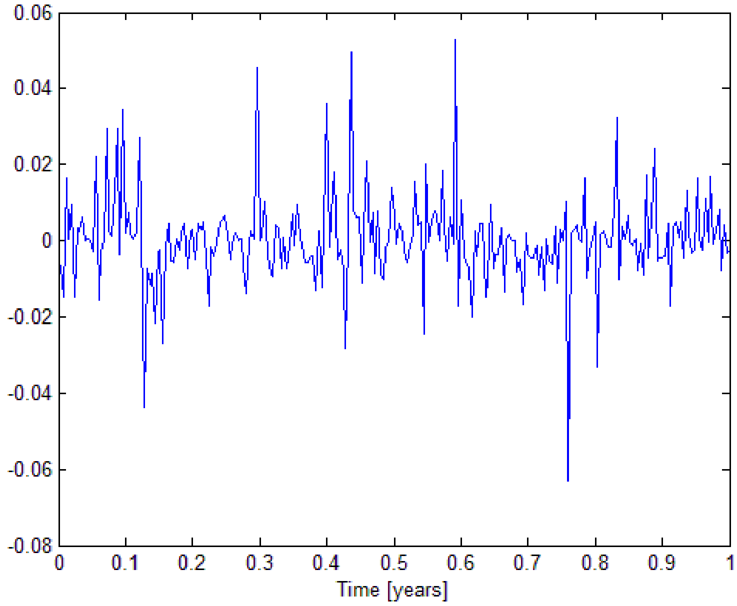

Figure 1: A NIG sample path. α = 50, β = -5, δ = 1.

where S_t is the stock price at time t. This model provides a NIG distribution for the log-returns of the stock. Although this gives us a much more realistic behavior of the stock than the geometric Brownian motion case, there are some drawbacks. For example, it is impossible to perform a continuous delta hedging strategy since the NIG model does not lead to a continuous asset price.

In the geometric Brownian motion case, moving from the real world to the risk-neutral world is easy; we just set the drift μ equal to the risk free rate r minus the continuous dividend yield q. However, for our model given in Eq. (5) there is no unique transformation from the real world to the risk neutral world. One possible way to achieve this is to use an Esscher transformation (Gerber and Shiu, 1994). Alternatively, Schoutens (2003) shows how it can be done by applying a mean-correcting measure change as follows,

$$S_t = S_0 \frac{e^{(r-q)t}}{\mathrm{E}[\exp(X_t^{NIG})]} \exp(X_t^{NIG}) \,. \tag{6}$$

It is clear that multiplication by the factor $e^{(r-q)t}/\mathrm{E}[\exp(N_t)]$ indeed provides a risk-neutral process. $\mathrm{E}[\exp(X_t^{NIG})]$ can be calculated by using the characteristic function for X_t^{NIG}, which is,

$$\mathrm{E}[\exp(\mathrm{i}uX_t^{NIG})] = \exp\left(-\delta t(\sqrt{\alpha^2-(\beta+\mathrm{i}u)^2}-\sqrt{\alpha^2-\beta^2})\right). \tag{7}$$

By putting $u = -\mathrm{i}$, it follows that

$$\mathrm{E}[\exp(X_t^{NIG})] = \exp\left(-\delta t(\sqrt{\alpha^2-(\beta+1)^2}-\sqrt{\alpha^2-\beta^2})\right). \tag{8}$$

Summing up, we may write

$$S_t = S_0 \exp\left(mt+\omega t+X_t^{NIG}\right), \tag{9}$$

where $m = r - q$ and $\omega = \delta(\sqrt{\alpha^2 - (\beta+1)^2} - \sqrt{\alpha^2 - \beta^2})$. While $m = r - q$ for the risk neutral case, in the real world, m must be estimated from historical data. It is important to notice that all parameters are changed when moving from the real world to the risk-neutral world in the NIG market model, unlike the geometric Brownian motion case, where the volatility stays the same under the Girsanov measure transformation. Using the NIG pdf from Eq. (1), the pdf for the log price $Z_t = \ln(S_t/S_0)$ of the NIG market model given in Eq. (9) is then,

$$f(z;t) = \frac{\alpha\delta t}{\pi} \exp\left(\delta t\sqrt{\alpha^2 - \beta^2} + \beta(z - \mu t)\right) \frac{K_1\left(\alpha\sqrt{\delta^2 t^2 + (z - \mu t)^2}\right)}{\sqrt{\delta^2 t^2 + (z - \mu t)^2}}, \quad (10)$$

where $\mu = m + \omega$. Depending on the use, m, α, β and δ will be either the risk-neutral or the real world parameters of the NIG process.

4 Data

The whole motivation for working with the more complex NIG process instead of a geometric Brownian motion is that the NIG process allows for both skewness and excess kurtosis, and should thus be able to fit historical stock market returns better. By way of verification, statistical analyses on historical data were carried out to make sure the NIG distribution provides a good fit for the stock indices under study.

4.1 Parameter estimation:

The three parameters of the NIG pdf were estimated using the maximum-likelihood method. In practice this is done numerically using optimization algorithms, e.g. the Nelder-Mead downhill simplex method, which is implemented in the built-in Matlab function *fminsearch*. As always with numerical optimization, one has to be careful when picking starting values for the algorithm. However, there is a nice way for automating the choice of starting values on a computer. This can be done by using the time series estimates of the different moments of the distribution, and then solve the corresponding equations in Table 1 numerically to obtain reasonable starting values.

4.2 Data analysis

The NIG process is calibrated to four well known stock indices. The length of the time series is different for the different indices, as reported in Table 2. The daily log returns are used for every trading day in the data set.

The calibration results are given in Table 3.

Although the NIG distribution generally provided a good fit to the data, it was clear from QQ-Plots of the S&P 500 and the FTSE 100 data that the NIG model is far from perfect. On both these plots one single data point clearly deviated from the straight line fit. This was due to the stock market crash on 19^{th} October 1987, and is an important observation. Even if the NIG-distribution

Index:	Data from:	Data to:
S&P 500	03.01.1950	07.01.2008
OMXS30	06.07.2004	28.12.2007
FTSE 100	02.04.1984	11.02.2008
OSEBX	25.05.2001	11.02.2008

Table 2: Short description of our data sets.

Index:	α	β	δ	m
S&P 500	96.2420	-5.3752	1.8435	0.0857
OMXS30	92.2307	-19.3533	2.4249	0.1400
FTSE 100	105.2935	-11.0490	2.6497	0.0808
OSEBX	79.6223	-17.3478	2.9745	0.1163

Table 3: NIG calibration results.

may give an overall good fit to stock index returns, it would generally underestimate the probability of an extreme event like the stock market crash of 1987.

P-values were calculated for the null-hypothesis that the given data sets were generated from a NIG distribution with parameters given in Table 3. This was done using a Pearson statistic. The obtained P-values for the indices given in Table 2 both for the NIG and the normal distributions, with their estimated parameter values are presented in Table 4. These values confirm what was observed graphically, the P-values for the NIG-distribution are very high and the null-hypothesis is in all cases accepted. The P-values for the normal distribution are as expected very small (less than $10^{-10} \approx 0$), and the null-hypothesis is rejected for all the indices. To conclude, the NIG-distribution seems to fit the stock index returns reasonably well, which has already been observed by Bølviken et al. (2000). The motivation for using the NIG-process on the indices at hand is thus strong.

Index:	NIG P-value:	Normal P-value:
S&P 500	0.3889	0
OMXS30	0.3151	0
FTSE 100	0.2366	0
OSEBX	0.6368	0

Table 4: P-values from the χ^2-tests.

5 Vanilla option pricing and calibration to market data

5.1 Pricing of a European call option

There is no closed form analytical solution for the European type option prices in the NIG market model as within the Black-Scholes framework. However, the option prices can be easily obtained by numerical integration over the density function. Although slower than an analytical solution, it is acceptable for any practical pricing purpose.

The formula for a European call option with strike X and time to maturity T is given by:

$$V(X,T) = \exp(-rT) \int_X^\infty (S-X) p_T(S|S_0) dS \tag{11}$$

Here $p_T(S|S_0)$ denotes the transition probability from S_0 to S over the time T. In the NIG market model case, these transition probabilities are found using equation Eq. (10) from the definition of the NIG process. In practice, we calculate this integral using Simpson's method. To ensure that the numerical implementation is correct, Monte Carlo simulation is used for verification. When the call prices are calculated, the corresponding put prices are obtained immediately by using put-call parity.

5.2 Calibration to market data

The NIG market model will now be calibrated to quotes observed in the vanilla option markets. The input to the calibration are all the market data, and the output are the corresponding model parameters. These parameters will be called the implied parameters. By doing this, insight is gained about the market's expectation of volatility, skewness and kurtosis. The implied parameters can be used to price vanilla options on the same underlying not quoted in the market (new strike and/or maturity), or, as will be done later, the implied parameters are used to price barrier options on the same asset. The calibration procedure can also be used to find mispriced options in the vanilla market.

The way the calibration is done in practice is to decide upon some error statistic, and then minimizing this error using some numerical optimization algorithm. The parameters that minimize the error statistic are then our implied parameters. There are several possible error statistics (Schoutens, 2003). The error statistic chosen for our calibration is a modified average relative percentage error given by:

$$\begin{aligned}
\text{ARPE}_{\text{bid-ask}} &= \frac{1}{\#\text{ options}} \sum_{options} \frac{|\text{market mid price - model price}|}{\text{market price} \cdot \frac{\text{ask-bid}}{\text{market price}}} \\
&= \frac{1}{\#\text{ options}} \sum_{options} \frac{|\text{market mid price - model price}|}{\text{ask-bid}}
\end{aligned} \tag{12}$$

This error statistic reflects the feature that the more liquid the option, the more information the option price contains, and it should be weighted higher in the calibration procedure. In general the spread (= ask - bid) is a measure of an option's liquidity, and a tighter spread means better liquidity. The relative spread can be measured as $\frac{\text{spread}}{\text{market mid price}}$.

The calibration is done by using a numerical optimization algorithm to find the parameters that minimizes the error statistic. We have used the Nelder-Mead algorithm in Matlab. When calibrating, one should give some thought to calibration risk, which is basically the risk of a wrong calibration, resulting in wrong option prices. One way to examine the calibration risk is to do the calibration with respect to several different error statistics, and check if the resulting parameter sets are approximately equal. See (Detlefsen and Hardle, 2007) for more on calibration risk.

The calibration is performed for some option chains found in different markets. The first data set consist of European call options on the Norwegian OBX index at closing time on 18th February 2008. The OBX index contains the 25 stocks on the Oslo Stock Exchange with the largest trading volume. The OBX index was at 356.68 at that time. The calibration gave the results shown in Table 5.

α :	60.883
β :	-44.5
δ :	1.371
ARPE:	4.90 %

Table 5: OBX calibration results.

As is seen, the average relative percentage error, ARPE, is measured to be 4.9% under optimal parameters. However, it is easier to understand the goodness of fit by looking at the fitted NIG option prices graphically, see Fig. 2.

Fig. 2 indicates that the NIG market model fits the option prices quoted in the market reasonably well, but not perfectly, as some of the NIG prices are slightly outside the bid-ask range in the market. We will discuss these results further, but first repeat the calibration process for a different set of options.

A lot of options are traded in the currency (FX) markets. We will now see how the NIG market model fits the option prices for the EUR/USD rate. Although equities and FX are two completely different markets, the modelling and option pricing techniques are very similar, unlike the fixed income market where the yield curve brings in another dimension. The only difference is that the dividend yield q is replaced with the risk free rate of the foreign currency, r_f. So for the EUR/USD rate, the Eurozone rate is denoted by r, and the US rate by r_f. The options in the FX market are in general more liquid than for most equities. While the OBX option chain contained 60 different options, our option chain for the EUR/USD rate contains 182 different options.

In the currency markets the shape of the volatility smiles can vary significantly. Both the curvature and the skewness can vary significantly for dif-

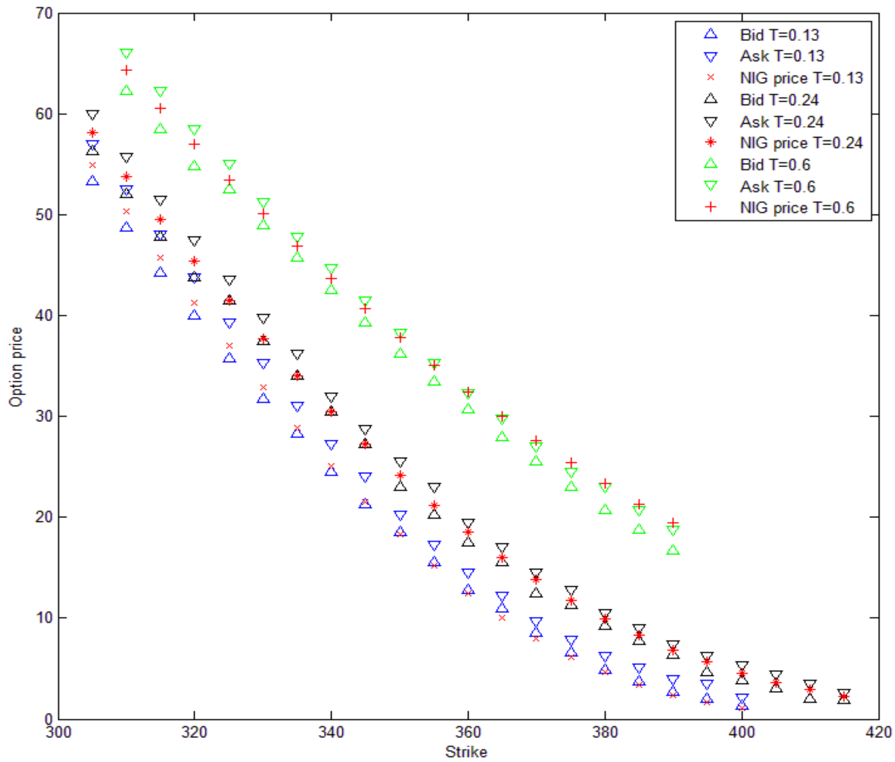

Figure 2: Fit of the OBX options.

ferent FX rates, while equity smiles in general have negative skewness and non-negative curvature. The implied volatility smiles for our option data set are shown in Fig. 3.

The calibration gives the NIG parameters listed in Table 6. The NIG market model fit looks good from the plot in Fig. 4. It displays very little deviation from the bid-ask interval. The ARPE at 4.20% is also a good indication of an accurate fit.

α :	28.521
β :	-8.947
δ :	0.236
ARPE:	4.20 %

Table 6: EUR/USD calibration results.

As we have seen from the two examples above, the NIG market model can give a very decent fit to market prices. It will fail to exactly fit the observed market mid prices, but we could never expect this to be achieved with just three time-invariant parameters to be calibrated. Since, in reality, we have no

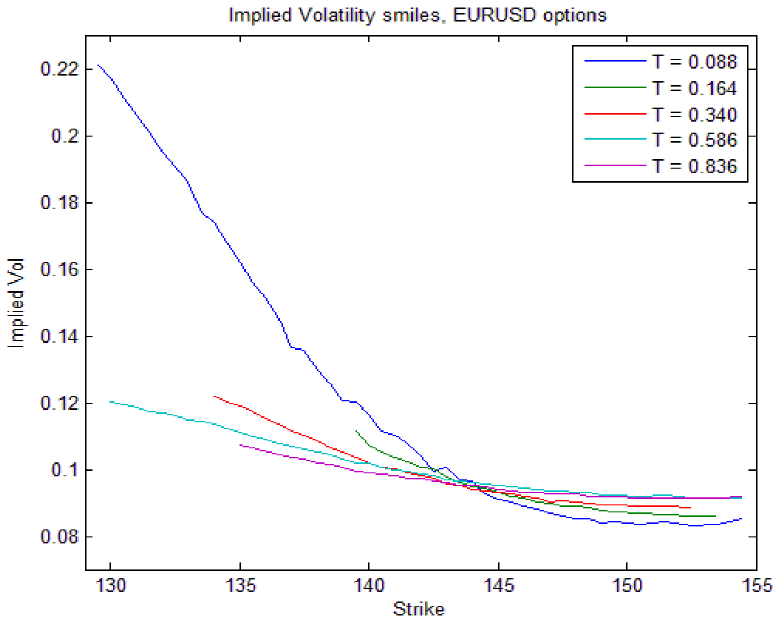

Figure 3: Implied volatility smiles. We see a strong skewness for the shortest maturity, then more moderate skewness for the longer maturities. The EUR/USD rate is quoted in (EUR/USD)*100.

exact observed market price, but rather a range between the bid price and the ask price, it seems like the NIG market model is doing quite well. (Schoutens, 2003) shows how we can get an exact fit for market mid prices, by including a stochastic volatility process. However, this comes with a serious cost of three to four extra parameters. This can complicate the calibration process, and definitely lead to a less intuitive understanding of the calibrated parameter values for the trader/user. If one can live with some mispricing with respect to the market mid price, the NIG market model seems to be adequate for option pricing. If a perfect fit to the market mid prices is desired, a more complicated model is required.

6 Barrier option pricing and the numerical path integration approach

We will now look at the pricing of several types of barrier options under the NIG market model. Barrier options are popular path-dependent options in the OTC markets, and certainly also a popular subject in the financial literature. Barrier options may be of both American and European type. However, in this paper only Europeans will be considered. Another important practical issue is how often the barrier is monitored. That is how often it is checked if the underlying is over/under the barrier. This monitoring time can be either discrete or continuous. In equities and commodities, it is said to be most common

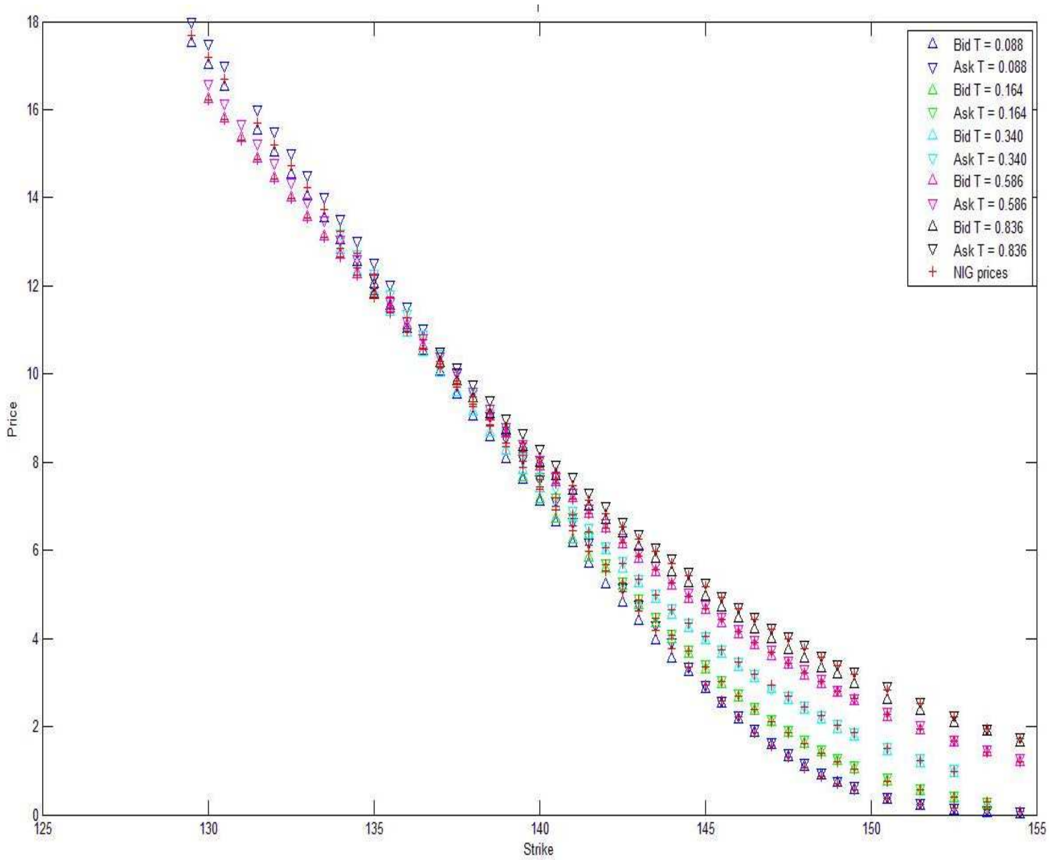

Figure 4: NIG prices fitted to the EUR/USD option prices.

with discrete monitoring, while in FX markets continuous monitoring is the standard. The discrete monitoring value is often set to the official daily closing price. Unless otherwise specified, all the barrier options will be considered to have daily monitoring. Also, no cash rebate K is paid out if the barrier option is knocked out/not knocked in.

Any given knock-out option can be constructed using its corresponding vanilla and knock-in option. The three option values are linked together by the in-out-parity: $Value_{vanilla} = Value_{knock-in} + Value_{knock-out}$.

Barrier options are known to be extremely model dependent, since the probability of hitting the barrier can be very model sensitive. Thus, the model risk for barrier options is also large. One should use several models when trading barrier options and their different prices will give some insights in the model risk. Schoutens et al. (2007) show in detail how model risk is significant for many exotic options including barrier options. There is no general analytical solution for the different types of barrier options under the NIG market model. Monte Carlo simulation can easily be applied, but for many trading purposes this is too slow. Instead of Monte Carlo simulations, it will be shown how the option prices can be calculated much more efficiently using numerical path integration. The method is based on the work in (Skaug and Naess, 2007).

As discussed in the previous section, vanilla call options can be priced numerically by calculating an integral, cf. Eq. (11). When dealing with path dependent options, numerical solutions become much more complicated. We propose to use numerical path integration for solving a wide range of such problems. It will be shown in detail how path integration works, how it may be implemented, and several numerical examples will be presented to confirm the accuracy and examine how much we gain in terms of computational time.

As an example, the price for an up-and-out-down-and-out call option will be calculated. Changing to another kind of barrier option is trivial. When we have calculated the price for an out barrier option, we can use in-out-parity to quickly find the price of the corresponding option.

As before, the strike price is denoted by X, the upper barrier by U, the lower barrier by L, and a constant interest rate r is assumed. Initial stock price is denoted S_0. The up-and-out-down-and-out option becomes worthless if the stock price S_t is equal to or bigger than the upper barrier U or equal to or smaller than the lower barrier L at any of the monitoring times. If the stock price stays within the barrier interval at every monitoring time until maturity, the call option value is $\max(0, S_T - X)$ at the maturity time T. The underlying stock is monitored at m different times $\tau_j, j = 1, ..., m$ such that $0 < \tau_1 < \tau_2 < ... < \tau_{m-1} < \tau_m = T$. For the initial conditions to make any sense $X < U$, and $L < S_0 < U$.

To find the price of the barrier option, we need the conditional probability distribution function for the stock price on the terminal day T, conditional on the property that the stock has stayed between the upper and lower barrier in the whole interval $[T_0, T]$. In the sequel we shall work with the process $Z_t = \ln(S_t/S_0)$ instead of S_t. Denoting this conditional probability density function expressed in terms of Z_t by $q_m(z)$, it is clear that the up-and-out-down-and-out barrier call value is given by:

$$c = e^{-rT} \int_{\ln(X/S_0)}^{\ln(U/S_0)} (S_0\, e^z - X) q_m(z)\, dz\,. \tag{13}$$

Similarly, the put value is given by:

$$p = e^{-rT} \int_{\ln(L/S_0)}^{\ln(X/S_0)} (X - S_0\, e^z) q_m(z)\, dz\,. \tag{14}$$

¿¿From this it follows that the only difference between pricing vanilla options and barrier options is the need to calculate the function $q_m(z)$. From the definition of $q_m(z)$, it is clear that it is given by

$$q_m(z) = \int_{\ln(L/S_0)}^{\ln(U/S_0)} \cdots \int_{\ln(L/S_0)}^{\ln(U/S_0)} p_{m|m-1}(z|z_{m-1}) ... p_{2|1}(z_2|z_1)\, p_{1|0}(z_1|0)\, dz_1 ... dz_{m-1}\,, \tag{15}$$

where $p_{i|i-1}(z|z_{i-1})$ denotes the transition probability density function of Z_{τ_i} from τ_{i-1} to τ_i given $Z_{\tau_{i-1}} = z_{i-1}$. When this multiple integral has been calculated, we can also calculate the option prices we need. In the case of an up-and-out option with a single barrier U, L is replaced with 0 in all of the above. Similarly, in the case of a down-and-out option with a single barrier L, U is replaced with ∞.

6.1 Implementation

Whenever the transition probability density $p_{i|i-1}(z|z')$, $i = 1, ..., m$ is known, Eq. (15) can be used recursively to obtain $q_m(z)$. When the NIG market model is used, the transition probability density can be derived from the pdf given in Eq. (10). Assuming equidistant monitoring times, with $\Delta\tau = \tau_i - \tau_{i-1}$, it is found that:

$$p_{1|0}(z|\,0) = q_1(z) = f(z; \Delta\tau)\,, \tag{16}$$

and

$$p_{i|i-1}(z|z') = f(z - z'; \Delta\tau)\,, \;\; i = 2, \ldots, m. \tag{17}$$

Hence, it follows that, for $i = 2, ..., m$:

$$q_i(z) = \int_{\ln(L/S_0)}^{\ln(U/S_0)} f(z - z'; \Delta\tau)\, q_{i-1}(z')\, dz'\,. \tag{18}$$

All the integrals are computed using numerical techniques. Since the NIG density function contains the modified Bessel function of the second kind, it contains a hidden integral, but both Matlab and Excel have this as a built-in function. Still, calculating the Bessel function is a significant part of the total computation time, when we apply the numerical path integration method on the NIG market model.

When the function $q_m(z)$ has been calculated for a given set of barrier(s), finding the option price $V(X,T)$ for any given strike X and maturity $T = \tau_m$, or any given binary option, or in fact for any given maturity, if the time to maturity is less than m days, is done just as fast as finding the value of a vanilla option,

$$V(X,T) = e^{-rT} \int_X^U (s - X)\, q_m\left(\ln\left(\frac{s}{S_0}\right)\right) \frac{ds}{s}\,. \tag{19}$$

Binary barrier options can also be found at trivial computational cost by calculating

$$\text{Binary} = e^{-rT} K \cdot \int_L^U q_m\left(\ln\left(\frac{s}{S_0}\right)\right) \frac{ds}{s}\,, \tag{20}$$

for any given value of K. For shorter maturities than m days, note that all $q_{m-i}(z)$, $i = 1, 2, ..., m-1$, have been calculated. So it follows that the calculation of all these barrier options can be done very efficiently using numerical path integration, if they are options on the same underlying asset.

As all the integration is done numerically, we have to discretize the stock price axis. Naturally, finer discretization will lead to better accuracy, but also higher computing time. In our calculations we have mainly used around 500 uniformly spaced grid points. This gives accurate answers and low computing time. We have experimented with non-uniformly spaced grid points, with finer discretization around the barriers and S_0, but we didn't gain much in terms of accuracy by doing this. We also tried a time-varying grid, but we didn't gain anything here either, except making the implementation complicated and unreadable. We believe it is important that the implementation of a numerical scheme is straightforward and easy to follow, and have thus chosen a simple uniformly spaced grid implementation. Since we will calculate prices for discretely monitored barrier options, the time axis discretization is automatically given. If we want to calculate prices for barrier options with continuous monitoring, we obviously have to pick a smaller time step. It is, however, hard to determine exactly how small time steps we should take to get an accurate solution. A way to determine this is to successively cut in half the length of the time step, and when the resulting changes in the option prices are small enough, we use this as our time step. A numerical example will be used to demonstrate this.

7 Numerical examples

Since it has been difficult to find any published numerical values for barrier-type options under the NIG framework, our results will be verified by Monte Carlo simulations. We will price several types of barrier options. First the accuracy of our numerical path integration implementation is verified for the down-and-out-up-and-out call option. The fixed parameters are given in Table 7, and the values for the upper and lower barrier are varied. Notice that the value for the first barrier option is almost the same as for the vanilla call.

S_0 :	100
X:	100
r:	0.05
α :	10
β :	-4
δ :	1
T:	0.2 (50 trading days)

Table 7: Fixed input parameters.

As we see in Table 8, the PI solution is very accurate, and takes less than a second to be calculated on a standard laptop. Of course in practice it doesn't make much sense to calculate barrier option prices with the accuracy of three decimals, since we know the model error will always be larger than 0.005, and a typical bid-ask spread can be as high as several percent. But because our purpose here is to verify the accuracy of our numerical implementation we do it anyway. The computation time for the path integration was around one second for the results in Table 8, with 400 grid points on the stock price axis. The Monte Carlo are derived from 6,000,000 simulations. The effect on the results of different number of grid points on the stock price axis are shown in Table 9. It is noticed that for $N = 400$ all results are accurate enough, and fast.

L	U	PI	MC
50	**150**	6.168	6.168
60	**140**	5.852	5.853
70	**130**	5.054	5.054
80	**120**	3.273	3.273
90	**110**	0.735	0.735
Vanilla price:		6.383	6.383

Table 8: Option prices for double knock-out options. PI denotes numerical path integration results, MC denotes Monte Carlo simulation results.

L	U	N=100	N=200	N=300	N=400	N=500	N=600	MC
50	**150**	86.308	4.559	5.968	6.164	6.169	6.169	6.168
60	**140**	8.883	5.185	5.836	5.852	5.852	5.852	5.853
70	**130**	3.434	4.987	5.055	5.054	5.054	5.054	5.054
80	**120**	2.843	3.272	3.272	3.273	3.273	3.273	3.273
90	**110**	0.734	0.735	0.735	0.735	0.735	0.735	0.735
Comp. time:		0.04 s	0.16 s	0.5 s	0.8 s	1.3 s	2.2 s	> 5000 s

Table 9: Results and computation times with different number N of discrete grid points on the stock price axis.

It has now been confirmed that the numerical path integration is both accurate and very fast for solving double barrier options. Another positive aspect is that the implementation is flexible, so we can easily change the type of barrier options. It is trivial to change the code from double barrier options to single barrier options. For the same scenario as in Table 7 we verify the accuracy for the numerical path integration on an up-and-out call, with different barriers. The results are listed in Table 8, and again it is seen that the numerical path integration is very accurate.

We shall also calculate prices of down-and-out options, and since the previous examples consisted of calls we now price puts. We again use the same scenario as before, given in Table 7. Results are shown in Table 11, and also for these options the numerical path integration is accurate.

As discussed earlier, barrier options can have different monitoring frequencies of the barrier, including continuous monitoring. In the numerical path integration implementation, we adjust for this by varying the time step Δt. Since all parameters are annualized, the time is measured in years. Assuming 250 trading days a year, daily time step corresponds to $\Delta t = \frac{1}{250} = 0.004$. To

Barrier:	**PI**	**MC**
105	0.114	0.114
110	0.794	0.794
115	2.000	2.001
120	3.277	3.279
125	4.317	4.319
130	5.054	5.056
135	5.541	5.541
140	5.852	5.850
145	6.046	6.042
150	6.168	6.164
155	6.245	6.246
Vanilla:	6.383	6.383

Table 10: Up-and-out call prices, with different barrier levels.

Barrier:	**PI**	**MC**
60	4.871	4.872
65	4.524	4.525
70	4.017	4.018
75	3.321	3.320
80	2.442	2.442
85	1.461	1.461
90	0.583	0.583
95	0.087	0.087
Vanilla:	5.388	5.388

Table 11: Down-and-out put option prices, with different barrier levels.

approximate continuous monitoring we must choose a shorter time step. Table 12 illustrates how the option prices changes with different time steps. As an example, we do this for the same down-and-out-up-and-out options we dealt with in Table 8, with the same parameters given in Table 7. Note that smaller time steps leads to correspondingly longer computing time.

		Time step in years:				
L	**U**	$\Delta t = 0.004$	$\Delta t = 0.002$	$\Delta t = 0.001$	$\Delta t = 0.0005$	$\Delta t = 0.00033$
50	**150**	6.168	6.166	6.172	6.150	6.149
60	**140**	5.852	5.856	5.843	5.841	5.838
70	**130**	5.054	5.035	5.036	5.033	5.029
80	**120**	3.273	3.241	3.229	3.218	3.217
90	**110**	0.735	0.712	0.698	0.690	0.685

Table 12: Barrier option prices with different monitoring times. We see the results converge to a given price. The same scenario as in Table 7 is used.

As we have seen through numerical examples, the numerical path integration approach is suitable for many kinds of barrier options. It has very good accuracy, and is a lot faster in general than Monte Carlo simulation.

8 Concluding remarks

It has been shown that the Normal Inverse Gaussian process provides a model that fits the financial data under study quite well. We have shown in detail how this statistical analysis is done, and how vanilla option pricing and market calibration can be performed.

It is clear that the numerical path integration approach is a very efficient method for pricing all sorts of barrier options under the NIG market model. Minimizing computing time is crucial for trading purposes. Numerical path integration is also very flexible; the restriction is that the chosen process has stationary, independent increments.

References

Abramowitz, M. and I. Stegun (1968). *Handbook of Mathematical Functions.* New York: Dover.

Barndorff-Nielsen, O. E. (1994). Normal inverse Gaussian distributions and the modelling of stock returns. Research report no. 300, Aarhus University, Aarhus, Denmark.

Black, F. and M. Scholes (1973). The pricing of options and corporate liabilities. *Journal of Political Economy 81*, 637–654.

Bølviken, E., , and F. E. Benth (2000). Quantification of risk in Norwegian stocks via the normal inverse Gaussian distribution. In *Proceedings AFIR*, pp. 87–98. Troms.

Detlefsen, K. and W. Hardle (2007). Calibration Risk for Exotic Options. SFB 649 Discussion paper. http://edoc.hu-berlin.de/series/sfb-649-papers/2006-1/PDF/1.pdf.

Gerber, H. U. and E. S. W. Shiu (1994). Option pricing by Esscher transforms. *Transactions of Society of Actuaries 46*, 99–191.

Haug, E. G. (2007). *The Complete Guide to Option Pricing Formulas* (2nd Edition). New York: McGraw-Hill, Inc.

Hull, J. C. (2006). *Options, Futures and Other Derivatives* (6th Edition). Upper Saddle River, NJ: Prentice Hall International Editions.

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

McDonald, R. L. (2002). *Derivatives Markets*. Upper Saddle River, NJ: Addison-Wesley.

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

Rebonato, R. (2004). *Volatility and Correlation. The Perfect Hedger and the Fox*. Chichester, UK: John Wiley & Sons, Ltd.

Schoutens, W. (2003). *Levy Processes in Finance - Pricing Financial Derivatives*. Chichester, UK: John Wiley & Sons, Ltd.

Schoutens, W., E. Simons, and J. Tistaert (2007). A Perfect Calibration! Now what? http://www.wilmott.com/detail.cfm?articleID=262.

Skaug, C. and A. Naess (2007). Fast and accurate pricing of discretely monitored barrier options by numerical path integration. *Computational Economics 30*, 143–151.