An Introduction to Implied Lévy Volatility with Applications [***]

Marc Cassagnol[†]

August 30, 2010

Abstract

The concept of implied volatility is one of the great successes of the Black-Scholes model. However, Black-Scholes relies on a normal distribution, which historical data shows is not adequate to describe real markets. This is often seen in implied volatility surfaces which "smile" and "smirk", taking on different values for different strike prices. In this paper we discuss a more empirically founded type of implied volatility based on Lévy distributions, first introduced by Corcuera, Guillaume, Leoni and Schoutens in 2009. We use implied Lévy volatility with real market prices to reduce the presence of smiles and smirks. We perform several delta and gamma hedging experiments to test the performance of the Black-Scholes model against Lévy models. Finally, we price barrier options using implied Black-Scholes and implied Lévy volatility and discuss how the change in volatility affects the prices.

[*]Department of Mathematics and Statistics, York University. Submitted in partial fulfillment of requirements of Master of Arts in Mathematics and Statistics, under the supervision of Prof. Hanna Jankowski.

[**]Schulich School of Business, York University. Submitted in partial fulfillment of requirements of Graduate Diploma in Financial Engineering.

[†]Email address: mjgc@mathstat.yorku.ca

Contents

1 Introduction

The Black-Scholes model has seen paramount success partially due to the concept of implied volatility which it introduced. Implied volatility allows traders to match market prices with theoretical model prices in a consistent way. It is often convenient to quote vanilla options by their implied Black-Scholes volatility as it does not require units of currency, and traders have gained an intuition for this quantity over the years. Implied volatility depends on strike price and time to maturity, and so it is usually reported as a curve or surface.

The drawback of implementing implied volatility using the Black-Scholes pricing model is that it is based on a normal distribution. Historical data indicates that return distributions on stocks are often skewed and heavier-tailed than can be described by the normal distribution. As a result of this, when implied Black-Scholes volatility is calculated for real option prices, the surfaces often "smile" or "smirk" with respect to strike price. The fact that different strikes yield different values for implied volatility indicates that the model does not accurately describe the behaviour of the stock. This creates a problem when pricing barrier options concerning which volatility to use as input; the volatility corresponding to the strike price, the barrier level, or perhaps some average of these two values.

In this paper we discuss Lévy distributions and the additional freedom they provide in calibrating skewness and kurtosis. We discuss a more generalized version of implied volatility which works for Lévy processes. Since Brownian motion is itself a Lévy process, the models discussed here reduce to the Black-Scholes model if a Brownian motion is used. The concept of implied Lévy volatility was first introduced in 2009 by Corcuera, Guillaume, Leoni and Schoutens in [6]. They discussed implied Lévy space and time volatility models, both of which will be covered here in great detail.

In addition, we discuss the Carr-Madan approach to option pricing using the fast Fourier transform (FFT). This method can be implemented for any Lévy process whose characteristic function is computable. This formulation will be used for implied Lévy volatility calculations. We also use the Carr-Madan formula to derive the delta and gamma of a call option under the Lévy framework. This will allow us to perform hedging experiments using Lévy models.

One of the benefits of using implied Lévy volatility is that the additional degrees of freedom provided by the model can be used to reduce the presence of smiling in the implied volatility surface. In Section 4 we examine options on AAPL stock that exhibit both smiling and smirking implied Black-Scholes volatility. We calibrate the corresponding implied Lévy space and time volatility models to yield flatter curves, suggesting that Lévy models are capable of more accurately describe the stock price.

In Section 5 we compare the performance of delta and gamma hedges when Black-Scholes and Lévy parameters are used. Using data which covers the 2007 financial crisis, we hedge a rolling at-the-money option on the S&P 500 and investigate the performance of the hedges under various models. Evaluating the hedging performance during the crisis works as a stress test for the models. We determine the optimal parameter settings which lead to the smallest hedging error for each model.

In Section 6 we price barrier options using Monte Carlo simulation and the implied Black-Scholes and Lévy volatility data from Section 4. We investigate how the flatter volatility curve affects the range of barrier option prices using both geometric Brownian

motion and a normal inverse Gaussian process to drive the simulations. We discuss several ways that volatility can directly or indirectly affect the price of barrier options. Finally we formulate our conclusions in Section 7 and discuss some areas where this work can be expanded upon in the future in Section 8.

2 Fundamental concepts

2.1 The Black-Scholes model

The Black-Scholes model assumes that the risk-neutral price process for the underlying stock follows a geometric Brownian motion:

$$S_t = S_0\, e^{\,(r-q-\sigma^2/2)t+\sigma W_t}$$

where $t \geq 0$, is the time remaining until maturity, $r \geq 0$ is the risk-free interest rate, $q \geq 0$ is the dividend yield, $\sigma > 0$ is the volatility parameter and W_t is a standard Brownian motion. To verify that this is a risk-neutral process, we must show that

$$E[S_t|\mathcal{F}_\tau] = S_t\, e^{\,(r-q)(t-\tau)}$$

for any $t > \tau$. We do this by showing that $X_t = S_0\, e^{\,\sigma W_t - \sigma^2 t/2}$ is a martingale:

$$\begin{aligned} dX_t &= \frac{\partial X}{\partial t} X_t\, dt + \frac{\partial X}{\partial w} X_t\, dW_t + \frac{1}{2}\frac{\partial^2 X}{\partial w^2} X_t\, (dW_t)^2 \\ &= -\frac{1}{2}\sigma^2 X_t\, dt + \sigma X_t\, dW_t + \frac{1}{2}\sigma^2 X_t (dt) \\ &= \sigma X_t\, dW_t. \end{aligned}$$

This shows that the drift of the process $S_t = X_t\, e^{\,(r-q)t}$ is equal to $(r-q)\, dt$ which implies that the expected continuous log-return of the stock is the risk-free rate (minus the dividend yield). Equipped with this model, one can calculate the price of a vanilla European call option to be the famous Black-Scholes formula:

$$C(S_0, K, T, r, \sigma, q) = S_0\, e^{-qT} N(d_1) - K\, e^{-rT} N(d_2) \tag{1}$$

where

$$d_1 = \frac{\log(S_0/K) + (r - q + \sigma^2/2)T}{\sigma\sqrt{T}},$$

$$d_2 = d_1 - \sigma\sqrt{T},$$

and $N(\cdot)$ is the standard normal cumulative distribution function. At any given time, the parameters S_0, r, σ and q are the same for every derivative written on the asset, so when pricing options it is convenient to abbreviate the call price function as $C(K, T)$ which we will do whenever possible.

2.2 Implied volatility

After the introduction of the Black-Scholes model during the 1970's, traders began using it to calibrate their models. Rather than calculating volatility by tracking the stock price, calculating daily returns and computing the standard deviation of the return distribution, they would find the value for volatility $\sigma(K,T)$ such that when plugged into the Black-Scholes formula, the resulting price equalled the current market price. This value is called the *implied Black-Scholes volatility*, or simply the *implied volatility* of the option, and it represents what volatility the market deems appropriate for the stock.

When the Black-Scholes model is used to calculate implied volatility one often obtains different numbers for different values of K and T. In particular, a "smile" or "smirk" shape is often observed in the plot of implied volatility versus strike price. Implied volatility tends to increase with maturity time, but is often larger for options with very short maturities. This is due to the increase in price that sometimes occurs when options are close to maturity as the price and the payoff converge.

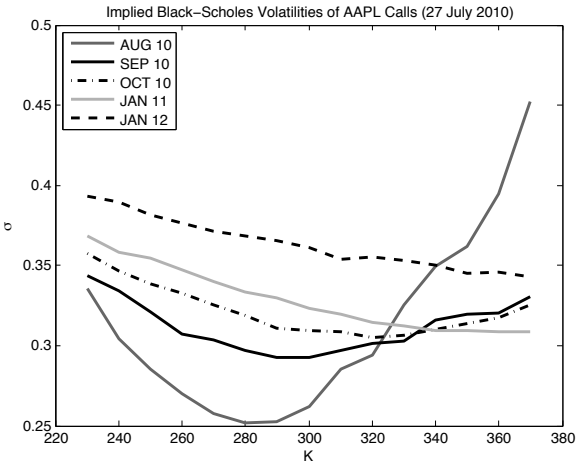

Figure 1: Implied Black-Scholes volatility for several AAPL calls; $S_0 = 264.08$. Data Source: Yahoo! Finance, July 27, 2010.

2.3 Hedging with the Greeks

The *Greeks* of an option measure the sensitivity of its price with respect to its input parameters S_0, K, T, r and σ. Knowing these sensitivities helps traders to hedge more effectively. The most common way to hedge a portfolio is to make it insensitive to changes in the underlying stock. The relevant Greeks for this are the first and second derivatives of the option price with respect to spot price; delta (Δ) and gamma (Γ) respectively. By differentiating the Black-Scholes formula with respect to S_0 we obtain the Black-Scholes delta:

$$\begin{aligned}
\Delta &= \frac{\partial C}{\partial S_0} \\
&= \frac{\partial}{\partial S_0}\left[S_0\, e^{-qT} N(d_1) - K\, e^{-rT} N(d_2)\right] \\
&= e^{-qT} N(d_1) + S_0\, e^{-qT}\, \frac{\partial N(d_1)}{\partial S_0} - K\, e^{-rT}\, \frac{\partial N(d_2)}{\partial S_0} \\
&= e^{-qT} N(d_1) + S_0\, e^{-qT}\, \frac{1}{S_0\,\sigma\sqrt{T}}\, n(d_1) - K\, e^{-rT}\, \frac{1}{S_0\,\sigma\sqrt{T}}\, n(d_2) \\
&= e^{-qT} N(d_1) + \frac{1}{S_0\,\sigma\sqrt{T}}\left[S_0\, e^{-qT}\, n(d_1) - K\, e^{-rT}\, n(d_2)\right] \\
&= e^{-qT} N(d_1)
\end{aligned}$$

where $n(\cdot)$ is the standard normal probability density function. We have used the identity $S_0\, e^{-qT}\, n(d_1) = K\, e^{-rT}\, n(d_2)$ in the last line. We differentiate delta with respect to S_0 to obtain gamma:

$$\begin{aligned}\Gamma &= e^{-qT}\,\frac{\partial}{\partial S_0}N(d_!) \\ &= e^{-qT}\,n(d_1)\frac{\partial d_1}{\partial S_0} \\ &= e^{-qT}\,\frac{n(d_1)}{S_0\,\sigma\sqrt{T}}.\end{aligned}$$

The ability to create a perfect continuous delta hedge is the basis for the arbitrage argument which gives rise to the Black-Scholes formula. A perfectly hedged portfolio can be created by holding a long call option and $-\Delta$ shares of the underlying asset or the reverse position; a short call option and Δ shares of the underlying. The purpose of holding the underlying is to neutralize the instantaneous sensitivity of the portfolio with respect to the spot price; the delta of the portfolio is simply the sum of the deltas of the option and underlying:

$$\Delta_{portfolio} = \pm\Delta_{call} \mp \Delta_{call}\,\frac{\partial S_0}{\partial S_0} = 0, \tag{2}$$

and so we call such a portfolio *delta neutral*. To maintain a perfect hedge, the portfolio must be rebalanced continuously, as the delta of the option changes due to movements in the stock price and and the passage of time. Therefore it is practically impossible to implement this strategy, as there will always be a measurable period of time between trades. The risk that arises by trading in discrete time is a form of *gap risk*; which is the risk associated with the value of the underlying changing in between trades.

An obvious way to manage gap risk is to rebalance more often. However, in practice this can get expensive due to the transaction costs associated with trading. A better approach is to further reduce the sensitivity of the portfolio in order to marginalize the change in value that could occur between trades. This can be achieved through *gamma hedging*. A portfolio is *gamma neutral* if the second derivative of its value with respect to the spot price equals zero. If the delta of the portfolio truly did not change over an interval of time, there would be no need to rebalance continuously, and this gamma hedge would have zero hedging error. However this is not always the case and thus one can not guarantee that a discretely rebalanced delta/gamma hedge will always outperform a discretely rebalanced delta hedge.

Purchasing the underlying asset has no effect on the gamma of the portfolio, as it has a gamma value of zero. So in order to neutralize the gamma of an option, one must use another derivative security written on the same underlying asset. Usually another option with a different strike price is used. Since this option will be sold and repurchased on occasion, one would in practice want to use a liquid option for this purpose, such as an option which is close to being at-the-money.

Let Δ and Γ denote the delta and gamma of the option we want to hedge, and Δ' and Γ' denote the delta and gamma of the option we are using to create the hedge. Let a denote the amount of underlying necessary to neutralize the delta of the options in the

portfolio, and b denote the amount of the second option we need to purchase to neutralize the gamma. To obtain a zero portfolio delta we then need

$$\Delta + a + b\,\Delta' = 0. \tag{3}$$

To obtain a zero portfolio gamma we need

$$\Gamma + b\,\Gamma' = 0. \tag{4}$$

Equation (4) shows that we must hold the opposite position

$$b = -\frac{\Gamma}{\Gamma'}$$

in the second option to neutralize the gamma of the first option. Therefore in order to delta-hedge the portfolio we must hold

$$a = -\Delta + \frac{\Gamma}{\Gamma'}\,\Delta'$$

units of the underlying stock.

In addition to hedging away the delta and gamma of the portfolio, it will be convenient for us to also maintain a portfolio with zero value. This can be achieved by taking a position in the money market (either borrowing or lending) equal to minus the value of the stocks and options in the portfolio. The benefit of constructing the portfolio this way is that the hedging error is simply equal to the portfolio value, with a perfect hedge having a portfolio value of zero. This makes is straight forward to calculate a hedging error distribution to evaluate the performance of the hedge.

As an example of how hedging error is calculated, consider the following delta hedge calculation. Beginning with a brand new call with price $C(K,T)$ and delta Δ we delta-hedge this option by purchasing $-\Delta$ shares of the underlying asset. The value of the portfolio is now

$$V_0 = C(K,T) - \Delta S_0.$$

In addition to this, we invest

$$-(C(K,T) - \Delta S_0)$$

in the money market which makes the value of the portfolio equal to zero. The next day (after δt time has passed) the value of the portfolio is:

$$V_{\delta t} = C(K, T-\delta t) - \Delta S_{\delta t} - (C(K,T) - \Delta S_0)e^{r\delta t}.$$

If our hedge was perfect, we would have $V_{\delta t} = 0$. Hence the hedging error (any deviation from zero) can be measured by the value of $V_{\delta t}$.

2.4 Barrier options

A *barrier option* is an exotic option which pays depending on whether or not the underlying stock price reaches a predetermined price called the *barrier level*, denoted by B, at some time before maturity. We will look at four varieties of barrier options:

Name	Barrier relative to spot price and call payoff
"Up and Out"	$S_0 < B$ $\text{payoff}_{call} = (S_T - K)^+ \mathbb{I}_{0 \le t \le T}(\max S_t < B)$
"Down and Out"	$S_0 > B$ $\text{payoff}_{call} = (S_T - K)^+ \mathbb{I}_{0 \le t \le T}(\min S_t > B)$
"Up and In"	$S_0 < B$ $\text{payoff}_{call} = (S_T - K)^+ \mathbb{I}_{0 \le t \le T}(\max S_t > B)$
"Down and In"	$S_0 > B$ $\text{payoff}_{call} = (S_T - K)^+ \mathbb{I}_{0 \le t \le T}(\min S_t < B)$

The "up" and "down" in the option name indicates that the barrier level is above below the spot price, respectively. The underlying must cross the barrier in order for "in" barrier options become activated (they must be *knocked in*). Conversely, "out" options are initially active and are *knocked out*, becoming void, if the underlying crosses the barrier.

Barrier options are always cheaper than a similar European option without a barrier, and so they can be used to make more specific bets than vanilla options. For example, if you are willing to bet that a stock price worth \$90 will increase, but not above \$100 you could purchase an up-and-out call and pay a smaller premium than a European call.

Since "in" barriers pay exactly when "out" barriers do not, it is possible to replicate a European option by purchasing an in-out pair with the same barrier level (i.e. a "down-in/down-out" pair or an "up-in/up-out" pair). Therefore it must be true that the sum of the prices of any in-out pair equals the price of a European call in order for there to be no arbitrage opportunities.

3 Lévy framework

3.1 Shortcomings of the Black-Scholes model

Although the Black-Scholes model has been widely adopted due to its simplicity, it is based on several unreasonable assumptions. Empirical evidence suggests that the Black-Scholes model does not adequately describe the statistical properties of the markets. In particular, historical data indicates that that log-returns typically do not follow a normal distribution. Table 1 shows the cumulants of the normal(μ, σ^2) distribution for reference.

Table 1: Cumulants of the Normal(μ, σ^2) distribution.

	Mean	SD	Skewness	Kurtosis
$N(\mu, \sigma^2)$	μ	σ	0	3

Table 2 shows the historical cumulants for several major indices and stocks from different industries. These datasets use daily market data from 1 January 2007 to 28 July 2010. Recall that skewness measures the asymmetry of a distribution and is defined as

$$\frac{E[(X - \mu_X)^3]}{\text{var}[X]^{3/2}}.$$

Table 2: Empirical mean, standard deviation, skewness and kurtosis of major indices and stocks (Apple Inc., Ford Motor Co., Potash Corp. of Saskatchewan Inc., Exxon Mobil Corp.).

Index	Mean	SD	Skewness	Kurtosis
Dow Jones	-0.00018	0.01644	0.05481	6.67717
Nasdaq	-0.00006	0.01854	-0.12699	4.83206
S&P 500	-0.00027	0.01807	-0.18269	6.31336
TSX	-0.00011	0.01689	-0.56274	6.04778
DAX	-0.00008	0.01736	0.21633	6.35851
FTSE 100	-0.00018	0.01658	-0.04907	5.57404
AAPL	0.00126	0.02711	-0.46031	4.79007
F	0.00060	0.04287	-0.04700	9.02968
POT	0.00081	0.03826	-1.12956	8.08810
XOM	-0.00014	0.02131	0.15629	11.32333

The empirical data is skewed in every case, suggesting that an asymmetrical distribution would be a better fit. Recall that kurtosis is defined as

$$\frac{E[(X-\mu_X)^4]}{\text{var}[X]^2}.$$

If the kurtosis is lower than 3, the distribution has a flatter top (*platykurtic*), and if the kurtosis is greater than 3, the distribution has a high peak and heavier tails (*leptokurtic*). We clearly see that the datasets give rise to return distributions which are more leptokurtic than the normal distribution, suggesting that we should use a distribution with a flexible kurtosis to create a better fit.

3.2 Lévy processes

Let X be a random variable with characteristic function $\phi(u) = E\left[e^{iuX}\right]$. If, for every positive integer n, there exist i.i.d. random variables $X_1, X_2, \ldots X_n$, such that $X = X_1 + X_2 + \ldots + X_n$ in distribution, we call X *infinitely divisible*.

For every infinitely divisible distribution, we can define a stochastic process $\{X_t : t \geq 0\}$ called a *Lévy process* with the following properties:

1. $X_0 = 0$,
2. stationary and independent increments,
3. $X_t - X_s$ has $(\phi(u))^{t-s}$ as a characteristic function (for any $0 \leq s \leq t$).

It is not difficult to show that geometric Brownian motion is a Lévy process, and hence the framework we build here has the Black-Scholes model as a special case. The Lévy process we will focus on in this paper is the normal inverse Gaussian (NIG) process. It is relatively straightforward to apply and can be simulated directly (as opposed to some Lévy processes which must be approximated), which will be useful when pricing barrier options. The NIG process $\{X_t : t \geq 0\}$ is a Lévy process such that X_t follows the $NIG(\alpha, \beta, \delta t, \mu t)$

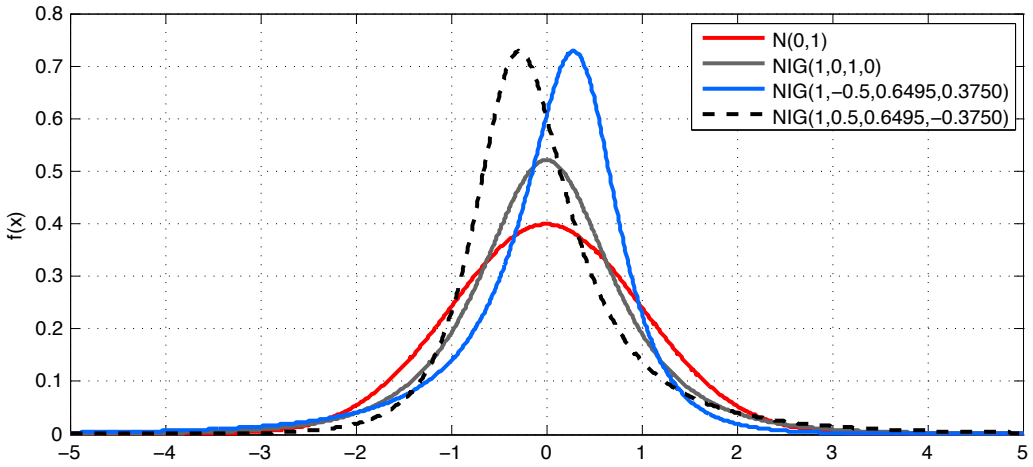

Figure 2: Density functions of standard normal and several NIG distributions with mean zero and variance one.

Table 3: Characteristics of the general (left) and symmetric (right) normal inverse Gaussian distributions

	$\text{NIG}(\alpha, \beta, \delta, \mu)$	$\text{NIG}(\alpha, 0, \delta, \mu)$
mean	$\mu + (\delta\beta/\sqrt{\alpha^2 - \beta^2})$	μ
variance	$\alpha^2\delta(\alpha^2 - \beta^2)^{-3/2}$	δ/α
skewness	$3\beta\alpha^{-1}\delta^{-1/2}(\alpha^2 - \beta^2)^{-1/4}$	0
kurtosis	$3\left(1 + \frac{\alpha^2+4\beta^2}{\delta\alpha^2\sqrt{\alpha^2-\beta^2}}\right)$	$3[1 + (1/\alpha\delta)]$

distribution, with tail-heaviness $\alpha > 0$, asymmetry parameter $\beta \in (-\alpha, \alpha)$, scale parameter $\delta > 0$, and location parameter $\mu \in \mathbb{R}$, for all t. The distribution scales as follows: if $X \sim NIG(\alpha, \beta, \delta, \mu)$ then $cX \sim NIG(\alpha/c, \beta/c, c\delta, c\mu)$. Figure 2 shows the density function for several NIG distributions with mean zero and variance one. The symmetric distribution has a kurtosis of 6. The asymmetric distributions have a kurtosis of 13.6667, and skewness values of -2 and 2.

The characteristic function of the NIG distribution is given by

$$\phi_{NIG}(u; \alpha, \beta, \delta, \mu) = \exp\left(iu\mu - \delta\left(\sqrt{\alpha^2 - (\beta + iu)^2} - \sqrt{\alpha^2 - \beta^2}\right)\right)$$

for $u \in \mathbb{R}$. We will replace β with $\kappa\alpha$ for $\kappa \in (-1, 1)$ in our numerical analysis. It will be necessary when modelling stock prices to choose a normal inverse Gaussian distribution such that $E[X_1] = 0$ and $\text{Var}[X_1] = 1$. This can be achieved by setting

$$\delta = \alpha\left(1 - k^2\right)^{3/2}$$

and

$$\mu = -\alpha\kappa\left(1 - \kappa^2\right).$$

3.3 Calculating option prices and Greeks under Lévy models

When we remove the assumption that the stock price follows a geometric Brownian motion in favour of a more general semi-martingale (of which Lévy processes are a subset), the fundamental theorem of asset pricing no longer holds. In particular, we lose market completeness as there is no longer a unique risk-neutral measure. This means we also lose the ability to perfectly hedge an option through continuous-time delta hedging. However, since trading is always performed discretely in practice, the non-existence of a perfect continuous hedge will not affect our ability to hedge under Lévy models. Most practitioners believe that the market is incomplete, so a pricing model which supports market incompleteness is desirable.

3.3.1 Call option pricing using characteristic functions

Following [4], we price options as follows. Let Q denote a measure such that the discounted stock price is a martingale. The price $C(K,T)$ of a European call option with strike K and maturity T can be defined as the expected price (under Q) at maturity, discounted by the risk-free rate:

$$C(K,T) = e^{-rT} E_Q[(S_T - K)^+]. \tag{5}$$

Under the Black-Scholes model, this is the unique arbitrage-free price. Under the general Lévy framework, this price lies in an interval of arbitrage-free prices corresponding to a range of martingale measures.

Whenever the risk-neutral density function is available, (5) can be evaluated analytically. However, it is usually a lengthy process. A faster approach is to use the fast Fourier transform (FFT) method developed by Carr and Madan in [4]. Their formula for the price of the European call option is

$$C(K,T) = \frac{e^{-\lambda \log(K)}}{\pi} \int_0^\infty e^{-iv \log(K)} \psi_T(v)\, dv \tag{6}$$

where

$$\psi_T(v) = \frac{e^{-rT}\varphi_T(v-(\lambda+1)i)}{(\lambda+iv)(\lambda+1+iv)},$$

and $\lambda > 0$ is a parameter which ensures the result $\psi_T(v)$ is finite. The only part of this formulation which depends on the model is the function $\varphi_T(u)$ which is (under Q) the characteristic function of the log-price process at maturity T,

$$\varphi_T(u) = E_Q[e^{iu \log(S_T)}].$$

The benefit of this approach to option pricing is that it can be implemented for any Lévy distribution whose characteristic function is available analytically. This gives us the freedom to use distributions with skewness and kurtosis which better match empirical values from the market. (6) can be evaluated numerically very quickly with the help of the inverse fast Fourier transform. The formulation can be used to price put options as well, through put-call parity.

3.3.2 Calculation of delta

The delta of an option is defined analogously in the Lévy case; it measures the sensitivity of the option with respect to changes in the underlying stock price. To calculate delta, we differentiate (6):

$$\begin{aligned}\Delta &= \frac{\partial}{\partial S_0}\left[\frac{e^{-\lambda \log(K)}}{\pi}\int_0^\infty e^{-iv\,\log(K)}\frac{e^{-rT}\varphi_T(v-(\lambda+1)i)}{\lambda^2+\lambda-v^2+i(2\lambda+1)v}\,dv\right] \\ &= \frac{e^{-\lambda\log(K)}}{\pi}\int_0^\infty \frac{e^{-iv\log(K)}e^{-rT}}{\lambda^2+\lambda-v^2+i(2\lambda+1)v}\frac{\partial\left[\varphi_T(v-(\lambda+1)i)\right]}{\partial S_0}\,dv. \qquad (7)\end{aligned}$$

In order to formulate this explicitly, we evaluate

$$\begin{aligned}\frac{\partial\left[\varphi_T(v-(\lambda+1)i)\right]}{\partial S_0} &= \frac{\partial}{\partial S_0}E_Q[e^{i(v-(\lambda+1)i)\,\log(S_T)}] \\ &= \frac{\partial}{\partial S_0}E_Q[e^{i(v-(\lambda+1)i)\left[\log(S_0)+\log(e^{(r-q+\omega)T+\sigma X_T})\right]}] \\ &= i(v-(\lambda+1)i)\frac{1}{S_0}E_Q[e^{i(v-(\lambda+1)i)\left[\log(S_0)+\log(e^{(r-q+\omega)T+\sigma X_T})\right]}] \\ &= \frac{\lambda+1+vi}{S_0}\varphi_T(v-(\lambda+1)i). \qquad (8)\end{aligned}$$

Combining this result with equation (7), and using the fact that $\lambda^2+\lambda-v^2+i(2\lambda+1)v = (\lambda+vi)(\lambda+1+vi)$ we obtain

$$\begin{aligned}\Delta &= \frac{e^{-\lambda\log(K)}}{\pi}\int_0^\infty \frac{e^{-iv\log(K)}e^{-rT}}{(\lambda+vi)(\lambda+1+vi)}\frac{\lambda+1+vi}{S_0}\varphi_T(v-(\lambda+1)i)\,dv \\ &= \frac{e^{-\lambda\log(K)}}{\pi}\int_0^\infty \frac{e^{-iv\log(K)}e^{-rT}}{S_0(\lambda+vi)}\varphi_T(v-(\lambda+1)i)\,dv. \qquad (9)\end{aligned}$$

3.3.3 Calculation of gamma

The gamma of an option measures the second order sensitivity of the option with respect to the underlying stock price.

$$\Gamma = \frac{\partial^2 C(K,T)}{\partial S_0^2} = \frac{\partial\Delta}{\partial S_0}.$$

We calculate gamma using the results of (8) and (9):

$$
\begin{aligned}
\Gamma &= \frac{\partial}{\partial S_0}\left[\frac{e^{-\lambda \log(K)}}{\pi}\int_0^\infty \frac{e^{-iv\log(K)}e^{-rT}}{S_0(\lambda+vi)}\varphi_T(v-(\lambda+1)i)\,dv\right] \\
&= \frac{e^{-\lambda\log(K)}}{\pi}\int_0^\infty \frac{e^{-iv\log(K)}e^{-rT}}{(\lambda+vi)}\frac{\partial}{\partial S_0}\left[\frac{\varphi_T(v-(\lambda+1)i)}{S_0}\right]dv \\
&= \frac{e^{-\lambda\log(K)}}{\pi}\int_0^\infty \frac{e^{-iv\log(K)}e^{-rT}}{(\lambda+vi)}\left[\frac{\frac{\lambda+1+vi}{S_0}\varphi_T(v-(\lambda+1)i)(S_0)-(1)\varphi_T(v-(\lambda+1)i)}{S_0^2}\right]dv \\
&= \frac{e^{-\lambda\log(K)}}{\pi}\int_0^\infty \frac{e^{-iv\log(K)}e^{-rT}}{(\lambda+vi)}\left[\frac{\varphi_T(v-(\lambda+1)i)(\lambda+1+vi-1)}{S_0^2}\right]dv \\
&= \frac{e^{-\lambda\log(K)}}{\pi}\int_0^\infty \frac{e^{-iv\log(K)}e^{-rT}}{S_0^2}\varphi_T(v-(\lambda+1)i)\,dv. \qquad (10)
\end{aligned}
$$

3.4 Simulating Lévy processes

If one wants to perform Monte Carlo calculations using Lévy processes, it is necessary to simulate from Lévy distributions. It is possible to simulate any Lévy process using an approximation of a compound Poisson process, but for specific Lévy processes, there are more sophisticated methods available. In particular, we focus on the NIG model, which can be simulated explicitly as a time-changed Brownian motion, with random time step sizes following an inverse Gaussian (IG) process. When used for option pricing, the location parameter μ of the NIG process has no effect on the price of the option, so for convenience we take $\mu = 0$.

The following algorithm can be used to simulate a sample path from a NIG$(\alpha, \beta, \delta, 0)$ distribution. In order to properly simulate the stock price, the NIG process must have a mean of zero and a variance of 1. Let T denote the time to maturity and N_t denote the number of time steps:

i. Set $dt = T/N_t$

ii. Set $a = dt$

iii. Set $b = \delta\sqrt{\alpha^2-\beta^2}$

iv. Generate N_t standard normal random variables $v_1, \ldots, v_{N_t}$

v. Set $y_i = v_i^2$ for $i = 1, \ldots, N_t$

vi. Set $x_i = \frac{a}{b} + \frac{y_i}{2b^2} - \frac{\sqrt{4aby_i + y_i^2}}{2b^2}$ for $i = 1, \ldots, N_t$

vii. Generate N_t standard uniform random variables $u_1, \ldots, u_{N_t}$

viii. If $u_i \leq \frac{a}{a+bx_i}$, set $\Delta IG_i = x_i$; Otherwise, set $\Delta IG_i = \frac{a^2}{x_i b^2}$ for $i = 1, \ldots, N_t$

ix. Generate N_t standard normal random variables $n_1, \ldots, n_{N_t}$

x. Set $X_i = \sum_{j=1}^{i} \beta\delta^2 \Delta IG_j + \delta n_j \sqrt{\Delta IG_j}$ for $i = 1, \ldots, N_t$

The X_i values follow the desired NIG process.

4 Implied Lévy volatility

We now discuss two Lévy models for a stock price process which, coupled with our option pricing model for Lévy processes, gives us the necessary tools to formulate the concept of implied Lévy volatility as introduced in [6]. We discuss Lévy space and time models, both of which can better match the heavy tails and skewness of log returns observed in Table 2.

4.1 Implied Lévy space volatility

Let $\{X_t : t \geq 0\}$ be a Lévy process. We denote the characteristic function of X_1 by

$$\phi_1(u) = E[\exp(iuX_1)]\,.$$

We additionally require that $E[X_t] = 0$ and $\mathrm{Var}[X_t] = t$, which ensures that $\mathrm{Var}[\sigma X_t] = \sigma^2 t$. Under the Lévy space model, the stock price process is modelled as

$$S_t = S_0\, e^{(r-q+\omega)t+\sigma X_t}, \qquad t \geq 0,$$

where

$$\omega = -\log(\phi_1(-\sigma i))$$

is the mean correcting term necessary to make the model risk-neutral, by making the discounted stock price a martingale. We now verify this fact.

Proposition 1. *$e^{\sigma X_t - \log(\phi_1(-\sigma i))t}$ is a martingale.*

Proof. Let $0 \leq s \leq t$.

$$\begin{aligned}
E\Big[e^{\sigma X_t - \log(\phi_1(-\sigma i))t} \mid \mathcal{F}_s\Big] &= e^{\sigma X_s - \log(\phi_1(-\sigma i))t} E\Big[e^{\sigma(X_t - X_s)} \mid \mathcal{F}_s\Big] \\
&\quad \downarrow \text{(independent increments)} \\
&= e^{\sigma X_s} e^{-\log(\phi_1(-\sigma i))t} E\Big[e^{\sigma(X_t - X_s)}\Big] \\
&\quad \downarrow \text{(stationary increments)} \\
&= e^{\sigma X_s - \log(\phi_1(-\sigma i))t} E\Big[e^{\sigma(tX_1 - sX_1)}\Big] \\
&= e^{\sigma X_s - \log(\phi_1(-\sigma i))t} E\Big[e^{-\log e^{\sigma X_1}(s-t)}\Big] \\
&= e^{\sigma X_s - \log(\phi_1(-\sigma i))t} E\Big[e^{-\log e^{i(-\sigma i)X_1}(s-t)}\Big] \\
&= e^{\sigma X_s - \log(\phi_1(-\sigma i))t} E\Big[e^{-\log E\left[e^{i(-\sigma i)X_1}\right](s-t)}\Big] \\
&= e^{\sigma X_s - \log(\phi_1(-\sigma i))t} E\Big[e^{-\log \phi_1(-\sigma i)(s-t)}\Big] \\
&= e^{\sigma X_s - \log(\phi_1(-\sigma i))s}.
\end{aligned}$$

□

In order to price options using the Carr-Madan formulation, we need to compute the characteristic function of $\log(S_T)$:

$$\begin{aligned}\varphi_T(u) &= E\Big[e^{iu\log S_T}\Big] \\ &= E\Big[e^{iu(\log S_0+(r-q+\omega)t+\sigma X_T)}\Big] \\ &= e^{iu(\log S_0+(r-q+\omega)t)}E\big[e^{iu\sigma X_T}\big] \\ &= e^{iu(\log S_0+(r-q+\omega)t)}\phi_T(\sigma u).\end{aligned}$$

The volatility parameter σ necessary to match the model price with a given market price is called the *implied Lévy space volatility* of the option.

4.2 Implied Lévy time volatility

We again begin with a Lévy process $\{X_t : t \geq 0\}$ such that $E[X_t] = 0$ and $\mathrm{Var}[X_t] = t$, and hence $\mathrm{Var}[X_{\sigma^2 t}] = \sigma^2 t$. We denote the characteristic function of X_1 by $\phi_1(u) = E[\exp(iuX_1)]$.

$$S_t = S_0\, e^{(r-q+\widetilde{\omega}\sigma^2)t+X_{\sigma^2 t}}, \qquad t \geq 0,$$

where

$$\widetilde{\omega} = -\log(\phi_1(-i))$$

is the mean correcting term necessary to make this new model risk-neutral. We now verify this fact with an expectation calculation similar to the previous section.

Proposition 2. $e^{\,X_{\sigma^2 t}-\log(\phi_1(-i))\sigma^2 t}$ *is a martingale.*

Proof. Let $0 \leq s \leq t$.

$$\begin{aligned}E\Big[e^{\,X_{\sigma^2 t}-\log(\phi_1(-i))\sigma^2 t} \mid \mathcal{F}_s\Big] &= e^{X_{\sigma^2 s}-\log(\phi_1(-i))\sigma^2 t}E\Big[e^{(X_{\sigma^2 t}-X_{\sigma^2 s})} \mid \mathcal{F}_s\Big] \\ &\quad \downarrow \text{(independent increments)} \\ &= e^{X_{\sigma^2 s}-\log(\phi_1(-i))\sigma^2 t}E\Big[e^{(X_{\sigma^2 t}-X_{\sigma^2 s})}\Big] \\ &\quad \downarrow \text{(stationary increments)} \\ &= e^{X_{\sigma^2 s}-\log(\phi_1(-i))\sigma^2 t}E\Big[e^{(\sigma^2 t X_1-\sigma^2 s X_1)}\Big] \\ &= e^{X_{\sigma^2 s}-\log(\phi_1(-i))\sigma^2 t}E\Big[e^{-\log e^{X_1}(\sigma^2 s-\sigma^2 t)}\Big] \\ &= e^{X_{\sigma^2 s}-\log(\phi_1(-i))\sigma^2 t}E\Big[e^{-\log e^{i(-iX_1)}(\sigma^2 s-\sigma^2 t)}\Big] \\ &= e^{X_{\sigma^2 s}-\log(\phi_1(-i))\sigma^2 t}E\Big[e^{-\log E\left[e^{i(-iX_1)}\right](\sigma^2 s-\sigma^2 t)}\Big] \\ &= e^{X_{\sigma^2 s}-\log(\phi_1(-i))\sigma^2 t}E\Big[e^{-\log \phi_1(-i)(\sigma^2 s-\sigma^2 t)}\Big] \\ &= e^{\,X_{\sigma^2 s}-\log(\phi_1(-i))\sigma^2 s}.\end{aligned}$$

□

In order to price options using the Carr-Madan formulation, we need to compute the characteristic function of $\log(S_T)$ under this model as well:

$$\begin{aligned}\varphi_T(u) &= E\left[e^{iu\log S_T}\right] \\ &= E\left[e^{iu(\log S_0+(r-q+\widetilde{\omega}\sigma^2)t+X_{\sigma^2 T})}\right] \\ &= e^{iu(\log S_0+(r-q+\widetilde{\omega}\sigma^2)t)}E\left[e^{iuX_{\sigma^2 T}}\right] \\ &= e^{iu(\log S_0+(r-q+\widetilde{\omega}\sigma^2)t)}\phi_{\sigma^2 T}(u).\end{aligned}$$

Similar to the Black-Scholes model, we call the volatility parameter σ necessary to match the model price with a given market price the *implied Lévy time volatility* of the option.

Notice that if we use a Brownian motion as our Lévy process, then σW_t and $W_{\sigma^2 t}$ follow the same $N(0,\sigma^2 t)$ distribution. Thus the space and time models coincide in the Black-Scholes setting and will yield the same option prices and implied volatilities. However this is not always the case for more general Lévy processes. Examples where the two models differ will be seen in the sections which follow.

4.3 Black-Scholes smiles and Lévy waves

In this section we calculate implied Lévy volatility curves for some options whose implied Black-Scholes volatility curve has a smile shape. We will consider how both symmetric and asymmetric NIG models perform at reducing the smile. For the data exhibiting the smile shape, we used the 27 July 2010 closing market prices of several AAPL (Apple Inc.) AUG 10 calls (maturing on 20 August 2010) with the following parameter settings:

$$S_0 = 264.08, \quad r = 2\%, \quad q = 0, \quad T = 24/365, \quad K = 230, 240, \ldots, 360, 370.$$

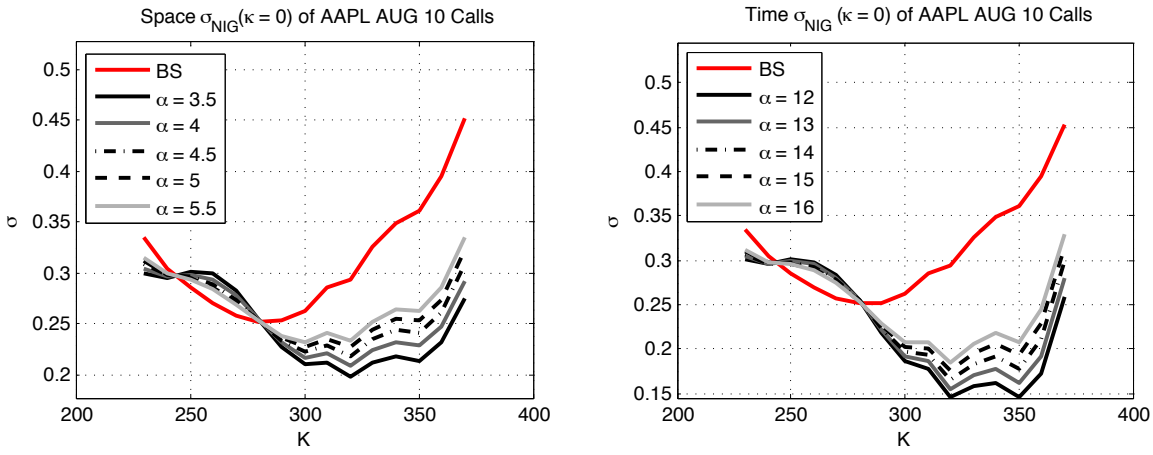

Figure 3: Implied volatility for the symmetric NIG space (left) and time (right) models.

Figure 3 shows the symmetric NIG implied Lévy space and time volatility curves for the range of values of α which performed the best. As we increase α, the shape of the implied

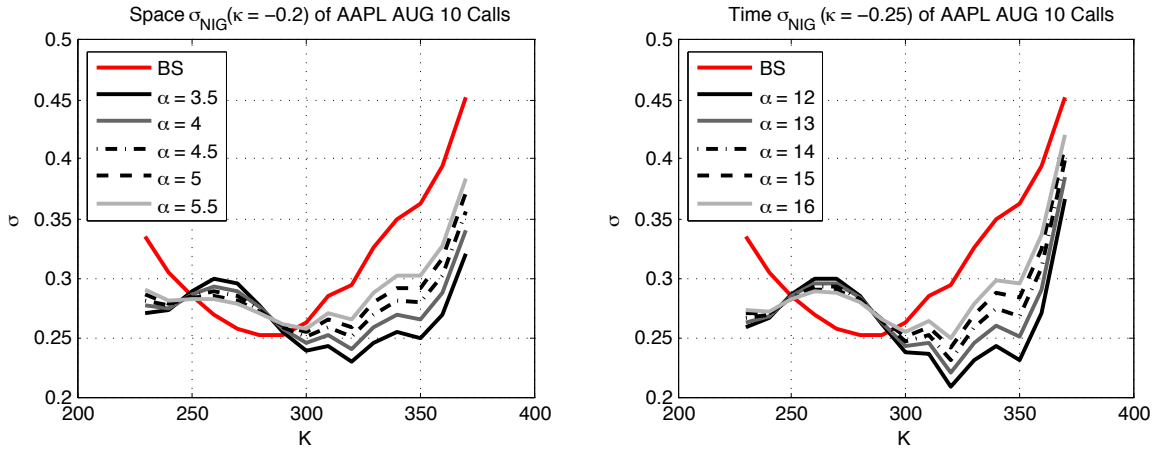

Figure 4: Implied volatility for asymmetric NIG space ($\kappa = -0.2$) (left) and time ($\kappa = -0.25$) (right) models.

Lévy volatility curves tend once again towards a smile shape. In order to obtain similar performance under the NIG space and time models, a much larger α is needed for the time model. Under the symmetric model, we see that the flattening of the smile is present but not remarkable; we attempt to improve this flattening by using an asymmetric model with carefully chosen parameters $\{\kappa, \alpha\}$.

Figure 4 shows asymmetric NIG implied Lévy space and time volatility curves for a range of α values which again were the best performing. By choosing NIG distributions with negative skewness, we were able to improve the flattening of the smile even further. We again see that in order to obtain similar performance in the space and time models, we must use a much larger α value for the time model. The NIG space model performed slightly better overall, although this is likely a result of our data set; given a different set of data or a different implied Lévy volatlilty shape, the time model could perform better.

Similar to the artifical market situation examined in [6], we see that in both the symmetric and asymmetric cases there appear to be two strike prices for which the implied Lévy volatility is invariant with respect to α. It may be possible to further improve the performance by using a different distribution, such as variance gamma, Meixner, CGMY, or any other infinitely divisible distribution for which the characteristic function can be computed.

4.4 Black-Scholes smirks and Lévy waves

We now calculate implied Lévy volatlilty curves for a different set of AAPL calls whose implied Black-Scholes volatility curve has a smirk shape. We again test both the symmetric and asymmetric model for their performance in flattening the smirk.

The options which form our data set are the 27 July 2010 closing prices of JAN 11 calls (maturing on 21 January 2011) with the following parameter settings:

$$S_0 = 264.08, \quad r = 2\%, \quad q = 0, \quad T = 178/365, \quad K = 230, 240, \ldots, 360, 370.$$

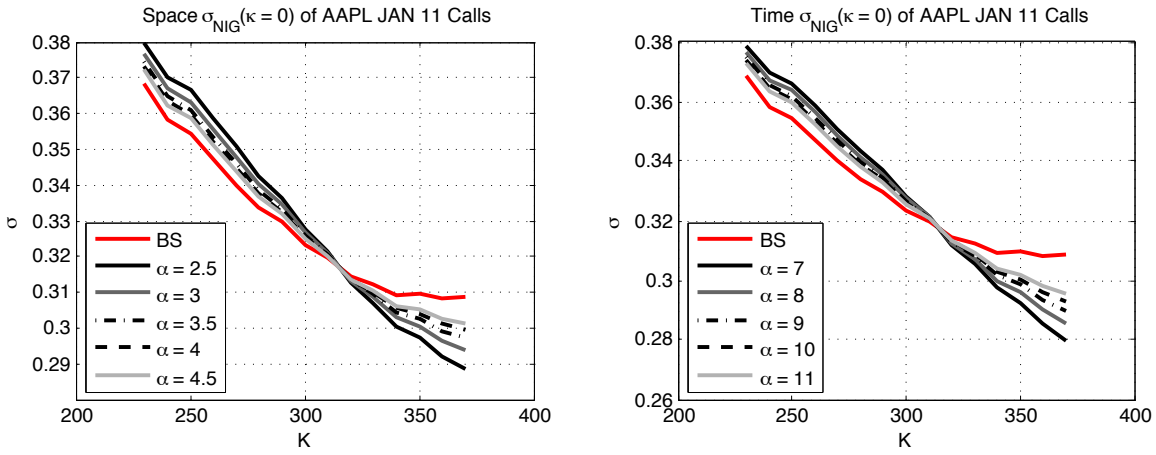

Figure 5: Implied volatility for symmetric NIG space (left) and time (right) models.

The implied volatility smirk directly captures the asymmetry of the true underlying distribution of the stock, so we expect to see the symmetric NIG distribution behave similarly to the normal distribution. This is exactly what was observed, as seen in Figure 5. Both the space and time models perform more or less the same although we required a larger α in the time case to achieve the same performance as the space model.

By moving to an asymmetric model, we attempt to obtain some form of flattening. Figure 6 shows implied volatility curves for some asymmetric NIG space and time models with $\kappa = -0.5$ for the range of values of α which were able to effectively flatten the smirk. For the ranges of α chosen, an increase in the value of α corresponded to a downward translation of the implied volatility curve. If α was increased further, the implied Lévy volatility curves tended towards the implied Black-Scholes volatility curve. If α was decreased further, the curves began to "frown" in slightly skewed manner. This behaviour is due to the inverse relationship between α and the skewness of the NIG distribution. This relationship was important when choosing our parameters as we were forced to strike a balance between tail heaviness and skewness.

We see that by moving away from a normal distribution and using a Lévy model with more parameters, we have the additional degrees of freedom necessary to calibrate our model in a way that better fits historical data.

5 Hedging performance

In order to further study how Lévy models perform when used with real market data, we compare their hedging performance with that of the Black-Scholes model. In each hedging strategy, we create a portfolio with an initial value of zero, so that the hedging error at a later date, $\text{HE}(\Delta t)$ is simply the value of the portfolio, with a value of zero corresponding to a perfect hedge.

For each model used, a different hedging error distribution will arise; as a measure of error we will calculate the absolute value of the mean plus the standard deviation and the variance of each distribution in order to compare performance. The data used in our

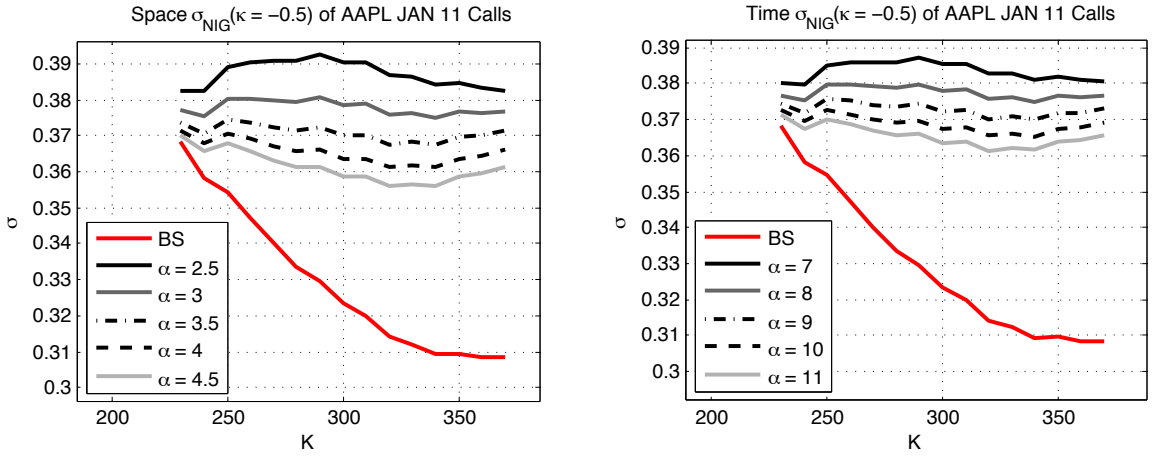

Figure 6: Implied volatility for asymmetric NIG space ($\kappa = -0.5$) (left) and time ($\kappa = -0.5$) (right) models.

analysis is composed of the VIX implied (Black-Scholes) volatility index of options on the S&P 500, as well S&P 500 price data ranging from 1 January 2007 to 27 July 2010. This data covers the financial crisis and the abnormal market dynamics that it caused. These dynamics may prove challenging for our hedges. Using this data, we first calculated the implied Lévy space and time volatilities for each day (Figure 7).

We observed that the implied Lévy models yielded consistently higher volatility values than the Black-Scholes model. During the large spike in volatility starting around day 400, the time volatility did not increase as sharply as the space volatility, suggesting that the space and time models behave similarly during a financial crisis.

Under the Black-Scholes model, the Black-Scholes formulas for the Greeks were used, with the VIX data as the input volatility. In the Lévy case, we used the formulas for the Greeks that were calculated in Section 3, coupled with the implied Lévy volatility values corresponding to the VIX data points (Figure 7).

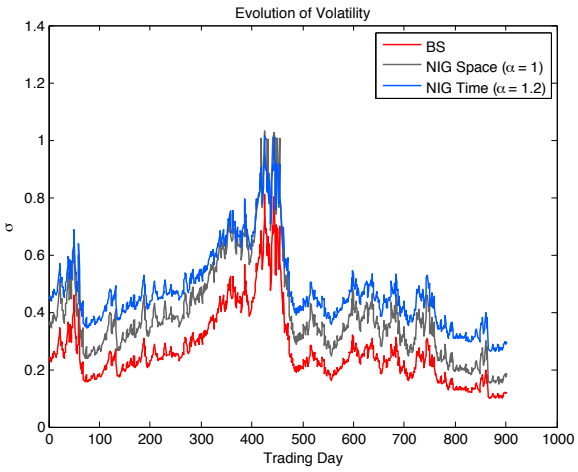

Figure 7: Evolution of implied volatility in the Black-Scholes (VIX), NIG space and NIG time models for some specific values of α.

5.1 Delta hedging

We consider a theoretical rolling at-the-money call option with one month to maturity. We delta-hedge this option for one day, and observe the hedging error. The next day we start with a fresh at-the-money call with one month to maturity and repeat this one-day hedge, producing many samples from the hedging error distribution. We use the symmetric NIG space and time models and asymmetric space and time models with κ values of 0.5 and

-0.5. for the hedging experiment. We begin by calculating the Black-Scholes, and Lévy delta values of an at-the-money option for each day.

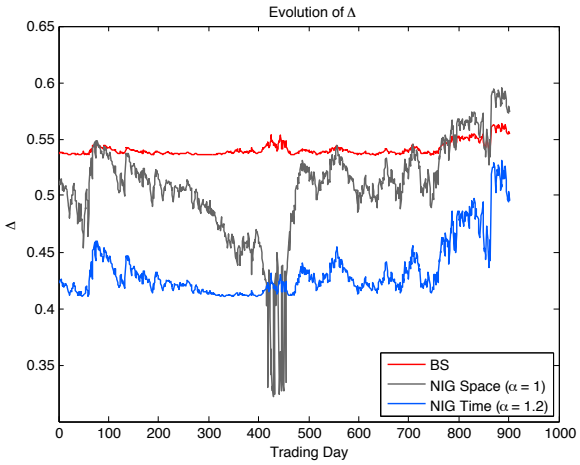

Figure 8: Evolution of Black-Scholes, and NIG space and time delta of at-the-money options on the S&P 500 for specific α vaules.

Figure 8 looks at the evolution of delta for the Black-Scholes, and NIG space and time models. It shows that the NIG time delta is hardly affected by the large spike in volatility seen around days 400-500 while the NIG space delta drops sharply. This could have a significant effect on the relative performance of the space and time hedges, as the delta dictates how much stock a trader should purchase to create a delta-neutral portfolio. It also suggests that the NIG space and time hedges are generally more conservative than Black-Scholes in the sense that they often dictate holding a smaller position in the underlying asset in order to neutralize the delta of the portfolio.

Figure 9 compares the delta hedging errors of the NIG space and time models with the Black-Scholes model. We see an improvement in hedging performance in the symmetric space and time models and the positively skewed space and time models, and decreased performance in the negatively skewed models, which is surprising given that the actual skewness of the S&P data was -0.1811.

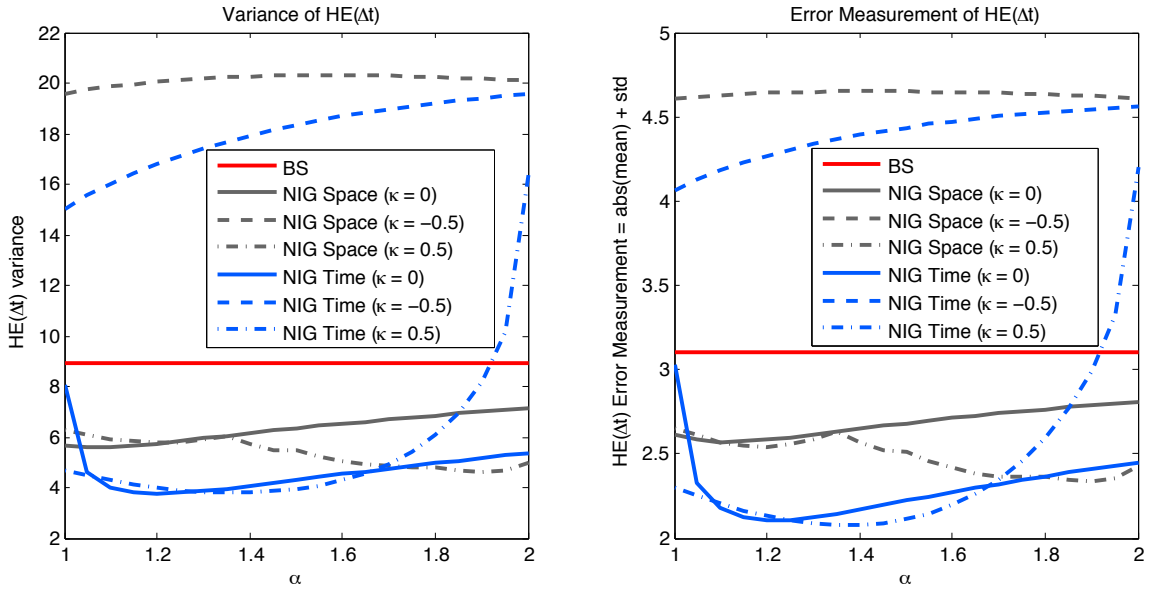

Figure 9: Variance (left) and error measurement (right) of the delta hedging error distribution of NIG space and time models.

The results are an indication of the improvement in hedging performance that can

be gained from Lévy models, when given the correct parameter settings. Comparing the variance plot with the plot of the error measurement it is clear that the standard deviation dominates the value of the error measurement.

5.2 Gamma hedging

The primary benefit of gamma-hedged portfolio is that it can be rebalanced less frequently than a delta-hedged portfolio without decreasing the performance of the hedge. This is beneficial to traders who wish to reduce transaction costs.

We consider the same rolling at-the-money call options on the S&P 500 for this analysis. Figure 10 shows the evolution of gamma under the three models over the time period of our experiment. During the large spike in volatility around days 400-500 the NIG space and time gamma values drop sharply down to similar levels to the Black-Scholes values. This means that during the peak of the financial crisis, all three gamma hedges would perform similarly. We performed three gamma hedging tests using 1 day, 2 weeks and 4 weeks between each trade in order to compare the performance. The option we used to hedge with was 0.5% out-of-the-money.

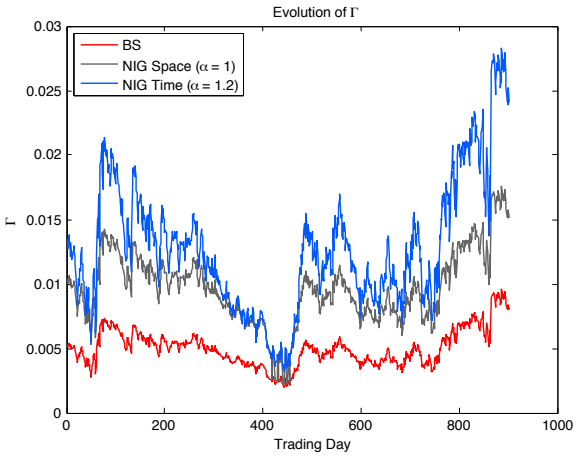

Figure 10: Evolution of Black-Scholes, and NIG space and time gamma of at-the-money options on the S&P 500 for specific α vaules.

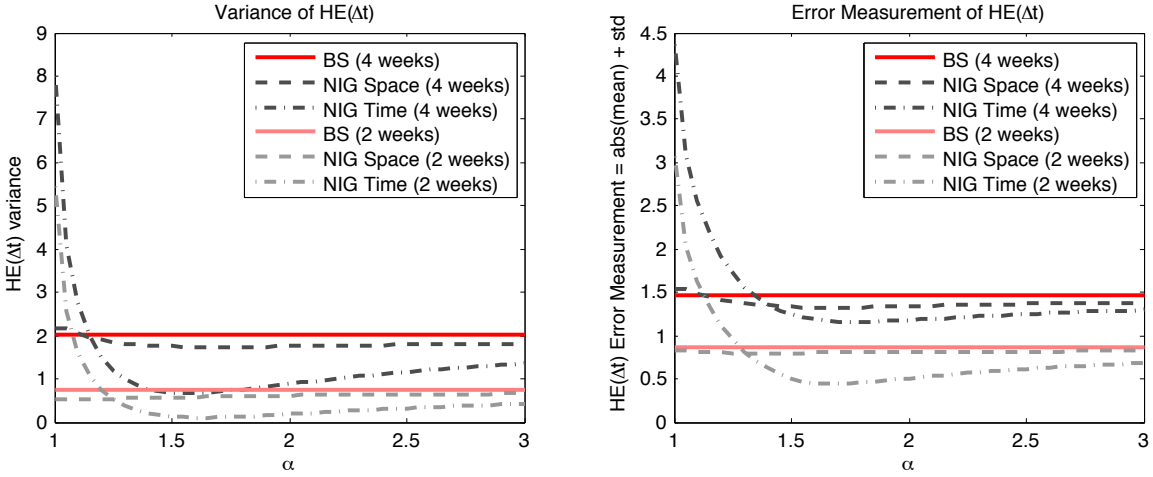

Figure 11: Variance (left) and error measurement (right) of the gamma hedging error distribution of NIG space and time models.

Figure 11 and Figure 12 show the hedging errors for the three trading frequencies examined. We see an improvement in gamma hedging performance in the NIG space and time models over Black-Scholes when certain α values are used. As expected, the more frequently rebalanced (daily) portfolio performed the best, though the shapes of the

hedging error curves are similar for all rebalancing frequencies. The optimal hedging error distributions corresponded to an α value of around 1.6 in both the space and time models. We again observe that the value of the error measurement is dominated by the standard deviation.

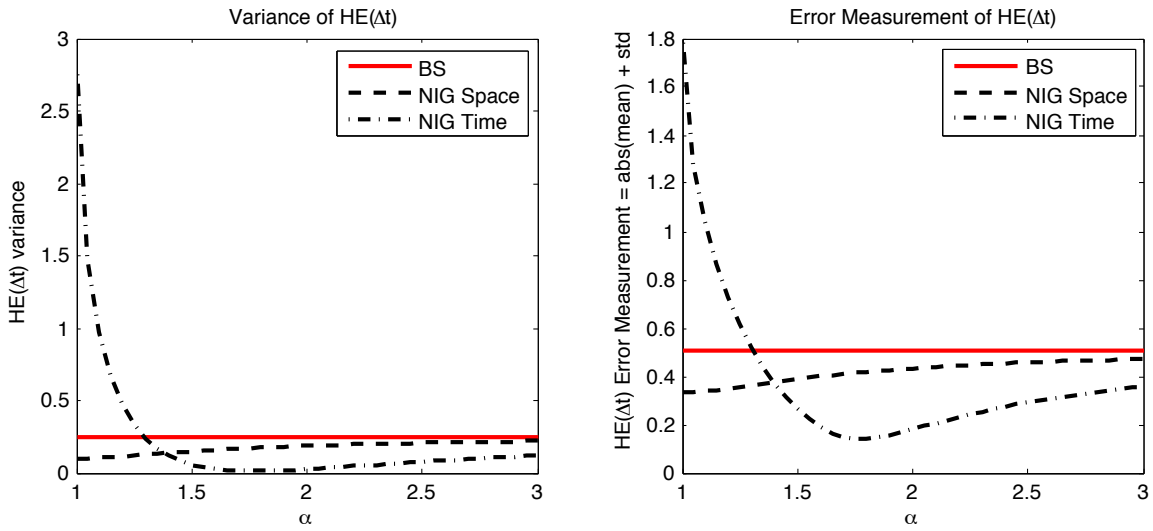

Figure 12: Variance (left) and error measurement (right) of the gamma hedging error distribution of NIG space and time models using daily hedging.

5.3 Delta and gamma hedging

When a portfolio is gamma-hedged in practice, it is typically delta-hedged as well. This type of hedge can theoretically outperform a portfolio that is either delta or gamma-hedged exclusively, for a carefully chosen rebalancing frequency. The option we used to gamma hedge with was again 0.5% out-of-the-money.

We include the results of delta/gamma hedging for 1 and 2 weeks between trades in Figure 13, and with daily trading in Figure 14. If we moved to 3 weeks or longer between trades, the Lévy models began failing to outperform the Black-Scholes model. Leaving such a large gap between trades allows the delta of the portfolio to depart considerably far from neutrality, and it seems the Lévy models are more sensitive to this change in trading frequency than the Black-Scholes model. This makes sense since Lévy deltas are generally smaller than the Black-Scholes delta values (Figure 8), suggesting that the smaller positions (in the underlying) would change by a larger percentage between trades.

Overall, we again see a performance improvement in the size of the hedging error of the Lévy models, but for a reduced range of α values; indicating that for a carefully chosen α, we can reduce the error measurement of the hedging error distribution. The improvement in the value of the error measurement was not as dramatic in this case though, possibly due to the poor performance of the delta hedge with such a large time gap between trades. However, the variance of the hedging error was reduced by almost 50% for some values of α which is a more promising result.

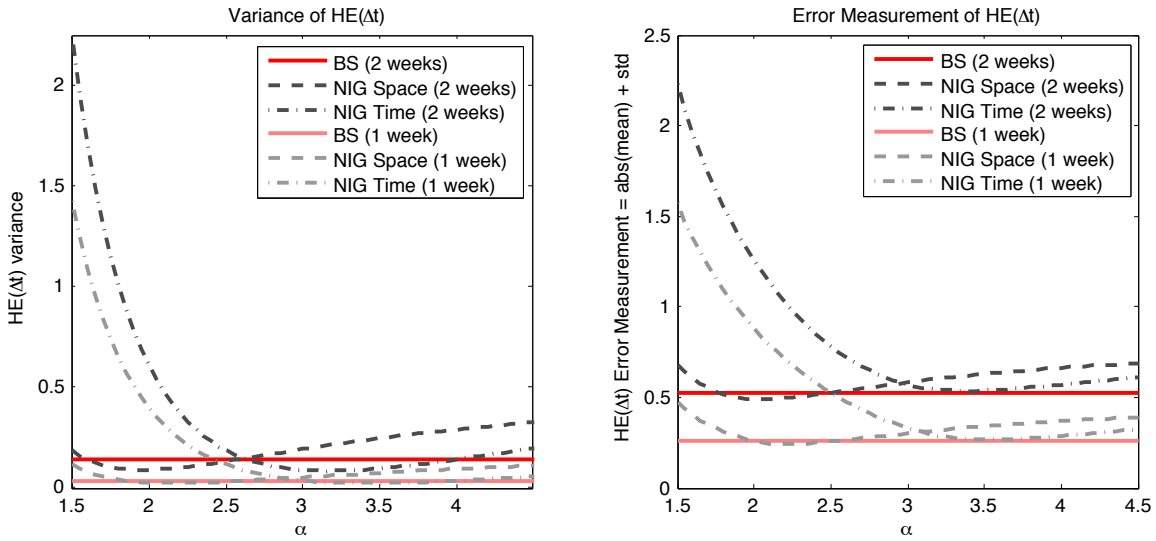

Figure 13: Variance (left) and error measurement (right) of the delta/gamma hedging error distribution of NIG space and time models.

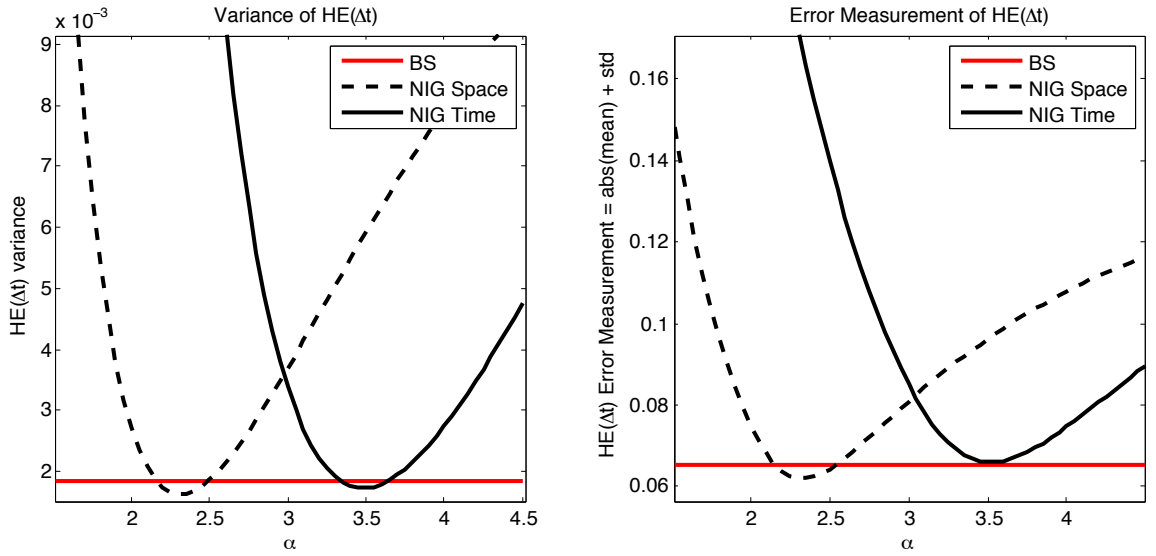

Figure 14: Variance (left) and error measurement (right) of the delta/gamma hedging error distribution of NIG space and time models using daily rebalancing.

6 Barrier pricing performance

As a test of how the prices of barrier options are affected by which pricing model and which implied volatility is used, we calculate the prices of theoretical barrier call options on AAPL stock using the implied volatility values from the asymmetric "smirk" NIG space case (Figure 6). We perform four sets of Monte Carlo simulations using vanilla call option prices as a control variate to reduce the variance of the sample payoffs. The four sets of simulations performed were configured as follows:

1. Geometric Brownian motion price process using $\sigma_{BS}(K)$
2. Geometric Brownian motion price process using $\sigma_{NIG}(K)$
3. NIG space price process using $\sigma_{BS}(K)$
4. NIG space price process using $\sigma_{NIG}(K)$

The parameter settings for the simulations were:

$$\begin{array}{llll} S_0 = 264.08, & r = 2\%, & q = 0, & T = 178/365, \\ K = 230, 240, \ldots, 360, 370, & B_{up} = 350, & B_{down} = 250. & \end{array}$$

The barrier levels B_{up} and B_{down} were chosen to ensure that the barrier option prices were not particularly one-sided. Recall the no-arbitrage condition that the sum of the prices of an "in" and "out" pair of options should equal the price of a vanilla option. This means that a decrease in the price of one member of the pair results in an increase in the price of the other member.

Our primary interest is in examining how moving from implied Black-Scholes volatility to implied Lévy volatility affects the barrier prices. Pricing the options in the four ways described will allow us to isolate the effects of the model (kurtosis, skewness, etc...) from the effects of the volatility. We create four separate pricing graphs for up-and-in, up-and-out, down-and-in, and down-and-out barrier call options. Figure 15 looks at the barrier prices obtained by using both a geometric Brownian motion price process and a Lévy price process, with both the Black-Scholes and Lévy implied volatilities. Table 4 contains the percentage values of barrier crossings that were observed in the simulations.

6.1 Geometric Brownian motion price process

We first discuss the prices obtained using the geometric Brownian motion price process. We see in Figure 15 that in the case of "up" barrier options, the increase in volatility from switching to implied Lévy volatility caused an increase in the price of up-and-in options and a slight decrease in the price of up-and-out options. This makes intuitive sense because a higher volatility makes the stock price more likely to cross the barrier, thus making "in" options more valuable and "out" options less valuable. The first two rows of Table 4 show an increase of approximately a 7–8% in the percentage of "up" barrier crossings observed after switching from Black-Scholes to Lévy volatility.

The results are slightly different in the case of the "down" options. The Lévy volatility prices were higher for both the down-and-in and the down-and-out calls. In this case the

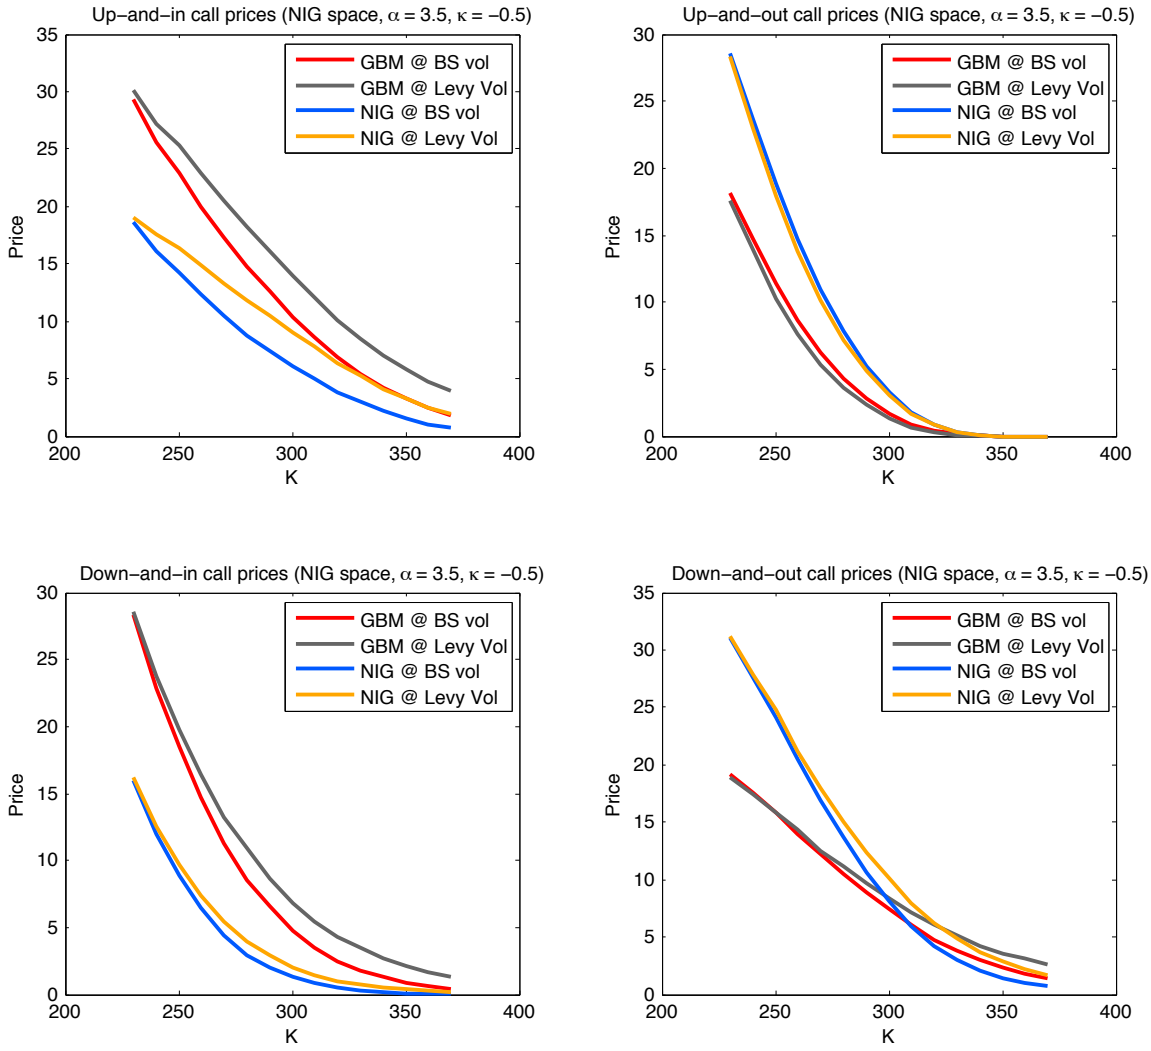

Figure 15: Barrier option prices for Black-Scholes and Lévy models, at implied Black-Scholes and Lévy volatilities.

results in Table 4 show that the number of "down" barrier crossings grew by 3 to 4 percent, which on its own would cause a decrease in the price of down-and-out options. The fact that we observed an increase in price must therefore be due to the option having matured deeper and more frequently in-the-money, causing higher observed payoffs and thus driving up the price.

6.2 Lévy price process

The NIG space model which yielded the flattest implied Lévy volatility curve in Figure 6 was chosen as the barrier pricing model. In this model, α was set to 3.5 and κ was set to -0.5. The skewness of this distribution is -0.5714; a relatively large negative skewness, which made upward movements in the stock price more likely than downward movements.

Figure 15 shows results similar to the previous section for barrier prices when a Lévy

Table 4: Percentage of Monte Carlo simulations where barrier was crossed for each type of barrier option and pricing configuration.

	Up-and-in	Up-and-out	Down-and-in	Down-and-out
GBM @ BS vol	17.07	16.97	79.58	79.70
GBM @ Levy vol	24.13	24.39	83.77	83.29
NIG @ BS vol	8.91	9.04	62.01	62.10
NIG @ Levy vol	14.98	14.79	66.51	66.27

price process is used. We see in the case of "up" options that the higher (implied Lévy) volatility increased the value of the up-and-in options and decreased slightly the value of the up-and-out options. This was coupled with a 5–6% increase in the number of barrier crossings. In the case of "down" options we observe that the negative skewness caused the number of barrier crossing events that occurred to increase by only 3–4%. So the increase in price when using Lévy volatility must again be caused by the options finishing deeper in-the-money on average. This was helped by the large negative skewness our distribution was calibrated with.

The effect on the prices by switching to a Lévy process was much more significant than the effect of the volatility. In Figure 15 we see in both the "up" and "down" cases that the switch to a Lévy process caused a decrease in the price of "in" options and an increase of equal value in the price of "out" options. This suggests a decrease in the number of barrier crossings upon switching to the Lévy process. Table 4 shows that in all 8 cases (4 types of barrier option, 2 types of volatility), the move to a Lévy process did dramatically decrease the number of observed barrier crossings.

The decrease in barrier crossings is surprising given that our Lévy process has a skewness of -0.5714 and a kurtosis of 3.807; slightly higher than the normal distribution. We would expect the move to a heavier tailed distribution to cause more barrier crossings. The fact that we observed less crossings implies that the NIG process we used tends to remain closer to S_0 over the course of the simulations compared to a geometric Brownian motion.

7 Conclusion

By using the concept of implied Lévy volatility first introduced by Corcuera et. al. in [6], we have shown that switching to a more flexible Lévy distribution allows us to perform several market functions more accurately than under the Black-Scholes model.

We first tested the performance of implied Lévy volatility with real market data. It was demonstrated that the curvature of an implied volatility plot could be minimized with a properly calibrated Lévy model. The additional degrees of freedom allowed us to use asymmetric distributions with heavier tails, which better fit the return distributions observed historically. This gave improvements over implied Black-Scholes volatility smiles and smirks, with the smirk case showing the greatest improvement by flattening the curve almost completely.

It was shown that one can calibrate a Lévy model such that the absolute mean and variance of the hedging error are lower than those values obtained by the Black-Scholes

model for delta, gamma and delta/gamma hedging. Even using data covering the 2007 financial crisis, we saw a reduction in the hedging error of as much as 50% in certain delta and gamma hedging models. The ability to reduce hedging error is desirable to any option trader as it can effectively minimize losses due to the errors that can arise from hedging discretely. This can in turn lead to lower economic capital requirements and thus higher margins for profits.

Lastly we used the flattened volatility curve obtained earlier to analyze the effect of volatility on the price of barrier options. A flatter volatility curve is desirable because it eliminates the problem of which implied volatility one should input into the pricing function; the volatility corresponding to the strike price, the barrier level or perhaps some average of the two. We saw changes in our Monte Carlo barrier prices caused by changes in the percentage of barrier crossings, and also how deep and how frequently the options finished in-the-money. If barrier options are priced using a model which is more accurately calibrated to the market, these changes in barrier crossings and moneyness should reflect more closely the type of behaviour that is observed historically.

8 Future Work

The results we have discussed here were for a normal inverse Gaussian process, but the procedures we used can be applied using any Lévy process for which the characteristic function is computable. This includes, for instance the Meixner, variance gamma, CGMY, and generalized tempered stable processes. We could possibly improve performance further by seeing how each of these models compare at flattening the implied volatility surface, especially in the case of a "smile", where the results we saw showed room for improvement.

There are other forms of hedging which can be expanded to Lévy models. These most notably include theta; sensitivity of the option value with respect to time, and vega; the sensitivity with respect to volatility. By formulating the theta and vega of an option under the Carr-Madan framework, we can compare the performance of a broader range of hedging strategies.

Lastly, the techniques we have for pricing barrier options using implied Lévy volatility can be applied to any other exotic option where there is some uncertainty in which volatility to use for pricing. The change in performance upon switching to implied Lévy volatility could be compared with what was seen in barrier options which could provide deeper understanding of how the volatility affects the value of exotic options.

References

[1] Back, K. (2005) *A Course in Derivative Securities*. Springer-Verlag Berlin Heidelberg.

[2] Barndorff-Nielsen, O.E. (1998) *Processes of Normal Inverse Gaussian Type*. Finance and Stochastics, **2**, 41-68.

[3] Black, F. and Scholes, M. (1973) *The Pricing of Options and Corporate Liabilities*. Journal of Political Economy, **81**, 637-654.

[4] Carr, P. and Madan, D. (1999) *Option Valuation using the Fast Fourier Transform*. Journal of Computational Finance, **2**, 61-73.

[5] Cont, R. and Tankov, P. (2003) *Financial Modelling with Jump Processes*. Chapman & Hall / CRC Press.

[6] Corcuera, J.M., Guillaume, F., Leoni, P., Schoutens, W. (2009) *Implied Lévy Volatility*. Journal of Quantitative Finance, **9**, 383-393.

[7] Schoutens, W. (2003) *Lévy Processes in Finance: Pricing Financial Derivatives*. John Wiley & Sons Ltd.

[8] Schoutens, W. and Symens, S. (2003) *The Pricing of Exotic Options by Monte Carlo Simulations in a Lévy Market with Stochastic Volatility*. Journal for Theoretical and Applied Finance, **6**, 839-864.

[9] Shreve, S. (2004) *Stochastic Calculus for Finance II*. Springer Science + Business Media, LLC.