# A Monte Carlo Method for the Normal Inverse Gaussian Option Valuation Model using an Inverse Gaussian Bridge

Claudia Ribeiro
City University
Cass Business School
London EC2Y 8HB
United Kingdom
Tel: +44 20 7040 8769
Email: C.A.G.C.Ribeiro@city.ac.uk

Nick Webber*
City University
Cass Business School
London EC2Y 8HB
United Kingdom
Tel: +44 20 7040 5171
Email: Nick.Webber@city.ac.uk

February 18, 2003

First version: November 2002

**Abstract**

The normal inverse Gaussian process has been used to model both stock returns and interest rate processes. Although several numerical methods are available to compute, for instance, VaR and derivatives values, these are in a relatively undeveloped state compared to the techniques available in the Gaussian case.

This paper shows how a Monte Carlo valuation method may be used with the NIG process, incorporating stratified sampling together with an inverse Gaussian bridge.

The method is illustrated by pricing average rate options. We find the method is up to around `200` times faster than plain Monte Carlo. These efficiency gains are similar to those found in a related paper, Ribeiro and Webber (02) [20], which develops an analogous method for the variance-gamma process.

*Corresponding author. Claudia Ribeiro gratefully acknowledges the support of Fundação para a Ciência e a Tecnologia and Faculdade de Economia, Universidade do Porto. We are grateful for helpful discussions with participants at QMF Sydney 2002, and to Lynda McCarthy for detailed comments on the manuscript.

# 1 Introduction

The normal inverse Gaussian model has been investigated by a number of authors and applied to to option valuation (for instance Eberlein and Keller (95) [13], Barndorff-Nielsen (95) [4], (98) [5], Rydberg (97) [21] and Carr, Geman, Madan and Yor (01) [10] amongst others). Analytical solutions are not usually available for European-style options, so valuation requires the use of numerical methods. These include Monte Carlo methods, (Rydberg (97) [21]) Fourier transform (FFT) methods (Carr and Madan (99) [11], Carr and Wu (02) [12]), and PDE approaches (the 'method of lines'; Albanese, Jaimungal and Rubisov (01a) [1], (01b) [2]).

This paper investigates the use of Monte Carlo methods with the normal inverse Gaussian (NIG) model. In particular we show how an inverse Gaussian bridge may be constructed and used in conjunction with stratified sampling. We apply the method to price average rate options, demonstrating that considerable efficiency gains are possible. The inverse Gaussian bridge can be used with other variance reduction techniques to achieve even greater speed-ups.

Other authors have exploited the standard Brownian bridge. Beaglehole, Dybvig and Zhou (97) [6] use a knowledge of the distribution of extremes of a Brownian bridge to significantly speed up a Monte Carlo method for pricing barrier options. In a related paper Ribeiro and Webber (02) [20] showed how to apply a bridge method to the variance-gamma process. They found large efficiency gains, particularly for average rate options. In the NIG case we find gains of around 200 for average rate options.

The second section of this paper recaps the normal inverse Gaussian process and its application to option pricing. We review how Monte Carlo methods may be applied by exploiting the subordinated Brownian motion representation of the normal inverse Gaussian process. In the third section we show how an inverse Gaussian bridge can be constructed and applied. The fourth section presents numerical results and the fifth section concludes.

# 2 The Normal Inverse Gaussian Framework

We review the normal inverse Gaussian process and its application to option pricing. We describe a 'plain' Monte Carlo method related to the subordinated Brownian motion representation of the NIG process.

## 2.1 The Normal Inverse Gaussian Process

The NIG distribution is a subclass of the generalized hyperbolic distributions introduced by Eberlein and Keller (95) [13] and Barndorff-Nielsen (95) [4], (97) [5], subsequently investigated by, for example, Bibby and Sørensen (01) [7]. Its use for financial modelling is discussed by Rydberg (97) [21], (99) [22], who also describes and applies a Monte Carlo method. She fits to time series of daily stock returns using a maxiumum likelihood method. Prause (97) [18] also esti-

mates parameter values from stock returns. Bølviken and Beth (00) [8] fit the NIG model to historical returns data by matching the first four moments of the returns process.

The normal inverse Gaussian process $L_t$ is a Lévy process where increments in $L_t$ are distributed according to the NIG distribution. It has parameters $\mu, \delta, \alpha, \beta \in \mathsf{R}$, with $\delta > 0$ and $0 \leq |\beta| \leq \alpha$. Set $\delta_t = \delta t$, $\mu_t = \mu t$ and $\gamma = \sqrt{\alpha^2 - \beta^2}$. The density of $L_t$ conditional on $L_0 = 0$ is

$$f_t^{\mathrm{NIG}}(l; \alpha, \beta, \delta_t, \mu_t) = \frac{\alpha}{\pi} \frac{K_1\left(\alpha\delta_t \sqrt{1 + \left(\frac{l-\mu_t}{\delta_t}\right)^2}\right)}{\sqrt{1 + \left(\frac{l-\mu_t}{\delta_t}\right)^2}} \exp\left(\delta_t\left(\gamma + \beta\left(\frac{l-\mu_t}{\delta_t}\right)\right)\right) \tag{1}$$

where $K_\lambda(z)$ is the modified Bessel function of the second kind,

$$K_\lambda(z) = \frac{1}{2}\int_0^\infty y^{\lambda-1} \exp\left(-\frac{1}{2}z\left(y + y^{-1}\right)\right) \mathrm{d}y. \tag{2}$$

The central moments of the distribution are

$$\mu_1 = \mu_t + \delta_t \beta \gamma^{-1}, \tag{3}$$
$$\mu_2 = \delta_t \alpha^2 \gamma^{-3}, \tag{4}$$
$$\mu_3 = 3\delta_t \beta \alpha^2 \gamma^{-5}, \tag{5}$$
$$\mu_4 = 3\delta_t \alpha^2 \left(\alpha^2 + 4\beta^2\right) \gamma^{-7}, \tag{6}$$

and skewness and kurtosis are

$$\text{skew:} \qquad \frac{\mu_3}{\mu_2^{3/2}} = 3\frac{\beta}{\alpha\sqrt{\delta_t \gamma}}, \tag{7}$$

$$\text{kurtosis:} \qquad \frac{\mu_4}{\mu_2^2} = 3\frac{\alpha^2 + 4\beta^2}{\delta_t \alpha^2 \gamma}. \tag{8}$$

The class of normal inverse Gaussian distributions is closed under convolution.

The characteristic triplet of the NIG distribution in the Lévy-Khintchine representation is $(a, 0, \ell^{\mathrm{NIG}})$ where

$$a = \mu + 2\frac{\delta\alpha}{\pi}\int_0^1 \sinh(\beta x) K_1(\alpha x) \mathrm{d}x \tag{9}$$

and the Lévy measure $\ell^{\mathrm{NIG}}(\mathrm{d}l) = k^{\mathrm{NIG}}(l)\,\mathrm{d}l$ has Lévy density

$$k^{\mathrm{NIG}}(l) = \pi^{-1}\delta\alpha |l|^{-1} K_1(\alpha |l|) \mathrm{e}^{\beta l}. \tag{10}$$

The NIG process $L_t$ can be represented as a subordinated Brownian motion, $L_t = \mu_t + w_{h(t)}$, where $w_t$ is Brownian motion with drift $\beta$ and variance 1 and

$h(t)$ is an inverse Gaussian process $h_t \sim \mathrm{IG}(\delta_t, \gamma)$. The density $f_t^{\mathrm{IG}}(x)$ of $h_t$ conditional on $h(0) = 0$ is

$$f_t^{\mathrm{IG}}(x) = \frac{\delta_t}{\sqrt{2\pi}} x^{-\frac{3}{2}} \exp\left(-\frac{1}{2}\frac{\gamma^2}{x}\left(x - \frac{\delta_t}{\gamma}\right)^2\right). \tag{11}$$

An alternative parameterisation is

$$f_t^{\mathrm{IG}}(x) = \sqrt{\frac{\lambda_t}{2\pi}} x^{-\frac{3}{2}} \exp\left(-\frac{\lambda_t}{2}\frac{(x-\mu_t)^2}{x\mu_t^2}\right) \tag{12}$$

where $\mu_t = \frac{\delta_t}{\gamma}$ and $\lambda_t = \delta_t^2$.

The NIG process is a special case of a set of generalised hyperbolic distributions. This set can be represented as a time changed Brownian motion where the time change $h_t$ is a member of a set of generalised inverse Gaussian distibutions indexed by $t$, $h_t \sim \mathrm{GIG}(\delta_t, \lambda, \gamma)$, with density $f_t^{\mathrm{GIG}}$,

$$f_t^{\mathrm{GIG}}(h; \delta_t, \lambda, \gamma) = \left(\frac{\gamma}{\delta_t}\right)^{\lambda} \frac{1}{2K_\lambda(\delta_t\gamma)} h^{\lambda-1} \exp\left(-\frac{1}{2}\left(\frac{\delta_t^2}{h} + \gamma^2 h\right)\right). \tag{13}$$

The set of distributions $\mathrm{GIG}(\delta_t, \lambda, \gamma)$ is not closed under convolution. The NIG process is a special case with an inverse Gaussian time change, a GIG process with $\lambda = -\frac{1}{2}$.[1] This subset is closed under convolution. The gamma process can be expressed as a limiting case of equation (13) with $\delta = 0$ and $\lambda = \lambda_t = \frac{t}{v}$, also closed under convolution.

## 2.2 Option Pricing and Subordinator Monte Carlo Methods

Let $S_t$ be the price at time $t$ of a non-dividend paying stock. We use the normal inverse Gaussian process $L_t$ to model returns to $S_t$. We take the state space $\Omega$ to be the path space of $L_t$ equipped with the filtration induced $L_t$. We model log-returns to the stock price process $S_t$ under the pricing measure $F$ as a Lévy process. Following Madan, Carr and Chang (98) [16] and Eberlein and Raible (99) [14] we set

$$S_t = S_0 \exp(rt + L_t - \varpi t) \tag{14}$$

where $L_t$ is a normal inverse Gaussian process, $r$ is the constant short rate, and

$$\varpi = \mu + \delta\gamma - \delta\sqrt{\alpha^2 - (1+\beta)^2} \tag{15}$$

is a compensator term, defined by $\mathrm{e}^{\varpi} = \mathsf{E}[\exp(L_1)]$, to ensure that $S_t\mathrm{e}^{-rt}$ is a martingale under $F$.[2]

[1] When $\lambda = -\frac{1}{2}$, $K_{-\frac{1}{2}}(z) = K_{\frac{1}{2}}(z)$ reduces to $K_{\frac{1}{2}}(z) = \sqrt{\frac{\pi}{2z}}\mathrm{e}^{-z}$.

[2] The focus of this paper is upon numerical solutions to options writen on $S_t$, and we are not concerned with the specification of $F$.

In the martingale framework the value $c_t$ at time $t < T$ of an option with payoff $H_T \equiv H_T(\omega)$ at time $T$ is

$$c_t = \mathsf{E}_t\left[H_T \mathrm{e}^{-r(T-t)}\right] \tag{16}$$

where we assume the use of the accumulator account numeraire and its associated measure. $H_T$ may depend on the state $\omega \in \Omega$. In this section we recall how (16) can be solved by Monte Carlo integration.

Construct a set $\{\widehat{\omega}^m\}_{m=1,\ldots,M}$ of discrete sample paths randomly selected under a measure $\widehat{F}$, a discrete time approximation to the measure $F$. Then the approximation $\widehat{c}_t$ to $c_t$ is

$$\widehat{c}_t = \mathrm{e}^{-r(T-t)} \frac{1}{M} \sum_{m=1}^{M} H_T(\widehat{\omega}^m). \tag{17}$$

Discrete sample paths for a subordinated Brownian motion, $L_t = w_{h(t)}$, can be constructed by first constructing discrete sample paths for the subordinator $h(t)$ and then sampling the process $w_t$ at times determined by the paths found for $h(t)$.[3]

Discretise time as $0 = t_0 < t_1 < \ldots < t_N = T$, and set $\Delta t_n = t_{n+1} - t_n$.[4] First construct a discrete sample path $\left\{\widehat{h}_n\right\}_{n=0,\ldots,N}$ for $h(t)$. Set $\widehat{h}_0 = 0$ and iteratively generate increments $\Delta\widehat{h}_n = \widehat{h}_{n+1} - \widehat{h}_n \sim \mathrm{IG}(\delta_{\Delta t_n}, \gamma)$. Now set $\widehat{w}_0 = 0$, and iteratively generate increments $\Delta\widehat{w}_n = \widehat{w}_{n+1} - \widehat{w}_n \sim \mathrm{N}\left(\beta_{\Delta\widehat{h}_n}, \Delta\widehat{h}_n\right)$. The path $\widehat{\omega} = \left\{\widehat{w}_n + \mu_{t_n}\right\}_{n=0,\ldots,N}$ is a discrete approximation to a sample path $\omega$ of $L_t$.

For the plain Monte Carlo method construct $M$ discrete sample paths $\{\widehat{\omega}^m\}_{m=1,\ldots,M}$, as above, then compute $H_T(\widehat{\omega}^m)$. The plain Monte Carlo estimate is given by (17).

Although the plain Monte Carlo method will give estimates $\widehat{c}_t$ that converge to the true option value $c_t$, convergence may be slow. The next section describes an improved Monte Carlo method exploiting stratified sampling and an inverse Gaussian bridge.

# 3 A Inverse Gaussian Bridge for the Normal Inverse Gaussian Process

Stratified sampling in conjunction with a Brownian bridge is often used to simulate processes driven by Wiener processes. This method is well known[5] to give effective speed-ups when valuing path dependent options. Ribeiro and

[3] This is the procedure adopted by Rydberg (97) [21]

[4] Later we may assume that $\Delta t = \frac{T}{N}$ is a constant. This assumption can easily be relaxed.

[5] For one application see Beaglehole, Dybvig and Zhou (97) [6]. See also Jackel (02) [15].

Webber (02) [20] outline the use of stratified sampling and review the theory of the Brownian bridge method. They go on to construct a bridge for the variance-gamma process. In this section we describe the construction of an inverse Gaussian bridge and its application to the normal inverse Gaussian process.

## 3.1 Stratified Sampling

A stratified sample permits superior sampling of an underlying distribution. Given a sample $\{\widehat{v}^m\}_{m=1,\ldots,M}$ drawn from the uniform distribution $\mathrm{U}[0,1]$, the set $\{\widehat{u}^m\}_{m=1,\ldots,M}$, where $\widehat{u}^m = \frac{m-1+\widehat{v}^m}{M}$, is a stratified sample of $\mathrm{U}[0,1]$. It has a single draw from each quantile band $\left[\frac{m-1}{M}, \frac{m}{M}\right]$, $1 \leq m \leq M$. Given a state variable $X_t$ with distribution function $F_t^X$ a stratified sample $\left\{\widehat{X}_t^m\right\}_{m=1,\ldots,M}$ of $X_t$ can be formed by the inverse transform method if the function $\left(F_t^X\right)^{-1} : [0,1] \to \mathsf{R}$ is known or can be constructed. Simply set $\widehat{X}_t^m = \left(F_t^X\right)^{-1}(\widehat{u}^m)$ where $\{\widehat{u}^m\}_{m=1,\ldots,M}$ is a stratified sample of $\mathrm{U}[0,1]$. If the payoff function $H_T$ depends solely on the value $X_T$ at time $T$, estimating $\widehat{c}_t$ using a stratified sample $\left\{\widehat{X}_T^m\right\}_{m=1,\ldots,M}$ will often result in an estimate with significantly reduced standard deviation.[6]

For many exotic options the payoff $H_T$ depends on a sample path for $X_t$ and not just upon its terminal value $X_T$. To exploit stratified sampling a bridge technique is required. Given a stratified sample $\left\{\widehat{X}_T^m\right\}_{m=1,\ldots,M}$ of $X_T$ a bridge method constructs a set of paths $0 = \widehat{X}_0^m < \widehat{X}_1^m < \ldots < \widehat{X}_N^m = \widehat{X}_T^m$ so that each $\widehat{X}_n^m$ has its correct conditional distribution. Intermediate points $\widehat{X}_n^m$ are constructed by sampling from a bridge distribution, defined and described in the next section. This sampling may also be stratified, leading to improved sampling at the intermediate times and of the path as a whole.

A bridge Monte Carlo algorithm for the NIG process $L_t$ proceeds as follows. First, construct a stratified sample $\left\{\widehat{h}_N^m\right\}_{m=1,\ldots,M}$ from $\mathrm{IG}(\delta_{t_N}, \gamma)$. Second, construct an inverse Gaussian bridge, $\widehat{h}^m = \left(\widehat{h}_0^m, \ldots, \widehat{h}_N^m\right)$, $m = 1, \ldots, M$, starting from $\widehat{h}_0^m = 0$. The bridge may be further stratified at intermediate times. Third, for each $m = 1, \ldots, M$ generate a sample point $\widehat{w}_N^m \sim \mathrm{N}\left(\beta_{\widehat{h}_N^m}, \widehat{h}_N^m\right)$ with mean rate $\beta$ and variance 1. Fourth, construct a bridge $\widehat{w}^m = (\widehat{w}_0^m, \ldots, \widehat{w}_N^m)$, at times $\widehat{h}_0^m, \ldots, \widehat{h}_N^m$, $m = 1, \ldots, M$, with $\widehat{w}_0^m = 0$. This is a standard Brownian bridge, as described in Ribeiro and Webber (02) [20], for instance. This bridge may also be stratified at intermediate times. Finally set $\widehat{L}^m = \left(\widehat{L}_0^m, \ldots, \widehat{L}_N^m\right)$ where $\widehat{L}_j^m = \widehat{w}_j^m + \mu_{t_j}$ for the NIG drift $\mu$. This is a stratified sample path for $L_t$.

[6] The true standard deviation of the stratified Monte Carlo estimate will often be significantly less than the reported standard error.

We now describe how the inverse Gaussian bridge is constructed and sampled.

## 3.2 A Bridge for an Inverse Gaussian Process

Suppose that $X \sim F_X$, $Y \sim F_Y$ and $Z \sim F_Z$ are random variables with densities $f_X$, $f_Y$ and $f_Z$, with $Z = X + Y$. We shall be interested in the case where $X$, $Y$ and $Z$ are increments in an inverse Gaussian process over intervals $[t_i, t_j]$, $[t_j, t_k]$ and $[t_i, t_k]$ respectively. Write $f_{X,Y}$ for the joint density function of $X$ and $Y$. Given a sample $z$ of $Z$ we are interested in sampling $X$ with the correct conditional distribution. The conditional density $f_{X|Z}$ of $X \mid Z$ is

$$f_{X|Z}(x) = \frac{f_{X,Y}(x, z-x)}{f_Z(z)} \tag{18}$$

$$= \frac{f_X(x) f_Y(z-x)}{f_Z(z)} \tag{19}$$

where the second equality follows if $X$ and $Y$ are independent, as they are in our case.

We first suppose that $X$, $Y$ and $Z$ are GIG variables. Let $h_t \sim \mathrm{GIG}(\delta_t, \lambda, \gamma)$ have a generalised inverse Gaussian distribution. Let $\tau_x = [t_i, t_j]$, $\tau_y = [t_j, t_k]$ and $\tau_z = [t_i, t_k]$ and suppose $X = h_{\tau_x} \sim \mathrm{GIG}(\delta_{\tau_x}, \lambda, \gamma)$, $Y = h_{\tau_y} \sim \mathrm{GIG}\left(\delta_{\tau_y}, \lambda, \gamma\right)$ and $Z = h_{\tau_x} \sim \mathrm{GIG}(\delta_{\tau_z}, \lambda, \gamma)$. The ratio of distributions denoted above by $f_{X|Z}(x)$ is

$$\begin{aligned}
f_{X|Z}(x) &= \left(\frac{\gamma}{\delta\tau_x}\right)^{\lambda} \frac{1}{2K_\lambda(\delta\tau_x\gamma)} x^{\lambda-1} \exp\left(-\frac{1}{2}\left(\frac{(\delta\tau_x)^2}{x} + \gamma^2 x\right)\right) \\
&\times \left(\frac{\gamma}{\delta\tau_y}\right)^{\lambda} \frac{1}{2K_\lambda(\delta\tau_y\gamma)} y^{\lambda-1} \exp\left(-\frac{1}{2}\left(\frac{(\delta\tau_y)^2}{y} + \gamma^2 y\right)\right) \\
&\times \left(\left(\frac{\gamma}{\delta\tau_z}\right)^{\lambda} \frac{1}{2K_\lambda(\delta\tau_z\gamma)} z^{\lambda-1} \exp\left(-\frac{1}{2}\left(\frac{(\delta\tau_z)^2}{z} + \gamma^2 z\right)\right)\right)^{-1} \\
&= \left(\frac{\gamma}{\delta}\right)^{\lambda} \left(\frac{\tau_x\tau_y}{\tau_z}\right)^{-\lambda} \frac{K_\lambda(\delta\tau_z\gamma)}{2K_\lambda(\delta\tau_x\gamma) K_\lambda(\delta\tau_y\gamma)} \left(\frac{xy}{z}\right)^{\lambda-1} \exp\left(-\frac{1}{2}\delta^2\left(\frac{\tau_x^2}{x} + \frac{\tau_y^2}{y} - \frac{\tau_z^2}{z}\right)\right),
\end{aligned} \tag{20}$$

where $y = z - x$. Since the set of GIG distributions is not closed under convolution, so that $h_{\tau_z}$ is not distributed as $h_{\tau_x} + h_{\tau_y}$, (20) does not represent a bridge distribution. However, in two special cases we do find that $h_{\tau_z} \sim h_{\tau_x} + h_{\tau_y}$. The first is the limiting case of $\delta = 0$ with $\lambda = \frac{t}{v}$ when we obtain gamma bridge density $f^{\Gamma}_{X|Z}(x)$,

$$f^{\Gamma}_{X|Z}(x) = \frac{1}{z} \frac{\Gamma\left(\frac{\tau_x}{\nu} + \frac{\tau_y}{\nu}\right)}{\Gamma\left(\frac{\tau_x}{\nu}\right)\Gamma\left(\frac{\tau_y}{\nu}\right)} \left(\frac{x}{z}\right)^{\frac{\tau_x}{\nu}-1} \left(1 - \frac{x}{z}\right)^{\frac{\tau_y}{\nu}-1} \tag{21}$$

(see Ribeiro and Webber (02) [20]). The second is when $\lambda = -\frac{1}{2}$ for the inverse Gaussian process; we obtain the inverse Gaussian bridge density $f^{\mathrm{IG}}_{X|Z}(x)$,

$$f^{\mathrm{IG}}_{X|Z}(x) = \frac{\delta}{\sqrt{2\pi}} \frac{\tau_x \tau_y}{\tau_z} \left(\frac{xy}{z}\right)^{-\frac{3}{2}} \exp\left(-\frac{1}{2}\delta^2 \left(\frac{\tau_x^2}{x} + \frac{\tau_y^2}{y} - \frac{\tau_z^2}{z}\right)\right), \tag{22}$$

where $y = z - x$. Note that this does not depend upon $\gamma$.

### 3.3 Sampling from the IG Bridge Distribution

To sample $X$ from the distribution in equation (22) we use Tweedie's theorem and a result of Michael, Schucany and Haas (76) [17] (MSH).

A version of Tweedie's theorem in given in Seshadri (93) [23]: Suppose $X \sim \mathrm{IG}\left(\delta_{\tau_x}, \delta^{-1}\right)$, $Y \sim \mathrm{IG}\left(\delta_{\tau_y}, \delta^{-1}\right)$ and $Z \sim \mathrm{IG}\left(\delta_{\tau_y}, \delta^{-1}\right)$ are inverse Gaussian with $Z = X + Y$, then

$$Q = \delta^2 \left(\frac{\tau_x^2}{X} + \frac{\tau_y^2}{Y} - \frac{\tau_z^2}{Z}\right) \tag{23}$$

is chi-squared with one degree of freedom, $Q \sim \chi^2_1$.

In our case $X$, $Y$ and $Z$ do not have these distributions posited in the theorem, but since $\gamma$ does not enter into (22) we are still able to apply a proof given in Seshadri (93) [23] to cover this case.[7] Hence, when $X$, $Y$ and $Z$ are increments to an inverse Gaussian process, as in our case, the variable $Q$ in equation (23) is $\chi^2_1$. In fact, it appears that this is precisely the most general case to which Seshadri's proof applies.

Let $q = \delta^2 \left(\frac{\tau_x^2}{x} + \frac{\tau_y^2}{y} - \frac{\tau_z^2}{z}\right)$ be the exponent in equation (22). Set $s = \frac{y}{x}$, $\lambda = \frac{\delta^2 \tau_y^2}{z}$ and $\mu = \frac{\tau_y}{\tau_x}$. Then

$$q = \lambda \frac{(s-\mu)^2}{s\mu^2} \equiv g(s). \tag{24}$$

For any $q$ there are exactly two solutions, $s_1$ and $s_2$, to equation (24). These are:

$$s_1 = \mu + \frac{\mu^2 q}{2\lambda} - \frac{\mu}{2\lambda}\sqrt{4\mu\lambda q + \mu^2 q^2}, \tag{25}$$

$$s_2 = \frac{\mu^2}{s_1}. \tag{26}$$

Changing variable in equation (22) to $S = \frac{Y}{X} \mid Z$, the density of $S$ is

$$f_S(s) = \sqrt{\frac{\lambda}{2\pi}} \frac{1}{1+\mu} s^{-\frac{3}{2}} (1+s) \exp\left(-\frac{1}{2}\lambda \frac{(s-\mu)^2}{s\mu^2}\right). \tag{27}$$

[7] The proof of Theorm 2.1 in Seshadri relies upon finding the conditional Laplace transform of $Q$ via equation (22). Our case follows with essentially the same proof since $f_{X|Z}(x)$ does not depend on $\gamma$.

This is a well defined distribution since $\mu$ is the mean of the inverse Gaussian density $\widehat{f}(s) = \sqrt{\frac{\lambda}{2\pi}} s^{-\frac{3}{2}} \exp\left(-\frac{1}{2}\lambda \frac{(s-\mu)^2}{s\mu^2}\right)$.

We can now apply the MSH result. Suppose $Q = g(S)$ where the first derivative $g'$ of $g$ exists, is continuous and is non-zero except on a closed set with probability zero. Suppose $S$ has density $f(s)$ and for a fixed $q$ suppose $q = g(s_i)$ for $i = 1, \ldots, N$. MSH show that a sample for $S$ can be drawn by first making a draw $q$ for $Q$ and then selecting the $i$th root $s_i$ with probability $p_i(q)$ where

$$p_i(q) = \left(1 + \sum_{j=1, j\neq i}^{N} \left|\frac{g'(s_i)}{g'(s_j)}\right| . \frac{f(s_j)}{f(s_i)}\right)^{-1}. \tag{28}$$

In our case $g(s)$ is given by equation (24) and $f(s)$ is given by equation (27), and the MSH method applies.[8] We have, using equation (26)

$$\frac{g'(s_1)}{g'(s_2)} = -\left(\frac{\mu}{s_1}\right)^2, \tag{29}$$

$$\frac{f(s_2)}{f(s_1)} = \frac{s_1^2}{\mu^3}\frac{\mu^2 + s_1}{1 + s_1}. \tag{30}$$

Hence, from (28), the smaller root, $s_1$, should be chosen with probability

$$p_1(q) = \frac{\mu(1 + s_1)}{(1 + \mu)(\mu + s_1)}. \tag{31}$$

### 3.4 Stratifying the Inverse Gaussian Bridge

For $t_i < t_k$, given $h_{t_i}$ and $h_{t_k}$, the value $h_{t_j}$ of an inverse Gaussian process at an intermediate time $t_j$ is generated by

1. Generating $q \sim \chi_1^2$ and computing the roots $s_1$ and $s_2$.
2. Set $s_{t_j} = s_1$ with probability $\frac{\mu(1+s_1)}{(1+\mu)(\mu+s_1)}$, else set $s_{t_j} = s_2$.
3. Set $h_{t_j} = h_{t_i} + \frac{h_{t_k} - h_{t_i}}{1+s_{t_j}}$.

A stratified sample for $s_{t_j}$ yields a stratified sample for $h_{t_j}$. It is straightforward to stratify both the distribution of $h_{t_N}$ at the terminal time and the distribution of $h_{t_j}$ at an intermediate time, by stratifying sampling from the $\chi_1^2$ and the Bernoulli distributions. We use the MHS algorithm to sample $h_{t_N}$ and our application of their result, as described above, to sample $h_{t_j}$.

It requires two draws to sample using either the MSH algorithm or using our algorithm for the bridge distribution, the first from the $\chi_1^2$ distribution and the second a Bernoulli draw. Let $\widehat{w}^m = (\widehat{u}^m, \widehat{v}^m)$, $m = 1, \ldots, M$, be a stratified

[8] This form of $g$ is one of the cases considered by MSH, but they consider a different density for $f$.

sample of the unit square. For each $m$, $\widehat{u}^m$ is used to construct the $\chi^2_1$ variate $v$ as the square of a normal variate by inverse transform. $\widehat{v}^m$ is then used to determine which root is chosen. If $\widehat{v}^m \leq \frac{\mu(1+s_1)}{(1+\mu)(\mu+s_1)}$ choose $s_1$, else choose $s_2$.

In fact, at each stratification time three uniform variates are required; the Brownian motion $w_t$ is also stratified. If there are $K$ stratification times a stratified sample from the unit hypercube of dimension $3K$ is required. Even for moderate $K$ this is only feasible using low discrepancy sampling. The algorithm we use restricts us to sampling from a unit hypercube of dimension less than 39. In the illustrations below we take $K = 2^P$ for some integer $L$, so we are restricted to $K \leq 8$.

We present in the next section comparisons of bridge Monte Carlo when stratified at different numbers of times. To stratify at $K$ times, where $K = 2^P$ and $N = QK$ for integer $P$ and $Q$, we first compute a stratified sample of points $h^m_N$ and $w^m_N$, $m = 1, \ldots, M$, at time $t_N$. We then stratify successively at times $\frac{1}{2}t_N$, $\frac{3}{4}t_N$, $\frac{1}{4}t_N$, $\frac{7}{8}t_N$, $\frac{5}{8}t_N$, $\frac{3}{8}t_N$, $\frac{1}{8}t_N$, and so on, until all times $\frac{k}{K}t_N$, $k = 1, \ldots, K$ have been stratified. The remaining intermediate points (when $Q > 1$) are constructed using ordinary random draws from chi-squared, Bernoulli and normal distributions.

# 4 Numerical Results

We use the bridge method to value average rate options in the NIG model and compare the results to those found with plain Monte Carlo. The performance of the bridge under various degrees of stratification is investigated. We investigate reset frequencies from quarterly to approximately daily.

We compare the efficiency of different Monte Carlo schemes using an efficiency gain measure. Suppose for some option plain Monte Carlo gives a standard deviation of $\sigma_P$ in time $t_P$ and an alternative Monte Carlo method gives $\sigma_A$ and $t_A$. The efficiency gain $E_{AP}$ of the alternative method to plain Monte Carlo is

$$E_{AP} = \frac{\sigma^2_P t_P}{\sigma^2_A t_A}. \tag{32}$$

Under the assumption that standard deviation scales inversely with the square root of the number of sample paths $M$, and that time taken is proportional to $M$, then $E_{AP}$ is the multiple of the time the plain method takes to achieve a particular standard deviation compared to the alternative method.

We find that the bridge method gives considerable efficiency gains. For average rate options with one year to maturity, daily resets and 8 stratification times we achieve gains of a factor of around 160.

A number of authors[9] fit the NIG process to daily stock returns data. For our illustrative purposes we use parameter values based on Rydberg (97) [21].[10]

[9] For example, Bølviken and Benth (00) [8], Prause (97) [18], Rydberg (97) [21], et cetera.

[10] The annualised parameter values we use are $\alpha = 75.49$, $\beta = -4.089$, $\delta = 3$, and $\mu = 0$. These are based on Rydberg's estimation from daily returns for Deutsche Bank. Rydberg's estimation procedure imposes $\mu = 0$. In fact in the stock process (14) the value of $\mu$ is

These values are not atypical of annualised values found by other authors. Parameter values implied from option prices would differ from these but since the emphasis in this paper is on a numerical algorithm we are content to proceed with these. These values give pronounced skews and kurtosis at short time horizons. At horizons of a year, because they vary inversely with time (equations (7) and (8)), skewness and kurtosis are very small.

## 4.1 Algorithm Issues

We require algorithms for generating uniform, normal, inverse Gaussian and chi-squared variates. Generating variates from the bridge density was discussed earlier.

Uniform variates are generated using a VBA version of ran2 from Numerical Recipes (92) [19]. All normal variates are generated by inverse transform. $N^{-1}$, the inverse of the normal distribution function, is computed using Applied Statistics Algorithm 111 [3] downloadable from lib.stat.cmu.edu/apstat/111.

To generate inverse Gaussian variates directly we use the MHS algorithm. This requires two uniform variates for each inverse Gaussian it generates, and is easily stratified by stratifying the uniform draws.

For low discrepancy sampling we use a Sobol' sequence based on Bratley and Bennett (88) [9]. Code is downloadable from www.netlib.org/toms/659. The code generates low discrepancy samples from a unit hypercube of dimension at most 39. Since bridge Monte Carlo uses three low discrepency coordinates at each stratified time, we are constrained to have at most 13 stratification times.

## 4.2 Valuing Average Rate Options

We value average rate options, maturing in one year, with various numbers of reset times up to final maturity, comparing the results to plain Monte Carlo. We use bridge Monte Carlo with various numbers of stratification times.

The number of reset times per year varies from 4 to 256, corresponding to quarterly up to approximately daily reset frequencies. The number of time steps is equal to the number of reset times. With $N$ reset times, resets are at times $\frac{1}{N}, \frac{2}{N}, \ldots, 1$. Bridge Monte Carlo is implemented with from 1 to 16 stratification times. With $K$ stratification times, stratifications are at times $\frac{1}{K}, \frac{2}{K}, \ldots, 1$. Stratification is by low discrepancy sampling. When the number of stratification times equals the number of reset times, the method is completely low discrepancy and non-stochastic. Results in this case are based on a single replication and no standard deviation is reported. For options with 4 and 8 resets we 'benchmark' by pricing using completely low discrepancy sampling with $M = 1,000,000$ sample paths. We note that convergence in $M$ for completely low discrepancy methods is not uniform.

Table 1 shows our results.[11] The payoff at time one is $H_T = \max(A - X, 0)$

compensated away.

[11] The initial asset value is $S_0 = 100$, exercise price $X = 100$ and riskless rate $r = 0.1$. Maturity time is 1 year.

| Average rate call options: Times and standard deviations. | | | | | | |
|---|---|---|---|---|---|---|
| **K** | **4 resets** | **8 resets** | **16 resets** | **32 resets** | **64 resets** | **256 resets** |
| 0, $M = 10^6$ | 8.5856 (0.0103) [71.6] | 7.7892 (0.0094) [141.4] | 7.4059 (0.0089) [278.6] | 7.2312 (0.0087) [551.8] | 7.1286 (0.0086) [1101.4] | 7.0698 (0.0086) [4390.1] |
| 1 | 8.6326 (0.044) [1.9] | 7.8205 (0.042) [2.4] | 7.3874 (0.048) [4.9] | 7.2347 (0.041) [7.7] | 7.0763 (0.047) [15.1] | 7.0656 (0.044) [58.8] |
| 2 | 8.5530 (0.021) [1.8] | 7.7695 (0.022) [2.1] | 7.4282 (0.022) [4.8] | 7.2457 (0.026) [7.7] | 7.0721 (0.026) [15.0] | 7.0497 (0.021) [58.8] |
| 4 | 8.5695 (−) [1.8] | 7.7963 (0.010) [2.0] | 7.4181 (0.011) [4.7] | 7.2105 (0.011) [7.6] | 7.1249 (0.011) [14.9] | 7.0430 (0.012) [58.6] |
| 8 | − | 7.7959 (−) [1.7] | 7.4045 (0.0048) [4.3] | 7.2121 (0.0059) [7.3] | 7.1296 (0.0051) [14.6] | 7.0519 (0.0059) [58.4] |
| **Bench-mark** | 8.5807 (−) [169] | 7.8072 (−) [169] | − | − | − | − |

Table 1: Average Rate Call Options: Comparison of Plain and Bridge Monte Carlo

where $A$ is the arithmetic average of the asset value at each reset time and $X = 100$ is the exercise price. In the table, the $K = 0$ row reports the plain Monte Carlo results. Plain Monte Carlo uses $M = 1,000,000$ sample paths. The results for bridge Monte Carlo are for $M = 10,000$. Actual standard deviations are shown in round brackets.[12] Times in seconds for a single replication are shown in square brackets.[13]

We see that the standard deviation decreases significantly with each additional level of stratification. Doubling the number of stratification times roughly halves the standard deviation but computation times remain similar. This means that each additional level of stratification is approximately quadrupling the efficiency gain. These are shown in Table 2.[14]

Efficiency gains are most pronounced for options with the least number of reset times, but even the daily reset option with two stratification times is 12 times faster than plain Monte Carlo. For the daily reset case ($N = 256$) using 8 stratification times we achieve an efficiency gain of a factor of 157 over plain Monte Carlo. We have no reason to suppose that efficiency gains would not

[12] For plain Monte Carlo the standard deviation is approximately equal to the standard error, so only the standard error is reported. For bridge Monte Carlo the true standard deviation is found from a hundred replications of the Monte Carlo procedure.

[13] All programmes were written in Visual Basic 6.0 and were run on a 900 Mhz PC.

[14] In fact efficiency gains decrease as $M$ decreases because of fixed set-up times in the implementation of the Monte Carlo algorithm.

| Average rate call options: Efficiency gains. | | | | | | |
|---|---|---|---|---|---|---|
| **K** | **4 resets** | **8 resets** | **16 resets** | **32 resets** | **64 resets** | **256 resets** |
| 1 | 2.1 | 3.0 | 2.0 | 3.2 | 2.4 | 2.9 |
| 2 | 9.8 | 12.4 | 9.9 | 8.2 | 8.2 | 12.2 |
| 4 | – | 58.8 | 42.3 | 42.8 | 45.7 | 40.1 |
| 8 | – | – | 219.4 | 167.1 | 216.2 | 157.0 |

Table 2: Average Rate Call Options: Efficiency Gains for Bridge Monte Carlo over Plain Monte Carlo

continue to increase with the introduction of further stratification times.

# 5 Conclusions

We have shown how an inverse Gaussian bridge may be used in conjunction with stratified sampling in the NIG model to give much improved Monte Carlo estimates of average rate option values. We find efficiency gains of a factor of around between 220 and 160 for 8 stratification times.

The use of the bridge Monte Carlo technique in the NIG framework should be considered whenever (i) Monte Carlo is used to value path dependent options or (ii) a single Monte Carlo run is used to price options with different maturities.

Bridge Monte Carlo may be used to maximum effect if an efficient algorithm is available to compute the inverse of the bridge distribution function. Such an algorithm exists in the NIG case. For the variance-gamma model Ribeiro and Webber (02) [20] found smaller efficiency gains then expected, which they attributed, in part, to the inefficiency of the available algorithm to compute the inverse of the bridge distribution function in that case.

In principle the bridge Monte Carlo method is widely applicable, but its ease of application depends upon the nature of the conditional distribution function at intermediate times, and on the efficiency of available algorithms to compute the inverse of that distribution function.

For the NIG process the use of the inverse Gaussian bridge is recommended for appropriate applications.

# References


[1] C. Albanese, S. Jaimungal, and D. Rubisov. Jumping in Line. Risk Magazine, LV(5):2259–2284, 2001.

[2] C. Albanese, S. Jaimungal, and D. Rubisov. The Model of Lines for Option Pricing with Jumps. 2001. Mimeo. http://www.math-point.com/research.htm.

[3] Algorithm 111. Inverse of the Normal Distribution Function. Applied Statistics, 26:118–121, 1977.

[4] O. E. Barndorff-Nielsen. Normal Inverse Gaussian Processes and the Modelling of Stock Returns. Research Report 300, University of Aarhus, 1995.

[5] O. E. Barndorff-Nielsen. Processes of normal inverse Gaussian type. Finance and Stochastics, 2(1):41–68, 1997.

[6] D. Beaglehole, P. Dybvig, and G. Zhou. Going to extreme: correcting simulation bias in exotic option valuation. Financial Analysts Journal, January/February:62–68, 1997.

[7] B. M. Bibby and M. Sorensen. Hyperbolic Processes in Finance. 2001.

[8] E. Bolviken and F. E. Benth. Quantification of Risk in Norwegian Stocks via the Normal Inverse Gaussian Distribution. Working paper, University of Oslo, pages 1–13, 2000.

[9] P Bratley and L.F Bennett. Algorithm 659 Implementing Sobol's Quasirandom Sequence Generator. Working Paper, University of Montreal, pages 88–100, 88.

[10] P. Carr, H. Geman, D. B. Madan, and M. Yor. Stochastic Volatility for Lévy Processes. April 2001.

[11] P. Carr and D. Madan. Option valuation using the fast Fourier transform. The Journal of Computational Finance, 2, 1999.

[12] P. Carr and L. Wu. Time-Changed Lévy processes and Option Prices. Working paper, Courant Institute, pages 1–41, 2002.

[13] E. Eberlein and U. Keller. Hyperbolic Distributions in Finance. 1995. Mimeo.

[14] E. Eberlein and S. Raible. Term structure models driven by general Lévy processes. Mathematical Finance, 9:31–35, 1999.

[15] P. Jackel. Monte carlo methods in finance. Wiley Finance, 2002.

[16] D. B. Madan, P. P. Carr, and E. C. Chang. The Variance Gamma Process and Option Pricing. European Finance Review, 2:79–105, 1998.

[17] J. Michael, W. Schucany, and R. Haas. Generating random variates using transformations with multiple roots. The American Statistician, 30:88–89, 1976.

[18] K. Prause. Modelling Financial Data Using Generalised Hyperbolic Distributions. Working paper, Universitat Freiburg i. Br., 48:1–14, 1997.

[19] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. Numerical Recipes in C. Cambridge University Press, 92.

[20] C. Ribeiro and N. Webber. Valuing Path Dependent Options in the Variance-Gamma Model by Monte Carlo with a Gamma Bridge. Working paper, City University Business School, 2002.

[21] T. H. Rydberg. The Normal Inverse Gaussian Lévy Process: Simulation and Approximation. Communications in Statistics: Stochastic Models, 13(4):887–910, 1997.

[22] T. H. Rydberg. Generalized hyperbolic diffusion processes with applications in Finance. Mathematical Finance, Vol.9, No. 2:183–201, 1999.

[23] V. Seshadri. The Inverse Gaussian Distribution. Clarendon Press, 1993.