Correcting for Simulation Bias in Monte Carlo methods to Value Exotic Options in Models Driven by Lévy Processes

Claudia Ribeiro
City University
Cass Business School
London EC1Y 8TZ
United Kingdom
Tel: +44 20 7040 8769
Fax: +44 20 7040 8546
Email: C.A.G.C.Ribeiro@city.ac.uk

Nick Webber*
City University
Cass Business School
London EC1Y 8TZ
United Kingdom
Tel: +44 20 7040 5171
Fax: +44 20 7040 8881
Email: Nick.Webber@city.ac.uk

July 28, 2003

First version: January 2003

Abstract

Lévy processes can be used to model asset return's distributions. Monte Carlo methods must frequently be used to value path dependent options in these models, but Monte Carlo methods can be prone to considerable simulation bias when valuing options with continuous reset conditions.

In this paper we show how to correct for this bias for a range of options by generating a sample from the extremes distribution of the Lévy process on subintervals.

We work with the variance-gamma and normal inverse Gaussian processes. We find the method gives considerable reductions in bias, so that it becomes feasible to apply variance reduction methods. The method seems to be a very fruitful approach in a framework in which many options do not have analytical solutions.

*Corresponding author. Claudia Ribeiro gratefully acknowledges the support of Fundação para a Ciência e a Tecnologia, CEMPRE - Centro de Estudos Macroeconómicos e Previsão and Faculdade de Economia, Universidade do Porto. We are grateful for insightful comments from Grace Kuan and Gianluca Fusai, and from participants at the Computational Economics conference, Seattle, 2003.

1 Introduction

There is a growing body of research on the use of Lévy processes to model asset returns. The two classes of processes that have been investigated extensively in the derivatives literature are the variance-gamma (VG) process due to Madan and Senata (1990), Madan and Milne (1991) and the normal inverse Gaussian (NIG) process due to Eberlein and Keller (1995) and Barndorff-Nielsen (1995). There are analytical solutions available in the first model for the values of European options (Madan, Carr and Chang (1998)), but most complex options require the use of numerical methods, including Monte Carlo. Monte Carlo methods have to be used to value path dependent options such as average rate options, lookback, and barrier options. Recently, exploiting the subordinator representation of a Lévy process, Ribeiro and Webber (2002a), (2002b) have shown how considerable speed-ups can be achieved for these options when the Lévy process is the VG process or the NIG process.

Ribeiro and Webber only consider path dependent options with discrete reset dates, so that, for instance, a barrier condition is tested only at a discrete number of times up to a final maturity date. It is well known that the values of options with continuous barriers are subject to considerable simulation bias if the barrier condition is only tested discretely.

In this paper we exploit the idea of Beaglehole, Dybvig and Zhou (1997) , and El Babsiri and Noel (1998), to correct for simulation bias by sampling from the distribution of the maximum or minimum of the bridge process connecting points along a simulated Monte Carlo path. These authors applied this idea to assets following geometric Brownian motion. We extend their results to enable us to accurately value barrier options and fixed and floating strike lookback options, when a subordinator representation of the underlying Lévy process is known and can be simulated. We apply our results to the VG process and the NIG process. We find that we are able to accurately and cheaply correct for simulation bias, and because the method allows us to price these options in far fewer steps we also achieve considerable efficiency gains.

Once simulation bias has been removed we are able to apply variance reduction techniques to achieve further efficiency gains.

We note, but do not explore here, that the method can be applied to options with time-dependent barriers by using a piecewise constant approximation to the barrier.

In the next section we describe our construction. In section three we present numerical results. Section four concludes.

2 The Modelling Framework

Let S_t be the price at time t of a non-dividend paying stock. We model relative log-returns to the stock price process $S = (S_t)_{t\geq 0}$ under the pricing measure Q as a Lévy process, $L = (L_t)_{t\geq 0}$. We take the state space Ω to be the path space of L equipped with the filtration $\mathcal{F} = \{\mathcal{F}_t\}$ induced L. We assume the existence of the accumulator account numeraire and use its associated measure.[1]

Following Madan, Carr and Chang (1998) and Eberlein and Raible (1999) we set

$$S_t = S_0 \exp\left(\left(r - \varpi\right) t + L_t\right) \tag{1}$$

where L is a Lévy process, r is the constant short interest rate, and ϖ is a compensator term, defined by

$$\mathrm{e}^{\varpi} = \mathbb{E}_0 \left[\exp\left(L_1\right)\right], \tag{2}$$

to ensure that $\left(S_t \mathrm{e}^{-rt}\right)_{t \geq 0}$ is a martingale under Q.

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

$$c_t = \mathbb{E}\left[\mathrm{e}^{-r(T-t)} H_T \mid \mathcal{F}_t\right] \equiv \mathbb{E}_t\left[\mathrm{e}^{-r(T-t)} H_T\right]. \tag{3}$$

H_T may depend on the state $\omega \in \Omega$, that is, on the path of S.

We consider Lévy processes that can be represented as a subordinated Brownian motion. For such a Lévy process L there is a Brownian motion $w = (w_t)_{t \geq 0}$ and a non-decreasing Lévy process $h = (h_t)_{t \geq 0}$, independent of w, such that $L_t = w_{h(t)}$. h is called the subordinator. We suppose the Brownian motion w has drift β and variance σ, $\mathrm{d}w_t = \beta \mathrm{d}t + \sigma \mathrm{d}z_t$ for a Wiener process z_t. The subordinator representations of the VG and NIG processes are given in section 3.

We recall how (3) can be solved by Monte Carlo integration, and describe the simulation bias correction.

2.1 The Monte Carlo Method

The Monte Carlo method constructs a set $\{\widehat{\omega}^m\}_{m=1,\dots,M}$ of discrete sample paths randomly selected under a measure $\widehat{Q}$, a discrete time approximation to the measure Q. Then the approximation $\widehat{c}_t$ to c_t is

$$\widehat{c}_t = \mathrm{e}^{-r(T-t)} \frac{1}{M} \sum_{m=1}^{M} \widehat{H}_T\left(\widehat{\omega}^m\right), \tag{4}$$

where $\widehat{H}_T$ is a discrete version of H_T. Sample paths are often computed by evolving forwards from S_0, taking successive draws of increments to the process L.

For $L = \left(w_{h(t)}\right)_{t \geq 0}$, since w_t and h_t are independent, the state space Ω decomposes into a direct product $\Omega = \mathcal{B} \times \mathcal{H}$ where $\mathcal{B}$ is generated by w and $\mathcal{H}$ by h, and the filtration decomposes similarly, $\mathcal{F}_t = \mathcal{F}_t^w \times \mathcal{F}_t^h$. Then

$$\begin{aligned} c_t &= \mathrm{e}^{-r(T-t)} \mathbb{E}_t\left[H_T\right] \equiv \mathrm{e}^{-r(T-t)} \mathbb{E}\left[H_T \mid \mathcal{F}_t\right] & (5) \\ &= \mathrm{e}^{-r(T-t)} \mathbb{E}\left[\mathbb{E}\left[H_T \mid \mathcal{F}_t^w \times \mathcal{F}_T^h\right] \mid \mathcal{F}_t\right]. & (6) \end{aligned}$$

[1] In this paper we are not concerned with change of measure to Q.

To find c_t one can first compute payoffs conditional upon realisations of h, and then take expectations over paths of h.

Discrete sample paths for a subordinated Brownian motion, $L = \left(w_{h(t)}\right)_{t\geq 0}$, can be constructed by first constructing discrete sample paths for the subordinator h and then sampling the process w at times determined by the paths found for h.[2] Ribeiro and Webber (2002a), (2002b) showed how a bridge simulation of h could lead to considerable variance reduction in the Monte Carlo method. They priced average rate, lookback and barrier options with discrete resets, finding large speed-ups.

We describe a method for pricing exotic options, such as barrier options and lookbacks, with continuous resets. The method is based upon simulating from an approximation to the extremes distribution of the Lévy bridge process.

Let time be discretised as $0 = t_0 < \ldots < t_N = T$, with the time step $\Delta t = t_j - t_{j-1}$, $j = 1, \ldots, N$, constant.[3] We first generate sample paths for the subordinator process h. Suppose h_t has the distribution function F_t^h. We obtain a discrete set $\widehat{h} = \left\{\widehat{h}_j\right\}_{j=0,\ldots,N}$ where $\widehat{h}_0 = 0$ and each $\Delta\widehat{h}_j = \widehat{h}_{j+1} - \widehat{h}_j \sim F_{\Delta t}^h$ is an increment over the period $[t_j, t_{j+1}]$.

Now set $\widehat{w}_{\widehat{h}(0)} = 0$, and iteratively generate increments $\Delta\widehat{w}_{\widehat{h}(j)} = \widehat{w}_{\widehat{h}(j+1)} - \widehat{w}_{\widehat{h}(j)} \sim \mathrm{N}\left(\beta\Delta\widehat{h}_j, \sigma^2\Delta\widehat{h}_j\right)$. Set $\widehat{L}_j = \widehat{w}_{\widehat{h}(j)}$. The path $\widehat{\omega} = \left\{\widehat{L}_j\right\}_{j=0,\ldots,N}$ is a discrete approximation to a sample path ω of L, and the set $\widehat{S} = \left\{\widehat{S}_j\right\}_{j=0,\ldots,N}$, with $\widehat{S}_j = S_0 \exp\left((r - \varpi)\, t_j + \widehat{L}_j\right)$, is a discrete sample path for S.

Now consider the problem of pricing a continuously reset barrier option by Monte Carlo, for instance an up-and-out barrier option. Let the barrier level be B. The barrier condition can be tested at each time step t_j, $j = 0, \ldots, N$. Generating M sample paths, $\{\widehat{\omega}^m\}_{m=1,\ldots,M}$, the option value at time 0 would be found to be

$$\widehat{c}_0 = \mathrm{e}^{-rT}\frac{1}{M}\sum_{m=1}^{M}\widehat{H}_T\left(\widehat{\omega}^m\right). \tag{7}$$

where for an up-and-out call with strike X

$$\widehat{H}_T\left(\widehat{\omega}^m\right) = \begin{cases} (S_T - X)_+, & B > \max_j\left\{\widehat{S}_j\right\}, \\ 0, & \text{otherwise.} \end{cases} \tag{8}$$

The estimate $\widehat{c}_0$ is a biased estimate of the barrier option value, since the sample path may have exceeded the barrier level, knocking out the option, in between times at which it has been observed. This is the source of simulation bias in the Monte Carlo method. It decreases only slowly as N increases.

When $L = z$ is a Wiener process Beaglehole, Dybvig and Zhou (1997) and El Babsiri and Noel (1998) have shown that simulation bias can be removed by

[2] This procedure was described by Rydberg (97).

[3] The assumption is for expositional simplicity only and can be relaxed with purely notational inconvenience.

performing at each time step along each sample path an additional simulation for the maximum of the Brownian bridge distribution in the interval $[t_j, t_{j+1}]$ between points $\widehat{L}_j$ and $\widehat{L}_{j+1}$. We show how this idea can be exploited to value barrier and other path dependent options when L is a Lévy process whose subordinator representation is known. This is the case for the VG and the NIG processes considered below.

For a process $Y = (Y_t)_{t \geq 0}$ set

$$M_{t,T}^Y \equiv M_{t,T}^Y(\omega) = \max_{s \in [t,T]} \{Y_s(\omega)\}, \tag{9}$$

$$m_{t,T}^Y \equiv m_{t,T}^Y(\omega) = \min_{s \in [t,T]} \{Y_s(\omega)\}. \tag{10}$$

Suppose we have evolved a sample path $\left\{\widehat{S}_j\right\}_{j=0,\ldots,N}$ and have generated values $\left\{\widehat{M}_{t_j,t_{j+1}}\right\}_{j=0,\ldots,N-1}$ of $\left\{M_{t_j,t_{j+1}}^S\right\}_{j=0,\ldots,N-1}$, conditional on $\left\{\widehat{S}_j\right\}_{j=0,\ldots,N}$, by some means. Write

$$\begin{array}{ll} M\left(\widehat{S}\right) = \max_{j=0,\ldots N}\left\{\widehat{S}_j\right\}, & M\left(\widehat{M}\right) = \max_{j=0,\ldots N-1}\left\{\widehat{M}_{t_j,t_{j+1}}\right\}, \\ m\left(\widehat{S}\right) = \min_{j=0,\ldots N}\left\{\widehat{S}_j\right\}, & m\left(\widehat{m}\right) = \min_{j=0,\ldots N-1}\left\{\widehat{m}_{t_j,t_{j+1}}\right\}. \end{array} \tag{11}$$

Table 1 contrasts the plain valuation method with the corrected valuation method for several exotic call options. X is an exercise price. The table gives the values of $\widehat{H}_T(\widehat{\omega})$ in each case. Over each subinterval we generate a value from the conditional maximum (minimum) distribution and use this either to test an exercise condition or to determine an exercise price.[4]

When exercise conditions do not change over the life of the option we may be able to use the corrected Monte Carlo method with a single time step up to maturity (a 'longstep' Monte Carlo). Ordinary Monte Carlo usually requires many time steps to begin to accurately approximate the option value. With a single time step the corrected Monte Carlo method requires three draws for each sample path, one for h_T, a second for $w_{h(T)}$, and the third for $M_{0,T}^S$.[5]

2.2 Using the Extremes Distribution

We need to be able to generate samples from the extremes distributions $M_{t,T}^S$ and $m_{t,T}^S$ of S.

Write

$$\Pr\left[\tau^Y \leq t \mid Y_0, Y_T; B_t\right] \tag{12}$$

for the probability that the process $Y = (Y_t)_{t \geq 0}$ hits a time varying barrier level B_t before time t, $0 < t < T$, conditional upon an initial value Y_0 and a final value Y_T of Y.

[4] To value a swing option, with payoff $(\max_t S_t - \min_t S_t - X)_+$, requires a knowledge of the joint distribution of the max and min of S_t.

[5] Each draw may require more than one uniform variate. For a small number of variates the draws may be fully stratified, leading to much greater speed-ups.

Set $R_t = \ln\left(\frac{S_t}{S_0}\right)$. Then, setting $\underline{B} = \ln\left(\frac{B}{S_0}\right)$, we have

$$\begin{aligned} \Pr\left[M_{0,T}^{S} \geq B \mid S_0, S_T\right] &= \Pr\left[\tau^{S} \leq T \mid S_0, S_T; B\right] && (13)\\ &= \Pr\left[\tau^{R} \leq T \mid R_0, R_T; \underline{B}\right] && (14)\\ &= \Pr\left[\tau^{R-(r-\omega)t} \leq T \mid R_0, R_T - (r-\omega)\,T; \underline{B} - (r-\omega)\,t\right] && (15)\\ &= \Pr\left[\tau^{L} \leq T \mid L_0, L_T; \underline{B} - (r-\omega)\,t\right] && (16) \end{aligned}$$

This is the hitting time of the bridge distribution of a Lévy process L to a linear non-constant barrier. For $L = \left(w_{h(t)}\right)_{t\geq 0}$ it can be computed exactly in several special cases, and in the general case there is an approximation available.

Special case: $h_t = t$ so L is a Brownian motion In this case the $(r-\omega)\,t$ term can be absorbed into the drift of w so that R is a Brownian motion and S is a geometric Brownian motion. Then

$$\begin{aligned} \Pr\left[M_{0,T}^{S} \leq B \mid S_0, S_T\right] &= \Pr\left[M_{0,T}^{R} \leq \underline{B} \mid R_0, R_T\right] && (17)\\ &= 1 - \exp\left(-2\frac{(\underline{B}-R_0)(\underline{B}-R_T)}{\sigma^2 T}\right), && (18) \end{aligned}$$

(see, for example, Karatzas and Shreve (1991) and El Babsiri and Noel (1998)).

To sample from the distribution (18) one may use inverse transform. Let $U \sim \mathrm{U}[0,1]$ be a draw of a standard uniform variate, then

$$\widehat{M} = \frac{R_0 + R_T + \sqrt{(R_0 - R_T)^2 - 2\sigma^2 T \ln(1-U)}}{2} \tag{19}$$

is a draw from the distribution of $M_{0,T}^{R}$. $S_0 \exp\left(\widehat{M}\right)$ is a draw from $M_{0,T}^{S}$.

The distribution of the conditional minimum is

$$\Pr\left[m_{0,T}^{R} \leq \underline{B} \mid R_0, R_T\right] = \exp\left(-2\frac{(\underline{B}-R_0)(\underline{B}-R_T)}{\sigma^2 T}\right) \tag{20}$$

and for $U \sim \mathrm{U}[0,1]$ the variate

$$\widehat{m} = \frac{R_0 + R_T - \sqrt{(R_0 - R_T)^2 - 2\sigma^2 T \ln(U)}}{2} \tag{21}$$

is a draw from $m_{0,T}^{R}$ and $S_0 \exp(\widehat{m})$ is a draw from $m_{0,T}^{S}$.

Special case: $r = \omega$ so $(r-\omega)\,t = 0$ First, suppose that h is deterministic. The previous results carry over directly and we have

$$\begin{aligned} \Pr\left[M_{0,T}^{R} \leq \underline{B} \mid R_0, R_T\right] &= \Pr\left[M_{0,T}^{L} \leq \underline{B} \mid L_0, L_T\right] && (22)\\ &= \Pr\left[M_{h(0),h(T)}^{w} \leq \underline{B} \mid w_{h(0)}, w_{h(T)}\right] && (23)\\ &= 1 - \exp\left(-2\frac{\left(\underline{B}-w_{h(0)}\right)\left(\underline{B}-w_{h(T)}\right)}{\sigma^2\left(h(T)-h(0)\right)}\right). && (24) \end{aligned}$$

To get a sample $\widehat{M}$ of $M^R_{0,T}$ one draws from the distribution (24), getting

$$\widehat{M} = \frac{w_{h(0)} + w_{h(T)} + \sqrt{\left(w_{h(0)} - w_{h(T)}\right)^2 - 2\sigma^2 \left(h\left(T\right) - h\left(0\right)\right) \ln\left(1 - U\right)}}{2} \tag{25}$$

with an analogous result for $m^S_{0,T}$.

When h is a stochastic subordinator one first generates a sample path for h. Conditional upon this sample path, $L = \left(w_{h(t)}\right)_{t \geq 0}$ is effectively a Brownian motion subordinated by a deterministic time-change, and one draws from (25) as above. Note that an entire path for h may not be required; a single draw for $h\left(T\right)$ may suffice.

General case: $r \neq \omega$ Finally, suppose $S_t = S_0 \exp\left(\left(r - \varpi\right) t + w_{h(t)}\right)$ for a stochastic time change h, where $r \neq \omega$. Then $R_t = \ln\left(S_t / S_0\right) = \left(r - \varpi\right) t + w_{h(t)}$ is no longer a Brownian motion time changed by h and we must make an approximation.

Suppose that T is small. Over the small interval $[0, T]$ one may suppose that hitting probability distribution is dominated not by non-linearity in the barrier level, but by the volatility of L. Then

$$\begin{aligned} \Pr\left[M^S_{0,T} \geq B \mid S_0, S_T\right] &= \Pr\left[\tau^L \leq T \mid L_0, L_T; \underline{B} - \left(r - \omega\right) t\right] && (26) \\ &\sim \Pr\left[\tau^L \leq T \mid L_0, L_T + \left(r - \varpi\right) T; \underline{B}\right] && (27) \\ &= \Pr\left[M^L_{0,T} \geq \underline{B} \mid L_0, L_T + \left(r - \varpi\right) T\right] && (28) \end{aligned}$$

and we approximate (26) using (28). In obtaining (28) we are assuming that in the distribution of $M^S_{0,T}$ the total volatility of L over the interval $[0, T]$ dominates non-linearities in h, which is reasonable for small T.

The general Monte Carlo method is thus to discretise the interval $[0, T]$ so that over each sub-interval $[t_j, t_{j+1}]$ the approximation (28) is valid.

2.3 Summary of the method

The method is thus

1. Generate a sample path $\widehat{h}_t$ for h, then generate a path $\widehat{L}_t = \widehat{w}_{\widehat{h}(t)}$ for L and $\widehat{S}_t = S_0 \exp\left(\left(r - \varpi\right) t + \widehat{L}_t\right)$ for S, ensuring that Δt is sufficiently small for (28) to be an acceptable approximation over each subinterval.

2. Generate approximate draws from the conditional maximum or minimum, using (28) or its analogue.

3. Compute the corrected payoff function, as in table 1.

4. Repeat steps 1 to 3 sufficient times and return the discounted average payoff as the option value.

The bias correction can be combined with the usual Monte Carlo speed-ups, such as stratifying the various draws.

For knock-in and knock-out options a variation of the method just described is possible. Instead of sampling random draws from (25) along each sample path, one computes the probability (28) that over each sub-interval the barrier is exceeded, and hence one obtains the probability that over a sample path for L the barrier level was exceeded. The payoff along the sample path can then be weighted by this probability. We see below that this refinement can considerably reduce the variance of the estimate since it requires fewer random draws.

3 Numerical Examples

We value a number of exotic options when the price process of the underlying asset is a geometric Lévy motion. We investigate simulation bias for an up-and-out and an up-and-in barrier option, and a fixed strike and a floating strike lookback option, for the VG and the NIG processes. We show that further variance reductions are possible once the correction has removed the simulation bias.

We suppose the initial stock price is $S_0 = 100$ and the riskless rate is $r = 0.1$. The final maturity time is $T = 0.25$ for all options and we assume there are 90 days in this three month period. For the fixed-strike lookback and the barrier options we take the exercise price to be $X = 105$. For the barrier options we take a barrier level of $B = 120$. Since there is bias in the usual standard error estimator, the true standard deviation of the Monte Carlo estimate is found from a hundred replications of the Monte Carlo procedure.

We find that the method reduces to relative insignificance the simulation bias for all four options, and that very significant speed-ups over a plain method are therefore possible. We also find it possible to reduce standard deviations using stratification methods.

We start by giving the subordinator representations of the VG and NIG processes and then briefly discuss algorithm issues, before presenting out numerical results.

3.1 Subordinator representations of the VG and NIG Processes

A VG process with parameters σ, ν and β can be represented (see, for example, Madan, Carr and Chang (1998)) as a Brownian motion, with drift β and volatility σ, subordinated to a Gamma process $h(t) \sim \Gamma(t/\nu, \nu)$. The density of $h(t)$ is

$$f_t^{\Gamma}(h) = \frac{h^{\frac{t}{\nu}-1}}{\nu^{\frac{t}{\nu}}\Gamma\left(\frac{t}{\nu}\right)} \exp\left(-\frac{h}{\nu}\right). \tag{29}$$

A VG process is then specified by $\mathrm{d}L_t = \beta \mathrm{d}h(t) + \sigma \mathrm{d}z_{h(t)}$. The compensator is

$$\varpi = -\frac{1}{\nu} \ln \left(1 - \mu\nu - \frac{1}{2}\sigma^2 \nu \right). \tag{30}$$

Similarly, a NIG process with parameters α, β, δ and μ can be represented (see, for example, Barndorff-Nielsen (1995)) as a Brownian motion, with drift β and volatility 1, subordinated to an Inverse Gaussian process, $h(t) \sim IG(\delta t, \gamma)$, whose density is

$$f_t^{IG}(h) = \frac{\delta t}{\sqrt{2\pi}} h^{-\frac{3}{2}} \exp\left(-\frac{\gamma^2}{2h} \left(h - \frac{\delta t}{\gamma} \right)^2 \right), \tag{31}$$

where $\gamma = \sqrt{\alpha^2 - \beta^2}$. A NIG process is then specified by $\mathrm{d}L_t = \mu \mathrm{d}t + \beta \mathrm{d}h(t) + \mathrm{d}z_{h(t)}$. The compensator for an NIG process is

$$\varpi = \mu + \delta \left(\gamma - \sqrt{\alpha^2 - (\beta + 1)^2} \right). \tag{32}$$

Note in the process S the contribution of the parameter μ in the drift of L is cancelled out by its appearance in the compensator ϖ.

In our numerical examples, parameters for the VG case are $\beta = -0.1436$, $\sigma = 0.12136$, $\nu = 0.3$ (based upon Madan, Carr and Chang (1998)) and for the NIG case are $\alpha = 75.49$, $\beta = -4.089$, $\delta = 3.0$, $\mu = 0$ (based upon Rydberg (1997)).

3.2 Algorithm Issues

We require algorithms for generating uniform, normal, gamma, beta random variates, inverse Gaussian and inverse Gaussian bridge variates.

Uniform variates are generated using a VBA version of ran2 from Numerical Recipes (1992). All normal variates were generated by inverse transform. N^{-1}, the inverse of the normal distribution function, is computed using Applied Statistics Algorithm 111 (1997) downloadable from lib.stat.cmu.edu/apstat/111.

To generate gamma variates directly we use the Best (1983) and Best (1978) algorithms as described in Devroye (1986). To compute the inverse of the gamma distribution function to use with the inverse transform method for stratified sampling, we use the algorithm of DiDonato and Morris (1987), downloadable from www.netlib.org/toms/654.

Beta variates with parameters α and β are generated directly by Cheng's method if $\min(\alpha, \beta) < 1$, Johnk's method if $\max(\alpha, \beta) < 1$, by Atkinson and Whittaker's method if $\min(\alpha, \beta) < 1 < \max(\alpha, \beta)$, and by ratio of gammas otherwise.[6] For stratified sampling, the inverse of the beta distribution function, B^{-1}, is computed using an algorithm due to Moshier (2000).

[6] See Dagpunar (1988) and Devroye (1986). Johnk's method fails at our level of machine precision in a small percentage of draws when both α and β are small. When this happens we resample. Experimentation leads us to believe that any bias introduced in our results by resampling is very small.

To generate inverse Gaussian variates directly we use the Michael, Schucany and Haas (1976) algorithm. Variates from the inverse Gaussian bridge distribution were computed by the algorithm of Ribeiro and Webber (2002b).

For low discrepancy sampling we use a Sobol' sequence based on Bratley and Bennett (1988). Code is downloadable from www.netlib.org/toms/659.

3.3 Simulation bias

Tables (2), (3), (4) and (5) show the effectiveness of the method in removing simulation bias in the VG and NIG models for our four example options. The discretisation intervals are three, two and one days, giving $N = 30$, 45, and 90, and also $N = 120$. These values are obtained using $M = 10,000$ sample paths. The tables show the simulated option values, their standard errors in round brackets, and computation times in square brackets.[7] All methods use bridge Monte Carlo as described by Ribeiro and Webber (2002a) , (2002b), with low discrepancy stratified sampling at 8 intermediate times including the final time. 'Corrected' uses either random draws from the maximum or minimum (Random) or the cumulative exceedence probability (Exceed), for barrier options. For the Random method draws from the extremes distribution are unstratified. For lookbacks only the Random correction can be used.

Figures (1), (2), (3) and (4) show graphically more detailed convergence for the four options in the VG case, plotting option values against computation times. Figures (5), (6), (7) and (8) show convergence in the NIG case. $M = 10,000$ sample paths were used here also.

We see that the price estimates of the Plain method are converging very slowly towards the prices provided by the corrected methods, and for these values of N the Plain price estimates are nowhere near their limiting values. The prices given by the Random and the Exceedence corrected methods differ by very little, each converging very rapidly in N. The bias of the Plain method would be reduced to insignificant levels only for very large N at the cost of very considerably higher computational times.

For barrier options, the Exeedence corrected method gives lower standard deviations than the Random method. It is also slightly faster, because the computation of the exceedence probability does not involve the additional draw of a uniform variate required by the Random method.

From the figures the corrected method applied with the NIG process converges more rapidly in all cases than when it is applied with the VG process. This is because with our parameter values the size of the biasing term $(r - \varpi)\,\Delta t$ is significantly smaller for the NIG process than the VG process.

3.4 Reduction in standard error

It only makes sense to apply variance reduction methods once simulation bias has been removed. We compare values and speed-ups achieved using a corrected

[7] The computations were performed in VBA on a 900 Mhz PC. All computation times are for a single replication of the Monte Carlo prodcedure.

method with $N = 20$ and $N = 50$ time steps, compared to the Plain method using $N = 100$ time steps. In each case we compare an unstratified version ('Unstrat.') with a stratified version. The Exceed correction method is used for the barrier options. Stratification is by low discrepancy sampling at 8 intermediate times, including the final time. For the Random method, at each stratification time both the sample path for S and the draw from the extremes distribution are stratified. In the tables below the number of sample paths is $M = 50,000$.

Tables (6) and (8) presents results for the four options. Each entry gives the simulated option value, its standard deviation in round brackets and the computation time in square brackets. Tables (7) and (9) give the efficiency gains[8] achieved over the plain-plain method. Efficiency gains reflect only the improvement is standard deviation. These gains are relatively small in comparison with the very considerable gains achieved through the removal of simulation bias.

The efficiency gains of the unstratified corrected methods reflect the reduced number of sample paths they use compared to the unstratified corrected method. Introducing stratification gives further gains ranging up to a factor of several hundred.

4 Conclusions

We have shown how to correct for simulation bias when using Monte Carlo methods to value options with continuous reset conditions when asset returns are modelled as Levy processes - specifically the VG and NIG processes. The correction works by either sampling from or computing values from an approxiation to the conditional distribution of the minimum or maximum, as relevant, of the bridge process of the underlying process. The correction can be applied to any Lévy process whose subordinator representation is known and whose subordinator bridge distribution can be sampled.

We apply the correction to barrier options and fixed and floating strike lookback options for VG and NIG processes. We found that the corrected method cheaply eliminates simulation bias, enabling us to value continuously reset barrier options in a relatively small number of time steps, and to use further variance reduction methods.

[8] Computed as the ratio of the square of the standard deviations times the computation time for each method. It represents the fraction of the time taken by the superior method to acheive the same standard error as the inferior method.

References

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

[2] M. El Babsiri and G. Noel. Simulating path dependent options: A new approach. *Journal of Derivatives*, Winter:65–83, 1998.

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

[4] D. R. Beaglehole, P. H. Dybvig, and G. Zhou. Going to extremes: Correcting simulation bias in exotic option valuation. *Financial Analysts Journal*, January/February:62–68, 1997.

[5] P Bratley and L.F Bennett. Algorithm 659 Implementing Sobol's Quasirandom Sequence Generator. *Working Paper, University of Montreal*, pages 88–100, 1988.

[6] J. Dagpunar. Principles of Random Variate Generation. *Clarendon Press*, 1988.

[7] L. Devroye. Non-Uniform Variate Generation. *Springer-Verlang, New York*, 1986.

[8] A.R. DiDonato and A.H. Morris. Incomplete gamma function ratios and their inverse. *ACM TOMS*, 13:318–319, 1987.

[9] E. Eberlein and U. Keller. Hyperbolic distributions in finance. *Bernoulli*, 1:281–299, 1995.

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

[11] I. Karatzas and S. Shreve. *Brownian Motion and Stochastic Calculus*. Springer Verlag, 2nd edition, 1997.

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

[13] D. B. Madan and F. Milne. Option Pricing with VG Martingale Components. *Mathematical Finance*, pages 39–55, 1(1991).

[14] D. B. Madan and E. Seneta. The Variance Gamma (V.G.) Model for Share Market Returns. *Journal of Business*, pages 511–524, 1990.

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

[16] S.L. Moshier. Inverse of the incomplete beta integral. *Cephes Math Library*, 2000.

[17] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. Numerical Recipes in C. *Cambridge University Press*, 1992.

[18] 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*, 2002a.

[19] C. Ribeiro and N. Webber. A Monte Carlo Method for the Normal Inverse Gaussian Option Valuation Model using an Inverse Gaussian Bridge. *Working paper, City University Business School*, 2002b.

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

Option		**Plain**	**Corrected**
Up-and-Out	$\widehat{H}_T =$	$\begin{cases} \left(\widehat{S}_T - X\right)_+, & M\left(\widehat{S}\right) < B, \\ 0, & \text{otherwise,} \end{cases}$	$\begin{cases} \left(\widehat{S}_T - X\right)_+, & M\left(\widehat{M}\right) < B, \\ 0, & \text{otherwise,} \end{cases}$
Up-and-In	$\widehat{H}_T =$	$\begin{cases} \left(\widehat{S}_T - X\right)_+, & M\left(\widehat{S}\right) \geq B, \\ 0, & \text{otherwise,} \end{cases}$	$\begin{cases} \left(\widehat{S}_T - X\right)_+, & M\left(\widehat{M}\right) \geq B, \\ 0, & \text{otherwise,} \end{cases}$
LB, Floating	$\widehat{H}_T =$	$\widehat{S}_N - m\left(\widehat{S}\right),$	$\widehat{S}_N - m\left(\widehat{m}\right),$
LB, Fixed	$\widehat{H}_T =$	$\left(M\left(\widehat{S}\right) - X\right)_+,$	$\left(M\left(\widehat{M}\right) - X\right)_+.$

Table 1: Computing the payoff: Comparison of Methods

Simulation bias, Up-and-out option						
N	VG process			NIG process		
	Plain	Corrected		Plain	Corrected	
		Random	Exceed		Random	Exceed
30	1.1977 (0.0017) [39.5]	1.1611 (0.0040) [41.4]	1.1621 (0.0013) [39.9]	1.6496 (0.0095) [7.3]	1.4465 (0.0100) [8.9]	1.4589 (0.0054) [7.9]
45	1.1941 (0.0015) [42.2]	1.1640 (0.0050) [44.8]	1.1637 (0.0012) [42.8]	1.6154 (0.0088) [10.8]	1.4588 (0.0103) [13.2]	1.4530 (0.0075) [11.5]
90	1.1944 (0.0017) [53.7]	1.1594 (0.0044) [54.0]	1.1611 (0.0012) [49.9]	1.5858 (0.0095) [21.3]	1.4705 (0.0094) [26.2]	1.4658 (0.0083) [22.6]
120	1.1922 (0.0016) [67.7]	1.1544 (0.0042) [59.3]	1.1603 (0.0013) [53.7]	1.5617 (0.0105) [28.3]	1.4694 (0.0118) [35.0]	1.4712 (0.0079) [30.0]

Table 2: Correcting simulation bias for up-and-out call options

Simulation bias, Up-and-in option						
N	VG process			NIG process		
	Plain	Corrected		Plain	Corrected	
		Random	Exceed		Random	Exceed
30	0.0693 (0.0013) [39.5]	0.1044 (0.0044) [41.0]	0.1019 (0.0012) [39.9]	1.2749 (0.0075) [7.0]	1.4586 (0.0096) [8.6]	1.4737 (0.0071) [7.5]
45	0.0701 (0.0015) [42.3]	0.0990 (0.0038) [44.6]	0.1016 (0.0012) [42.8]	1.3064 (0.0099) [10.6]	1.4596 (0.0115) [13.0]	1.4671 (0.0068) [11.3]
90	0.0717 (0.0016) [48.9]	0.1069 (0.0042) [53.9]	0.1029 (0.0012) [49.9]	1.3373 (0.0109) [21.2]	1.4552 (0.0097) [26.2]	1.4715 (0.0088) [22.5]
120	0.0742 (0.0018) [52.4]	0.1008 (0.0041) [59.1]	0.1018 (0.0012) [53.7]	1.3432 (0.0101) [28.2]	1.4612 (0.0095) [34.8]	1.4786 (0.0076) [30.0]

Table 3: Correcting simulation bias for up-and-in call options

Simulation bias, Fixed Strike Lookback				
N	VG process		NIG process	
	Plain	Corrected	Plain	Corrected
30	1.4934 (0.0024) [39.7]	2.1199 (0.0108) [41.5]	4.5211 (0.0073) [7.2]	5.2938 (0.0077) [9.1]
45	1.4987 (0.0025) [42.5]	2.1098 (0.0098) [45.2]	4.6341 (0.0070) [10.9]	5.2887 (0.0094) [13.8]
90	1.5052 (0.0026) [49.5]	2.1089 (0.0102) [55.1]	4.7711 (0.0074) [21.8]	5.2917 (0.0092) [28.9]
120	1.5120 (0.0025) [53.2]	2.1252 (0.0097) [60.8]	4.8096 (0.0088) [29.1]	5.3128 (0.0093) [37.1]

Table 4: Correcting simulation bias for fixed-strike lookback call options

Simulation bias, Floating Strike Lookback				
N	VG process		NIG process	
	Plain	Corrected	Plain	Corrected
30	5.1518 (0.0028) [39.7]	6.1327 (0.0092) [41.4]	7.9538 (0.0077) [7.2]	8.9360 (0.0080) [9.1]
45	5.1685 (0.0032) [42.5]	6.1334 (0.0092) [45.2]	8.0860 (0.0081) [10.9]	8.9211 (0.0093) [13.8]
90	5.1842 (0.0032) [49.4]	6.1345 (0.0103) [55.1]	8.2781 (0.0081) [21.7]	8.9372 (0.0088) [27.7]
120	5.1947 (0.0035) [53.1]	6.1391 (0.0095) [60.7]	8.3134 (0.0081) [28.9]	8.9382 (0.0096) [37.0]

Table 5: Correcting simulation bias for floating-strike lookback call options

NIG process: Comparison of error, standard deviations						
	Plain		Corrected			
	$N = 100$		$N = 20$		$N = 50$	
Option	Unstrat.	Strat.	Unstrat.	Strat.	Unstrat.	Strat.
Up-and-out	1.452 (0.013) [127.8]	1.580 (0.004) [118.9]	1.475 (0.013) [28.6]	1.460 (0.003) [25.1]	1.448 (0.013) [65.5]	1.461 (0.003) [62.7]
Up-and-in	1.347 (0.022) [120.9]	1.342 (0.004) [118.4]	1.458 (0.022) [28.7]	1.471 (0.003) [26.3]	1.480 (0.022) [65.7]	1.465 (0.004) [63.4]
LB, fixed	4.782 (0.027) [124.5]	4.784 (0.004) [121.7]	5.288 (0.028) [34.7]	5.306 (0.003) [30.7]	5.329 (0.028) [80.9]	5.299 (0.004) [79.1]
LB, floating	8.265 (0.031) [124.2]	8.289 (0.004) [121.5]	8.924 (0.031) [34.5]	8.926 (0.004) [30.6]	8.965 (0.031) [80.6]	8.925 (0.004) [77.1]

Table 6: NIG case: Conparison of error, stardard deviations

NIG process: Efficiency gains (s.d. only)					
	Plain	Corrected			
	$N = 100$	$N = 20$		$N = 50$	
Option	Strat.	Unstrat.	Strat.	Unstrat.	Strat.
Up-and-out	9.3	4.6	130	2.0	40
Up-and-in	25	4.3	345	1.8	71
LB, fixed	50	3.3	247	1.4	87
LB, floating	66	3.5	284	1.5	89

Table 7: Efficiency gains in the NIG case

VG process: Comparison of error, standard deviations						
	Plain		Corrected			
	$N = 100$		$N = 20$		$N = 50$	
Option	Unstrat.	Strat.	Unstrat.	Strat.	Unstrat.	Strat.
Up-and-out	1.188 (0.010) [96.6]	1.189 (0.001) [257.5]	1.160 (0.009) [30.6]	1.1591 (0.0004) [185.0]	1.160 (0.009) [59.2]	1.1590 (0.0005) [214.9]
Up-and-in	0.081 (0.0054) [96.3]	0.0764 (0.0006) [248.2]	0.103 (0.0055) [32.3]	0.1071 (0.0004) [184.7]	0.106 (0.0057) [62.6]	0.1068 (0.0005) [215.5]
LB, fixed	1.504 (0.012) [100.2]	1.509 (0.001) [254.2]	2.116 (0.013) [38.3]	2.111 (0.004) [189.3]	2.133 (0.014) [78.2]	2.119 (0.004) [236.9]
LB, floating	5.219 (0.016) [99.5]	5.195 (0.001) [259.5]	6.126 (0.016) [38.1]	6.124 (0.005) [190.9]	6.124 (0.016) [77.8]	6.135 (0.005) [230.9]

Table 8: VG case: Conparison of error, stardard deviations

VG process: Efficiency gains (s.d. only)					
	Plain	Corrected			
	$N = 100$	$N = 20$		$N = 50$	
Option	Strat.	Unstrat.	Strat.	Unstrat.	Strat.
Up-and-out	95	3.5	351	1.8	158
Up-and-in	30	2.9	97	1.4	58
LB, fixed	48	2.0	3.7	0.9	3.7
LB, floating	43	2.6	6.0	1.3	4.1

Table 9: Efficiency gains in the VG case

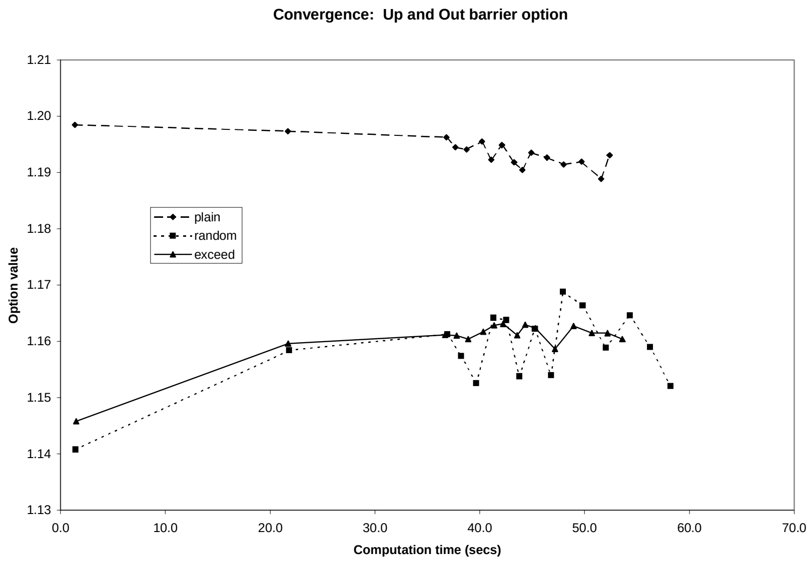

Figure 1: VG case, convergence: Up and Out barrier option.

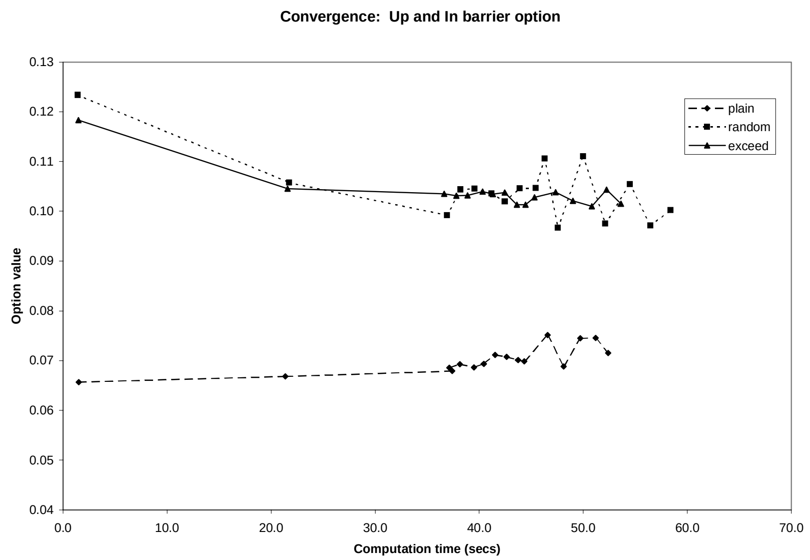

Figure 2: VG case, convergence: Up and In barrier option.

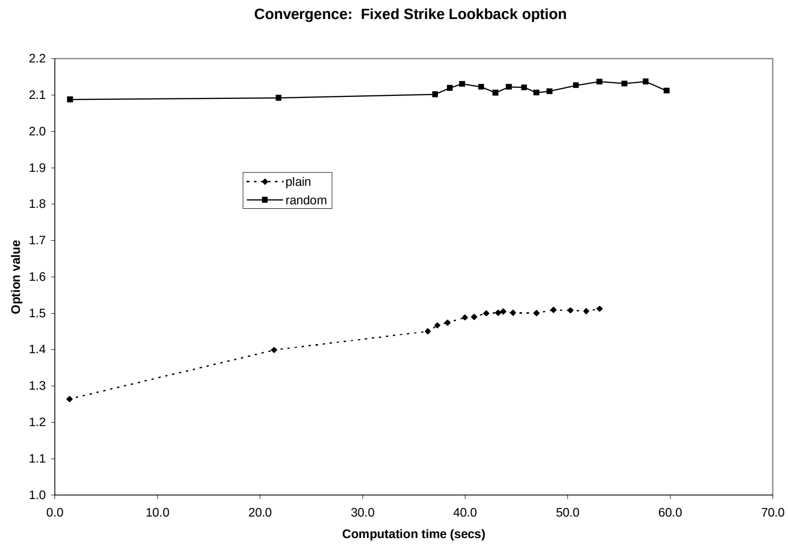

Figure 3: VG case, convergence: Fixed strike lookback

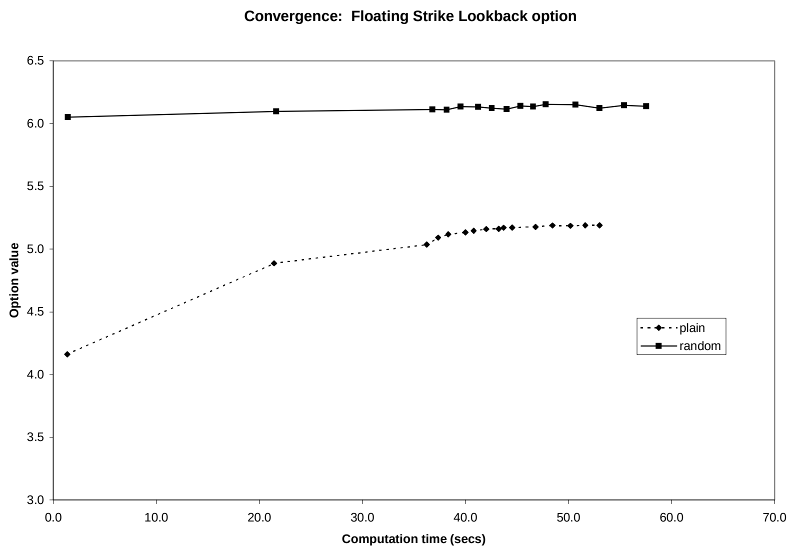

Figure 4: VG case, convergence: Floating strike lookback

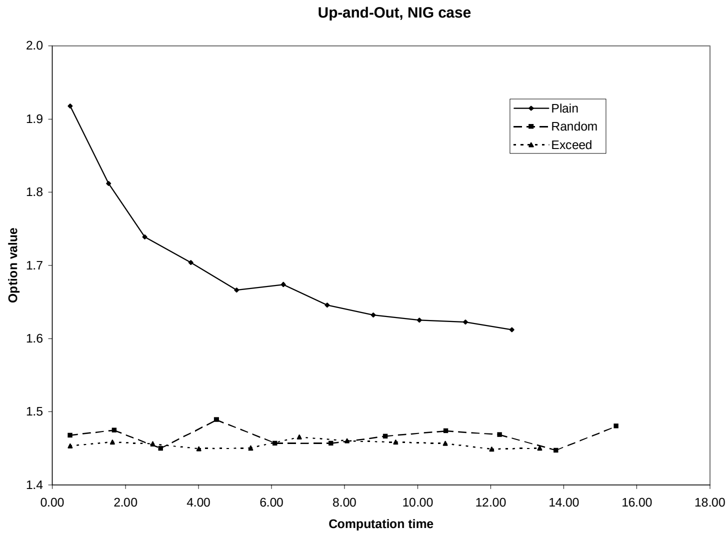

Figure 5: NIG case, convergence: Up and Out barrier option.

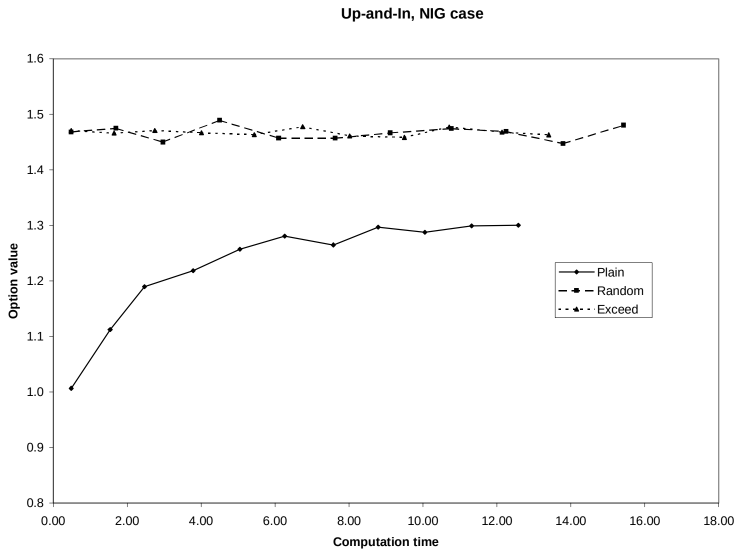

Figure 6: NIG case, convergence: Up and In barrier option.

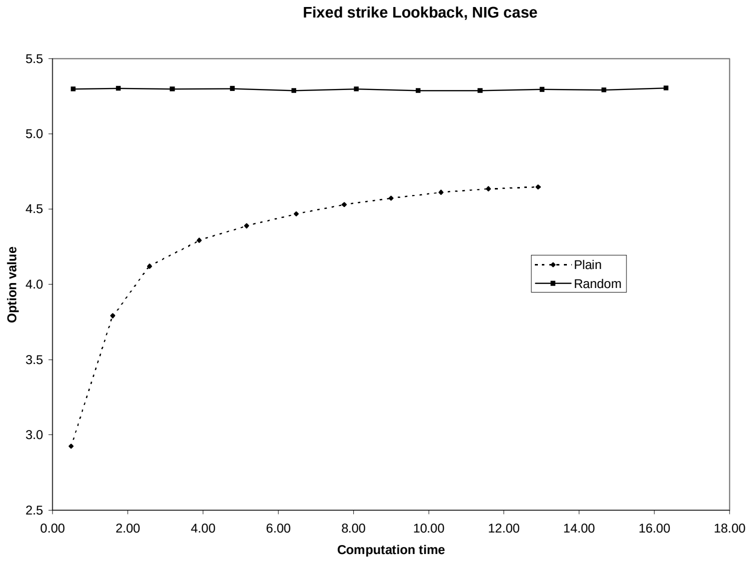

Figure 7: NIG case, convergence: Fixed Strike Lookback.

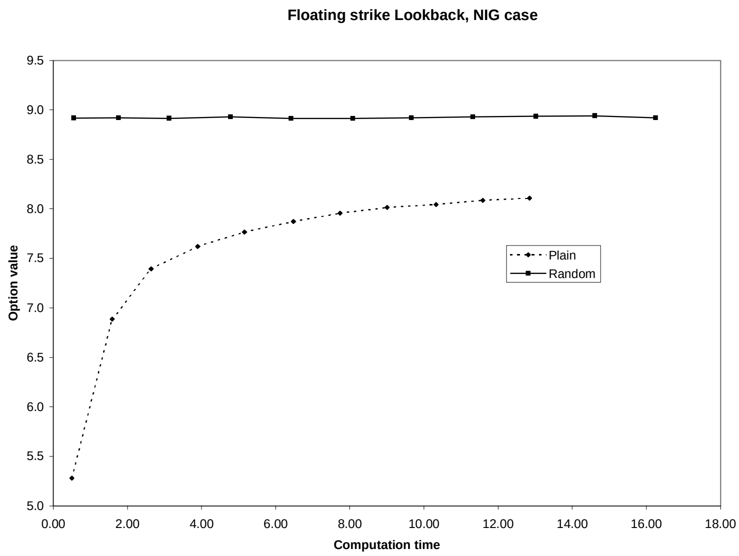

Figure 8: NIG case, convergence: Floating Strike Lookback.