Unbiased Monte Carlo valuation of lookback, swing and barrier options under variance gamma models

Martin Becker*
Saarland University, Saarbrücken, Germany

May 19, 2008

We introduce fast and unbiased methods for Monte Carlo valuation of lookback, swing, and barrier options under variance gamma models. For the valuation of lookback and swing options, a procedure to draw samples of final value, infimum, and supremum of variance gamma processes with additional drift term up to arbitrary precision is developed, for barrier options, a separate method is designed to improve the performance by exploiting the particular payoff structure. Both algorithms substantially rely on adaptive difference-of-gammas brige sampling, a newly introduced enhancement of the truncated difference-of-gammas bridge sampling method, which was originally presented in Avramidis et al. (2003) and successfully applied to lookback and barrier option valuation in Avramidis and L'Ecuyer (2006). By applying our algorithms to examples treated in Avramidis and L'Ecuyer (2006) we observe considerable reductions in computational effort and memory requirements. Furthermore, our algorithms feature a priori precision bounds, whereas Avramidis and L'Ecuyer (2006) are restricted to a posteriori bounds.

Keywords: Monte-Carlo simulation, variance gamma process, gamma bridge, barrier options, lookback options

*Tel.: +49 681 302 3571, Fax.: +49 681 302 3551, Address: Saarland University, Campus C3.1, Room 206, 66123 Saarbrücken, email: martin.becker@mx.uni-saarland.de

1 Introduction

The family of variance gamma (VG) processes was introduced by Madan and Seneta (1990) and later extended by Madan et al. (1998) as a continuous time model for log-prices of securities. Originally established as Brownian motions subordinated by independent gamma processes, VG processes are pure-jump, infinite activity, finite variation Lévy-processes. Although there is empirical evidence that stock (log-)prices do not follow Lévy-processes (see e.g. Klößner (2006)), the VG model provides a better fit to market option prices than the Black-Scholes model, particularly mitigating the volatiliy smile effect. While closed-form expressions for European (vanilla) option prices in the VG model are available, the valuation of path-dependent options poses a greater challenge. Recently, two different approaches for efficient Monte Carlo and Quasi-Monte Carlo option pricing under the VG model were proposed, one by Avramidis and L'Ecuyer (2006), the other one by Ribeiro and Webber (2006). Both approaches make use of gamma bridge sampling techniques, but in different fashions and with different success, see Becker (2007).

A fundamental tool for the Monte Carlo methods introduced in this paper is a newly developed, fast and exact simulation method for final, minimal and maximal values of VG processes (including an additional drift term). Applications for final and extremal values of price processes include volatility/covariance estimation (see e.g. Rogers et al. (1994), Yang and Zhang (2000), Brandt and Diebold (2006), Rogers and Zhou (2008)), specification tests (see e.g. Becker et al. (2007), Klößner (2006), Klößner (2007)) and model calibration (see e.g. Lildholdt (2002), Venter et al. (2005)). In these applications, simulated triplets of final, minimal and maximal values can be used for performance measurements of estimators and power studies for specification tests. Moreover, the final, minimal and maximal values of an asset determine the payoffs of special path-dependent options, in particular single/double barrier options (with continuous reset condition), lookback options and swing options, which leads to the Monte Carlo option pricing applications this paper focuses upon.

The remainder is organized as follows: section 2 introduces some basic notation and describes the difference-of-gammas bridge sampling method (DGBS), which was introduced by Avramidis et al. (2003) and successfully applied by Avramidis and L'Ecuyer (2006). In section 3, an adaptive generalization of the DGBS method is developed to achieve a fast and exact (up to given precision bounds) simulation procedure for final, minimal and maximal values of VG processes with additional drift term. In section 4, the special payoff structure of barrier options is exploited to design a fast and completely unbiased MC option pricing method for single and double barrier options with continuous reset conditions. Section 5 compares the performance of the newly introduced methods to the MC methods developed in Avramidis and L'Ecuyer (2006) and traditional methods based on full dimensional path sampling. Section 6 concludes with a short summary.

2 Preliminaries

For $\theta \in \mathbb{R}$, $\sigma > 0$ let $W^{(\theta,\sigma)} = (W_t^{(\theta,\sigma)})_{t\geq 0}$ denote a Brownian motion with drift θ and volatility σ, i.e. $W_t^{(\theta,\sigma)} = \theta t + \sigma W_t$, $t \geq 0$, for a standard Brownian motion $W = (W_t)_{t\geq 0}$. For $a, b > 0$ let $Gamma(a, b)$ denote the gamma distribution with shape parameter a and scale parameter

b, i.e. the distribution given by the probability density function

$$f_{Gamma(a,b)}(x) = x^{a-1}\frac{e^{-x/b}}{b^a\Gamma(a)}\mathbb{1}_{x>0}\ .$$

For $\mu,\nu > 0$ let $G^{(\mu,\nu)} = (G_t^{(\mu,\nu)})_{t\geq 0}$ denote a gamma process with drift μ and volatility ν, i.e. a process with $G_0^{(\mu,\nu)} = 0$ and independent stationary increments $G_{t+h}^{(\mu,\nu)} - G_t^{(\mu,\nu)}$ following a $Gamma(h\frac{\mu^2}{\nu},\frac{\nu}{\mu})$ distribution ($t \geq 0, h > 0$). A VG process $\widetilde{X} = (\widetilde{X}_t)_{t\geq 0}$ with parameters (θ,σ,ν) is obtained by subordinating a Brownian motion $W^{(\theta,\sigma)}$ with drift θ and volatility σ by a gamma process $G^{(1,\nu)}$with drift $\mu = 1$ and volatility ν, which is independent of $W^{(\theta,\sigma)}$, i.e.

$$\widetilde{X}_t = \theta G_t^{(1,\nu)} + \sigma W_{G_t^{(1,\nu)}},\ t \geq 0,$$

for a standard Brownian motion W independent of $G^{(1,\nu)}$.
Apart from the representation as a subordinated Brownian motion, every VG process can be constructed as the difference of two independent gamma processes. In particular, a VG process $\widetilde{X} = (\widetilde{X}_t)_{t\geq 0}$ with parameters (θ,σ,ν) can be built via $\widetilde{X}_t = G_t^+ - G_t^-$, $t \geq 0$, where $G^+ = (G_t^+)_{t\geq 0}$ is a gamma process with drift $\mu_p = (\sqrt{\theta^2 + 2\sigma^2/\nu}+\theta)/2$ and volatility $\nu_p = \mu_p^2\nu$ and $G^- = (G_t^-)_{t\geq 0}$ is a gamma process with drift $\mu_n = (\sqrt{\theta^2 + 2\sigma^2/\nu} - \theta)/2$ and volatility $\nu_n = \mu_n^2\nu$, which is independent of G^+ (see Madan et al. (1998)).
For risk-neutral modeling purposes, every VG processes $\widetilde{X} = (\widetilde{X}_t)_{t\geq 0}$ with parameter set (θ,σ,ν) may be endowed with an additional drift $\mu \in \mathbb{R}$, resulting in a process $X = (X_t)_{t\geq 0}$ given by $X_t = \mu t + \widetilde{X}_t$, $t \geq 0$. For the remainder of this paper, we will only consider these extended VG processes, determined by their parameter set (μ,θ,σ,ν).
Sampling the path of a VG process involves sampling gamma processes in either representation. The standard method for sampling a gamma process $G^{(\mu,\nu)}$ on a discrete time grid $0 = t_0 < t_1 < \cdots < t_n$ is to independently draw the increments $G_{t_i}^{(\mu,\nu)} - G_{t_{i-1}}^{(\mu,\nu)}$ from the corresponding $Gamma((t_i - t_{i-1})\frac{\mu^2}{\nu},\frac{\nu}{\mu})$ distribution ($i \in \{1,\ldots,n\}$) and accumulate. A second approach, which was already employed by Dufresne et al. (1991), makes use of the conditional distribution of $G_t^{(\mu,\nu)}$ given $G_{\tau_1}^{(\mu,\nu)}$ and $G_{\tau_2}^{(\mu,\nu)}$, the so called "gamma bridge", for $0 \leq \tau_1 < t < \tau_2$. Since $G_t^{(\mu,\nu)} \overset{d}{=} G_{\tau_1}^{(\mu,\nu)} + (G_{\tau_2}^{(\mu,\nu)} - G_{\tau_1}^{(\mu,\nu)})Y$, where $Y \sim Beta((t-\tau_1)\frac{\mu^2}{\nu},(\tau_2 - t)\frac{\mu^2}{\nu})$ and $Beta(\alpha,\beta)$ denotes the beta distribution, given by the probability density function

$$f_{Beta(\alpha,\beta)}(x) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{\int_0^1 y^{\alpha-1}(1-y)^{\beta-1}dy}\mathbb{1}_{0<x<1}\ ,$$

sampling of the gamma process $G^{(\mu,\nu)}$ on a time grid $0 = t_0 < t_1 < \cdots < t_n$ can be done in arbitrary order, see Avramidis and L'Ecuyer (2006).
Originally, Avramidis et al. (2003) combined the representation of a VG process as difference of gamma processes with gamma bridge sampling methods in order to improve efficiency of Quasi-Monte Carlo techniques for the valuation of Asian options under VG models. The paths of both gamma processes, G^+ and G^-, were sequentially generated on a dyadic partition of length 2^m, $m \in \mathbb{N}$, for the considered time range $[0,T]$, i.e. on the discrete time grid $(t_i)_{i=0}^{2^m-1}$ with $t_0 = T$, $t_1 = \frac{1}{2}T$, $t_2 = \frac{1}{4}T$, $t_3 = \frac{3}{4}T$, $t_4 = \frac{1}{8}T$, $t_5 = \frac{3}{8}T$, …, $t_{2^m-1} = \frac{2^m-1}{2^m}T$ or equivalently $t_0 = T$, $t_i = \frac{2(i-2^{\lceil\log_2(1+i)-1\rceil})+1}{2^{\lceil\log_2(1+i)\rceil}}T$, $i \in \{0,\ldots,2^m-1\}$. Later, this simulation method was extended to arbitrary time sequences $(t_i)_{i=0}^{\infty}$, which are dense in $[0,T]$, and dynamic truncation

of the simulation procedure after a finite number of steps. Nevertheless, in all numerical experiments, a dyadic partition was retained, so that fast simulation methods for symmetric Beta distributions were applicable.
The so called "difference-of-gammas bridge sampling" (DGBS) provides some interesting bounds on the simulated VG path: let $X = (X_t)_{t\geq 0}$ be an arbitrary VG process given by the difference-of-gammas representation $X_t = \mu t + G_t^+ - G_t^-$, $t \geq 0$, for $\mu \in \mathbb{R}$ and independent gamma processes $G^+ = (G_t^+)_{t\geq 0}$, $G^- = (G_t^-)_{t\geq 0}$, let $\mu^+ = \max\{0,\mu\}$ and $\mu^- = \min\{0,\mu\}$. Due to the monotonicity of G^+ and G^-, it is obvious that

$$L_{\tau_1}^{\tau_2} := \mu\tau_1 + G_{\tau_1}^+ - G_{\tau_2}^- + (\tau_2 - \tau_1)\mu^- \leq X_t \leq \mu\tau_1 + G_{\tau_2}^+ - G_{\tau_1}^- + (\tau_2 - \tau_1)\mu^+ =: U_{\tau_1}^{\tau_2} \qquad (1)$$

holds for all $t \in [\tau_1, \tau_2]$, $0 \leq \tau_1 < \tau_2$, and that these bounds are narrowing with increasing τ_1 and decreasing τ_2 (cf. Corollary 1, Avramidis and L'Ecuyer (2006)).
In the following section, we will introduce the adaptive DGBS method, which enhances the (truncated) DGBS method by dynamic modification of the time grid $(t_i)_{i=0}^{\infty}$ for the path discretization, particularly by an adaptive selection of a subgrid $(t_{j_i})_{i=0}^{\infty}$, $j_i < j_{i+1} \forall i$, in order to skip dispensable simulation steps. The dynamic selection in the adaptive DGBS method will substantially rely on the bounds provided by (1).

3 Simulating final and extremal values

The truncated DGBS method of Avramidis and L'Ecuyer (2006) is already feasible for the simulation of minima (infima) and maxima (suprema) of VG processes X with additional drift term on time intervals $[0, T]$, $T > 0$, when neglecting time and memory requirements: given a tolerance level $\varepsilon > 0$, one could truncate the DGBS procedure when the bounds for X differ by at most 2ε in all (or at least all "relevant") intervals and return the smallest and largest midpoints of these intervals as minimal and maximal value, diverging from the "true" minimal and maximal values by at most ε. In section 5 it can be seen that, apart from speed issues, huge memory requirements may arise for common values of ε.
The main idea of the adaptive DGBS method (implemented in figure 1 based on a dyadic partition) is to use the information in the bounds provided by (1) to detect intervals, which cannot include either minimal or maximal values, and exclude them from further subdivisions. For this purpose, the algorithm holds a variable bl containing the biggest lower bound for the supremum and a variable su containing the smallest upper bound for the infimum which have shown up so far. If $su < L_{\tau_1}^{\tau_2} \leq U_{\tau_1}^{\tau_2} < bl$ holds for the current interval $[\tau_1, \tau_2]$, obviously $\inf_{t\in[\tau_1,\tau_2]} X_t > \underline{X}_T := \inf_{t\in[0,T]} X_t$ and $\sup_{t\in[\tau_1,\tau_2]} X_t < \overline{X}_T := \sup_{t\in[0,T]} X_t$, and the current interval can therefore be excluded from further considerations. Otherwise, the interval will be included in subsequent subdivions, furthermore bl and su are updated: obviously, $\max\{X_{\tau_1}, X_{\tau_2}\}$ is a lower bound for $\overline{X}_T$ and $\min\{X_{\tau_1}, X_{\tau_2}\}$ an upper bound for $\underline{X}_T$.
To aquire a feasible stopping rule the algorithm holds in similar fashion to bl and su the biggest upper bound, bu, and the smallest lower bound, sl, seen so far in the current stage (or depth) of the dyadic partition. At the end of each stage, it is checked whether the difference of the smallest upper bound su and the smallest lower bound sl as well as the difference between biggest upper bound bu and biggest lower bound bl stays below 2ε. In this case, the estimators $(sl + su)/2$ for the infimum and $(bl + bu)/2$ for the supremum fulfill the error constraints and the algorithm terminates. Instead of these estimates, slight modifications are calculated in the last line of figure 1 to ensure that the estimated infimum will not exceed zero or the final value,

Input: $\sigma, \nu > 0$, $\mu, \theta \in \mathbb{R}$, $T > 0$, $\varepsilon > 0$, $m_{\max} \in \mathbb{N}$
Output: Realisation $(x, \underline{x}, \overline{x})$ of $(X_T, \underline{X}_T, \overline{X}_T)$ for VG process X

$a \leftarrow \nu^{-1}$; $b_1 \leftarrow (\frac{1}{2}\nu(\sqrt{\theta^2 + 2\sigma^2/\nu} + \theta))^{-1}$; $b_2 \leftarrow (\frac{1}{2}\nu(\sqrt{\theta^2 + 2\sigma^2/\nu} - \theta))^{-1}$
$G_0^+ \leftarrow 0$; $G_0^- \leftarrow 0$; $\mu^+ \leftarrow \max\{0, \mu\}$; $\mu^- \leftarrow \min\{0, \mu\}$
$G_1^+ \leftarrow Gamma(Ta, b_1)$ rn; $G_1^- \leftarrow Gamma(Ta, b_2)$ rn; $t_0 \leftarrow 0$; $t_1 \leftarrow T$
$su \leftarrow \min\{0, G_1^+ - G_1^- + t_1\mu\}$; $bl \leftarrow \max\{0, G_1^+ - G_1^- + t_1\mu\}$
$lint_1 \leftarrow 0$; $rint_1 \leftarrow 1$; $nint \leftarrow 1$; $pos \leftarrow 1$
for $i \in \{1, \ldots, m_{\max}\}$ **do**
  $lint^o \leftarrow lint$; $rint^o \leftarrow rint$; $nint^o \leftarrow nint$; $nint \leftarrow 0$
  $sl \leftarrow bl$; $bu \leftarrow su$
  for $j \in \{1, \ldots, nint^o\}$ **do**
    $pos \leftarrow pos + 1$
    $i_l \leftarrow lint_j^o$; $i_m \leftarrow pos$; $i_r \leftarrow rint_j^o$; $t_{i_m} \leftarrow \frac{t_{i_l} + t_{i_r}}{2}$; $\delta \leftarrow \frac{t_{i_r} - t_{i_l}}{2\nu}$
    $Y^+ \leftarrow Beta(\delta, \delta)$ rn; $Y^- \leftarrow Beta(\delta, \delta)$ rn
    $G_{i_m}^+ \leftarrow G_{i_l}^+ + (G_{i_r}^+ - G_{i_l}^+)Y^+$; $G_{i_m}^- \leftarrow G_{i_l}^- + (G_{i_r}^- - G_{i_l}^-)Y^-$
    for $(i_-, i_+) \in \{(i_l, i_m), (i_m, i_r)\}$ **do**
      $lb \leftarrow G_{i_-}^+ - G_{i_+}^- + t_{i_-}\mu + (t_{i_+} - t_{i_-})\mu^-$
      $ub \leftarrow G_{i_+}^+ - G_{i_-}^- + t_{i_-}\mu + (t_{i_+} - t_{i_-})\mu^+$
      if $(lb < su)$ or $(ub > bl)$ **then**
        $nint \leftarrow nint + 1$; $lint_{nint} \leftarrow i_-$; $rint_{nint} \leftarrow i_+$
        $sl \leftarrow \min\{sl, lb\}$; $bu \leftarrow \max\{bu, ub\}$
        $bl \leftarrow \max\{bl, G_{i_-}^+ - G_{i_-}^- + t_{i_-}\mu, G_{i_+}^+ - G_{i_+}^- + t_{i_+}\mu\}$
        $su \leftarrow \min\{su, G_{i_-}^+ - G_{i_-}^- + t_{i_-}\mu, G_{i_+}^+ - G_{i_+}^- + t_{i_+}\mu\}$
  if $\max\{su - sl, bu - bl\} < 2\varepsilon$ **then break**
$x \leftarrow G_1^+ - G_1^- + T\mu$
$\underline{x} \leftarrow \min\{0, \frac{1}{2}(su + sl), x\}$; $\overline{x} \leftarrow \max\{0, \frac{1}{2}(bu + bl), x\}$

Figure 1: Simulation of $(X_T, \underline{X}_T, \overline{X}_T)$ for VG processes with drift

and that the supremum will not fall below zero or the final value. Obviously, the correction of these simulation artifacts may only decrease the approximation error.

Figure 2 illustrates a typical run of the simulation algorithm of figure 1 (for the process parameters, see section 5). The maximum depth of the dyadic partion, which is usually set large enough to never come into effect for real applications, is set to 9 here for illustrational purposes. The shaded regions visualize the bounds for the process subpaths, particularly for the infimum and the supremum. Process values are generated with bridge sampling only at the positions marked by filled circles. The filled regions of the bar at the lower border denote the subintervals which have to be divided in the next stage of the dyadic partition. It can be seen that most of the subintervals can be excluded from subdivisions very early, so the workload is reduced to a small fraction of the workload for full dimensional sampling. Details on the performance of the algorithm are provided in section 5.

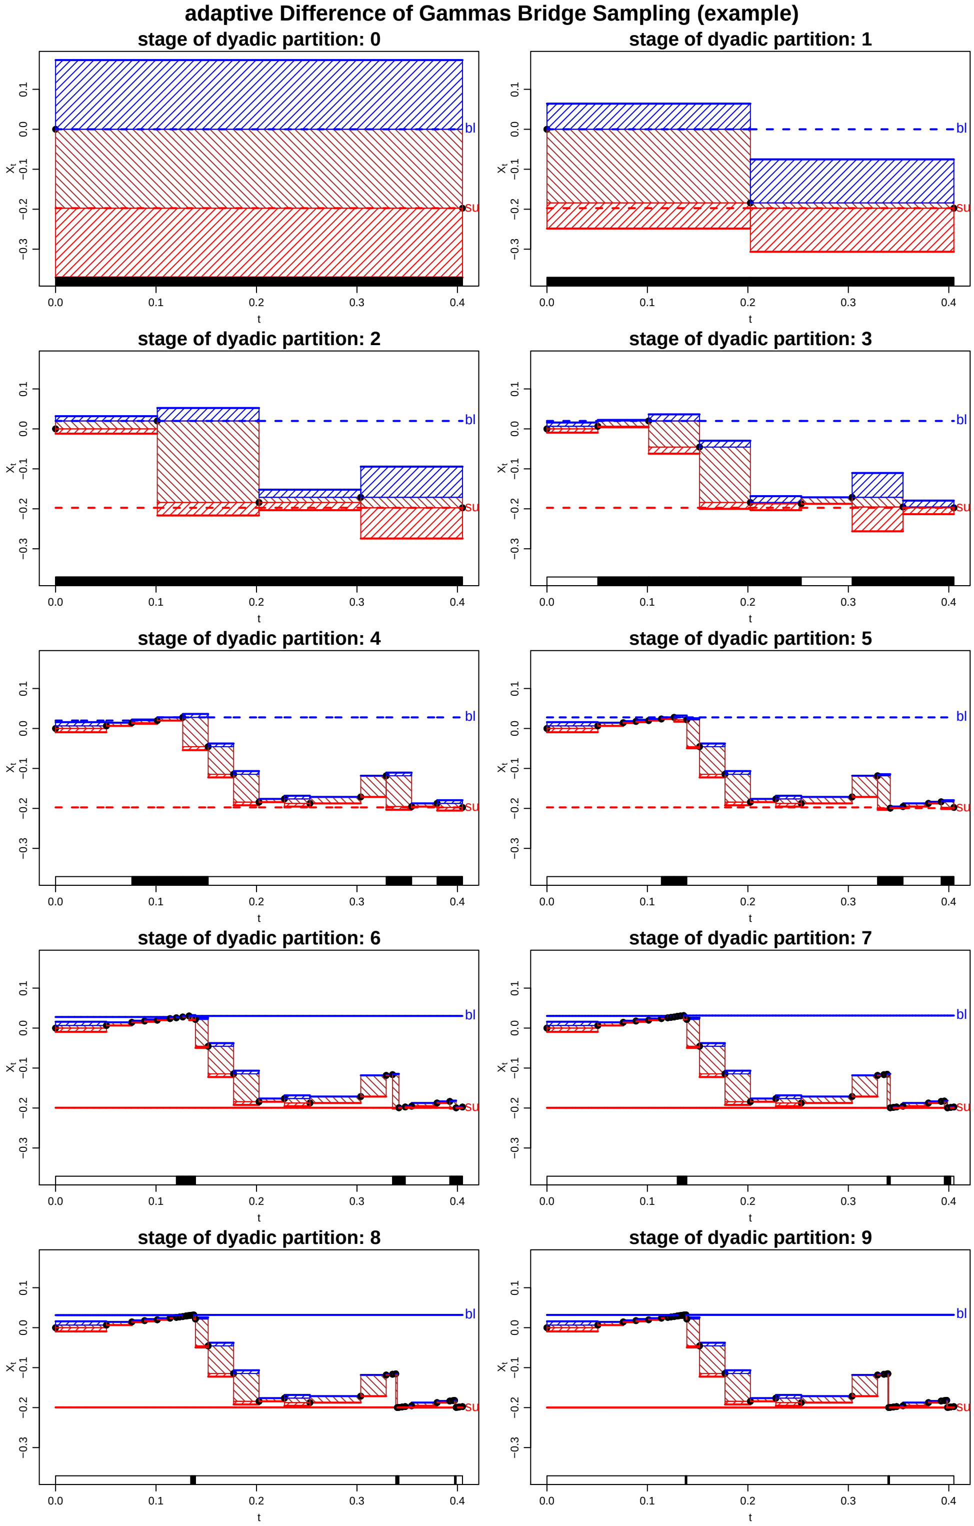

Figure 2: Simulation steps for VG process with drift

4 Application to option pricing

For option valuation, we rely on a risk-neutral approach based on general equilibrium arguments, which was already used by Madan et al. (1998) to derive valuation formulas for vanilla options under the VG model. In this setting, the risk-neutral asset price process $S = (S_t)_{t\geq 0}$ is defined by

$$S_t = S_0 e^{X_t} \ , \tag{2}$$

where $X = (X_t)_{t\geq 0}$ follows a VG process with parameter set $(\mu, \theta, \sigma, \nu)$. The drift μ is determined by $\mu = \omega + r - q$, where r is the risk-free interest rate, q the continuously compounded dividend yield, and $\omega = \ln(1 - \theta\nu - \sigma^2\nu/2)/\nu$ the compensator, which assures that the discounted value of the asset is a martingale, in particular $E(S_t) = S_0 e^{(r-q)t}$ (see Avramidis and L'Ecuyer (2006)).
The simulation procedure for final, minimal and maximal values of VG processes with drift from section 3 provides in principle a Monte Carlo valuation method for options with payoffs depending only on final, minimal and maximal values of the underlying: given a sample of n triplets $(x_i, \underline{x}_i, \overline{x}_i)$ of the final value x_i at time $T > 0$, the minimal value $\underline{x}_i$ and maximal value $\overline{x}_i$ up to time T of the VG process, the corresponding sample $(s_i, \underline{s}_i, \overline{s}_i)$ from the price process is easily calculated using (2), and the fair price $\widehat{C}$ of the option can be estimated as mean discounted payoff via

$$\widehat{C} = e^{-rT} \cdot \frac{1}{n} \sum_{i=1}^{n} \texttt{Payoff}(s_i, \underline{s}_i, \overline{s}_i) \ . \tag{3}$$

For lookback options, Avramidis and L'Ecuyer (2006), who focus on efficiency gains through applications of Quasi-MC methods, truncate their DGBS procedure quite early after a fixed number of steps and correct for a big amount of the thereby catched simulation bias with extrapolation methods. With the adaptive DGBS method from section 3, early truncation is not required, in particular, lookback and swing options can be valuated with negligible bias by setting the tolerance level of the simulation procedure sufficiently small.
Avramidis and L'Ecuyer (2006) exploit the special payoff structure of barrier options to truncate the DGBS method as early as possible, obtaining huge efficiency gains in comparison to full path discretization with $m^* = 2^{k^*}$ equally spaced points while obtaining the same accuracy. The computational effort of their method grows only linearly in $k^* = \log_2(m^*)$ compared to m^* for the full discretization method, but a (small) bias still appears by occasionally reaching m^* before the corresponding payoff is determined. Avramidis and L'Ecuyer again use some extrapolation techniques to reduce this bias. Applying the adaptive DGBS method of section 3 to similar truncation criteria results in a completely unbiased estimation, because there is no need to restrict the algorithm to a practically relevant maximum number of discretization points m^*, instead the algorithm continues until the corresponding (exact) payoff is determined.
The algorithm in figure 3 implements the adaptive DGBS method for valuation of general barrier options with continuous reset conditions under VG models (for the efficient calculation of *mean* and *se* in lines 29–30, see Knuth (1998), p. 232). Some parts of the algorithm depend on particular option characteristics, e.g. put/call, single/double barrier, down-in/down-out, up-in/up-out. Table 1 summarizes the different fragments which have to be pasted into figure 3. In a first step, it is checked whether a positive payoff is still possible after sampling the final process value and obtaining the first lower and upper process bounds (line 8). In this case, the adaptive DGBS procedure is started. After each bridge sampling step, the knock-in and knock-out conditions (depending on the option type) are inspected, and the algorithm

terminates if these conditions are fulfilled. The bounds for the process subpaths now provide the information on which of the subintervals may contain barrier crossings and therefore have to be considered for further bridge sampling steps. Examples for the application to up-and-out calls are given in figures 4 (knock-out case) and 5 (no knock-out case).

Input: $\sigma, \nu, T > 0$, $\theta \in \mathbb{R}$, $r, q \in \mathbb{R}$, $n, m_{\max} \in \mathbb{N}$, option parameter set Θ
Output: Estimator (*mean*) with standard error (*se*) for price of (single or double) barrier option under VG process

$\mu \leftarrow r + \frac{1}{\nu}\ln(1 - \theta\nu - \frac{1}{2}\sigma^2\nu) - q$; $mean \leftarrow 0$; $squares \leftarrow 0$
$a \leftarrow \nu^{-1}$; $b_1 \leftarrow (\frac{1}{2}\nu(\sqrt{\theta^2 + 2\sigma^2/\nu} + \theta))^{-1}$; $b_2 \leftarrow (\frac{1}{2}\nu(\sqrt{\theta^2 + 2\sigma^2/\nu} - \theta))^{-1}$
$G_0^+ \leftarrow 0$; $G_0^- \leftarrow 0$; $\mu^+ \leftarrow \max\{0, \mu\}$; $\mu^- \leftarrow \min\{0, \mu\}$
for $k \in \{1, \ldots, n\}$ **do**
    $G_1^+ \leftarrow Gamma(Ta, b_1)$ rn; $G_1^- \leftarrow Gamma(Ta, b_2)$ rn; $t_0 \leftarrow 0$; $t_1 \leftarrow T$
    $x \leftarrow G_1^+ - G_1^- + T\mu$
    $lint_1 \leftarrow 0$; $rint_1 \leftarrow 1$; $nint \leftarrow 1$; $pos \leftarrow 1$
    if *initial conditions met for* Θ **then**
        $in \leftarrow$ true; $done \leftarrow$ false
        for $i \in \{1, \ldots, m_{\max}\}$ **do**
            $lint^o \leftarrow lint$; $rint^o \leftarrow rint$; $nint^o \leftarrow nint$; $nint \leftarrow 0$
            for $j \in \{1, \ldots, nint^o\}$ **do**
                $pos \leftarrow pos + 1$
                $i_l \leftarrow lint_j^o$; $i_m \leftarrow pos$; $i_r \leftarrow rint_j^o$; $t_{i_m} \leftarrow \frac{t_{i_l} + t_{i_r}}{2}$; $\delta \leftarrow \frac{t_{i_r} - t_{i_l}}{2\nu}$
                $Y^+ \leftarrow Beta(\delta, \delta)$ rn; $Y^- \leftarrow Beta(\delta, \delta)$ rn
                $G_{i_m}^+ \leftarrow G_{i_l}^+ + (G_{i_r}^+ - G_{i_l}^+)Y^+$; $G_{i_m}^- \leftarrow G_{i_l}^- + (G_{i_r}^- - G_{i_l}^-)Y^-$
                for $(i_-, i_+) \in \{(i_l, i_m), (i_m, i_r)\}$ **do**
                    $lb \leftarrow G_{i_-}^+ - G_{i_+}^- + t_{i_-}\mu + (t_{i_+} - t_{i_-})\mu^-$
                    $ub \leftarrow G_{i_+}^+ - G_{i_-}^- + t_{i_-}\mu + (t_{i_+} - t_{i_-})\mu^+$
                    if *knock-out conditions met for* Θ **then**
                        $done \leftarrow$ true; $in \leftarrow$ false; **break**
                    if *knock-in conditions met for* Θ **then**
                        $done \leftarrow$ true; **break**
                    if *current interval relevant for* Θ **then**
                        $nint \leftarrow nint + 1$; $lint_{nint} \leftarrow i_-$; $rint_{nint} \leftarrow i_+$
            if ($done =$ true) or ($nint = 0$) **then break**
    if *knock-in conditions not met for* Θ **then** $in \leftarrow$ false
    if $in =$ true **then** $po \leftarrow \texttt{Payoff}(x, \Theta)$ **else** $po \leftarrow 0$
    $delta \leftarrow e^{-rT} \cdot po - mean$; $mean \leftarrow mean + delta/k$
    $squares \leftarrow squares + delta \cdot (e^{-rT} \cdot po - mean)$
$se \leftarrow \sqrt{squares/(N(N-1))}$

Figure 3: MC valuation of barrier options under VG processes

line	condition type/payoff	option characteristic	condition/value
8	initial conditions (every relevant condition must hold)	put call up-in / up-out down-in / down-out	$S_0e^x \leq K$ $S_0e^x \geq K$ $G_1^+ + T\mu^+ > \ln\frac{B_u}{S_0}$ $G_1^+ - T\mu^- < \ln\frac{B_l}{S_0}$
19	knock-in conditions (every relevant condition must hold)	up-in down-in	$lb > \ln\frac{B_u}{S_0}$ $ub < \ln\frac{B_l}{S_0}$
21, 26	knock-out conditions (at least one relevant condition must hold)	up-out down-out	$lb > \ln\frac{B_u}{S_0}$ $ub < \ln\frac{B_l}{S_0}$
23	relevance conditions (at least one relevant condition must hold)	up-in / up-out down-in / down-out	$ub > \ln\frac{B_u}{S_0}$ $lb < \ln\frac{B_l}{S_0}$
28, 29	`Payoff`(x, Θ)	put call	$\max\{0, K - S_0e^x\}$ $\max\{0, S_0e^x - K\}$

Table 1: Conditions/payoffs for figure 3

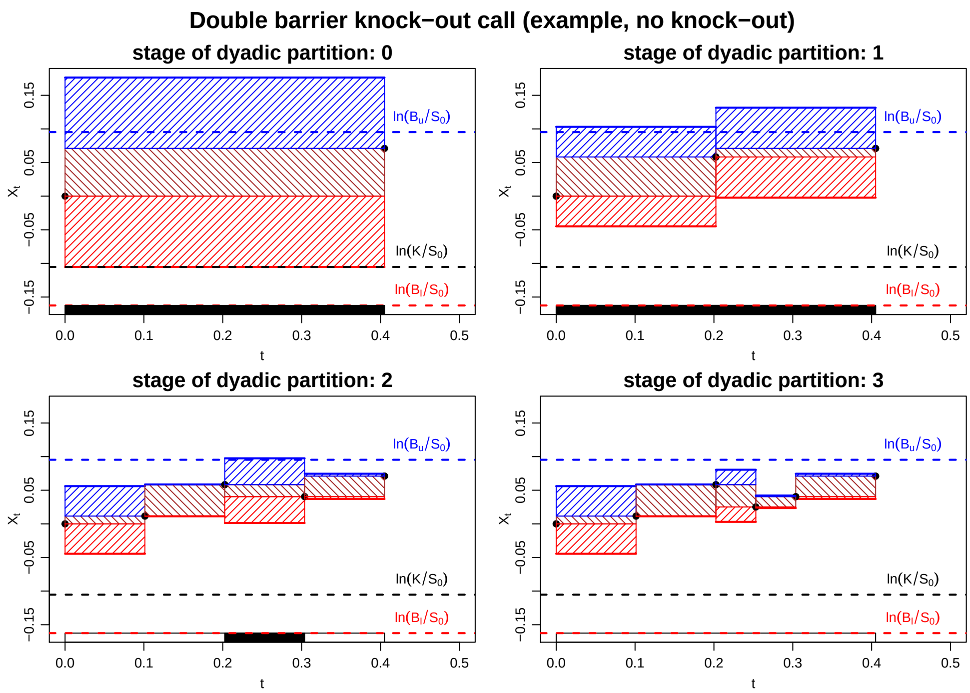

Figure 4: Simulation steps for barrier option payoff (no knock-out)

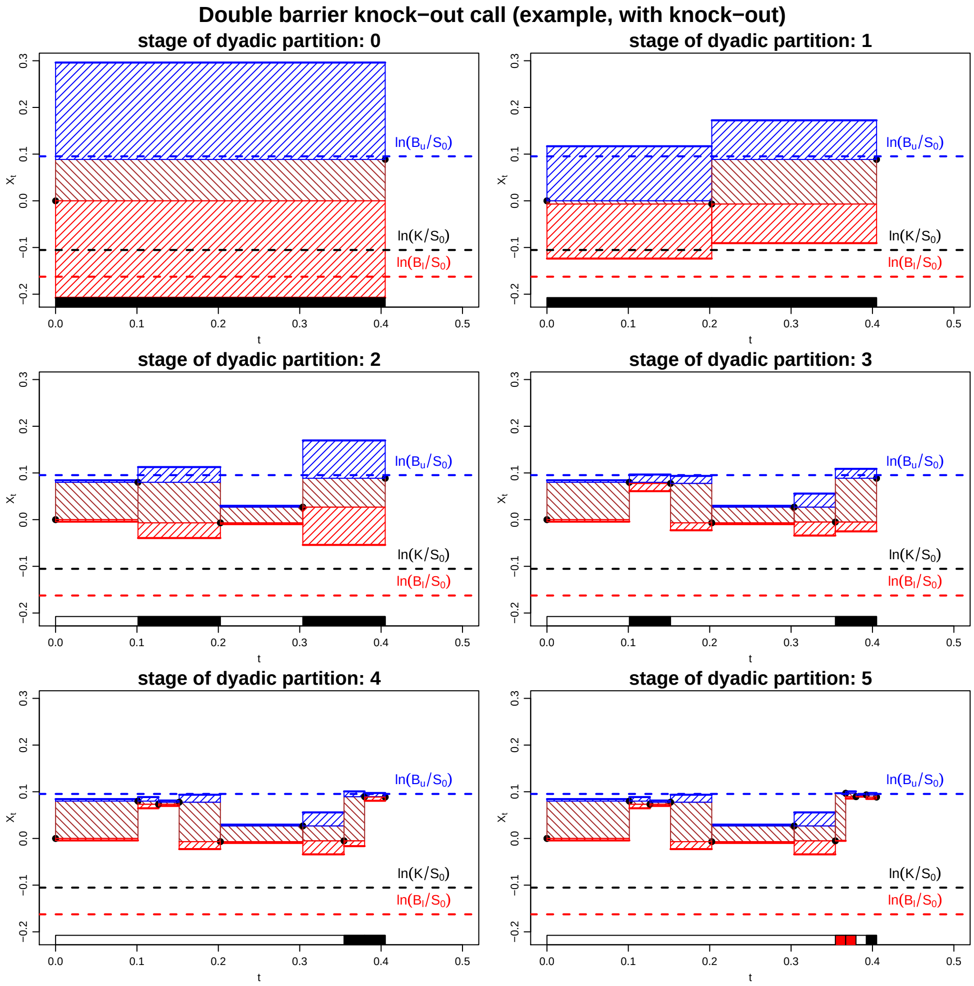

Figure 5: Simulation steps for barrier option payoff (knock-out)

5 Performance

In this section, the performance of the algorithms in figure 1 and figure 3 will be analysed and compared to related methods, including some of the MC methods introduced in Avramidis and L'Ecuyer (2006), using the same VG model/option parameters, in particular $\sigma = 0.1927$, $\nu = 0.2505$, $\theta = -0.2859$, $T = 0.40504$, $r = 0.0548$, $q = 0$, $S_0 = 100$, furthermore for the up-and-in call $K = 100$ and $B_u = 120$. The corresponding additional drift term of the VG process results in $\mu = 0.31356$.

In a first analysis, we examine the efficiency gain of the adaptive DGBS method for the simulation of final, minimal and maximal values of VG processes (figure 1) compared to (truncated) full dimensional path sampling. For this purpose, we simulate $N = 10^7$ triplets for tolerance levels $\varepsilon = 10^{-i}$, $i \in \{2, 6, 10, 14\}$, with adaptive DGBS. We calculate the average number of simulated points per path for adaptive DGBS as well as the average number of simulated points for the truncated full dimensional sampling with exactly the same accuracy, i.e. $\frac{1}{N}\sum_{j=1}^{N} 2^{k_j^{(i)}}$, where $k_j^{(i)}$ is the depth of the dyadic partition where adaptive DGBS stops in the jth simulation run for tolerance level 10^{-i}. The results, including the efficiency gain for the adaptive DGBS method (calculated as the ratio of average number for truncated full dimensional path sampling to average number for adaptive DGBS), are summarized in table 2. It can be seen that large efficiency gains are obtained even for moderate error bounds ε.

tolerance level ε	10^{-2}	10^{-6}	10^{-10}	10^{-14}
adaptive DGBS	16.32	50.44	76.40	102.30
truncated full dim. sampling	105.44	$2.47 \cdot 10^5$	$2.01 \cdot 10^9$	$1.64 \cdot 10^{13}$
efficiency gain	6.46	$4.90 \cdot 10^3$	$2.63 \cdot 10^7$	$1.60 \cdot 10^{11}$

Table 2: Comparison of average number of simulated points per path (adaptive DGBS vs. truncated full dimensional sampling)

Obviously, the bias of Monte Carlo option valuation for floating strike lookback call options with discounted payoff $e^{-rT}(S_T - \inf_{0\le t\le T} S_t)$ based on adaptive DGBS with tolerance level ε can not exceed $S_0 e^{-rT}\varepsilon$, since $S_t = S_0 e^{X_t}$ and $\inf_{0\le t\le T} X_t$ is simulated with error smaller than ε, furthermore $\inf_{0\le t\le T} X_t < 0$ and therefore applying the exponential function reduces the error in $e^{\inf_{0\le t\le T} X_t}$. Probably the bias is much smaller, because there are no apparent reasons for the error to be systematic. Avramidis and L'Ecuyer (2006) applied extrapolation techniques to reduce the bias for their lookback option pricing method to about 10^{-4} with DGBS and 256 simulated points per path. As table 2 shows, the corresponding tolerance level of about 10^{-6} for the VG process can be reached with 50.44 simulated points in average when applying adaptive DGBS. Furthermore, if a workload of 102.30 points per path (in average) is feasible, the simulation bias for the option value reduces to less then 10^{-12} when applying adaptive DGBS.

In a second analysis, we compare the expected truncation value (which corresponds to the number of simulated points per path) for the truncated DGBS method when applied to barrier options (Avramidis and L'Ecuyer (2006), table 1), to the expected truncation number for the adaptive DGBS method (figure 3). While the computational effort of the truncated DGBS method grows roughly linear in $k^* = \log_2(m^*)$, where m^* is the maximum number of points which may be simulated per path, the computational effort of the adaptive DGBS method stabilizes for $k^* \geq 8$. The results are illustrated in figure 6.

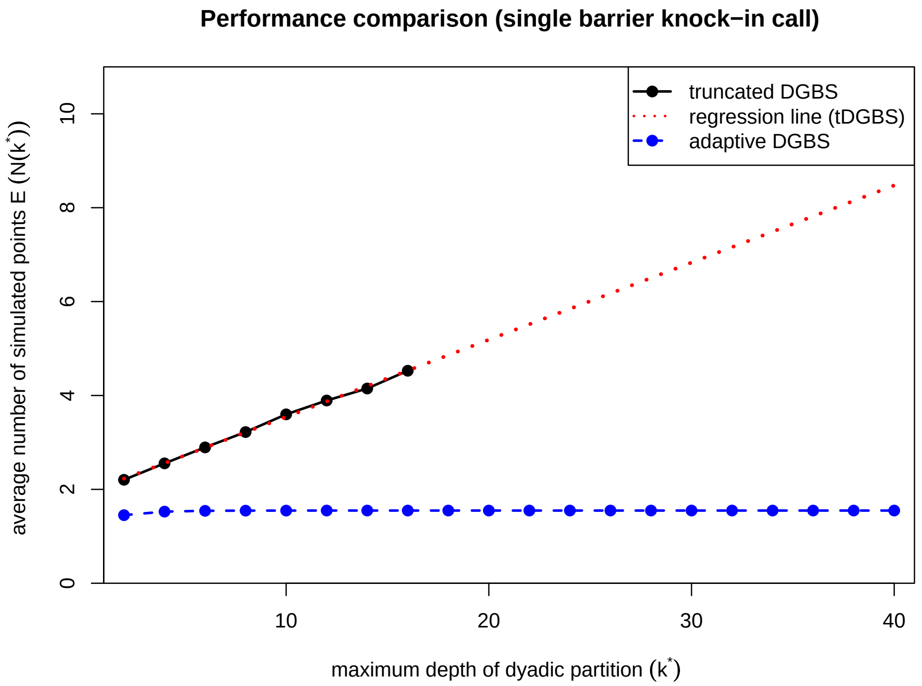

Figure 6: Estimated expected truncation values (up-and-in call) for adaptive and truncated DGBS

To measure the speed and to verify the correctness of the proposed techniques, the simulation method for final and extremal values (figure 1) and the Monte Carlo valuation method for barrier options (figure 3) were implemented in C, using the random number facilities and the data management of the statistical computing language R (R Development Core Team (2007)). For the generation of $Beta(a, a)$ random numbers, a combination of two methods is used, more precisely the method of Devroye (1986, p. 437) for $a > \frac{1}{2}$ and the method of Devroye (1996) for $0 < a \leq \frac{1}{2}$.

For the valuation of lookback options, the simulation of either the infimum or supremum is dispensible. For this application, the algorithm of figure 1 can be simplified in an obvious manner for speed improvements. When applying our algorithms to the examples in Avramidis and L'Ecuyer (2006), we are able to fully reproduce the reference prices. Table 3 summarizes the results and the speed of the proposed methods (Pentium D 3.0GHz, 1 core, Windows XP). Apart from the saving of computation time, memory requirements are an issue both for lookback and barrier options when applying the truncated DGBS method. Since m^* values of the two gamma processes G^+ and G^- must be stored for the regular DGBS method, memory allocation of at least $16 \cdot m^*$ bytes is necessary. The dyadic partitioning in the adaptive DGBS method for completely unbiased barrier option valuation (figure 3) occasionally reached depths greater than 30, which corresponds to memory requirements of more than 16 GB for the truncated DGBS method, which is not feasible on standard PCs. In comparison, the memory

option type	# samples	price estimate	est. std. deviation	payoffs/sec.
lookback call ($\varepsilon = 10^{-6}$)	$5 \cdot 10^8$	9.39827	0.0003244	11,350
up-and-in call	$1 \cdot 10^{10}$	2.15705	0.0000709	323,562
swing ($\varepsilon = 10^{-6}$)	$1 \cdot 10^7$	17.07974	0.0025971	6,506

Table 3: Results and speed of adaptive DGBS for option valuation

requirement of the adaptive DGBS method was about 14 KB. Due to the small memory and time requirements of the adaptive DGBS method, the maximum truncation number can be set to values that are very unlikely to be ever reached, e.g. 2^{256}.

We didn't investigate in performance comparisons to the methods introduced by Ribeiro and Webber (2006), who claimed to correct for simulation bias in the Monte Carlo valuation of barrier options by applying gamma bridge sampling to the subordinating gamma process and exploiting well-known results of the distribution of the maximum of a Brownian bridge. Because of a mistake in the derivation of their results, they overlook that the proposed method is not applicable to VG processes (nor to NIG processes, which they also considered). Experimentally applying their method to the examples considered in Ribeiro and Webber (2006) results in an 'overcorrection' and amplifies the bias to a $15 - 142$ times higher bias in the opposite direction, see Becker (2007).

6 Summary

In this paper, we proposed fast and unbiased Monte Carlo methods for the valuation of lookback, swing, and barrier options under variance gamma models. As a by-product, a fast method was developed for the simulation of final value, infimum, and supremum of variance gamma processes including an additional drift term, which is exact up to arbitrarily small a priori precision bounds.

The key ingredient of our algorithms is the newly introduced adaptive difference-of-gammas bridge sampling method, an enhanced version of the (truncated) difference-of-gammas bridge sampling procedure, which was originally established by Avramidis et al. (2003). For Monte Carlo lookback and barrier option valuation, our performance measurements showed considerable reductions in computional effort and memory requirements, in particular for small precision bounds ε. Therefore, a priori bounds ε close to machine precision level are perfectly feasible, which constitutes the fundament of our effectively unbiased valuation methods.

Avramidis and L'Ecuyer (2006) focused on further efficiency improvements by applying Quasi-Monte Carlo techniques for variance reduction. They observed considerable efficiency gains compared to pure Monte Carlo methods for their valuation procedures based on truncated difference-of-gammas bridge sampling. Variance reducing techniques, in particular importance sampling and Quasi-Monte Carlo methods, should in principle be applicable for our methods based on adaptive difference-of-gammas bridge sampling as well. A detailed analysis of possible efficiency gains is left to further studies.

References

Avramidis, A. N. and P. L'Ecuyer (2006) Efficient Monte Carlo and Quasi-Monte Carlo Option Pricing Under the Variance Gamma Model, *Management Science*, 52(12), pp. 1930–1944.

Avramidis, A. N., P. L'Ecuyer, and P.-A. Tremblay (2003) Efficient simulation of gamma and variance-gamma processes, in *Proceedings of the 2003 Winter Simulation Conference*, edited by Chick, S., P. J. Sánchez, D. Ferrin, and D. J. Morrice, pp. 319–326, IEEE Press.

Becker, M. (2007) Comment on "Correcting for Simulation Bias in Monte Carlo Methods to Value Exotic Options in Models Driven by Lévy Processes" by C. Ribeiro and N. Webber, `http://www.wiwi.uni-saarland.de/friedmann/papers/CommentOnRW2006.pdf`, submitted.

Becker, M., R. Friedmann, S. Klößner, and W. Sanddorf-Köhle (2007) A Hausman test for Brownian motion, *AStA - Advances in Statistical Analysis*, 91(1), pp. 3–21.

Brandt, M. W. and F. X. Diebold (2006) A No-Arbitrage Approach to Range-Based Estimation of Return Covariances and Correlations, *Journal of Business*, 79(1), pp. 61–74.

Devroye, L. (1986) *Non-Uniform Random Variate Generation*, Springer, New York.

Devroye, L. (1996) Random variate generation in one line of code, in *Proceedings of the 1996 Winter Simulation Conference*, edited by Charnes, J. M., D. J. Morrice, D. T. Brunner, and J. J. Swain, pp. 265–272, IEEE Press.

Dufresne, F., H. U. Gerber, and E. S. W. Shiu (1991) Risk Theory with the Gamma Process, *ASTIN Bulletin*, 21(2), pp. 177–192.

Klößner, S. (2006) Empirical Evidence: Intraday Returns are neither Symmetric nor Lévy Processes, Paper presented at Statistische Woche, Dresden, 2006, September 18-21.

Klößner, S. (2007) On Intraday Time-Reversibility of Return Processes, Paper presented at Statistics under one umbrella, Bielefeld, 2007, March 27-30.

Knuth, D. E. (1998) *The Art of Computer Programming, Volume 2: Seminumerical Algorithms*, Addison-Wesley Longman, Amsterdam, 3. edition.

Lildholdt, M. P. (2002) Estimation of GARCH models based on open, close, high, and low prices, Centre for Analytical Finance Working Paper Series No. 128, University of Aarhus, `http://www.cls.dk/caf/wp/wp-128.pdf`.

Madan, D. B., P. P. Carr, and E. C. Chang (1998) The Variance Gamma Process and Option Pricing, *European Finance Review*, 2(1), pp. 79–105.

Madan, D. B. and E. Seneta (1990) The Variance Gamma (V.G.) Model for Share Market Returns, *Journal of Business*, 63(4), pp. 511–524.

R Development Core Team (2007) *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.

Ribeiro, C. and N. Webber (2006) Correcting for Simulation Bias in Monte Carlo Methods to Value Exotic Options in Models Driven by Lévy Processes, *Applied Mathematical Finance*, 13(4), pp. 333–352.

Rogers, L. C. G., S. E. Satchell, and Y. Yoon (1994) Estimating the volatility of stock prices: A comparison of methods that use high and low prices, *Applied Financial Economics*, 4(3), pp. 241–247.

Rogers, L. C. G. and F. Zhou (2008) Estimating correlation from high, low, opening and closing prices, To appear in: The Annals of Applied Probability.

Venter, J. H., P. J. de Jongh, and G. Griebenow (2005) NIG-GARCH models based on open, close, high and low prices, *South African Statistical Journal*, 39(2), pp. 79–101.

Yang, D. and Q. Zhang (2000) Drift-Independent Volatility Estimation Based on High, Low, Open, and Close Prices, *Journal of Business*, 73(3), pp. 477–491.