Maastricht University

Fast Drift Approximated Pricing in the Bgm Model

Citation for published version (APA):

Pietersz, R., Pelsser, A., & van Regenmortel, M. (2004). Fast Drift Approximated Pricing in the Bgm Model. *Journal of Computational Finance*, *8*(1), 93-124. https://doi.org/10.21314/JCF.2004.113

Document status and date:
Published: 01/01/2004

DOI:
10.21314/JCF.2004.113

Document Version:
Publisher's PDF, also known as Version of record

Document license:
Taverne

Please check the document version of this publication:

• A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
• The final author version and the galley proof are versions of the publication after peer review.
• The final published version features the final layout of the paper including the volume, issue and page numbers.

Link to publication

General rights
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

• Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
• You may not further distribute the material or use it for any profit-making activity or commercial gain
• You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy
If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Download date: 23 Apr. 2024

Fast drift-approximated pricing in the BGM model

Raoul Pietersz

Erasmus Research Institute of Management, Erasmus University Rotterdam, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands

Antoon Pelsser

Econometric Institute, Faculty of Economics, Erasmus University Rotterdam, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands

Marcel van Regenmortel

Product Development Group (HQ7011), ABN AMRO Bank, P.O. Box 283, 1000 EA Amsterdam, The Netherlands

It is demonstrated that the forward rates process discretized by a single time step together with a separability assumption on the volatility function allows for representation by a low-dimensional Markov process. This in turn leads to efficient pricing by, for example, finite differences. We then develop a discretization based on the Brownian bridge that is especially designed to have high accuracy for single time stepping. The scheme is proven to converge weakly with order one. We compare the single time step method for pricing on a grid with multi-step Monte Carlo simulation for a Bermudan swaption, reporting a computational speed increase by a factor 10, yet maintaining sufficiently accurate pricing.

1 Introduction

The BGM framework, developed by Brace, Gątarek and Musiela (1997), Miltersen, Sandmann and Sondermann (1997) and Jamshidian (1996, 1997), is now one of the most popular models for pricing interest rate derivatives. In the BGM framework almost all prices are computed using Monte Carlo simulation. An advantage of Monte Carlo is its applicability to almost any product. However, it has the drawback of being computationally rather slow. In an attempt to limit the computational time, Hunter, Jäckel and Joshi (2001a,b), Jäckel (2002, Section 12.5) and Kurbanmuradov, Sabelfeld and Schoenmakers (1999, 2002) introduced

We are grateful for the comments of Tomas Björk, Dick Boswinkel, Marije Elkenbracht, Dariusz Gątarek, Lane Hughston, Dimitri Neumann, Martin Martens, Ton Vorst, Lars Wulf, and seminar participants at ABN AMRO Bank, ECFR, Mathematics in Finance 2002, South Africa, and Quantitative Methods in Finance 2002, Australia. We thank Glyn Baker for the use of his hopscotch implementation. The first author is grateful for the financial support of the Erasmus Center of Financial Research.

predictor-corrector drift approximations, which reduce the Monte Carlo stage to single time-step simulation.

This paper presents a significant addition to the single time step pricing method. We show that much more efficient numerical methods (either numerical integration or finite differences) may be used at the cost of a minor additional assumption, *separability*. The latter is a non-restrictive requirement on the form of the volatility function. The single time step together with separability renders the state of the BGM model completely determined by a low-dimensional Markov process. This enables efficient implementation.

We give an example of the fast single time step pricing framework for Bermudan swaptions. A comparison is made with prices obtained by least-squares multi-time step Monte Carlo simulation in the BGM model. This includes the use of the Longstaff and Schwartz (2001) method.

The computational speed increase achieved with the use of finite differences for BGM single time step pricing is the main result. This paper also contains two other results:

- The first result is a new time discretization using a Brownian bridge, as introduced in Section 3, which is proven to have least-squares error in a certain sense (to be defined) for single time step discretizations. In Section 4 it is shown numerically that the Brownian bridge scheme outperforms (in the case of single time steps) various other discretizations for the Libor-in-arrears density test. In the first part of Section 5, we prove theoretically that the Brownian bridge scheme converges weakly with order one when used for multi-time step Monte Carlo. In the second part of Section 5, we compare the Brownian bridge scheme numerically with other discretizations for multi-time steps.
- The second result is a method for measuring the accuracy of single time stepping. This is the timing inconsistency test as outlined in Section 8.

A further application of the Brownian bridge drift approximation is its use in the likelihood ratio method. This method, introduced by Broadie and Glasserman (1996), efficiently estimates risk sensitivities for Monte Carlo pricing. The particular application of the likelihood ratio method to the Libor market model has been developed by Glasserman and Zhao (1999), who proposed the use of drift approximations.

The outline of this paper is as follows. After setting out some basic notation and the most important formulas for the BGM model, the single time step pricing framework is developed, various discretization schemes are discussed and the Brownian bridge scheme is introduced. The Brownian bridge scheme is then investigated theoretically and numerically for both single and multi-time steps, respectively. Next, the proposed framework is worked out for the one-factor case. This is followed by an example of the pricing of Bermudan swaptions, both for a one- and a two-factor model. A test is then developed to assess the quality of single time steps. Finally, conclusions are drawn.

2 Notation for BGM model

In this section our notation of the BGM model is introduced.

Consider a BGM model, $\mathcal{M}$.[1] Such a model features N forward rates, L_i, $i = 1,\ldots, N$, where forward i accrues from time T_i to time T_{i+1}, $0 < T_1 < \ldots < T_{N+1}$. Define the accrual factor, δ_i, to be $T_{i+1} - T_i$. Denote by $B_i(t)$ the time-t price of a discount bond that expires at time T_i. Bond prices and forward rates are linked by the relation

$$1+\delta_i L_i(t) = \frac{B_i(t)}{B_{i+1}(t)}$$

Each forward rate is driven by a d-dimensional Brownian motion (where d is the number of stochastic factors in the BGM model), $\mathbf{W}$, as follows:

$$\frac{\mathrm{d}L_i(t)}{L_i(t)} = \tilde{\mu}_i(t)\,\mathrm{d}t + \sigma_i(t)\cdot \mathrm{d}\mathbf{W}(t) \tag{1}$$

Here σ_i is the d-dimensional volatility vector, and $\tilde{\mu}_i$ is the drift term, whose form will in general depend on the choice of probability measure. Throughout, we use the numeraire probability measure associated with the bond maturing at time T_{N+1}, the so called *terminal measure*. There is a specific reason why we use the terminal measure, and this is explained in Remark 2 of Section 3. For the terminal measure, the drift term will have the following form for $i < N$:

$$\tilde{\mu}_i\big(t, L_{i+1},\ldots, L_N\big) = -\sum_{k=i+1}^{N} \frac{\delta_k L_k \sigma_k(t)\cdot\sigma_i(t)}{1+\delta_k L_k} \tag{2}$$

For $i = N$ the drift term is zero. This simply expresses the well-known fact that a forward rate is a martingale under its associated forward measure.

For the remainder of this paper it will be useful to have stochastic differential equation (SDE) (1) in logarithmic form:

$$\begin{aligned} \mathrm{d}\log L_i(t) &= \mu_i(t)\,\mathrm{d}t + \sigma_i(t)\cdot \mathrm{d}\mathbf{W}^{N+1}(t), \\ \mu_i(t) &= \tilde{\mu}_i(t) - \frac{1}{2}\,\|\sigma_i(t)\|^2 \end{aligned} \tag{3}$$

Last, we introduce the notion of all available forward rates at a given point in time. Define $i(t)$ to be the smallest integer i such that $t \le T_i$. Define $\mathbf{L}$ to consist of all forward rates that have not yet expired at time t, ie,

$$\mathbf{L}(t) = \big(L_{i(t)}(t),\ldots, L_N(t)\big) \tag{4}$$

[1] The construction of such a model may be found in, eg, Musiela and Rutkowski (1997), Pelsser (2000) or Brigo and Mercurio (2001).

3 Single time step method for pricing on a grid

The two key elements in the development of a method to price interest rate derivatives in the BGM model by low-dimensional finite differences are:

- the forward rates process should be discretized by a single time step scheme; and
- the volatility structure should be separable, which permits the dynamics of the single time step forward rates process to be represented by a low-dimensional Markov process.

3.1 Justification of the above assumptions

Because the forward rates are approximated by a single-step scheme, the model will in general no longer be arbitrage-free. This timing inconsistency is addressed in Section 8, where it is shown that its impact is negligible for most cases. The single-step approximation is accurate enough for the pricing of derivatives, as shown numerically in Section 7. At the end of this section we introduce a novel discretization scheme based on the Brownian bridge that is especially designed for single time stepping. Its superiority (for single time steps only) over other discretizations is established in Section 4.

We proceed by first introducing notation for the single step-approximated forward rates process. This is followed by a statement of the separability assumption, after which we establish the low-dimensional Markov representation result. Single time step discretizations are then discussed, and we end by considering methods for pricing American style options with Monte Carlo methods.

3.2 Notation

We assume as given a time discretization $\tau_1 < \dots < \tau_J$. Define $Z_i(u,v) = \int_u^v \sigma_i(t) \cdot \mathrm{d}W^{N+1}(t)$. Given a scheme for the log rates

$$\log L_i(\tau_{j+1}) = \log L_i(\tau_j) + D_i\big(\tau_j, \tau_{j+1}, \mathbf{L}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1})\big) + Z_i(\tau_j, \tau_{j+1}) \qquad (5)$$

then denote by

$$L_i^{\mathrm{A}}(t) = L_i(0) \exp\big\{D_i\big(0, t, \mathbf{L}(0), \mathbf{Z}(0,t)\big) + Z_i(0,t)\big\}$$

its single time step-approximated equivalent. Here D stands for "drift approximation" and it is determined by the scheme applied, which may be the Euler, the predictor-corrector or the Brownian bridge scheme. These schemes will be elaborated on at the end of this section. The A in L^{A} stands for "approximated". The vector $\mathbf{Z}$ is defined by analogy with $\mathbf{L}$ in Equation (4).

3.3 Separability

Definition 1 (Separability) *A collection of instantaneous volatility functions $\sigma_i : [0, T_i] \to \mathbb{R}^d$, $i = 1, \dots, N$, is called "separable" if there exists a vector-valued* *function* $\sigma : [0, T] \to \mathbb{R}^d$ *and vectors* $\mathbf{v}_i \in \mathbb{R}^d$, $i = 1, \ldots, N$, *such that*

$$\sigma_i(t) = \mathbf{v}_i \sigma(t) \tag{6}$$

(no vector product; entry-by-entry multiplication) for $0 \leq t \leq T_i$, $i = 1, \ldots, N$.

Separability appears regularly in the context of requiring a process to be Markov. We mention three examples. First, we mention Ritchken and Sankarasubramanian (1995, Proposition 2.1). Working in the HJM model (Heath, Jarrow and Morton, 1992), they show that separability is a necessary and sufficient condition on the volatility structure such that the dynamics of the term structure may be represented by a two-dimensional Markov process. Second, we mention the Wiener chaos expansion framework of Hughston and Rafailidis (2002). In this framework any interest rate model is completely characterized by its so-called *Wiener chaos expansion*. The nth chaos expansion is represented by a function $\phi_n : \mathbb{R}_+^n \to \mathbb{R}$ that satisfies certain integrability conditions. If all ϕ_n are separable, the resulting interest rate model turns out to be Markov. Third, we mention the finite-dimensional Markov realizations for stochastic volatility forward rate models (see Björk, Landén and Svensson, 2002). Here a necessary condition for a stochastic volatility model to have a finite-dimensional Markov realization is that the drift term and each component of the volatility term in the Stratonovich representation of the short rate SDE should be a sum of functions that are separable in time to expiry and the stochastic volatility driver.

We give an example of a separable volatility function in the case of a one-factor model ($d = 1$).

Example (mean-reversion) Following De Jong, Driessen and Pelsser (2002), the instantaneous volatility may be specified as

$$\sigma_i(t) = \gamma_i \mathrm{e}^{-\kappa(T_i - t)} \tag{7}$$

The constant κ is usually referred to as the *mean-reversion parameter*.

3.4 Single time step method

The following proposition shows that a single time step plus separability yields low-dimensional representability.

PROPOSITION 1 *Suppose that* $\mathcal{M}$ *is a d-factor BGM model, for which the instantaneous volatility structure is separable. Then the single time step discretized forward rates process may be represented by a d-dimensional Markov process.*

PROOF Define the Markov process $\mathbf{X} : [0, T] \to \mathbb{R}^d$ by

$$\mathbf{X}(t) = \int_0^t \sigma(s)\, \mathrm{d}\mathbf{W}^{N+1}(s)$$

(entry-by-entry multiplication) where σ is as in Definition 1. Then the single time step process $\mathbf{L}^A : [0, T] \to (0, \infty)^{n-i(t)+1}$ at time t satisfies

$$L_i^A(t) = L_i(0) \exp\left\{D_i(0, t, \mathbf{L}(0), \mathbf{v}\mathbf{X}(t)) + \mathbf{v}_i \cdot \mathbf{X}(t)\right\} \tag{8}$$

Here D_i is defined implicitly by Equation (5) and $\mathbf{v}$ is a matrix of which row i is $\mathbf{v}_i$. The claim follows, bar a clarifying remark:

The second term in the exponent of Equation (8) is exactly equal to the stochastic part occurring in the BGM SDE (1), in virtue of the separability of the volatility structure:

$$\begin{aligned}\int_0^t \sigma_i(s) \cdot \mathrm{d}\mathbf{W}^{N+1}(s) &= \int_0^t \big(v_i \sigma(s)\big) \cdot \mathrm{d}\mathbf{W}^{N+1}(s) \\ &= \mathbf{v}_i \cdot \mathbf{X}(t)\end{aligned}$$

where the notation of Definition 1 has been used. □

Remark 1 The vector of single time-stepped rates may be considered (if separability holds) to be a time-dependent function of the Markov process $\mathbf{X}$, ie,

$$\mathbf{L}^A(t) = \mathbf{f}\big(t, \mathbf{X}(t)\big)$$

for some function $\mathbf{f}$. Hunt, Kennedy and Pelsser (2000, Theorem 1) showed that this is impossible to achieve for the true BGM forward rates themselves in the case when $\mathbf{X}$ is one-dimensional and under some technical restrictions.

Another essential building block for the fast single time step pricing framework is use of the terminal measure. This is explained in the following remark.

Remark 2 (Choice of numeraire) For the workings of the fast single time step pricing algorithm it is essential that the terminal measure be used. This is explained as follows. As proven in Proposition 1, the time-t single time-stepped forward rates are fully determined by $\mathbf{X}(t)$. This result holds for any choice of measure or numeraire. However, for the terminal numeraire, the value of the numeraire at time t is fully determined by the forward rate values at time t, but this does not hold in the case of, for example, the spot numeraire, in that the latter is generally determined by bond values observed at earlier times. The spot numeraire B_0 rolls its holdings over by the spot Libor account. Its time-T_i value is

$$B_0(T_i) = \frac{1}{\prod_{j=1}^{i} B_j(T_{j-1})}, \qquad T_0 := 0$$

Put in another way, the value of the spot numeraire is path-dependent, whereas that of the terminal numeraire is not. For pricing on a grid it is essential that the numeraire value is known given the value of $\mathbf{X}(t)$. Therefore the fast single time step framework requires the use of the terminal numeraire.

3.5 Valuation of interest rate derivatives with the single time step method

Interest rate derivatives with mild path-dependency may be valued by numerical integration, by a lattice/tree or by finite differences, provided that the single time-stepped rates are used and the separability assumption holds. The derivatives that may be valued include, but are not restricted to: caps, floors, European and Bermudan swaptions, trigger swaps and discrete barrier caps.

3.6 Discretizations

We discuss four time-discrete approximation schemes of the log BGM SDE (3):

❑ Euler;
❑ predictor-corrector;
❑ Milstein second-order scheme; and
❑ Brownian bridge.

The notation (Equation (5)) for a discretization of SDE (3) is recalled here:

$$\log L_i(\tau_{j+1}) = \log L_i(\tau_j) + D_i\big(\tau_j, \tau_{j+1}, \mathbf{L}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1})\big) + Z_i(\tau_j, \tau_{j+1})$$

We implicitly define $\tilde{D}$ by

$$D_i\big(\tau_j, \tau_{j+1}, \mathbf{L}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1})\big) = \tilde{D}_i\big(\tau_j, \tau_{j+1}, \mathbf{L}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1})\big) - \frac{1}{2}\int_{\tau_j}^{\tau_{j+1}} \|\sigma_i(s)\|^2 \, \mathrm{d}s$$

so as to remove the term common to the Euler, predictor-corrector and Brownian bridge discretizations.

3.6.1 Euler discretization

The Euler discretization (see, for example, Kloeden and Platen (1999, Equation (9.3.1))) sets

$$\tilde{D}_i\big(\tau_j, \tau_{j+1}, \mathbf{L}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1})\big) = -\left\{\sum_{k=i+1}^{N} \frac{\delta_k L_k(\tau_j)\,\sigma_k(\tau_j)\cdot\sigma_i(\tau_j)}{1+\delta_k L_k(\tau_j)}\right\}\big(\tau_{j+1}-\tau_j\big)$$

3.6.2 Predictor-corrector discretization

The predictor-corrector discretization was introduced to the setting of Libor market models by Hunter, Jäckel and Joshi (2001a). The key idea is to use predicted information to more accurately estimate the contribution of the drift to the increment of the log rate. For the terminal measure, an iterative procedure may be applied that loops from the terminal forward rate, N, to the spot Libor rate, $i(t)$. Initially, we set $\tilde{D}_N(\tau_j, \tau_{j+1}, \mathbf{L}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1})) = 0$. Then, for $i = N - 1, \ldots, i(t)$,

$$\tilde{D}_i\big(\tau_j, \tau_{j+1}, \mathbf{L}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1})\big) =$$

$$-\left\{\frac{1}{2}\sum_{k=i+1}^{N}\frac{\delta_k L_k(\tau_j)\,\sigma_k(\tau_j)\cdot\sigma_i(\tau_j)}{1+\delta_k L_k(\tau_j)} + \frac{1}{2}\sum_{k=i+1}^{N}\frac{\delta_k L_k(\tau_{j+1})\,\sigma_k(\tau_{j+1})\cdot\sigma_i(\tau_{j+1})}{1+\delta_k L_k(\tau_{j+1})}\right\}\times$$

$$\big(\tau_{j+1} - \tau_j\big)$$

with $L_k(\tau_{j+1})$ dependent on $L_m(\tau_j)$ and $Z_m(\tau_j, \tau_{j+1})$, $m = k + 1, \ldots, N$.

3.6.3 Milstein discretization

The second-order Milstein scheme (see, for example, Kloeden and Platen (1999, Equation (14.2.1))) was introduced to the setting of Libor market models in the series of papers by Glasserman and Merener (2003a,b and 2004). Moreover, these papers extended the convergence results to the case of jump–diffusion with thinning, which is key to the development of the jump–diffusion Libor market model. Also, these papers considered discretizations in various different sets of state variables, such as forward rates, log-forward rates, relative discount bond prices and log-relative discount bond prices. In Glasserman and Merener (2003b, 2004) it is shown numerically that the time-discretization bias of the log-Euler scheme is less than the bias of other discretizations, for example, in terms of the bonds. The results of Glasserman and Merener thus justify the log-type discretization (5) used in the present work.

The Milstein scheme can indeed be used to obtain a single time step discretization of the forward rates process – and hence it may be applied to the single time step pricing framework – but it is not particularly suited to single large time steps, as shown in the numerical comparisons for single time step accuracy in Section 4. Therefore we omit here the exact form of the scheme.

3.6.4 Brownian bridge discretization

Here we develop a novel discretization for the drift term. The idea is to calculate the expectation of the drift integral given the (time-changed) Wiener increment.

$$\tilde{D}_i\big(\tau_j, \tau_{j+1}, \mathbf{L}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1})\big) =$$

$$-\mathbb{E}^{N+1}\left[\int_{\tau_j}^{\tau_{j+1}}\sum_{k=i+1}^{N}\frac{\delta_k L_k(s)\,\sigma_k(s)\cdot\sigma_i(s)}{1+\delta_k L_k(s)}\,\middle|\,\mathcal{F}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1})\right] \tag{9}$$

The Brownian bridge discretization is superior when a single time step is applied.

This is shown theoretically and numerically in Section 4. Viewed as a numerical scheme for multi-step discretizations, it converges weakly with order one, as will be shown in the first part of Section 5. In the multi-step Monte Carlo numerical experiments of the second part of Section 5, we show that the bias is significantly less than for the Euler discretization.

In the remainder of this section, we first show how expression (9) can be calculated in practice, and, second, we establish that the Brownian bridge scheme has least-squares error (in a yet to be defined sense).

REMARK 3 (CALCULATION OF EXPRESSION (9)) In practice, expression (9) can be approximated with high accuracy. The calculation proceeds in four steps (it is indicated when a step contains an approximation):

Step 1 To calculate expression (9), the first step is to note that the order of the expectation and integral may be interchanged.

$$-\mathbb{E}^{N+1}\left[\int_{\tau_j}^{\tau_{j+1}} \sum_{k=i+1}^{N} \frac{\delta_k L_k(s)\,\sigma_k(s)\cdot\sigma_i(s)}{1+\delta_k L_k(s)}\,\mathrm{d}s \;\middle|\; \mathcal{F}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1})\right]$$

$$= -\int_{\tau_j}^{\tau_{j+1}} \mathbb{E}^{N+1}\left[\sum_{k=i+1}^{N} \frac{\delta_k L_k(s)\,\sigma_k(s)\cdot\sigma_i(s)}{1+\delta_k L_k(s)} \;\middle|\; \mathcal{F}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1})\right]\mathrm{d}s$$

This is a straightforward application of Fubini's theorem (see, for example, Williams (1991, Section 8.2)).

Step 2 (approximation) For the purposes of calculating the conditional expected value of expressions of the form $L/(1+\delta L)$, the forward rates are approximated with a single-step Euler discretization. Note that once this assumption has been made, the drift no longer affects the calculation. This stems from a property of the Brownian bridge: a Wiener process with deterministic drift conditioned to pass through a given point at some future time is always a Brownian bridge, independently of its drift prior to conditioning. Thus the estimation of the drift integral (9) is the same whether it is assumed that the forward rates are driftless or whether these follow a single time step Euler approximation.

$$-\int_{\tau_j}^{\tau_{j+1}} \mathbb{E}^{N+1}\left[\sum_{k=i+1}^{N} \frac{\delta_k L_k(s)\,\sigma_k(s)\cdot\sigma_i(s)}{1+\delta_k L_k(s)} \;\middle|\; \mathcal{F}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1})\right]\mathrm{d}s$$

$$\approx -\int_{\tau_j}^{\tau_{j+1}} \mathbb{E}^{N+1}\left[\sum_{k=i+1}^{N} \frac{\delta_k L_k^{\mathrm{BB}}\,\sigma_k\cdot\sigma_i}{1+\delta_k L_k^{\mathrm{BB}}} \;\middle|\; \mathcal{F}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1})\right]\mathrm{d}s$$

where BB indicates the use of the Brownian bridge, and where we have suppressed the dependence of time s.

Note that the assumption of singe-step Euler discretization for the calculation of expression (9) renders this calculation an approximation. In principle, the approximation could affect the quality of the discretization. We show numerically that this is not the case in the Libor-in-arrears case considered in Section 4.

Step 3 The conditional mean and conditional variance of the log forward rates are calculated. See Appendix A for details.

Step 4 (approximation) The drift expression (9) may be approximated by a single numerical integration over time; the expectation term is approximated by inserting the conditional mean of the forward rates process:[2]

$$-\int_{\tau_j}^{\tau_{j+1}} \mathbb{E}^{N+1}\left[\sum_{k=i+1}^{N} \frac{\delta_k L_k^{\mathrm{BB}} \sigma_k \cdot \sigma_i}{1+\delta_k L_k^{\mathrm{BB}}} \,\middle|\, \mathcal{F}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1}) \right] \mathrm{d}s$$

$$\approx -\int_{\tau_j}^{\tau_{j+1}} \sum_{k=i+1}^{N} \frac{\delta_k \mathbb{E}^{N+1}\left[L_k^{\mathrm{BB}} \middle| \mathcal{F}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1})\right] \sigma_k \cdot \sigma_i}{1+\delta_k \mathbb{E}^{N+1}\left[L_k^{\mathrm{BB}} \middle| \mathcal{F}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1})\right]} \mathrm{d}s$$

REMARK 4 If a two-point trapezoidal rule (ie, the average of the begin and end points) is used to evaluate the time integral in expression (9), the Brownian bridge reduces to the predictor-corrector scheme. In this sense, the predictor-corrector scheme is a special case of the Brownian bridge scheme.

We end this section with a discussion of the method used in this paper for pricing American-style options with Monte Carlo. The method used is the regression-based method of Longstaff and Schwartz (2001), which is a method of stochastic mesh type (see Broadie and Glasserman (2004)). Convergence of the method to the correct price follows generically from the asymptotic convergence property of stochastic mesh methods, as shown by Avramidis and Matzinger (2004).

4 The Brownian bridge scheme for single time steps

In this section, we establish theoretically and numerically that the Brownian bridge scheme has superior accuracy for single time steps.

[2] Alternatively, the expectation term could be evaluated by numerical integration as well, but this is computationally expensive. The full numerical integration ("BB alternative") has been compared numerically in Section 4 with the mean-insertion approximation ("BB"); the loss in accuracy is negligible on an absolute level. A theoretical error analysis of the mean-insertion approximation is given in Appendix B.

4.1 Theoretical result

Consider a stochastic differential equation of the form

$$\mathrm{d}X(t) = \mu\big(t, X(t)\big)\,\mathrm{d}t + \sigma(t)\,\mathrm{d}W(t) \tag{10}$$

Note that the BGM log SDE (3) is of the above form. We consider a certain class of discretizations:

Definition 2 *Let the function* $\bar{\mu}(\cdot,\cdot,\cdot)$ *denote a single time step discretization of SDE (10) with the following form:*

$$Y(\tau_{j+1}) = Y(\tau_j) + \bar{\mu}\big(\tau_j, Y(\tau_j), Z(\tau_j, \tau_{j+1})\big) + Z(\tau_j, \tau_{j+1}) \tag{11}$$

Here $Z(\tau_j, \tau_{j+1}) = \int_{\tau_i}^{\tau_{i+1}} \sigma(s)\,\mathrm{d}W(s)$. *Any such discretization is said to use information about the Gaussian increment to estimate the drift term.*

Note that Euler, predictor-corrector and Brownian bridge are such schemes. The next theorem states that, for the BGM setting, the Brownian bridge scheme (9) has least-squares error for a single time step over all discretizations that use information about the Gaussian increment for the drift term.

Lemma 1 *Let* $\{Y\}$ *be a single time step discretization of SDE (10) that uses information about the Gaussian increment for the drift term. Consider the discretization expected squared error*

$$S^2(\{Y\}) := \mathbb{E}\left[\Big(Y(\tau_{j+1}) - X_{\{\tau_j, Y(\tau_j)\}}\Big)^2 \,\Big|\, \mathcal{F}(\tau_j)\right]$$

Here $X_{\{t,x\}}$ *denotes the solution of SDE (10) starting from* (t, x). *Then the discretization* $\{Y^*\}$ *that yields least squared error,* S^2, *over all possible discretizations that use information about the Gaussian increment to estimate the drift term is defined by*

$$\begin{aligned}
&\bar{\mu}^*\Big(\tau_j, Y(\tau_j), Z(\tau_j, \tau_{j+1})\Big) \\
&\quad = \mathbb{E}\left[\int_{\tau_j}^{\tau_{j+1}} \mu\Big(s, X_{\{\tau_j, Y(\tau_j)\}}(s)\Big)\,\mathrm{d}s \,\middle|\, \mathcal{F}(\tau_j), \mathbf{Z}(\tau_j, \tau_{j+1})\right]
\end{aligned} \tag{12}$$

Proof Define

$$I := \int_{\tau_j}^{\tau_{j+1}} \mu\Big(s, X_{\{\tau_j, Y(\tau_j)\}}(s)\Big)\,\mathrm{d}s$$

For ease of exposition we write $Z = Z(\tau_j, \tau_{j+1})$ and $\bar{\mu} = \bar{\mu}(\tau_j, Y(\tau_j), Z)$, but we keep in mind that $\bar{\mu}$ is $\{\mathcal{F}(\tau_j), Z\}$-measurable. Also write $\mathbb{E}_t[\cdot] := \mathbb{E}[\cdot \mid \mathcal{F}(t)]$. Then let

$\{Y'\}$ with drift term $\bar{\mu}'$ be a discretization of the form of Definition 2. First, we condition on Z:

$$\begin{aligned}\mathbb{E}_{\tau_j}\big[\{\bar{\mu}' - I\}^2 \mid Z\big] &\geq \mathbb{E}_{\tau_j}\big[\{\mathbb{E}_{\tau_j}[I \mid Z] - I\}^2 \mid Z\big] \\ &= \mathbb{E}_{\tau_j}\big[\{\bar{\mu}^* - I\}^2 \mid Z\big]\end{aligned}$$

The inequality holds since expectation equals projection, and the latter has, by definition, least squared error over all possible $\{\mathcal{F}(\tau_j), Z\}$-measurable drift terms. Continuing, we find

$$\begin{aligned}S^2(\{Y'\}) = \mathbb{E}_{\tau_j}\big[\{\bar{\mu}' - I\}^2\big] &= \mathbb{E}_{\tau_j}\Big[\mathbb{E}_{\tau_j}\big[\{\bar{\mu}' - I\}^2 \mid Z\big]\Big] \\ &\geq \mathbb{E}_{\tau_j}\Big[\mathbb{E}_{\tau_j}\big[\{\bar{\mu}^* - I\}^2 \mid Z\big]\Big] = S^2(\{Y^*\})\end{aligned}$$

ie, Y^* has less squared error than Y'. As Y' was an arbitrary discretization of the form of Definition 2, the result follows. □

4.2 Libor-in-arrears case

We estimate numerically the accuracy in the Libor-in-arrears test of the various schemes of Section 3. We extend here the Libor-in-arrears test of Hunter, Jäckel and Joshi (2001a) by including the Milstein and Brownian bridge schemes. The test is designed to measure the accuracy of a single time step discretization. The idea of the test is briefly described here; for details the reader is referred to Hunter, Jäckel and Joshi (2001a).

Consider the distribution of a forward rate under the measure associated with the numeraire of a discount bond maturing at the *fixing* time of the forward. Note that the forward rate is not a martingale under such a measure as the natural payment time of the forward is not the same as its fixing time. An analytical formula for the associated density, however, is known. We can thus compare the density obtained from a single time step discretization with the analytical formula for the density. The results of this test are displayed in Figure 1. It is shown (for the particular set-up) that the Brownian bridge scheme reduces the maximum error in the density by a factor 100 over the predictor-corrector scheme.

5 The Brownian bridge scheme for multi-time step Monte Carlo

This section consists of two parts. First, we show theoretically that the Brownian bridge scheme converges weakly with order one. Second, we estimate numerically the convergence behavior of the various schemes of Section 3.

In a financial context, the interest lies in calculating the prices of derivatives, which are in certain cases expectations of payoff functions. Therefore we are interested mainly in *weak convergence* of Monte Carlo simulations. The definition is recalled here and may be found in, for example, Kloeden and Platen (1999, Section 9.7).

FIGURE 1 Plots of the estimated densities and error in densities of various single time step discretizations.

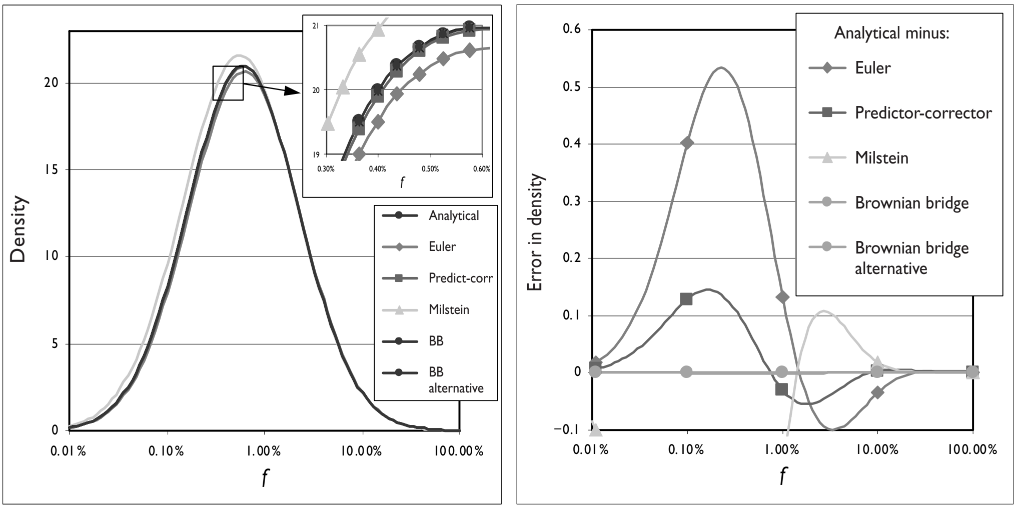

The deal set-up is the same as in Hunter, Jäckel and Joshi (2001a); the three-month forward rate fixing 30 years from today is set initially to 8% and its volatility to 24%. The legend "BB" denotes Brownian bridge and "BB alternative" denotes full numerical integration of the expectation term. Note that three densities have been added to the above figures compared with Figure 1 of Hunter, Jäckel and Joshi (2001a): Milstein and the two Brownian bridge schemes. On both figures, however, the differences between the analytical and Brownian bridge densities are indiscernible to the eye. The most notable addition is the Milstein density. Outside of the error graph, the Milstein scheme reaches a maximum absolute error that is around twice the maximum absolute error for the Euler scheme. The maximum absolute error in the density for the Brownian bridge and its alternative are 10^{-3} and 6×10^{-4}, respectively. In this particular test the Brownian bridge scheme thus achieves a reduction by a factor of 100 in the maximum absolute error over the predictor-corrector scheme, the latter being the second best scheme.

Definition 3 (Weak convergence) *A scheme* $\{Y^{\varepsilon}(\tau_j)\}$ *with maximum step size* ε *is said to convergence weakly with order* β *to X if, for each function g with* $2(\beta + 1)$ *polynomially bounded derivatives, there exists a constant C such that, for sufficiently small* ε,

$$\left| \mathbb{E}\left[g(X(T)) \right] - \mathbb{E}\left[g(Y^{\varepsilon}(T)) \right] \right| \leq C \cdot \varepsilon^{\beta} \tag{13}$$

A criterion that is easier to verify than the above definition is the concept of *weak consistency*, and under quite natural conditions it follows that weak consistency implies weak convergence. The definition of weak consistency is recalled here, and may be found for example on page 327 of Kloeden and Platen (1999). Here we develop the remainder of the theory in terms of approximating an autonomous one-dimensional SDE, say,

$$\mathrm{d}X(t) = a\big(X(t)\big)\,\mathrm{d}t + b\big(X(t)\big)\,\mathrm{d}W(t), \qquad X(0) \text{ deterministic} \tag{14}$$

However, the theory holds in more general cases too.

DEFINITION 4 (WEAK CONSISTENCY) *A scheme* $\{Y^{\varepsilon}(\tau_j)\}$ *with maximum step size* ε *is weakly consistent if there exists a function* $c = c(\varepsilon)$ *with*

$$\lim_{\varepsilon \downarrow 0} c(\varepsilon) = 0 \tag{15}$$

such that

$$\mathbb{E}\left[\left|\mathbb{E}\left[\frac{Y^{\varepsilon}(\tau_{j+1}) - Y^{\varepsilon}(\tau_j)}{\Delta\tau_j} \,\Big|\, \mathcal{F}(\tau_j)\right] - a\big(Y^{\varepsilon}(\tau_j)\big)\right|^2\right] \le c(\varepsilon) \tag{16}$$

and

$$\mathbb{E}\left[\left|\mathbb{E}\left[\frac{1}{\Delta\tau_j}\left\{Y^{\varepsilon}(\tau_{j+1}) - Y^{\varepsilon}(\tau_j)\right\}\left\{Y^{\varepsilon}(\tau_{j+1}) - Y^{\varepsilon}(\tau_j)\right\} \,\Big|\, \mathcal{F}(\tau_j)\right] - b\big(Y^{\varepsilon}(\tau_j)\big)\, b\,\big(Y^{\varepsilon}(\tau_j)\big)\right|^2\right] \le c(\varepsilon) \tag{17}$$

Here $\{\mathcal{F}(t)\}$ *is the filtration generated by the Brownian motion driving SDE (14).*

Kloeden and Platen prove the following theorem (see Theorem 9.7.4 of Kloeden and Platen (1999)) linking weak consistency to weak convergence.

THEOREM 1 (LINKING WEAK CONSISTENCY TO WEAK CONVERGENCE) *Suppose that* a *and* b *of Equation (14) are four times continuously differentiable with polynomial growth and uniformly bounded derivatives. Let* $\{Y^{\varepsilon}(\tau_j)\}$ *be a weakly consistent scheme with equitemporal steps* $\Delta\tau_j = \varepsilon$ *and initial value* $Y^{\varepsilon}(0) = X(0)$ *which satisfies the moment bounds*

$$\mathbb{E}\left[\max_j \left|Y^{\varepsilon}(\tau_j)\right|^{2q}\right] \le K\left(1 + \left|X(0)\right|^{2q}\right), \qquad q = 1, 2, \ldots$$

and

$$\mathbb{E}\left[\frac{1}{\varepsilon}\left|Y^{\varepsilon}(\tau_{j+1}) - Y^{\varepsilon}(\tau_j)\right|^6\right] \le c(\varepsilon) \tag{18}$$

where $c(\varepsilon)$ *is as in Definition 4. Then* Y^{ε} *converges weakly to* X.

In the proposition below we show that the Brownian bridge scheme with the proposed calculation method is weakly consistent. The above theorem then allows us to deduce that the Brownian bridge scheme converges weakly.

PROPOSITION 2 (BROWNIAN BRIDGE SCHEME IS WEAKLY CONSISTENT) *Assume that the volatility functions* $\sigma_i(\cdot)$ *are piecewise analytical on the model horizon* $[0, T]$. *Then the Brownian bridge scheme defined by Equation (9) and by the four-step calculation method described in Remark 3 is weakly consistent with the forward rates process defined in Equation (3).*

PROOF Without loss of generality, we may assume that the volatility functions are analytical. Otherwise, due to the piecewise property of the volatility functions, we can break up the problem into sub-problems for which each has analytical volatility functions. Note also that all derivatives of the volatility functions are bounded because the interval $[0, T]$ is compact.

We need only verify the consistency Equation (16) for the drift term. To achieve this, define for i and for all $\tau \in [0, T]$ and for all $\mathbf{L}$ the function $f_{\{i,\tau,\mathbf{L}\}} : [0, T - \tau] \to \mathbb{R}$:

$$f_{\{i,\tau,\mathbf{L}\}}(t) = -\sum_{k=i+1}^{N} \frac{\delta_k L_k}{1 + \delta_k L_k} \int_0^t \sigma_k(\tau + s) \cdot \sigma_i(\tau + s)\,\mathrm{d}s$$

Due to the assumption that the volatility functions are analytical, it follows that the function $f_{\{i,\tau,\mathbf{L}\}}$ is analytical in t. Taylor's formula states that there exists an error term $E_{\{i,\tau,\mathbf{L}\}}(\cdot)$ depending on i, τ and $\mathbf{L}$ such that

$$f_{\{i,\tau,\mathbf{L}\}}(t) = f_{\{i,\tau,\mathbf{L}\}}(0) + t\frac{\partial f_{\{i,\tau,\mathbf{L}\}}}{\partial t}(0) + E_{\{i,\tau,\mathbf{L}\}}(t) \tag{19}$$

with

$$\lim_{t \downarrow 0} \frac{\left|E_{\{i,\tau,\mathbf{L}\}}(t)\right|}{t^2} < \infty \tag{20}$$

Due to the analyticity, boundedness and limiting behavior of the function $h(x) = x/(1 + x)$, namely $h \uparrow 1$ $(h \downarrow 0)$ as $x \to \infty$ $(x \to -\infty$, respectively), we have that all its derivatives are bounded. Viewed as a function $[0, T] \times [0, T] \times \mathbb{R}^N \to \mathbb{R}$,

$$(t, \tau, \mathbf{L}) \mapsto f_{\{i,\tau,\mathbf{L}\}}(t)$$

We can thus find a bound on the second derivative, $\partial^2 f_{\{i,\tau,\mathbf{L}\}}/\partial t^2$, independent of $(\tau, \mathbf{L})$. Theorem 7.7 of Apostol (1967) then states that the error term of Equation (19) may be chosen independently of τ and $\mathbf{L}$. Hence we find that

$$f_{\{i,\tau,\mathbf{L}\}}(t) = t\left\{-\sum_{k=i+1}^{N} \sigma_k(\tau)\cdot\sigma_i(\tau)\frac{\delta_k L_k}{1+\delta_k L_k}\right\} + E(t)$$

with E satisfying the second-order Equation (20). Here we have used

$$f_{\{i,\tau,\mathbf{L}\}}(0) = 0 \quad \text{and}$$

$$\left.\frac{\partial f_{\{i,\tau,\mathbf{L}\}}}{\partial t}\right|_{t=0} = \left\{-\sum_{k=i+1}^{N} \sigma_k(\tau)\cdot\sigma_i(\tau)\frac{\delta_k L_k}{1+\delta_k L_k}\right\}$$

If $\mathbf{Y}^{\varepsilon}$ denotes the Brownian bridge scheme, then

$$\begin{aligned}\mathbb{E}\left[Y_i^{\varepsilon}(\tau_{j+1}) - Y_i^{\varepsilon}(\tau_j)\,\middle|\,\mathcal{F}(\tau_j)\right] &= f_{\{i,\tau_j,\mathbf{Y}^{\varepsilon}(\tau_j)\}}(\varepsilon)\\ &= \varepsilon\left\{-\sum_{k=i+1}^{N}\frac{\delta_k Y_k^{\varepsilon}(\tau_j)\,\sigma_k(\tau_j)\cdot\sigma_i(\tau_j)}{1+\delta_k Y_k^{\varepsilon}(\tau_j)}\right\} + E(\varepsilon)\end{aligned}$$

Note that the term within braces is exactly drift term i evaluated at $(\tau_j, \mathbf{Y}^{\varepsilon}(\tau_j))$. It follows that consistency Equation (16) holds with $c(\varepsilon)$ equal to $(E(\varepsilon)/\varepsilon)^2$. The function $c(\cdot)$ is then quadratic in ε. □

Corollary 1 (Brownian bridge scheme converges weakly with order one) *Under the assumptions of Proposition 2, the Brownian bridge scheme defined by Equation (9) and by the four-step calculation method described in Remark 3 converges weakly to the forward rates process defined in Equation (3). It has order of convergence one.*

Proof We only need verify the claim with regards to the order of convergence. In the proof of Theorem 1 in Kloeden and Platen (1999), it is shown that the error term in the weak convergence criterion (13) is less than $\sqrt{c(\varepsilon)}$, with $c(\cdot)$ satisfying the requirements (15), (16), (17) and (18). All these requirements can be met for the Brownian bridge scheme with a quadratic function c. Taking the square root then yields first-order weak convergence for the Brownian bridge scheme. □

5.1 Numerical results

We now turn to the second part of Section 5, in which the various discretization schemes are compared numerically. A floating leg and a cap were valued with 10 million simulation paths. This large number of paths was used because the time discretization bias for the log rates is small compared to the standard error often observed with 10,000 paths. For example, the Euler one-step-per-accrual

FIGURE 2 Plots of the estimated biases for a floating leg and a cap for the Euler, predictor-corrector, Milstein and Brownian bridge schemes.

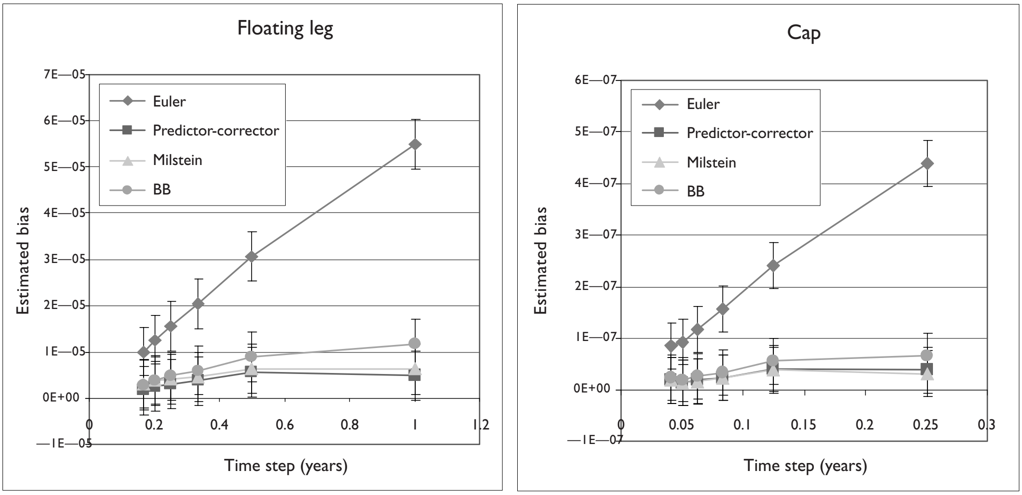

A single-factor model was applied. The floating leg is a six-year deal, with the fixings at 1,…, 5 years, and payments of annual Libor at 2,…, 6 years. The cap is a 1.5-year deal, with the fixings at 0.25, 0.5,…, 1.25 years, and payments of quarterly Libor above 5% (if at all) at 0.5, 0.75,…, 1.5 years. The market conditions are the same for both deals: all initial forward rates are 6%, and all volatility is constant at 20%. The net present values of the floating leg and cap are 0.24 and 0.013, respectively, on a notional of one unit of currency. The error bars denote a 95% confidence bound based on twice the sample standard error.

discretization relative bias for the floating leg and the cap was estimated at 0.02% and 0.003%, whereas twice the standard error at 10,000 paths is 0.07% and 0.01%, respectively.

To filter out the time discretization bias from the simulation standard error we reduce the latter by simultaneously simulating the prices under the respective forward measures. Under the forward measure, there is no drift term and the Euler log-scheme solves the stochastic differential equation without time discretization error; in such a way unbiased prices are obtained. The standard error of the simulated bias is then a measure of its accuracy. Because the correlation between the discounted payoff under the terminal and the forward measure is high, the standard error will be lower than the analytical value of the contract.

The results are presented in Figure 2. They show that the predictor-corrector, Milstein and Brownian bridge schemes have a time discretization bias that is hardly distinguishable from the standard error of the estimate. The Euler scheme, however, has a clear time discretization bias for larger time steps. We classify the schemes from best suited to worst suited (for the particular numerical cases under consideration) using the criterion of the minimal computational time required to achieve a bias that is indistinguishable from the standard error at 10,000,000 paths. As Milstein is slightly faster than predictor-corrector, which in turn is faster than the Brownian bridge, we obtain: first, Milstein; second, predictor-corrector; third, Brownian bridge; and fourth, Euler. We stress here that this classification might be particular to the numerical cases that we considered. We also stress that the strength of the Brownian bridge lies in single time steps rather than in multi-time steps.

6 Example: one-factor drift-approximated BGM framework

This section illustrates the framework for fast single time step pricing in BGM by setting it up in the special case of a one-factor model with a volatility structure as in the example in Section 3.3. This structure may be written as follows:

$$\sigma_i(t) = \tilde{\gamma}_i \, \mathrm{e}^{\kappa t}$$

for certain constants $\tilde{\gamma}_i$. The corresponding Markov factor, X, is then defined as and characterized by

$$X(t) = \int_0^t \mathrm{e}^{\kappa s} \mathrm{d}W(s), \qquad X(t) \sim \mathcal{N}\big(0, \Sigma^2(t)\big)$$

where

$$\Sigma^2(t) = \int_0^t \mathrm{e}^{2\kappa s} \mathrm{d}s = \begin{cases} \dfrac{\mathrm{e}^{2\kappa t} - 1}{2\kappa}, & k \neq 0 \\ t, & k = 0 \end{cases}$$

Prices may now be computed by either numerical integration or finite differences. In the case of numerical integration, if $\Pi(t, X)$ denotes the numeraire-deflated value of the contingent claim, we have

$$\Pi\big(0, X(0)\big) = \int_{-\infty}^{\infty} \Pi(t, x) \, p\big(x; 0, \Sigma^2(t)\big) \, \mathrm{d}x$$

where t denotes the expiry of the contingent claim and $p(\cdot\,; \mu, \Sigma^2)$ denotes the Gaussian density with mean μ and standard deviation Σ. In case of finite differences, Feynman–Kac yields the following PDE for the price relative to the terminal bond:

$$\frac{\partial \Pi}{\partial t} + \frac{1}{2} \mathrm{e}^{2\kappa t} \frac{\partial^2 \Pi}{\partial X^2} = 0 \tag{21}$$

with use of appropriate boundary conditions. For example, for a Bermudan payer swaption we have $\Pi(\cdot, -\infty) \equiv 0$, zero convexity $\partial^2 \Pi / \partial X^2 \equiv 0$ at $X = \infty$, and exercise boundary conditions at the exercise times.

TABLE 1 A simple numerical example.

i	(I) $L_i(0)$	(II) $\mu_i(0)$	(III) $-\frac{1}{2}\tilde{\gamma}_i^2\Sigma^2(1)$	(IV) $\tilde{\gamma}_i X(1)$	(V) Drift frozen $L_i(1)$	(VI) Equation $(9)_{i-1}-(9)_i$	(VII) Brownian bridge $L_i(1)$
5	7.00%	0.00000	−0.03644	0.25	8.67%	−0.00569	8.67%
4	7.00%	−0.00409	−0.03644	0.25	8.63%	−0.00567	8.62%
3	7.00%	−0.00818	−0.03644	0.25	8.60%	−0.00564	8.57%
2	7.00%	−0.01227	−0.03644	0.25	8.56%	−0.00562	8.53%
1	7.00%	−0.01636	−0.03644	0.25	8.53%		8.47%

6.1 A simple numerical example

We will evolve five annual ($\delta_i = 1$) forward rates over a one-year period. Forward rate i accrues from year i until year $i + 1$, $i = 1,\dots, 5$. Take $L_i(0) = 7\%$, $\tilde{\gamma}_i = 25\%$ and $\kappa = 15\%$; then $\Sigma^2(1) \approx 1.166196$. Suppose that, after one year, the process X jumps to 1; thus $X(1) = 1$. All computations are displayed in Table 1. Column (II) is determined by Equation (2). To evaluate the effect of the Brownian bridge scheme over the Euler scheme, the "drift-frozen" forward rates (where the drift is evaluated at time zero) are displayed in column (V), using the equation (V) = (I) exp ((II) + (III) + (IV)). Then, we start with computing the Brownian bridge scheme forward rate 5 and work back to forward rate 1. Forward rate 5 is easily computed as no drift terms are involved. To compute the drift term integral at time 1 for forward rate 4, we compute the drift term integral of Equation (9) for forward rate 5. The result is displayed in column (VI). This we may then use to compute the Brownian bridge scheme forward rate 4 (see column (VII)), where we use the equation $(\mathrm{VII})_i = (\mathrm{I})\exp(\{\sum_{j=i+1}^{N}(\mathrm{VI})_j\} + (\mathrm{III}) + (\mathrm{IV}))$. Continuing, we compute the drift for forward rate 3 using only the Brownian bridge forward rates 4 and 5. And so on until all forward rates have been computed.

7 Example: Bermudan swaption

As an example of the single time step pricing framework, an analysis is made for Bermudan swaptions in comparison with a BGM model combined with the least-squares Monte Carlo method introduced by Longstaff and Schwartz (2001). The one-factor set-up introduced in the previous section was used with zero mean-reversion.

Callable Bermudan and European payer swaptions were priced in a one-factor BGM model for various tenors and non-call periods. The zero rates were taken to be flat at 5%, and the volatility of the forwards was set flat at 15%. The Bermudans were priced on a grid, the Europeans through numerical integration. The PDE was solved using an explicit finite-difference scheme. The explanatory

TABLE 2 Specification of the Bermudan swaption comparison deal.

Callable Bermudan swaption	
Market data	
Zero rates	Flat at 5%
Volatility	Flat at 15%
Product specification	
Tenor	Variable (2–8 years)
Non-call period	Variable
Call dates	Semi-annual
Pay/receive	Pay fixed
Fixed leg properties	
Frequency	Semi-annual
Date roll	None
Day count	Half year = 0.5
Fixed rate	5.06978% (ATM)
Floating leg properties	
Frequency	Semi-annual
Date roll	None
Day count	Half year = 0.5
Margin	0%
Numerics	
Simulation paths	10,000
Finite-difference scheme	Explicit
Longstaff–Schwartz	
Explanatory variable	Swap NPV
Basis function type	Monomials
No. of basis functions	Two (constant and linear)

variable in the least-squares Monte Carlo was taken to be the NPV of the underlying swap. This was regressed on to a constant and a linear term. These two basis functions yield sufficiently accurate results because the value of a Bermudan swaption increases almost linearly with the value of the underlying swap.

Problems may possibly occur for American-style derivatives in the single time step framework. Since the framework is not arbitrage-free, spurious early or delayed exercise may take place to collect the arbitrage opportunity. The effects of these phenomena have been analyzed by comparing the exercise boundaries[3] and risk sensitivities of Longstaff–Schwartz and single time step BGM. In both

[3] In the Longstaff–Schwartz case, the future discounted cashflows are regressed against the NPV of the underlying swap with a constant and linear term – say, with coefficients a and b. So the option is exercised whenever $S > a + bS \Leftrightarrow S > a/(1 - b) =: S^*$, where it is assumed that $b < 1$, which turns out to hold in practice. Hence the exercise boundary S^* may be computed from the regression coefficients by the above formula.

TABLE 3 Results of the Bermudan swaption comparison deal. The notation $XNCY$ in the first column denotes an X-year underlying swap with a non-call period of Y years. In the case of a European swaption, it means that the swaption is exercisable after Y years exactly. All prices and standard errors are in basis points.

	Drift-approx. BGM	Bermudan Longstaff–Schwartz	Standard error	Drift-approx. BGM	European Monte Carlo BGM	Standard error
2NC1	29.40	28.85	0.42	27.36	26.88	0.43
3NC1	64.33	62.78	0.83	53.78	52.92	0.83
4NC1	101.66	101.51	1.29	78.04	78.77	1.24
4NC3	44.09	43.59	0.70	42.93	42.55	0.71
5NC1	141.22	137.95	1.68	100.85	99.31	1.55
5NC3	89.25	86.75	1.34	83.08	80.83	1.36
6NC1	182.16	179.48	2.22	122.27	123.36	1.92
6NC3	134.88	136.43	2.01	120.60	123.06	2.03
6NC5	50.93	50.79	0.86	50.07	50.09	0.87
7NC1	224.40	221.38	2.61	142.93	140.66	2.19
7NC3	181.20	177.11	2.53	156.15	153.71	2.53
7NC5	101.84	100.59	1.64	97.28	96.57	1.65
8NC1	266.63	266.35	3.15	159.38	161.00	2.50
8NC3	226.55	226.94	3.14	185.20	190.98	3.08
8NC5	151.23	151.13	2.38	137.73	140.95	2.41
8NC7	54.20	53.70	0.96	52.38	53.12	0.96

models the exercise rule turned out to be of the following form: exercise whenever the NPV of the underlying swap, S, is larger than a certain value S^*, which is then defined to be the *exercise boundary*.

For a full description of the deal see Table 2. Results have been summarized in Table 3. Computational times may be found in Table 4. Exercise boundaries for the 8 year deal are displayed in Figure 3, including confidence bounds on the Longstaff–Schwartz boundaries.[4] We looked at exercise boundaries for other deals

[4] The empirical covariance matrix of the regression-estimated coefficients a and b may be used to obtain the empirical variance of S^*. Denote random errors in a and b by ϵ_a and ϵ_b, respectively. If it is assumed that these errors are relatively small, a Taylor expansion yields (ignoring second-order terms)

$$S^* \approx \frac{a}{1-b}\left(1 + \frac{\epsilon_a}{a} + \frac{\epsilon_b}{1-b}\right)$$

We thus obtain the empirical variance of S^* (as well as its standard error). Assuming that S^* is normally distributed, a 95% confidence interval is given by plus and minus twice the standard error.

TABLE 4 Computational times for the Bermudan swaption comparison deal on a computer with a 700MHz processor. The notation XNCY in the first column denotes an X-year underlying swap with a non-call period of Y years. In the single time step framework Bermudans are priced on a grid and Europeans are priced through numerical integration. All computational times are in seconds.

	Bermudan		European	
	Drift-approximated BGM	Longstaff Schwartz	Drift-approximated BGM	Monte Carlo BGM
2NC1	0.4	3.0	0.0	1.9
3NC1	0.4	6.6	0.1	3.7
4NC1	0.7	11.1	0.2	6.1
4NC3	0.2	4.5	0.1	3.4
5NC1	1.4	17.3	0.6	9.1
5NC3	0.3	9.0	0.1	6.2
6NC1	2.4	24.5	0.6	12.8
6NC3	0.7	14.6	0.2	9.8
6NC5	0.2	5.8	0.0	4.8
7NC1	4.0	33.1	0.8	16.8
7NC3	1.4	21.2	0.4	13.5
7NC5	0.3	11.4	0.2	8.6
8NC1	5.6	45.9	1.2	23.9
8NC3	2.2	30.2	0.6	18.8
8NC5	0.6	18.4	0.2	13.5
8NC7	0.1	7.4	0.0	7.8

TABLE 5 BGM pricing simulation re-run for 500,000 paths using pre-computed exercise boundaries. The standard errors for both prices were virtually the same in all cases, so only a single standard error is reported. All prices and standard errors are in basis points.

	BGM simulation price		
	LS pre-computed exercise boundaries	DA pre-computed exercise boundaries	Standard error
2NC1	28.63	28.62	0.06
3NC1	62.80	62.77	0.12
4NC1	99.51	99.58	0.18
5NC1	138.38	138.55	0.24
6NC1	178.08	179.41	0.30
7NC1	221.51	222.49	0.36
8NC1	263.05	265.27	0.42

FIGURE 3 Exercise boundaries for the eight-year deal.

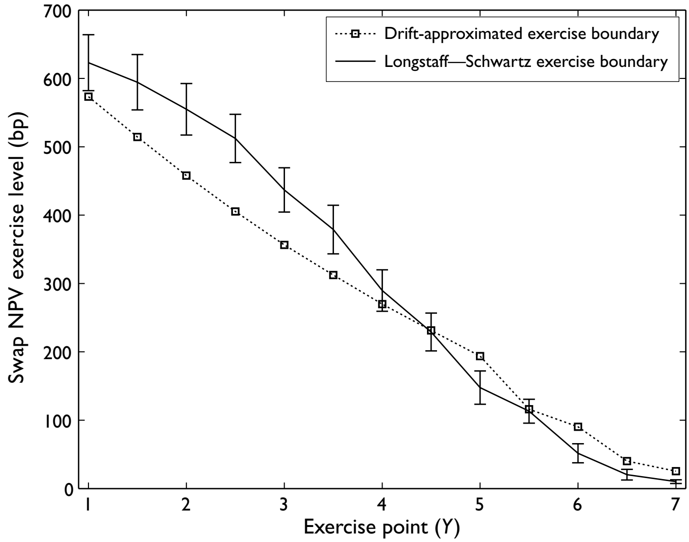

as well and these revealed a similar picture. Risk sensitivities for the various deals are displayed in Figure 4.

The results show that the single time step BGM pricing framework indeed prices the Bermudan swaptions close to Longstaff–Schwartz, including correct estimates of risk sensitivities for shorter-maturity deals. In all cases the price difference is within twice the standard error of the simulation. Moreover, the computational time involved is less by a factor 10. Note that the exercise boundary is calculated slightly differently by the Longstaff–Schwartz and drift-approximated (DA) approach. Also, risk sensitivities for longer-maturity deals (seven to eight years) can be outside of the two-standard-error confidence bound. The Brownian bridge drift approximation thus becomes worse for longer-maturity deals, as also explained in Section 8. To determine which approach computed the best exercise boundaries, the BGM pricing simulation was re-run for 500,000 paths using the pre-computed exercise boundaries. The results, given in Table 5, show that the drift-approximated exercise boundaries are not worse than their Longstaff–Schwartz counterparts and are even slightly better.[5] Hence there is no problem with the spurious early exercise opportunities arising from the absence of no-arbitrage in the fast single time step framework. The non-arbitrage-free issue is investigated further in the next section. This section ends with the results for a two-factor model.

[5] This does not necessarily mean that the DA framework outperforms Longstaff–Schwartz because we only regress on the NPV of the underlying swap. Longstaff–Schwartz may possibly yield better exercise boundaries when it is regressed on to more explanatory variables.

FIGURE 4 Risk sensitivities: deltas and vegas with respect to a parallel shift in the zero rates and caplet volatilities, respectively.

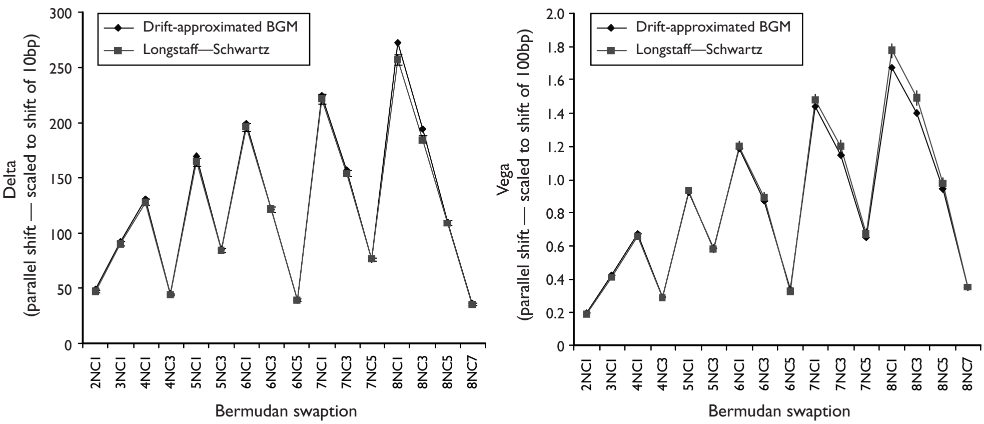

The error bars for the Longstaff–Schwartz prices represent a 95% confidence bound based on twice the empirical standard error.

7.1 Two-factor model

We consider a two-factor model with the same set-up as above with the exception of the volatility structure, which we now take to be

$$\frac{\mathrm{d}L_i(t)}{L_i(t)} = v_{i,1}\mathrm{d}W_1^{i+1}(t) + v_{i,2}\mathrm{d}W_2^{i+1}(t)$$

Here $|\mathbf{v}_i| = 15\%$. For a model with forward expiry structure $T_1 < \ldots < T_N$, we take the $\mathbf{v}_i \in \mathbb{R}^2$ to be

$$\mathbf{v}_i = (15\%)\left(a_i, \sqrt{1-a_i^2}\right), \quad a_i = \frac{T_i - T_1}{T_N - T_1}$$

This instantaneous volatility structure is purely hypothetical. It has the property that correlation steadily drops between more separated forward rates. To solve the two-dimensional PDE version of Equation (21) we used the hopscotch method (see paragraph 48.5 of Wilmott (1998)). Results for the two-factor model are displayed in Table 6. In a two-factor model (with de-correlation) the exercise decision no longer depends only on the NPV of the underlying swap but also on all forward swap rates. We therefore take the results with regression on all forward swap rates to be the benchmark. Indeed, the drift-approximated prices agree more with the benchmark than with prices obtained when Longstaff–Schwartz regresses on the NPV of a single swap. The computational time for the fast drift-approximated pricing two-dimensional grid was, on average, only a quarter of the computational time for Monte Carlo.

TABLE 6 Two-factor model comparison. 50,000 paths were used for the Longstaff–Schwartz simulation. "Swap NPV only" and "All forward rates" indicate that Longstaff–Schwartz regressed on only the NPV of the swap and on all forward swap rates, respectively. All prices and standard errors are in basis points.

	Fast drift approximation	Longstaff–Schwartz Swap NPV only	Longstaff–Schwartz All forward rates (benchmark)	Standard error
2NC1	25.45	23.27	24.64	0.2
3NC1	59.22	55.79	58.08	0.3
4NC1	94.67	89.54	93.00	0.5
5NC1	132.35	124.79	129.42	0.7
6NC1	171.41	162.89	169.76	0.9
7NC1	212.15	202.97	210.89	1.1
8NC1	252.49	242.59	251.88	1.3
9NC1	292.62	283.89	294.68	1.5

8 Test of accuracy of drift approximation

Besides the approximation of the drift, the framework (Proposition 1) contains a timing inconsistency. The inconsistency is best described by an example (see Figure 5). Suppose that the underlying Markov process $\mathbf{X}$ jumps to $\mathbf{X}(2)$, say, in two years. Consider computing the value of the forwards at year 2. We could jump immediately to year 2 and calculate the forwards there. Alternatively, we could consider first calculating the forwards at time 1 (under the assumption that $\mathbf{X}$ jumps to some value $\mathbf{X}(1)$) and from this point calculate the forwards at time 2 (assuming that $\mathbf{X}$ then jumps to the very same $\mathbf{X}(2)$). In general, the so computed forwards at time 2 will be different.

In a way, any low-dimensional approximation of BGM will exhibit this timing inconsistency. Consider the following. Given the value of $\mathbf{X}(t)$, we cannot determine all time-t forward rates. We do, however, know the value of $L_N(t)$ because

FIGURE 5 Timing inconsistency in the single time step framework for BGM.

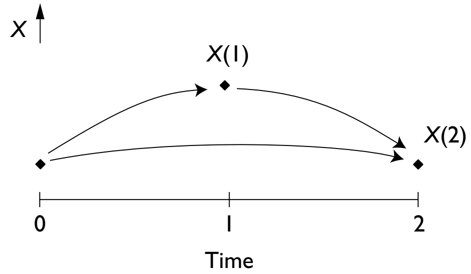

L_N has zero drift under the terminal measure $N+1$. The value of any other forward rate $L_i(t)$ does not depend solely on the value of $\mathbf{X}(t)$ but is dependent on the whole path that $\mathbf{X}$ traversed on the interval $[0,t]$. The framework for fast single time step pricing simply calculates the most likely value of $L_i(t)$ given the value of $\mathbf{X}(t)$. If we start from a different initial model state (for example, if we start from the state determined by $\mathbf{X}(1)$), then almost surely our guess for the most likely value of $L_i(t)$ will be different. In this way, it is not really fair to consider this timing inconsistency, but we will nonetheless investigate it. In the following, a test will be proposed to evaluate the size of the inconsistency error.

8.1 Test of accuracy of drift approximation based on no-arbitrage

The accuracy test is described by an example. Consider some time T at which forwards $i,\dots,N$ have not yet expired. The framework for fast drift-approximated pricing yields time-T forward rates as a function of $\mathbf{X}(T)$. Under the assumptions that the model state is determined by the Markov process $\mathbf{X}$, and that the framework is arbitrage-free, the fundamental arbitrage-free pricing formula will yield values of forward rates at time $t<T$ *as a function of* $\mathbf{X}(t)$ given by the following formula:[6]

$$\begin{aligned} L_i^{\mathrm{AF}}(t,\mathbf{x}) &= \frac{1}{\delta_i}\left\{\frac{B_i^{\mathrm{AF}}(t)\big/B_{N+1}^{\mathrm{AF}}(t)}{B_{i+1}^{\mathrm{AF}}(t)\big/B_{N+1}^{\mathrm{AF}}(t)}-1\right\} \\ &= \frac{1}{\delta_i}\left\{\frac{\mathbb{E}^{N+1}\left[\dfrac{B_i^{\mathrm{DA}}(T)}{B_{N+1}^{\mathrm{DA}}(T)}\,\middle|\,\mathbf{X}(t)=\mathbf{x}\right]}{\mathbb{E}^{N+1}\left[\dfrac{B_{i+1}^{\mathrm{DA}}(T)}{B_{N+1}^{\mathrm{DA}}(T)}\,\middle|\,\mathbf{X}(t)=\mathbf{x}\right]}-1\right\} \end{aligned} \tag{22}$$

where each of the above-stated T random variables should be evaluated at $(T,\mathbf{X}(T))$. The second equality follows from $B_i^{\mathrm{AF}}/B_{N+1}^{\mathrm{AF}}$ being a martingale by the assumption of no arbitrage. The "arbitrage-free" forward rates $L_i^{\mathrm{AF}}(t,\mathbf{x})$ obtained in this way may then be compared with forward rates $L_i^{\mathrm{DA}}(t,\mathbf{x})$ obtained by single time stepping.

8.2 Numerical results for single time step test

The inconsistency test was performed under the following set-up. Ten annual forward rates were considered where forward rate i accrued from year i to $i+1$, for $i=20,\dots,29$. Under the notation of the previous section, t was taken to be 10 years, T was taken to be 20 years and T_{N+1} was taken to be 30 years. See also Figure 6. $L_i(0)$ was taken to be 5%, and mean-reversion, κ, was varied at 0%, 5%

[6] Here the notations "AF" and "DA" indicate "arbitrage-free" and "drift-approximated", respectively.

FIGURE 6 Set-up for inconsistency test.

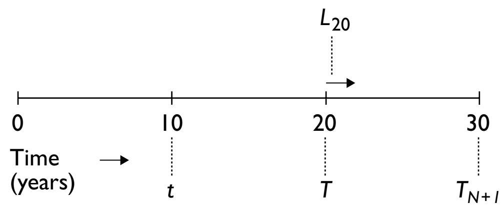

and 10%. The $\tilde{\gamma}_i$ were chosen such that the volatility of the corresponding caplet was equal to some general volatility level v, which was varied at 10%, 15% and 20%. Let SD denote the standard deviation of $X(10)$. $X(10)$ moves were considered for 0, ±SD/2, and ±SD. Results for the volatility/mean-reversion scenario 15%/10% are given in Table 7. The comparison is only reported for L_{20} because this forward rate contains the most drift terms, and therefore its corresponding error is the largest among $i = 20, \ldots, 29$. Note that the error for L_{29} is always zero as it is fully determined by X. In Table 8 the maximum error (over the five considered $X(10)$ moves) between $L_{20}^{\mathrm{AF}}(10)$ and $L_{20}^{\mathrm{DA}}(10)$ is reported.

TABLE 7 Quality of drift approximations: comparison of $L_{20}^{\mathrm{AF}}(10)$ and $L_{20}^{\mathrm{DA}}(10)$ under different $X(10)$ moves for the volatility/mean-reversion scenario 15%/10%. SD denotes the standard deviation of $X(10)$. All variables are evaluated at time $t = 10$.

	Brownian bridge			Predictor-corrector		
$X(10)$	L_{20}^{AF} (%)	L_{20}^{DA} (%)	$L_{20}^{\mathrm{DA}} - L_{20}^{\mathrm{AF}}$ (bp)	L_{20}^{AF} (%)	L_{20}^{DA} (%)	$L_{20}^{\mathrm{DA}} - L_{20}^{\mathrm{AF}}$ (bp)
−SD	3.75	3.81	5.11	3.74	3.81	7.17
−SD/2	4.23	4.27	4.03	4.19	4.27	7.94
0	4.77	4.79	2.37	4.70	4.79	8.81
+SD/2	5.38	5.38	−0.05	5.28	5.38	9.79
+SD	6.07	6.03	−3.47	5.92	6.03	10.91

TABLE 8 Quality of drift approximations: maximum of $|L_{20}^{\mathrm{AF}}(10) - L_{20}^{\mathrm{DA}}(10)|$ over $X(10)$ moves 0, ±SD/2, ±SD, for different volatility/mean-reversion scenarios. SD denotes the standard deviation of $X(10)$. Differences are denoted in basis points.

	Brownian bridge			Predictor-corrector		
Mean-	Volatility level (v)			Volatility level (v)		
reversion	10%	15%	20%	10%	15%	20%
0%	2.97	9.34	28.73	2.86	8.60	37.45
5%	2.56	8.21	19.46	2.32	12.29	53.85
10%	1.46	5.11	12.56	1.69	10.91	44.59

The test was performed for both the Brownian bridge and predictor-corrector schemes. The results show that the former outperforms the latter in the timing inconsistency test.

The inconsistency test results show that, for less volatile market scenarios, the single time step framework performs very accurately, with errors only up to a few basis points. For more volatile market scenarios the approximation deteriorates. But for realistic yield curve and forward volatility scenarios there are no problems with respect to pricing (see Section 7). The worsening of the approximation for more volatile scenarios is what may be expected from the nature of the drift approximations: as the model dimensions increase, the single time step approximation will break up. By "model dimensions" we mean the volatility level, the tenor of the deal, the difference between the forward index i and N, or time zero forward rates, etc. Care should be taken in applying the single time step framework for BGM that the market scenario does not violate the realm where the single time step approximation is reasonably valid.

9 Conclusions

We have introduced a fast approximate pricing framework as an addition to the predictor-corrector drift approximation developed by Hunter, Jäckel and Joshi (2001a). These authors used the drift approximation only to speed up their Monte Carlo by reducing it to single time step simulation. We have shown that, at a slight cost, much faster computational methods may be used, such as numerical integration or finite differences. The additional cost is a non-restrictive assumption, namely, separability of the volatility function. The proposed drift approximation framework was applied to the pricing of Bermudan swaptions, for which it yielded very accurate prices with much lower computation times.

Appendix A Mean of generalized geometric Brownian bridge

In this appendix, the time-t mean of the process L_k defined in Equation (9) is determined. Equivalently, we may determine the time-t mean of the process Y, given by

$$\frac{\mathrm{d}Y(t)}{Y(t)} = \sigma(t) \cdot \mathrm{d}\mathbf{W}(t), \quad Y(0) = y_0, \quad Y(t^*) = y^*$$

(Compare with Equation (9).) The solution of Y (unconditional of time-t^*) is given by

$$Y(t) = y_0 \mathrm{e}^{X(t) - \frac{1}{2}\Sigma^2(t)}$$

where

$$X(t) := \int_0^t \sigma(s) \cdot \mathrm{d}\mathbf{W}(s), \qquad \Sigma^2(t) := \int_0^t \| \sigma(s) \|^2 \mathrm{d}s$$

Note that

$$\left\{ \omega \in \Omega; Y(t^*) = y^* \right\} = \left\{ \omega \in \Omega; X(t^*) = \log(y^*/y_0) + \tfrac{1}{2}\Sigma^2(t^*) =: x^* \right\}$$

According to the martingale time change theorem (for example Theorem 4.6 of Karatzas and Shreve (1991)), we have that $X(\tau(\cdot))$ is a Brownian motion, where the time change τ is defined by

$$\tau(t) = \inf\left\{ s \geq 0; \Sigma^2(t) > s \right\}$$

Working in the time-changed time coordinates, $X(\cdot) \,|\, X(\tau^*) = x^*$ is a standard Brownian bridge, and so, according to Section 5.6.B of Karatzas and Shreve (1991),

$$X(\tau) \,\big|\, X(\tau^*) = x^* \sim \mathcal{N}\left(\frac{\tau}{\tau^*} x^*, \, \tau - \frac{\tau^2}{\tau^*} \right)$$

Back in the original time coordinates, this translates to

$$X(t) \,\big|\, X(t^*) = x^* \sim \mathcal{N}\left(\frac{\Sigma^2(t)}{\Sigma^2(t^*)} x^*, \, \Sigma^2(t) - \frac{\left(\Sigma^2(t)\right)^2}{\Sigma^2(t^*)} \right)$$

With this, we may evaluate the mean of $Y(t) \,|\, Y(t^*) = y^*$ to be

$$\mathbb{E}\left[Y(t) \,\big|\, Y(t^*) = y^* \right] = y_0 \left(\frac{y^*}{y_0} \right)^{\frac{\Sigma^2(t)}{\Sigma^2(T)}} \exp\left\{ \frac{1}{2} \frac{\Sigma^2(t)}{\Sigma^2(T)} \left(\Sigma^2(T) - \Sigma^2(t) \right) \right\}$$

where the following simple rule has been used: $\mathbb{E}[\mathrm{e}^Z] = \mathrm{e}^{\beta + \tau^2/2}$ whenever Z is normally distributed, $Z \sim \mathcal{N}(\beta, \tau^2)$.

Appendix B Approximation of substituting the mean in the expectation of expression (9)

In Section 3 a four-step method for the calculation of expression (9) is described. An approximating fourth step is proposed that evaluates the expectation of the BGM drift inserting the mean. In this appendix an error bound for this approximation is derived, and it is shown that the approximation is of order two in volatility in the neighbourhood of zero.

The expectation term can always be rewritten as

$$f(\mu,\sigma) = \mathbb{E}\left[\frac{\exp\{\mu+\sigma Z\}}{1+\exp\{\mu+\sigma Z\}}\right]$$

where Z is distributed standard normally. It is straightforward to verify that the above function $f\colon \mathbb{R}^2 \to \mathbb{R}$ is infinitely differentiable at every point of the whole real plane. Note that approximating the above expectation at the mean signifies that the above function is approximated as

$$f(\mu,\sigma) \approx f(\mu,0) = \frac{\exp\{\mu\}}{1+\exp\{\mu\}}$$

Fix μ and calculate the derivative of f with respect to σ. The interchange of differentiation and expectation is a subtle argument that may, for example, be found in Williams (1991, paragraph A.16.1). We carefully verified that in the above case all the requirements for interchange are satisfied. We then find

$$\frac{\partial f}{\partial \sigma}(\mu,\sigma) = \mathbb{E}\left[Z\frac{\exp\{\mu+\sigma Z\}}{\left(1+\exp\{\mu+\sigma Z\}\right)^2}\right]$$

Due to the odd nature of the above integrand at the point $\sigma = 0$, we find that

$$\frac{\partial f}{\partial \sigma}(\mu,0) = 0$$

Taylor's formula then states that there exists $C \geq 0$ (possibly depending on μ) such that

$$\left| f(\mu,\sigma) - \frac{\exp\{\mu\}}{1+\exp\{\mu\}} \right| \leq C\sigma^2$$

Because a bound on the second derivative of $\sigma \mapsto f(\mu,\sigma)$ may be found independently of μ on some interval $[0,\bar{\sigma}]$, it follows from Theorem 7.7 of Apostol (1967) that the constant C may then be chosen independently of μ for all $\sigma \in [0,\bar{\sigma}]$.

REFERENCES

Apostol, T. M. (1967). *Calculus*, Volume 1, Second edition. John Wiley & Sons, Chichester.

Avramidis, A. N., and Matzinger, H. (2004). Convergence of the stochastic mesh estimator for pricing Bermudan options. *Journal of Computational Finance* **7**(4), 73–91.

Björk, T., Landén, C., and Svensson, L. (2002). Finite dimensional Markovian realizations for stochastic volatility forward rate models. Working paper.

Brace, A., Gątarek, D., and Musiela, M. (1997). The market model of interest rate dynamics. *Mathematical Finance* **7**(2), 127–55.

Brigo, D., and Mercurio, F. (2001). *Interest rate models: theory and practice.* Springer-Verlag, Berlin.

Broadie, M., and Glasserman, P. (1996). Estimating security price derivatives using simulation. *Management Science* **42**(2), 269–85.

Broadie, M., and Glasserman, P. (2004). A stochastic mesh method for pricing high-dimensional American options. *Journal of Computational Finance* **7**(4), 35–72.

De Jong, F., Driessen, J., and Pelsser, A. A. J. (2002). LIBOR market models versus swap market models for pricing of interest rate derivatives: an empirical analysis. *European Finance Review* **5**(3), 201–37.

Glasserman, P., and Merener, N. (2003a). Cap and swaption approximations in LIBOR market models with jumps. *Journal of Computational Finance* **7**(1), 1–36.

Glasserman, P., and Merener, N. (2003b). Numerical solution of jump–diffusion LIBOR market models. *Finance and Stochastics* **7**(1), 1–27.

Glasserman, P., and Merener, N. (2004). Convergence of a discretization scheme for jump–diffusion processes with state-dependent intensities. *Proceedings of the Royal Society* **460**(2041), 111–27.

Glasserman, P., and Zhao, X. (1999). Fast Greeks by simulation in forward LIBOR models. *Journal of Computational Finance* **3**(1), 5–39.

Heath, D., Jarrow, R., and Morton, A. (1992). Bond pricing and the term structure of interest rates: a new methodology for contingent claims valuation. *Econometrica* **60**(1), 77–105.

Hughston, L. P., and Rafailidis, A. (2002). A chaotic approach to interest rate modelling. To appear in *Finance and Stochastics*.

Hunt, P., Kennedy, J., and Pelsser, A. A. J. (2000). Markov-functional interest rate models. *Finance and Stochastics* **4**(4), 391–408.

Hunter, C. J., Jäckel, P., and Joshi, M. S. (2001a). Drift approximations in a forward-rate-based LIBOR market model. Working paper, www.rebonato.com.

Hunter, C. J., Jäckel, P., and Joshi, M. S. (2001b). Getting the drift. *Risk*, July.

Jäckel, P. (2002). *Monte Carlo methods in finance.* J. Wiley & Sons, Chichester.

Jamshidian, F. (1996). LIBOR and swap market models and measures. Working paper, Sakura Global Capital, London.

Jamshidian, F. (1997). LIBOR and swap market models and measures. *Finance and Stochastics* **1**(4), 293–330.

Karatzas, I., and Shreve, S. E. (1991). *Brownian motion and stochastic calculus*, Second edition. Springer-Verlag, Berlin.

Kloeden, P. E., and Platen, E. (1999). *Numerical solution of stochastic differential equations*, Volume 23 of *Applications of mathematics*. Springer-Verlag, Berlin.

Kurbanmuradov, O., Sabelfeld, K., and Schoenmakers, J. (1999). Lognormal random field approximations to LIBOR market models. Working paper, Weierstrass Institute, Berlin (www.wias-berlin.de).

Kurbanmuradov, O., Sabelfeld, K., and Schoenmakers, J. (2002). Lognormal approximations to LIBOR market models. *Journal of Computational Finance* **6**(1), 69–100.

Longstaff, F. A., and Schwartz, E. S. (2001). American options by simulation: a simple least-squares approach. *Review of Financial Studies* **14**(1), 113–47.

Miltersen, K. R., Sandmann, K., and Sondermann, D. (1997). Closed form solutions for term structure derivatives with log-normal interest rates. *Journal of Finance* **52**(1), 409–30.

Musiela, M., and Rutkowski, M. (1997). Continuous-time term structure models: forward measure approach. *Finance and Stochastics* **1**(4), 261–91.

Pelsser, A. A. J. (2000). *Efficient methods for valuing interest rate derivatives.* Springer-Verlag, Berlin.

Ritchken, P., and Sankarasubramanian, L. (1995). Volatility structures of the forward rates and the dynamics of the term structure. *Mathematical Finance* **5**(1), 55–72.

Williams, D. (1991). *Probability with martingales.* Cambridge University Press, Cambridge.

Wilmott, P. (1998). *Derivatives: the theory and practice of financial engineering.* John Wiley & Sons, Chichester.