NEW AND ROBUST DRIFT APPROXIMATIONS FOR THE LIBOR MARKET MODEL

MARK JOSHI AND ALAN STACEY

ABSTRACT. We present four new methods for approximating the drift in the LIBOR market model. These are compared to a variety of existing methods including PPR, Glasserman-Zhao and predictor-corrector. We see that two of them which use correlation adjustments to better approximate the drift are more effective than existing methods.

1. INTRODUCTION

The LIBOR market model has become an important, and possibly standard, model for the pricing of interest rate derivatives in recent years. Despite this, various details of its implementation are still being worked out. These include how to handle derivatives with early exercise or callability features, the choice and appropriate use of calibration instruments, and how to approximate the evolution of the underlying stochastic differential equation. The last of these is important for two reasons, firstly better approximations will lead to more accurate and efficient Monte Carlo simulations, and secondly, it is useful when trying to develop Markov functional approaches to implementation. In this article, we focus on improving Monte Carlo simulations via better drift approximations, however, we believe the techniques will also have relevance to alternative implementation methodologies.

We recall the set-up of the LIBOR market model. For more detail, we refer the reader to the fundamental papers [2],[8], [12] and [14], or to the books [3], [5], [10] [13] and [16]. The basic idea is to evolve discrete market-observable forward rates, rather than hidden unobservable factors. We have tenor dates $T_0, T_1, \ldots T_n$, and corresponding forward rates $f_0, \ldots f_{n-1}$. Set $\tau_i = T_{i+1} - T_i$.

Let $P(T)$ denote the zero-coupon bond which pays 1 at time T and let its value at time $t \leq T$ be denoted $P_t(T)$. Taking $P(T_n)$ as numeraire, the rates f_i have the following evolution

$$\frac{df_i}{f_i} = \sigma_i dW_i + \mu_i dt,$$

1

where σ_i is a deterministic function of time (typically either piecewise constant or of abcd-form as in [16]), (W_i) is an n-dimensional Brownian motion whose projection on to each of the coordinate axes is a standard Brownian motion, and the drift term is given by

$$\mu_i = - \sum_{j>i} \frac{f_j \tau_j}{1 + f_j \tau_j} \rho_{ij} \sigma_i \sigma_j,$$

where ρ represents the (instantaneous) correlation between W_i and W_j. Note that for convenience we consider $\sigma_i(t)$ to be equal to 0 for $t \geq T_i$

Many pure LIBOR products for example a trigger swap or an auto-cap can be priced purely from knowing the distribution of the forward rates at their own reset times. Such products are typically products that depend only on LIBOR rates and do not have callability features. To price these products, we can integrate their deflated pay-offs against the distribution of forward rates. Since typically the number of forward rates will be large (≥ 20,) this will be done by Monte Carlo simulation.

Since rates do not move after their reset times, we wish to simulate from the distribution of their values at time T_{n-1} (or T_n). Provided we can do so accurately, we would then like to just carry out one Monte-Carlo step from 0 to T_{n-1}, a practice known as *long-stepping*. It is this case which we should bear in mind throughout. For this reason, we will always take the final bond $P(T_n)$ as numeraire.

In what follows, we will discuss evolution from a time $S \geq 0$ to at a time $T \leq T_n$. We will work with the logs of the rates since this yields constant volatility terms. We then have to evolve

$$d \log f_j = \left(\mu_j(f, t) - \frac{1}{2} \sigma_j^2(t) \right) dt + \sigma_j(t) dW_j. \qquad (1.1)$$

If drifts were not state-dependent this would be easy and have an exact solution. Since they are not, most work on the topic, [1], [7], [15], has focussed on ways to approximate the drift. An alternate approach is to work with the bonds instead, [6] at the gain of eliminating drift problems whilst introducing volatility problems.

Here we introduce four new drift approximation methods, truncated-predictor-corrector (tpc), mean-setting-predictor-corrector (mpc), correlation-adjusted predictor-corrector (capc), and correlation-adjusted numerical-integration predictor-corrector (cani). All of these rely on the observation that if one wishes to estimate the value of the drift of f_j at the end of a time-step ending at $T > T_j$, it is the values of $f_k(T_j)$ for $k > j$, that are important rather than the values of $f_k(T_k)$, since f_j does not move after T_j. The methods tpc and mpc are based on doing a log-linear interpolation to obtain an estimate of the value of $f_k(T_j)$. They

do not, however, perform particularly well. We solve their deficiencies with capc and cani, via the observation that the correlation between rates can be used to obtain better estimates of $f_k(T_j)$ than by log-linear interpolation, and crucially that this can be done with only a moderate increase in the computing time required. We will see that these two new methods are substantially better than existing methods.

We structure the paper as follows. In Section 2, we review the basics of how we do the evolution and set up some notation. In Section 3, we develop a class of methods we call *iterative*, including all our new methods. In Section 4, we review the predictor-corrector method and the Glasserman-Zhao method. We present numerical results in Section 5. We conclude in Section 6.

Note that the true distribution of the rates is not log-normal in the pricing measure, and all of these methods will lead to non-lognormal rates. An alternate approach is to try and find the best log-normal approximation and this has been pursued in [4] and [11]. Note that whilst such methods have the virtue of ease of implementation, they will never accurately capture the full nature of a non-lognormal distribution.

2. The evolution

The non-drift part of the evolution is straightforward: we obtain a pseudo-square-root A of the covariance matrix of the distribution $(\int_S^T \sigma_i dW_i)_{i<n}$, and then set

$$Y = AZ$$

where Z_i $(0 \leq i < n)$ are independent $N(0,1)$; then $(Y_i)_{i<n}$ has the same joint distribution as $(\int_S^T \sigma_i dW_i)_{i<n}$. Note that since we are typically doing one long step, the covariance matrix will be of full rank, even if there is perfect instantaneous correlation and flat volatility since each forward rate will terminate at a different time.

Now we have

$$\log f_i(T) = \log f_i(S) + \int_S^T \sigma_i dW_i - \frac{1}{2} \int_S^T \sigma_i^2(t)dt + \int_S^T \mu_i(t)dt. \quad (2.1)$$

Let I_i denote the middle of these three integrals. We use $(\hat{f}_i)$ to denote our attempted draw from the distribution of $f_i(T)$.

3. Iterative Methods

When P_n is taken to be numeraire, the drifts have the interesting property that the drift of f_j only depends on the values of f_k for $k > j$. Iterative methods make use of this information by working backwards

and doing each rate one by one. Noting $\mu_{n-1} \equiv 0$, all of the iterative methods start by setting

$$\log \hat{f}_{n-1} = \log f_{n-1}(S) + Y_{n-1} + I_{n-1}.$$

The various methods then — successively for $i = n - 2, \ldots, 0$ — use information already calculated to determine an estimate, which we denote $\hat{\mu}_i$, of the difficult state-dependent term $\int_S^T \mu_i dt$, and set

$$\log \hat{f}_i = \log f_i(S) + Y_i + I_i + \hat{\mu}_i. \tag{3.1}$$

Euler stepping. Take

$$\hat{\mu}_i = -\sum_{j>i} \frac{f_j(S)\tau_j}{1 + f_j(S)\tau_j} \int_S^T \rho_{ij}(t)\sigma_i(t)\sigma_j(t)dt.$$

In other words, calculate the drift term as though the f_is across the interval $[S, T]$ remained equal to the values $f_i(S)$. This method is iterative in the trivial sense that it makes no use of any information after time S.

Iterative Predictor-Corrector (ipc). Take

$$\hat{\mu}_i = -\frac{1}{2}\sum_{j>i} \left(\frac{f_j(S)\tau_j}{1 + f_j(S)\tau_j} + \frac{\hat{f}_j\tau_j}{1 + \hat{f}_j\tau_j} \right) \int_S^T \rho_{ij}(t)\sigma_i(t)\sigma_j(t)dt. \tag{3.2}$$

Note that this computation is possible because $\hat{\mu}_i$ is needed only after all the $\hat{f}_j$ for $j > i$ have already been calculated.

Truncated Predictor-Corrector (tpc). For each i, let the *local setting time* of f_i be given by

$$\tilde{T}_i = \mathrm{mid}(S, T_i, T).$$

Of course, if we are long-stepping, then this will be the same as the setting time T_i.

Under the ipc method, the drift of f_i is calculated using the $\hat{f}_j$s $(j > i)$, our estimate for the values of f_j at times $\tilde{T}_j$, despite the fact that f_i sets at time $\tilde{T}_i$ which is generally earlier. It would seem to make more sense to use an estimate of $f_j(\tilde{T}_i)$ instead of using $\hat{f}_j$.

Having determined $\hat{f}_j$ (for all $j > i$) which we also denote $\hat{f}_j(\tilde{T}_j)$, we define $\hat{f}_j(t)$ for $S \leq t \leq \tilde{T}_j$ by log-linear interpolation between $f_j(S)$ and $\hat{f}_j(\tilde{T}_j)$, i.e.,

$$\log \hat{f}_j(t) = \frac{\tilde{T}_j - t}{\tilde{T}_j - S} \log f_j(S) + \frac{t - S}{\tilde{T}_j - S} \log \hat{f}_j(\tilde{T}_j). \tag{3.3}$$

(Note that if $\tilde{T}_j = S$ it is clear what is meant since in that case $\hat{f}_j = f_j(S)$.)

In fact we really want to adjust the convexity coefficients in (3.3) to take into account any non-constant volatility. To see how this is done, see the PPR method below. When the volatility is flat, which holds in the numerical examples we give here, it is not an issue.

Based on the crude observation that the drift of f_i is zero after T_i, we estimate the overall drift by averaging the f_j-dependent term at S and our estimate for that term at $\tilde{T}_i$. To do this, we replace the $\hat{f}_j$ appearing in (3.2) by $\hat{f}_j(\tilde{T}_i)$. So we have

$$\hat{\mu}_i = -\frac{1}{2} \sum_{j>i} \left(\frac{f_j(S)\tau_j}{1 + f_j(S)\tau_j} + \frac{\hat{f}_j(\tilde{T}_i)\tau_j}{1 + \hat{f}_j(\tilde{T}_i)\tau_j} \right) \int_S^T \rho_{ij}(t)\sigma_i(t)\sigma_j(t)dt. \quad (3.4)$$

Mean-setting Predictor-Corrector (mpc). This method is a variation of truncated predictor-corrector based on the following observations.

- If $\log f \sim N(\mu, \sigma^2)$ then $\mathbb{E}f = e^{\mu + \frac{1}{2}\sigma^2}$.
- If W_t is a Brownian motion with constant drift and constant volatility σ, then, for $a < b < c$, the distribution of W_b conditional on W_a and W_c is normal with mean

$$\frac{c-b}{c-a}W_a + \frac{b-a}{c-a}W_b$$

and variance

$$\frac{(c-b)(b-a)}{c-a}\sigma^2$$

See [9] for more information. We proceed as for truncated predictor-corrector except that we replace (3.3) by

$$\log \hat{f}_j(t) = \frac{\tilde{T}_j - t}{\tilde{T}_j - S} \log f_j(S) + \frac{t - S}{\tilde{T}_j - S} \log \hat{f}_j(\tilde{T}_j) + \frac{1}{2} \frac{(\tilde{T}_j - t)(t - S)}{(\tilde{T}_j - S)^2} \int_S^T \sigma_i^2 \quad (3.5)$$

Balland. This variant of predictor-corrector is used in [1]; the author was more interested in Markov functional models than Monte Carlo and therefore does not make any claims of its superiority over other methods. However, we include it for completeness. It takes a geometric mean of the forward rates (or equivalently the arithmetic mean of their logarithms) and uses that in the drift computation. Set

$$\bar{f}_j = \sqrt{f_j(S)\hat{f}_j}$$

and then put

$$\hat{\mu}_i = -\sum_{j>i} \frac{\bar{f}_j \tau_j}{1 + \bar{f}_j \tau_j} \int_S^T \rho_{ij}(t)\sigma_i(t)\sigma_j(t)dt.$$

Pietersz-Pelsser-Regenmortel (PPR). This method is taken from [15]. Like tpc and mpc it relies on the estimation of values of $f_j(t)$ for intermediate values of t given $f_j(S)$ and $\hat{f}_j$. One wishes the log-linear interpolation to take appropriate account of time-dependent volatility. Set

$$V_j(t) = \int_S^t \sigma_i^2$$

for $S \leq t \leq T$. One variant of PPR defines $\hat{f}_j(t)$ (given the value of $\hat{f}_j = \hat{f}_j(\tilde{T}_j) = \hat{f}_j(T)$) by

$$\log \hat{f}_j(t) = \frac{V_j(T) - V_j(t)}{V_j(T)} \log f_j(S) + \frac{V_j(t)}{V_j(T)} \log \hat{f}_j. \qquad (3.6)$$

The other variant of PPR differs at precisely this point and incorporates an adjustment similar to that used in mean-setting predictor-corrector described above. For this variant, one has

$$\log \hat{f}_j(t) = \frac{V_j(T) - V_j(t)}{V_j(T)} \log f_j(S) + \frac{V_j(t)}{V_j(T)} \log \hat{f}_j + \frac{1}{2} \frac{V_j(t)(V_j(T) - V_j(t))}{V_j(T)}.$$
$$(3.7)$$

We shall use the second method in this paper, since it appears to be the superior choice.

In calculating the drift one wishes to estimate terms of the form

$$\int_S^T \frac{f_j(t)\tau_j}{1 + f_j(t)\tau_j} \rho_{ij}(t)\sigma_i(t)\sigma_j(t) \, dt. \qquad (3.8)$$

In order to do this one fixes a number of subintervals for the application of a version of the trapezoid rule. Let $S = t_0 < t_1 < \ldots t_m = T$ be evenly spaced. Then we estimate terms of the kind in (3.8) by integrating out the $\rho_{ij}\sigma_i\sigma_j$ term over each subinterval and the average of the estimates of the f_j-term at the two endpoints. So we obtain

$$\hat{\mu}_i = -\sum_{j>i}\sum_{k=0}^{m-1} \frac{1}{2}\left(\frac{\hat{f}_j(t_k)\tau_j}{1 + \hat{f}_j(t_k)\tau_j} + \frac{\hat{f}_j(t_{k+1})\tau_j}{1 + \hat{f}_j(t_{k+1})\tau_j}\right)\int_{t_k}^{t_{k+1}} \rho_{ij}\sigma_i\sigma_j.$$

where $\hat{f}_j(t)$ is as defined in (3.7).

Correlation-Adjusted Predictor-Corrector (capc). Three of the above methods involve making estimates of $f_j(t)$ for $S < t < T$ from the value of $f_j(S)$ and the randomly generated $\hat{f}_j(T)$. However, even conditional on $f_j(T)$, the value of $f_j(t)$ is strongly correlated with the values of other f_is. In particular, if $S < \tilde{T}_i < \tilde{T}_j$ then the value of $f_j(\tilde{T}_i)$ (which the truncated and mean-setting predictor-corrector methods attempt to estimate) is strongly correlated with the value of $f_i(\tilde{T}_i)$ — indeed in a model with high instantaneous correlation between rates it may be more strongly correlated with $f_i(\tilde{T}_i)$ than it is with $f_j(\tilde{T}_j)$.

The capc method attempts to use the additional information contained in the term Y_i — as in (3.1) — to obtain a better estimate of $f_j(\tilde{T}_i)$ than those obtained by other methods. Of course an even better estimate could be obtained by use of the other Ys but this appears to be too computationally intensive and would probably not make much additional improvement.

Fixing $i < j$, let

$$X_i = \int_S^T \sigma_i dW_i \qquad X_j = \int_S^T \sigma_j dW_j \qquad X_j^i = \int_S^{\tilde{T}_i} \sigma_j dW_j.$$

From the covariance structure of these three random variables, using for example Cholesky decomposition, it is not difficult to obtain constants a, b and c such that

$$X_j^i = aX_i + bX_j + c\aleph_{ij}$$

where $\aleph_{ij}$ is standard Gaussian and independent of (X_i, X_j).

We then define an approximation, $\hat{f}_j^i$ to the value of $f_j(\tilde{T}_i)$ by

$$\log \hat{f}_j^i = \log f_j(S) + aY_i + bY_j + \frac{1}{2}c^2 + \frac{V_j(\tilde{T}_i)}{V_j(\tilde{T}_j)}(I_j + \hat{\mu}_j).$$

Note that although it is correct to multiply I_j by $V_j(\tilde{T}_i)/V_j(\tilde{T}_j)$, the multiplication of $\hat{\mu}_j$ by this quantity is just a rough approximation to the drift of f_j between S and $\tilde{T}_i$.

One then uses the iterative predictor-corrector method with $\hat{f}_j^i$ in the place of $\hat{f}_j$, so we have

$$\hat{\mu}_i = -\frac{1}{2}\sum_{j>i}\left(\frac{f_j(S)\tau_j}{1 + f_j(S)\tau_j} + \frac{\hat{f}_j^i \tau_j}{1 + \hat{f}_j^i \tau_j}\right)\int_S^T \rho_{ij}(t)\sigma_i(t)\sigma_j(t)dt. \quad (3.9)$$

Correlation-Adjusted Numerical Integration (cani). We highlight one systematic source of error in capc which is exemplified as follows. Suppose the early rates are low at their setting times, but

that the later rates, say $f_{10}, f_{11}, \ldots$ are high. Then, particularly in the strongly-correlated cases, the later rates will tend to be lower at the earlier times. This will make the drift terms involving the rates $f_{10} \ldots$ smaller at earlier times, and hence the overall drift of all the rates will tend to be smaller, including the drift of the later rates. That is to say, lower earlier rates implies smaller drift even for the later rates. This effect is not picked up by capc.

Correlation-adjusted numerical integration makes precisely the same estimates as capc for the intermediate values of $f_j(\tilde{T}_i)$, but makes more extensive use of this information. In calculating the drift term, a numerical integration is carried out (using just the trapezoid method). Defining, for convenience of notation, $\tilde{T}_{-1} = S$ and $\hat{f}_j^{-1} = f_j(S)$, we set

$$\hat{\mu}_i = -\frac{1}{2} \sum_{j>i} \sum_{k=0}^{i} \left(\frac{\hat{f}_j^{k-1}\tau_j}{1 + \hat{f}_j^{k-1}\tau_j} + \frac{\hat{f}_j^k\tau_j}{1 + \hat{f}_j^k\tau_j} \right) \int_{\tilde{T}_{k-1}}^{\tilde{T}_k} \rho_{ij}(t)\sigma_i(t)\sigma_j(t)dt.$$

(3.10)

Note that some of these integrals may vanish because $\tilde{T}_{i-1} = \tilde{T}_i$ and so the summation becomes smaller, although if we are evolving with one long step they will generally all be non-zero.

Although it does not take much longer than other methods for a modest number of rates (20, say) the use of (3.10) becomes too computationally intensive when the number of rates is large. However, virtually identical results can be obtained by carrying out the numerical integration using only a handful of points, i.e., taking an inner sum with fewer terms.

4. OTHER METHODS

We also consider the performance of two other methods, which seem currently to be the most widely used for long-stepping.

Predictor-Corrector (pc). This method does not proceed iteratively, and is therefore applicable in greater generality: it does not rely upon the fact that the drift term for f_i only involves rates f_k with $k > i$. It uses Euler stepping to calculate an estimate $\tilde{f}_i$ for each $f_i(T)$. It then uses these values to estimate the drift of each rate at time T by setting

$$\hat{\mu}_i = -\frac{1}{2} \sum_{j>i} \left(\frac{f_j(S)\tau_j}{1 + f_j(S)\tau_j} + \frac{\tilde{f}_j\tau_j}{1 + \tilde{f}_j\tau_j} \right) \int_S^T \rho_{ij}(t)\sigma_i(t)\sigma_j(t)dt, \quad (4.1)$$

and applying (3.1). This is the method suggested in [7], although the subtle distinction between it and iterative predictor corrector is often missed.
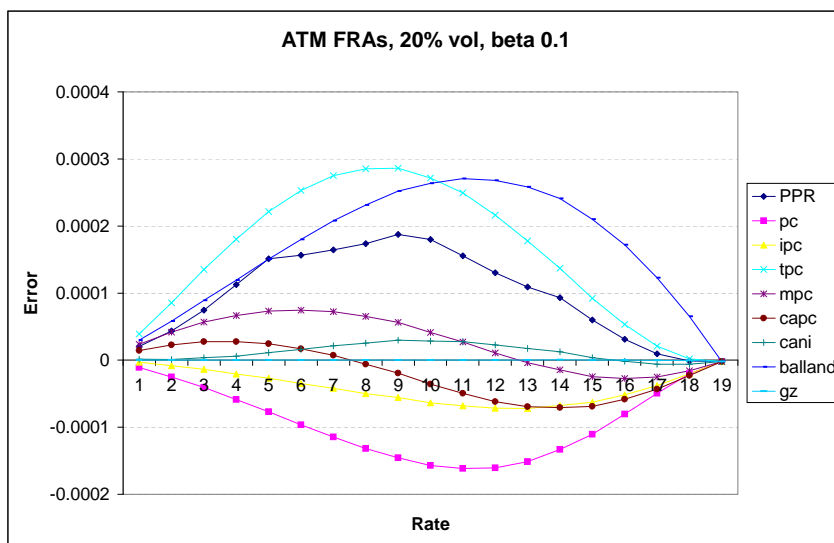
FIGURE 5.1. Pricing errors for at-the-money forward rates agreements in the base case

Glasserman-Zhao (gz). This method uses Euler stepping after a change of coordinates. Instead of discretizing the equations for the forwards, one uses these to derive equations for the evolution of the discounted zero-coupon bonds, $P(T_j)/P(T_n)$. As the bonds are tradables, the drift term is zero, although the volatility is now state-dependent. There is a number of variants of this approach, the best of which seems to be to consider the logarithm of difference between adjacent bonds, $\log((P(T_{j-1}) - P(T_j)/P(T_n))$; it is this best method which we consider. This method has the virtue that deflated bond prices are martingales even in the discretized measure which implies that there is no discretization error at all for the pricing of zero-coupon bonds and forward rate agreements (the latter being a linear combination of the former). For more details see [6] or [5].

5. NUMERICAL COMPARISONS OF THE VARIOUS METHODS

We suppose for concreteness in our discussions that we are dealing with rates $f_0, \ldots, f_{19}$ with $T_i = i + 1$. The correlation between rate i
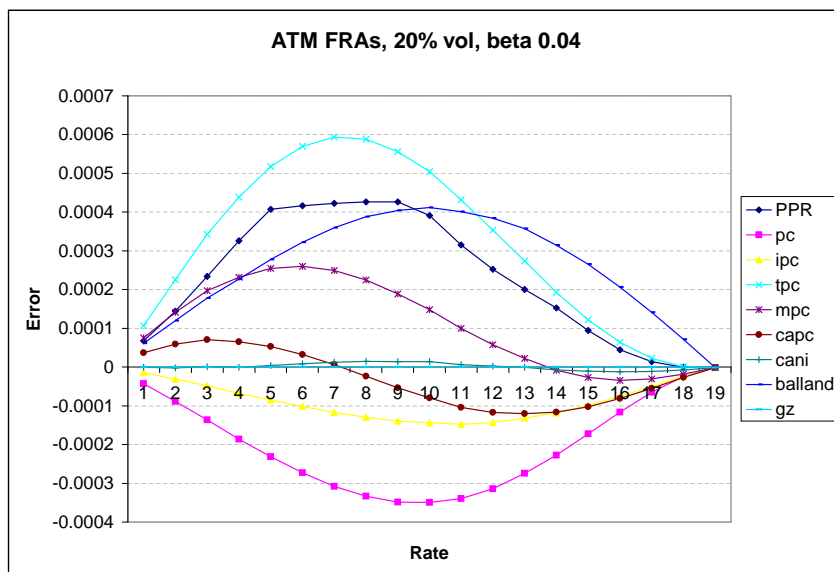
FIGURE 5.2. Pricing errors for at-the-money forward
rates agreements in the base case except beta=0.04.

and rate j is

$$e^{-\beta|T_i - T_j|}.$$

The *base case* we consider has a flat instantaneous volatility of 20%
and $\beta = 0.1$ but we will also consider high correlation (smaller β); we
will note deviations from the base case in our graphs. Since drifts are
proportional to correlation, the higher the correlation the tougher the
test. We only consider flat volatility, which leads to higher terminal
correlation and is therefore a more stringent test; non-flat volatility
gives broadly similar results to those obtained with lower flat volatility.
The base case has initial forward rates equal to 5%.

We run simulations with 2^{22} paths using Sobol numbers which en-
sures that errors due to convergence are below 0.1 basis points. When
pricing a product that involves cash-flows before the final time, the
value of the cash-flows is used to purchase one unit of the discretely
compounding money market account which is rolled up to time T_n and
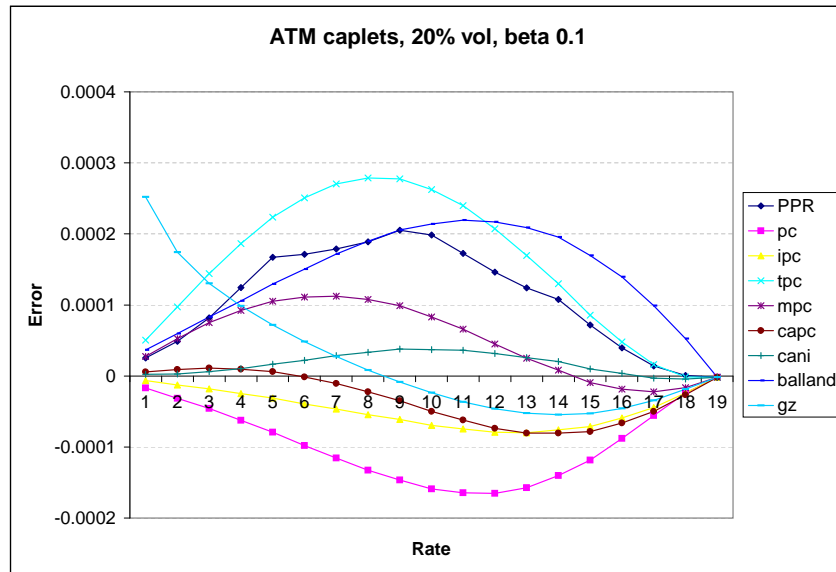then used to purchase units of the numeraire bond.

FIGURE 5.3. Pricing errors for at-the-money caplets in the base case.

When using the PPR method we typically take $m = 4$, although using more trapezoids does not seem to make a very significant difference, nor does it necessarily improve the results.

In general, when we do not know analytically the quantity we are trying to estimate (e.g. the price of an exotic interest rate derivative) we run the simulation with decreasing step-sizes where we can see all the methods converging to the same result. When we do know the price analytically (e.g. with FRAs or caplets) we usually just use one long step and measure the errors. We do not present results for the Euler method as these demonstrate huge errors which swamp the other methods.

Vanilla options. In figures 5.1 and 5.2, we show pricing errors for at-the-money forward rates agreements. As expected, Glasserman-Zhao prices these perfectly. However, cani is almost perfect and capc does better than the remaining methods. The methods pc and tpc both do badly but in different directions. Despite ipc's similarity to pc, it does a lot better. The mean-setting variant of tpc, that is mpc, performs much
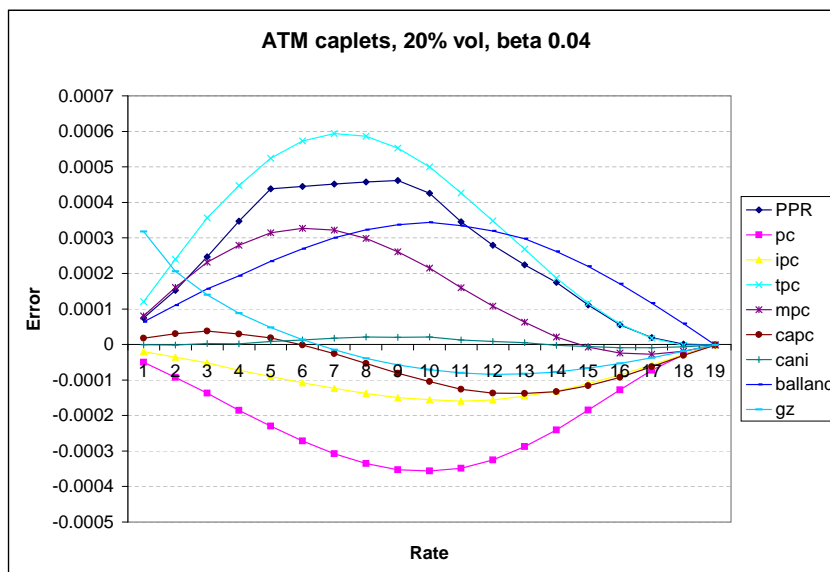
FIGURE 5.4. Pricing errors for at-the-money caplets in the base case, except beta 0.04.

better than it. PPR does surprisingly poorly. The Balland method is not impressive either.

We perform the same analysis for at-the-money caplets in figures 5.3 and 5.4. Glasserman-Zhao is no longer so good, and does quite poorly for short-dated options. For the remaining methods, errors are qualitatively very similar to the FRAs case. We see that cani is still very good with maximum error less than half a basis point, and capc again does better than the remaining methods. As before, the methods pc and tpc both do badly but in different directions. The ipc variant of pc again does a lot better than pc. The mpc method is again bad. PPR does poorly and the Balland method is again unimpressive.

We perform the same analysis for caplets struck at 8%; in figures 5.5 and 5.6, we show pricing errors. The true prices of these caplets run from 36 bps to 72 bps. So even 1bp is a substantial error in fractional terms. Glasserman-Zhao is now the worst method. For the remaining methods, errors are qualitatively very similar to the at-the-money case. We see that cani is still very good with maximum error less than half a basis point, and capc again does better than the remaining methods. As before, the methods pc and tpc both do badly but in different
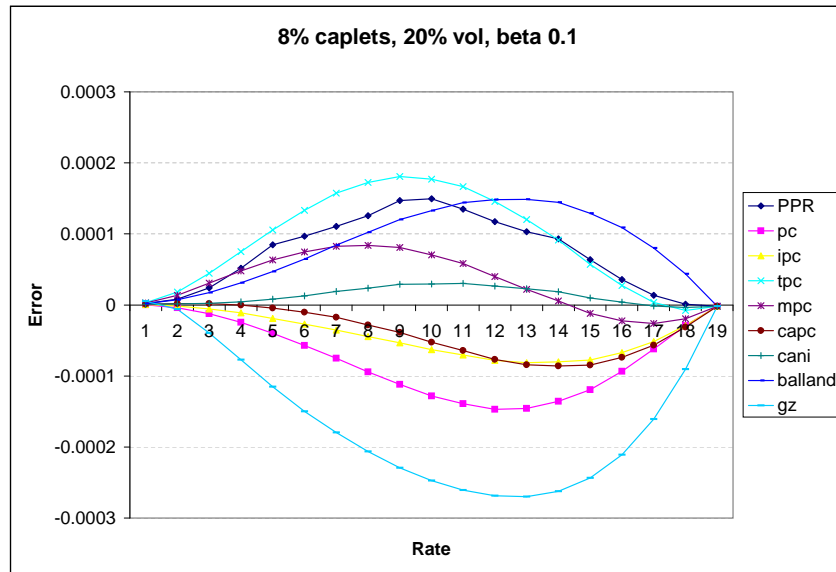
FIGURE 5.5. Pricing errors for caplets struck at 8 % in the base case

directions. The ipc variant of pc again does a lot better than pc. The mpc method is again bad. PPR does poorly and the Balland method is again unimpressive.

Auto-Caps. Ultimately, we are interested in drift methods for the pricing of exotic options. It is capturing the co-dependence of the rates that is subtle. We consider here a popular product which depends on the joint distribution of all the rates in an interesting way. The particular product we consider consists of at most seven caplets each with strike 5%. The caplets apply to the first seven rates which exceed the trigger level of 7%, or all such rates if there are fewer than seven in total.

We can see that although capc does not perform especially well compared with other good methods, cani is clearly extremely good. To keep the graph readable, we have not included pc, mpc and tpc, since the other graphs have illustrated that these should not be used.

6. CONCLUSION

We have presented some new methods for approximating the drift in the LIBOR market model. The method cani is very good in a range
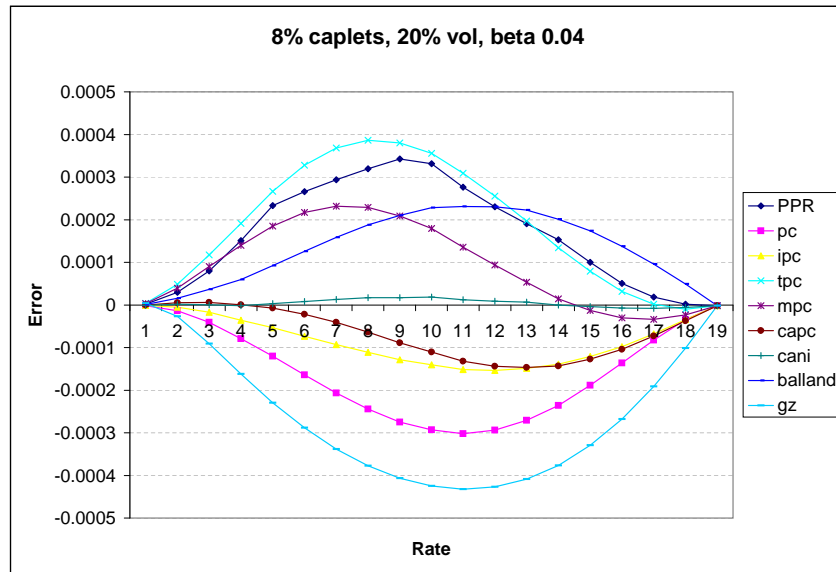
FIGURE 5.6. Pricing errors for caplets struck at 8 % the in base case, except beta =0.04

of cases, even where other well-known methods fail. The method capc is pretty good in general but sometimes fails. The methods tpc and mpc do surprisingly poorly but this has helped us understand that simply using the value of the rate at the end points to interpolate is not sufficient.

The original predictor-corrector method became popular because of its ability to price across long time horizons. However, we have seen here that the iterative variant is more effective and that the small additional effort to use capc or cani is well worthwhile.

References

[1] P. Balland Stoch-Vol for Libor Model, presentation at *International Centre for Business Information* (ICBI) conference, Madrid, 2004.

[2] A. Brace, D. Gatarek, M. Musiela, The market model of interest-rate dynamics, *Mathematical Finance* **7**, 127–155, 1997

[3] D. Brigo, F. Mercurio, *Interest Rate Models – Theory and Practice*, Springer Verlag, 2001

[4] A. Daniluk, D. Gatarek, A fully lognormal LIBOR market model, *Risk magazine*, September 2005, 115–118
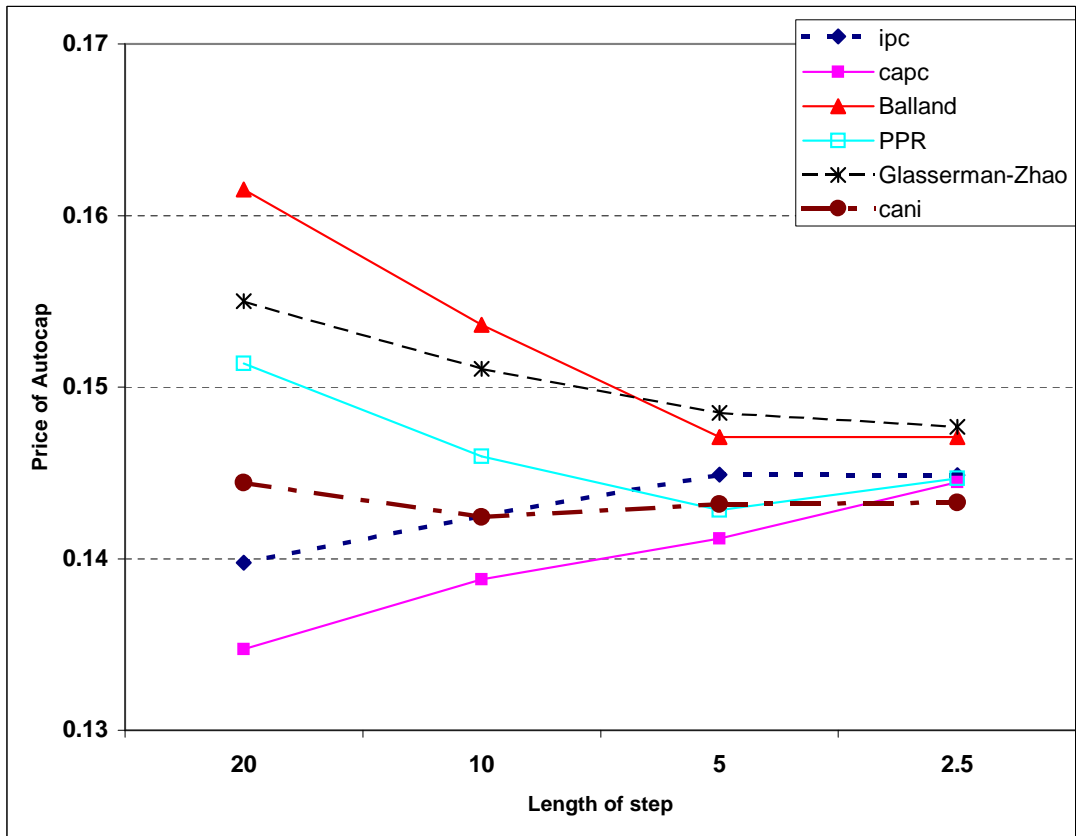
FIGURE 5.7. An autocap with 7 caplets, strike 5% and trigger level 7%. Base case parameters except volatility increased to 30%.

[5] P. Glasserman, *Monte Carlo Methods in Financial Engineering*. Springer-Verlag, New York, 2004.

[6] P. Glasserman and X. Zhao, Arbitrage-free discretization of lognormal forward Libor and swap rate models, *Finance and Stochastics* **4** 35–68, 2000.

[7] C. Hunter, P. Jäckel and M. Joshi, Getting the drift, *Risk*, **14**(7) 81–84, 2001.

[8] F. Jamshidian, LIBOR and swap market models and measures, *Finance and Stochastics* **1**, 293–330, 1997

[9] I. Karatzas, E. Shreve, *Brownian Motion and Stochastic Calculus*, Second edition, Springer Verlag, 1997

[10] M. Joshi, *The Concepts and Practice of Mathematical Finance*, Cambridge University Press, 2003.

[11] O. Kurbanmuradov, K. Sabelfeld, J. Schoenmakers, Lognormal approximations to LIBOR market models, *Journal of Computational Finance*, Volume 10, Number 1, 2002

[12] M. Musisela, M. Rutkowski, Continuous-time term structure models: forward measure approach. *Finance and Stochastics,* **1** 261–292, 1997

[13] M. Musiela, M. Rutkowski, *Martingale Methods in Financial Modelling*, Springer Verlag, 1997.

[14] K. R. Miltersen, K. Sandmann, D. Sondermann, Closed form solutions for term structure derivatives with log-normal interest rates, *Journal of Finance,* **52**, (1) 409–430, 1997

[15] R. Pietersz, A. Pelsser, M. van Regenmortel (2004). Fast drift-approximated pricing in the BGM model, *Journal of Computational Finance,* Volume 8, Number 1, 2004

[16] R.Rebonato *Modern Pricing of Interest-Rate Derivatives*, Princeton University Press, 2002.

CENTRE FOR ACTUARIAL SCIENCES, DEPARTMENT OF ECONOMICS, UNIVERSITY OF MELBOURNE, VICTORIA 3010, AUSTRALIA
   E-mail address: mark@markjoshi.com

LEHMAN BROTHERS INTERNATIONAL (EUROPE), 25 BANK STREET, LONDON, E14 5LE, UNITED KINGDOM
   E-mail address: astacey@lehman.com