Iterating cancelable snowballs and related exotics in a many-factor Libor model

Christian Bender, Anastasia Kolodko, John Schoenmakers

May 22, 2006

We propose a valuation method for exotic cancelable and callable structures in a multi-factor Libor model which are path-dependent in the sense that after canceling or calling, one cancels a sequence of cash-flows or receives a sequence of cash-flows in the future, respectively. The method, which is based on a Monte Carlo procedure for standard Bermudans recently developed and further extended in Kolodko and Schoenmakers [8], Bender and Schoenmakers [2], is compared to and also combined with popular known approaches by Andersen [1], Longstaff and Schwartz [9], and Piterbarg [10].

As a main example we consider the (cancelable) snowball swap, a highly sensitive interest rate product with growing popularity, in a full-blown Libor market model. From the treatment of this example it will be clear how to design Monte Carlo valuation algorithms for related cancelable or callable path-dependent products. For our example it turns out that price lower bounds obtained by the regression approach using explanatory variables as in [10] may be significantly off. Even an enhancement of this approach by an Andersen like modification of the corresponding exercise boundary does not lead to an acceptable small gap between the lower price and an upper bound price obtained via the dual method by Rogers [12], Haugh and Kogan [5] (see [7] for an alternative dual and [3] for upper bounds via consumption processes). However, via improving the stopping rule entailed by this enhancement using the iteration procedure in [8] we end up with acceptable prices.

The proposed approach is quite generic, as in principle it only requires a Monte Carlo simulation mechanism for an underlying Markovian system, for instance a Markovian system of SDEs. In particular, it can be used to improve upon popular methods, such as [1], [9], and [10] to get satisfactorily accurate target results (in this respect the method outperforms the standard policy iteration described in [11], see [8] for details).

Straightforward application of the policy iteration procedure based on [8] requires just like the duality approach a nested Monte Carlo simulation and is thus rather slow. Therefore we include a method of variance reduction which has a flavor of stratified sampling. Moreover, we underline that the improved stopping rule which is important for the buyer of the product for instance, can already be obtained at the cost of a standard (not nested) Monte Carlo simulation.

Recap of the Libor market model

The Libor market model is a popular and advanced tool for modelling interest rates and pricing of interest rate products. Let us first recall the Libor market model with respect to a tenor structure $0 = T_0 < T_1 < \ldots < T_n$ in the spot Libor measure P^*. For $1 \leq i < n$, the forward Libor $L_i(t)$, i.e. the annualized forward rate on a loan over period $[T_i, T_{i+1}]$ contracted at date t, which is to be settled at T_{i+1}, is governed, for $0 \leq t \leq T_i$, by the SDE (e.g., see Jamshidian [6])

$$dL_i = \sum_{j=\kappa(t)}^{i} \frac{\delta_j L_i L_j \, \gamma_i \cdot \gamma_j}{1 + \delta_j L_j} \, dt + L_i \, \gamma_i \cdot dW^*, \quad (1)$$

where $\delta_i = T_{i+1} - T_i$ are day count fractions, $t \to \gamma_i(t) = (\gamma_{i,1}(t), \ldots, \gamma_{i,d}(t))$ are deterministic volatility vector functions defined in $[0, T_i]$ (called factor loadings), and $\kappa(t) := \min\{m : T_m \geq t\}$ denotes the next reset date at time t. In (1), $(W^*(t) \mid 0 \leq t \leq T_{n-1})$ is a standard d-dimensional Wiener process under the spot Libor measure P^* with d, $1 \leq d < n$, being the number of driving factors.

This measure is induced by the numeraire

$$\begin{aligned} B_*(0) &:= 1 \\ B_*(t) &:= \frac{B_{\kappa(t)}(t)}{B_1(0)} \prod_{i=1}^{\kappa(t)-1} (1+\delta_i L_i(T_i)), \; t>0, \end{aligned}$$

with $\prod_{i=1}^{0} := 1$, and where $B_i(t)$ is the value of a zero coupon bond with face value \$1 at time $t \le T_i$. For further use we here also introduce the filtration $(\mathcal{F}_t)_{t\ge 0}$ (history information process) connected with the Libor process.

Path dependent cancelable and callable products

Consider a subset of tenor dates $\mathcal{T}_1 < \cdots < \mathcal{T}_k$, hence $\{\mathcal{T}_1, \ldots, \mathcal{T}_k\} \subset \{T_0, \ldots, T_n\}$, and adapted cash-flows C_i defined at $\mathcal{T}_i$ for $i = 1, \ldots, k$. Let us consider a path dependent contract which involves the right to cancel a sequence of (possibly negative) cash-flows $C_1, \ldots, C_\tau$, at a date τ to be decided by the option holder. The cash-flows of this contract are equivalent to an aggregated cash-flow $B_*(\mathcal{T}_\tau)\mathcal{Z}_\tau := B_*(\mathcal{T}_\tau)\sum_{j=1}^{\tau} Z_j$ at the cancellation date, with $Z_i := C_i/B_*(\mathcal{T}_i)$ being discounted cash-flows with respect to B_*. Indeed, it is equivalent to invest each cash-flow C_i, $i \le \tau$, in the numeraire B_*, yielding an amount $C_i/B_*(\mathcal{T}_i) = Z_i$ which is worth $B_*(\mathcal{T}_\tau)Z_i$ at date τ. By general arguments (see Duffie [4]) it follows that the value of the cancelable product at time zero is given by

$$V_0^{cancel} := \sup_{\tau\in\{1,\ldots,k\}} E^0 \mathcal{Z}_\tau = \sup_{\tau\in\{1,\ldots,k\}} E^0 \sum_{j=1}^{\tau} Z_j, \tag{2}$$

where the supremum is taken over all stopping indices with values in the set $\{1, \ldots, k\}$. Note that $\tau = k$ may be interpreted as "not canceled".

Naturally, in contrast to cancelable products we may consider callable ones.

$$\begin{aligned} V_0^{call} &:= \sup_{\tau\in\{1,\ldots,k\}} E^0 \sum_{j=\tau+1}^{k} Z_j \\ &= E^0 \sum_{j=1}^{k} Z_j + \sup_{\tau\in\{1,\ldots,k\}} E^0 \sum_{j=1}^{\tau} (-Z_j) \end{aligned}$$

is the price of a callable product which generates cash-flows $C_{\tau+1}, \ldots, C_k$, when called at date τ. Hence a callable product can be seen as the sum of a non-callable and a cancelable one (and vice versa). However, the virtual cash-flow in the callable representation is in the form of a conditional expectation, hence not explicit at hand. So, it is advantageous to write a callable product as sum of non-callable and a cancelable. We therefore concentrate on cancelable products throughout the paper.

Iterating path-dependent cancelables

The path dependent cancelable product introduced above can be seen as a standard Bermudan product with respect to a (virtual) cash-flow $\mathcal{Z}_i$. Therefore it can be evaluated by the iterative method developed in [8], which is studied further concerning numerical stability and extended to multiple stopping in [2]. Although here the cash-flow $\mathcal{Z}_i$ can be negative, it is easy to see that this does not provide any additional difficulties.

Let us briefly recall the iterative method. Suppose we are given some (generally suboptimal) exercise policy τ_i, $i = 1, \ldots, k$ for a Bermudan product with cash-flow process $\mathcal{Z}$; τ_i is the stopping rule according to which the option should be exercised, provided the option has not been exercised before $\mathcal{T}_i$. We assume that the exercise policy τ has the following properties,

$$\begin{aligned} i \le \tau_i \le k, \; \tau_k = k, \\ \tau_i > i \Rightarrow \tau_i = \tau_{i+1}, \qquad 0 \le i < k. \end{aligned} \tag{3}$$

This policy provides a lower bound process Y_i for the discounted Bermudan prices Y_i^*, also called Snell envelope,

$$Y_i^* \ge Y_i := E^i \mathcal{Z}_{\tau_i}, \qquad i = 1, \ldots, k,$$

where E^i denotes conditional expectation with respect to $\mathcal{F}_{\mathcal{T}_i}$. We next construct a new exercise policy $(0 \le i \le k)$,

$$\widehat{\tau}_i := \inf\{j \ge i : \mathcal{Z}_j \ge \max_{p:\, j\le p\le k} E^j \mathcal{Z}_{\tau_p}\} \tag{4}$$

which clearly satisfies (3) also, and consider the new lower bound process

$$\widehat{Y}_i := E^i \mathcal{Z}_{\widehat{\tau}_i}, \qquad i = 1, \ldots, k,$$

which is generally an improvement of Y,

$$Y_i \le \widehat{Y}_i \le Y_i^*, \qquad i = 1, \ldots, k.$$

Naturally, we may iterate the above procedure, i.e. improve $\widehat{\tau}$ in the same way and so forth. It is shown that after iterating this procedure $k-1$ times, the Snell envelope is attained independently of the choice of the starting stopping family.

Based on each of these lower bound processes one may construct upper bounds by the dual approach of [12] and [5]. This sequence of upper bounds terminates at the Bermudan price after finitely many steps as well.

For the path-dependent cancelable, introduced in the previous section, the improved policy (4) reads $(0 \leq i \leq k)$,

$$\widehat{\tau}_i = \inf\{j \geq i : 0 \geq \max_{p:j+1 \leq p \leq k} E^j \sum_{q=j+1}^{\tau_p} Z_q\}. \quad (5)$$

In most cases, both the cash-flow Z_i and the event $\{\tau_i = i\}$ are determined by the state of an underlying Markovian process (for example the Libor process (1)) at date i. In such a situation the conditional expectations involved in the iterative procedure can be estimated by Monte Carlo simulation, which thus leads to a Monte Carlo algorithm in a natural way. We refer to [8] for a detailed description of the general algorithm and to [13], Section 5.4.3, for a linear implementation of one improvement step. The numerical stability of these procedures is proved in [2].

To reduce the computational cost we recommend the following variance reduction for practical implementation. It is based upon the intuition that, for a good input stopping family τ_i, $E^j \mathcal{Z}_{\tau_{i+1}}$ is a fair (lower) approximation of $E^j \mathcal{Z}_{\hat{\tau}_{i+1}}$. Writing, by the consistency property (3),

$$\mathcal{Z}_{\hat{\tau}_0} - E\mathcal{Z}_{\hat{\tau}_0} = \sum_{i=0}^{\hat{\tau}_0} E^{i+1}\mathcal{Z}_{\hat{\tau}_i} - E^i \mathcal{Z}_{\hat{\tau}_i},$$

we hence expect that,

$$\mathcal{Z}_{\hat{\tau}_0} - \sum_{i=0}^{k} \mathbf{1}_{\{\hat{\tau}_0 > i-1\}} \left(E^{i+1}\mathcal{Z}_{\tau_i} - E^i \mathcal{Z}_{\tau_i}\right) \quad (6)$$

has a much lower variance than $\mathcal{Z}_{\hat{\tau}_0}$. The conditional expectations in (6) can be calculated for any outer path via plain Monte-Carlo with the already simulated inner paths, i.e. without additional simulation cost. Since the simulated version of $\mathbf{1}_{\{\hat{\tau}_0 > i-1\}}$ does only depend on the inner paths simulated at times $j = 0, \ldots, i-1$ and the simulated version of $E^{i+1}\mathcal{Z}_{\tau_i} - E^i \mathcal{Z}_{\tau_i}$ depends only on the inner paths simulated at times $j = i, i+1$, one can easily see that the expectation of the simulated version of $\sum_{i=0}^{k} \mathbf{1}_{\{\hat{\tau}_0 > i-1\}} \left(E^{i+1}\mathcal{Z}_{\tau_i} - E^i \mathcal{Z}_{\tau_i}\right)$ is zero. Therefore we average over the simulated versions of (6) instead of the simulated versions of $\mathcal{Z}_{\hat{\tau}_0}$ to estimate $E\mathcal{Z}_{\hat{\tau}_0}$ more efficiently in our numerical experiments. In our experiments we typically achieve a variance reduction of factor 20-60 this way (depending on the initial stopping family).

Generic construction of a good input stopping family

An important issue is the choice of the input stopping family τ. By choosing the trivial family $\tau_i \equiv i$ we are faced with the evaluation of the conditional expectations $E^j Z_q$ for $q > j$ in (5). When these are available in closed form we may compute (estimate) $\widehat{Y}_i := E^i \mathcal{Z}_{\hat{\tau}_i}$ via (standard) Monte Carlo simulation. After improving $\widehat{\tau}$ in turn via (5) again, we obtain a next improved stopping family $\widehat{\widehat{\tau}}$ via Monte Carlo simulation along each simulated Libor trajectory. We so arrive at a next improved estimation $\widehat{\widehat{Y}}_i$ via a nested Monte Carlo simulation.

Not in all situations closed form solutions (or close approximations) for the Europeans $E^j Z_q$, $q > j$, are known. For such cases it is usually better not to start (5) with the trivial stopping family above. For our applications we will consider three kinds of input families, respectively, due to an Andersen [1] like method, Piterbarg's [10] version of the Longstaff-Schwartz [9] regression method, and a backward optimization of the exercise boundary resulting from the latter approach.

Andersen-like method

As a starting policy one could take

$$\tau_i^A := \inf\{j \geq i : H_j \geq Z_j\}, \quad (7)$$

where the deterministic sequence H is precomputed via a standard optimization procedure as studied in [1] for Bermudan swaptions (see also [13]). The family $\widehat{\tau}$ obtained via (5) is then an improved exercise policy in the sense that $\widehat{Y}_i$, which requires nested Monte Carlo simulation, is generally closer to Y^*. We note that for interest rate contracts the cash-flow Z_{j+1} is typically already known at T_j (e.g. for the snowball). In such situation we replace Z_j by Z_{j+1} in (7).

Longstaff-Schwartz a-lá Piterbarg

We here formulate a version of the Longstaff-Schwartz algorithm designed for the cancelable path dependent product.

Let $(\beta_q)_{q\geq 1}$ be a system of base functions on the state space. Suppose we have a sample $L^{(m)}$, $m = 1, \dots, M$, of Libor trajectories. Set $\tau_k^{(m)} :\equiv k$, and $C_k^{(m)} :\equiv -1$, $m = 1, \dots, M$. For $j < k$ we recursively construct $\tau_j^{(m)}$, $C_j^{(m)}$, $m = 1, \dots, M$, from $\tau_{j+1}^{(m)}$, $C_{j+1}^{(m)}$, $m = 1, \dots, M$ as follows. Via a standard least squares minimization we compute a system of regression coefficients $(\widehat{c}_{jq})_{q\geq 1}$,

$$(\widehat{c}_{jq})_{q\geq 1} :=$$

$$\underset{(c_q), q\geq 1}{\operatorname{argmin}} \sum_{m=1}^{M} \left(\sum_{q\geq 1} c_q \beta_q (L^{(m)}(\mathcal{T}_j)) - \sum_{p=j+1}^{\tau_{j+1}^{(m)}} Z_p^{(m)} \right)^2,$$

and set

$$C_j^{(m)} := \sum_{q\geq 1} \widehat{c}_{jq} \beta_q (L^{(m)}(\mathcal{T}_j)), \quad m = 1, \dots, M.$$

We then define for $m = 1, \dots, M$,

$$\tau_j^{(m)} = j \quad \text{if} \quad C_j^{(m)} < 0, \quad \text{else} \quad \tau_j^{(m)} = \tau_{j+1}^{(m)}.$$

Thus working backwards, we end up with an (approximate) continuation rest-value process

$$C_j := \sum_{q\geq 1} \widehat{c}_{jq} \beta_q (L(\mathcal{T}_j)), \quad j = 1 \leq j < k, \quad C_k = -1.$$

By this we may obtain a lower biased approximation of the Bermudan price by an independent Monte Carlo re-simulation and using the stopping rule

$$\tau_i^{LS} = \inf\{j \geq i : C_j \leq 0\}. \tag{8}$$

For the typically high dimensional Libor process the choice of base functions and their number is in general a problematic issue. In order to keep the above regression method robust [10] suggests to consider base functions which are only defined on a small set of explanatory variables, though accepting a priori a bias in this way. As a generic choice he proposes the spot Libor $L_j(\mathcal{T}_j)$ and the swap rate over the period $[\mathcal{T}_j, \mathcal{T}_k]$. In our experiments we will see that we may so obtain relatively close lower bounds for the cancelable snowball swap, but, particularly in more factor cases these are not close enough.

Backward optimization of a given exercise boundary

Similar to the optimization procedure for (7) we may improve the exercise criterion obtained by the regression method once again to

$$\tau_i^{LS,A} = \inf\{j \geq i : C_j + \alpha_j \leq 0\}, \tag{9}$$

by backward optimization of the deterministic sequence α.

Our experiments show that the input stopping family (8) is generally insufficient in the sense that the gap due to the improved lower bound $\widehat{\tau}^{LS}$ and the dual upper bound corresponding to τ^{LS} is still too large. In all our cases, however, the gap due to $\widehat{\tau}^{LS,A}$ and the dual corresponding to $\tau^{LS,A}$ is acceptably small (though the lower bound due to $\tau^{LS,A}$ may be not sufficiently close). Therefore, we recommend (9) as a generically 'good' input stopping family.

Specification and valuation of the cancelable snowball swap

Starting (5) with one of the generic stopping families constructed above, a wide range of Libor exotics can be priced. As an example, let us consider a *snowball swap* contract on a \$1 nominal loan. According to this contract one receives floating Libor and has to pay so called Snowball coupons which follow the following term sheet. One pays on a semi-annual base a constant rate I over the first year and in the forthcoming years (Previous Coupon+A-Libor)$^+$, where A is specified in the contract. A *cancelable snowball swap* is a snowball swap which may be canceled after first year. We here consider this cancelable snowball product in a (semi-annual) Libor model (1). The snowball coupons K_i, settled at T_{i+1} ($i = 0, \dots, n-1$), are thus specified by

$$\begin{aligned} K_i &:= I, \quad i = 0, 1, \\ K_i &:= (K_{i-1} + A_i - L_i(T_i))^+ \quad i = 2, \dots, n-1. \end{aligned}$$

We consider a contract where A increases on an annual base according to $A_2 := S$, $A_{i+1} = A_i$ if i is even, and $A_{i+1} = A_i + s$ if i is odd, where S and s are given in the contract. The value V_0 of the cancelable snowball swap at $t = T_0 = 0$ is given by (2) with

$$Z_q := \frac{L_{q-1}(T_{q-1})\delta_{q-1} - K_{q-1}\delta_{q-1}}{B_*(T_q)}, \quad q = 1, \ldots, n.$$

Note that $Z_1, .., Z_n$ is an adapted (even predictable) sequence of cash-flows.

We now present different methods for computing V_0. To this end we consider the Markov process $(B_*(T_j), L(T_j), K_j)$ and evaluate V_0 via a Monte Carlo algorithm for the iteration procedure (5) using different starting policies described above.

Numerical results

We carry out simulation experiments for a 10yr Snowball with

$$I = 7\%, \quad S = 3\%, \quad s = 0.25\%,$$

and 19 semi-annual exercise possibilities starting at 1yr. In the Libor model (1) we take $\delta_i \equiv 0.5$, flat 3.5% initial Libor curve and constant volatility loadings

$$\gamma_i(t) \equiv 0.2e_i,$$

where e_i are d-dimensional unit vectors decomposing an input correlation matrix of rank d. We take as basis an endogenously full-rank correlation structure of the form

$$\rho_{ij} = \exp\left[\frac{|j-i|}{n-2}\ln\rho_\infty\right], \quad 1 \le i, j \le n-1. \quad (10)$$

with $n > 2$ and $\rho_\infty = 0.3$ (for more general correlation structures we refer to [13]). Then, for a particular choice of d we deduce from ρ in (10) a rank-d correlation matrix ρ^d with decomposition $\rho^d_{ij} = e_i \cdot e_j$, $1 \le i, j < n$, by principal component analysis.

We now investigate the stopping families τ^A, τ^{LS} and $\tau^{LS,A}$. For these stopping families, we construct lower bounds, their corresponding dual upper bounds, and the via (5) improved lower bounds.

From Table 1 we conclude that the stopping family τ^A leads to a very crude approximation of the Bermudan price, even for a 1-factor model. The improved stopping family $\widehat{\tau}^A$ provides a much better lower bound. However, the gap between $\widehat{Y}_0^A$ and the dual upper bound $Y^A_{up,0}$ is still rather large, it varies from 2% to 20% (relative to the upper bound prices). Here we use $2 \cdot 10^7$ Monte Carlo trajectories for Y_0^A and $6 \cdot 10^4$ Monte Carlo trajectories (with 500 inner simulations) for $Y^A_{up,0} - Y_0^A$ to keep the standard deviation within 0.5% relative. Further, we compute $\widehat{Y}_0^A$ using the variance reduction technique (6) by 10^5 outer and 500 inner simulations to keep the standard deviation within 1% relative.

We now consider price estimations via τ^{LS} obtained by Piterbarg's version of the Longstaff-Schwartz method. As basis functions we use quadratic polynomials in the explanatory variables $L_i(T_i)$, the swap rate over $[T_i, T_{20}]$, and the coupon K_i which is needed for Markovianity of the underlying process. In Table 2 we see that, compared to τ^A, the stopping family τ^{LS} provides better lower bounds and dual upper bounds (except for the 1-factor case). However, the gaps between Y_0^{LS} and $Y^{LS}_{up,0}$ are still large, they may even exceed 30% relative (in the 19-factor model). Again, the improved stopping family $\widehat{\tau}^{LS}$ leads to much better lower bounds. For this table we used 10^7 Monte Carlo trajectories for Y_0^{LS} and $2.5 \cdot 10^4$ Monte Carlo trajectories (with 500 inner simulations) for $Y^{LS}_{up,0} - Y_0^{LS}$ to keep the standard deviation within 0.5% relative. The improved lower bound $\widehat{Y}_0^{LS}$ is computed using variance reduction method (6) with $5 \cdot 10^4$ outer Monte Carlo trajectories and 500 inner simulations, to keep the standard deviation within 1% relative.

The optimization (9) of the Longstaff-Schwartz exercise boundary provides better lower bounds (6%-20% relative to Y_0^{LS}), see Table 3, but, the gaps between $Y_0^{LS,A}$ and $Y^{LS,A}_{up,0}$ are still rather large. The improved stopping family $\tau^{LS,A}$, however, leads to acceptable estimations of product prices: the gaps between $\widehat{Y}_0^{LS,A}$ and $Y^{LS,A}_{up,0}$ do not exceed 4% relative, and are overall less than 4 base points in absolute sense. We use 10^7 Monte Carlo trajectories for $Y_0^{LS,A}$ and 10^4 outer Monte Carlo trajectories with 500 inner simulations for $Y^A_{up,0} - Y_0^A$, in order to keep the standard deviation within 0.5% relative. Further, we construct $\widehat{Y}_0^A$ using variance reduction (6) by $5 \cdot 10^4$ outer and 500 inner simulations to keep the standard deviation within 1% relative.

Figure 1 shows the exercise frequency of the example product at different exercise dates for different policies in a full factor model $d = 19$. This picture gives an impression of the sensitivity (in a sense) of this product; we see that the snowball swap is mostly canceled quite early or very late. Although the exercise profile does not differ dramatically for the different stopping times, the comparably few scenarios, in which the τ^{LS}, respectively $\tau^{LS,A}$ exercise rule cancels much too late, has a significant influence on the price.

In Figure 2 we compare a conservative estimation of the duality gap $(Y_{up,0}^{LS,A} + 2(\mathrm{SD})) - (\widehat{Y}_0^{LS,A} - 2(\mathrm{SD}))$ with $(Y_{up,0}^{LS,A} + 2(\mathrm{SD})) - (Y_0^{LS,A} - 2(\mathrm{SD}))$ as a function of the number of simulated outer trajectories in the full factor case ($d = 19$). We see that with only 5000 trajectories, hence with much less computation time, we still have a major improvement upon the gap due to policy (9). On our Pentium-III computer 5000 outer simulations with 500 inner simulations took 15 minutes.

We now conclude with the slogan:

Iterate the stopping strategy obtained via an Andersen like enhanced Piterbarg version of Longstaff–Schwartz, and compute its dual due to Rogers, Haugh and Kogan.

The authors are with the Weierstrass Institute for Applied Analysis and Stochastics, Mohrenstr. 39, D-10117 Berlin, Germany.

Acknowledgment

The authors are grateful for interesting discussions with the Quantitative Analysis Dept. at Bankgesellschaft Berlin AG, on this subject. The paper strongly benefited from the constructive remarks of two anonymous referees. This work is supported by the DFG Research Center "Matheon" (FZT 86) in Berlin.

References

[1] Andersen, L. (1999) A simple approach to the pricing of Bermudan swaptions in a multifactor Libor market model. *Journal of Computational Finance*, **3**, 5-32

[2] Bender, C., Schoenmakers, J. (2006) An iterative algorithm for multiple stopping: Convergence and stability. *Advances in Appl. Prob.*, to appear

[3] Belomestny, D., Milstein, G.N. (2006) Monte Carlo evaluation of American options using consumption processes. *Int. Journal of Theoretical and Applied Finance*, to appear

[4] Duffie, D., *Dynamic Asset Pricing Theory.* Princeton: Princeton University Press 2001

[5] Haugh, M. B., Kogan, L. (2004) Pricing American options: A duality approach. *Operations Research*, **52**, 258-270

[6] Jamshidian, F. (1997) LIBOR and swap market models and measures. *Finance and Stochastics*, **1**, 293-330

[7] Jamshidian, F. (2004) Numeraire-invariant option pricing and American, Bermudan, and trigger stream rollover. Working paper, URL: http://wwwhome.math.utwente.nl/

[8] Kolodko, A., Schoenmakers, J. (2006) Iterative construction of the optimal Bermudan stopping time. *Finance and Stochastics*, **10**, 27-49

[9] Longstaff, F. A., Schwartz, R. S. (2001) Valuing American options by simulation: A simple least-square approach. *Review of Financial Studies*, **14**, 113-147

[10] Piterbarg, V.V. (2004) Pricing and hedging callable Libor exotics in forward Libor models, *Journal of Computational Finance,* **8** (2), 65-117

[11] Puterman, M. Markov decision processes. New York: Wiley 1994

[12] Rogers, L. C. G. (2002) Monte Carlo valuation of American options. *Math. Finance*, **12**, 271-286

[13] Schoenmakers, J. (2005) *Robust Libor Modelling and Pricing of Derivative Products.* Boca Raton, FL: Chapman & Hall – CRC Press

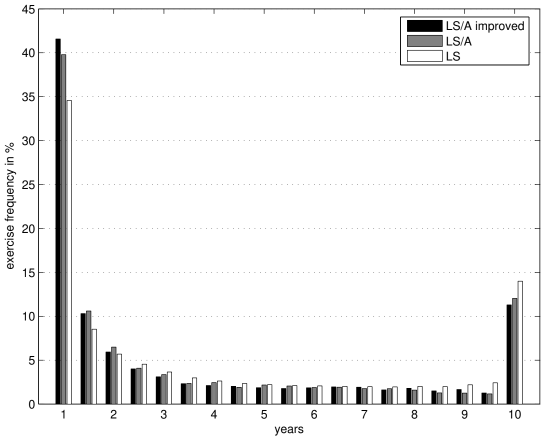

Figure 1: Simulated exercise frequencies of τ^{LS}, $\tau^{LS,A}$, and $\widehat{\tau}^{LS,A}$

d	Y_0^A (SD)	$\widehat{Y}_0^A$ (SD)	$Y_{up,0}^A$ (SD)
1	206.89(0.29)	215.66(0.44)	219.76(0.31)
2	122.79(0.26)	152.37(0.76)	169.44(0.52)
5	76.067(0.25)	116.19(0.84)	140.28(0.65)
10	64.32(0.24)	106.32(0.88)	129.60(0.65)
19	57.91(0.24)	97.81(0.87)	124.07(0.66)

Table 1: Price estimations via τ^A

d	Y_0^{LS} (SD)	$\widehat{Y}_0^{LS}$ (SD)	$Y_{up,0}^{LS}$ (SD)
1	202.35(0.41)	214.94(0.68)	221.63(0.47)
2	130.32(0.38)	152.81(0.76)	167.93(0.56)
5	94.35(0.36)	117.70(0.74)	135.77(0.58)
10	83.93(0.36)	109.46(0.75)	125.05(0.57)
19	77.54(0.36)	103.28(0.74)	119.78(0.58)

Table 2: Price estimations via τ^{LS}

d	$Y_0^{LS,A}$ (SD)	$\widehat{Y}_0^{LS,A}$ (SD)	$Y_{up,0}^{LS,A}$ (SD)
1	215.00(0.40)	216.78(0.70)	218.12(0.42)
2	150.26(0.37)	156.79(0.74)	159.02(0.45)
5	111.62(0.35)	123.22(0.87)	126.63(0.51)
10	100.27(0.34)	112.97(0.86)	116.23(0.54)
19	93.52(0.34)	106.47(0.84)	110.22(0.55)

Table 3: Backward optimization of the Longstaff-Schwartz exercise boundary

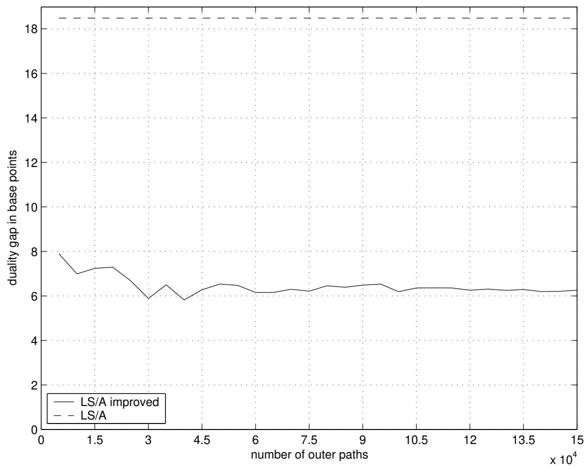

Figure 2: Estimated duality gap $(Y_{up,0}^{LS,A}+2(\mathrm{SD}))-(\widehat{Y}_0^{LS,A}-2(\mathrm{SD}))$ as a function of the number of simulated outer path compared with $(Y_{up,0}^{LS,A}+2(\mathrm{SD}))-(Y_0^{LS,A}-2(\mathrm{SD}))$