

Working Paper No. 22

# An Empirical Analysis of Valuation Algorithms for Pricing Callable Snowball Floaters

Damir Filipović, Nils Friewald, and Stefan Pichler

First version: September 2009

Current version: September 2009



# An Empirical Analysis of Valuation Algorithms for Pricing Callable Snowball Floaters

Damir Filipović\*, Nils Friewald†, Stefan Pichler‡

September 15, 2009

## Abstract

In this paper we value a callable snowball floater, a complex interest rate instrument with variable coupon payments, which depend on the prevailing interest rates in arrears and recursively on previous coupon payments. The embedded option requires solving an optimal stopping problem using the dynamic programming principle. A well-known and widely used algorithm to estimate conditional expectations is a specific form of least squares Monte Carlo simulation introduced by Longstaff and Schwartz (2001), which we refer to as the *LSM approach*. Contrary to the *standard approach*, where discounted option values of the subsequent period are regressed on the current state variables, Longstaff and Schwartz (2001) use the ex post realized payoffs of in-the-money option scenarios from continuation instead. They argue that, in doing so, they get values less than or equal to the value implied by the optimal stopping rule, which provides an objective convergence criterion.

We compare the LSM approach with the standard approach and use the price estimate from an elaborate nested Monte Carlo simulation as a benchmark. We empirically find that the LSM estimate of the embedded option might be considerably downward biased, whereas the standard estimate is much closer to the benchmark price. Moreover, we find that there is no optimal type of basis function that can generally be recommended for pricing interest rate instruments. Instead, we suggest using the LSM approach to determine the optimal type of basis function that results in the largest option value and rely on the standard approach to price the instrument. These are important issues to consider when pricing complex interest rate instruments, in general, and callable snowball floaters, in particular.

**JEL-Classification:** C15, G12, G13

**Keywords:** snowball floater, bermudan option, least squares Monte Carlo, nested Monte Carlo simulation

---

\*Vienna Institute of Finance, University of Vienna, and Vienna University of Economics and Business, Heiligenstädter Straße 46-48, 1190 Vienna, Austria; *email:* damir.filipovic@vif.ac.at

Damir Filipović gratefully acknowledges the financial support from WWTF (Vienna Science and Technology Fund).

†Vienna University of Economics and Business, Department of Finance and Accounting, Heiligenstädter Straße 46-48, 1190 Vienna, Austria; *email:* nils.friewald@wu.ac.at

‡Vienna University of Economics and Business, Department of Finance and Accounting, Vienna Graduate School of Finance, Heiligenstädter Straße 46-48, 1190 Vienna, Austria; *email:* stefan.pichler@wu.ac.at

# 1 Introduction

Snowball floaters are complex debt instruments with variable coupon payments, which depend on the prevailing interest rates (e.g., EURIBOR) in arrears and recursively also on previous coupon payments. A standard snowball floater promises high initial coupons, which then might get even larger if interest rates fall, but may also quickly melt down if interest rates rise. Coupons are floored at zero, thereby guaranteeing only positive payments to the investor. Furthermore, snowball floaters are often endowed with an option that allows the issuer to buy back the floater at par whenever it is favorable for her. Hence, a callable snowball floater is an early-exercise as well as a path-dependent instrument.

In general, path-dependent products can easily be priced using Monte Carlo techniques, which are based on forward propagation in time, whereas the early-exercise feature relies on backward propagation of the option value, which is typically implemented with binomial or trinomial trees. However, trees are not really suitable for underlying variables of high dimensionality such as interest rate instruments. Monte Carlo methods are much more convenient in this respect, but the early exercise feature must additionally be accounted for.

A number of studies address this issue in the literature: An early contribution is by Bossaerts (1989) who solves the optimal exercise strategy by simulation where he maximizes the option value. Other contributions are Tilley (1993), Broadie and Glasserman (1997), Broadie et al. (1997), Boyle et al. (1997), Rogers (2002), Andersen and Broadie (2004), Bally et al. (2005), etc. However, probably the most important papers on this topic are by Tsitsiklis and Van Roy (1999), Tsitsiklis and Van Roy (2001), Carriere (1996) and Longstaff and Schwartz (2001) which propose an approximation technique that relies on Monte Carlo simulation and accounts for the early exercise feature. The difficulty in pricing American options, for which there are no closed-form solutions, is to appraise the accuracy of the estimate. Various studies address this issue by producing lower and upper price bounds of the option value (e.g., Broadie and Detemple (1996), Jamshidian (2003), Haugh and Kogan (2004), Kolodko and Schoenmakers (2006), and Belomestny et al. (2008)). On the basis of these bounds we obtain convergence criteria, which provide guidance to assess the accuracy.

Interest rate instruments, such a snowball floaters, have gained increasing popularity in the recent past when the interest rate level was rather low. Institutional investors, such as corporate clients or hedge funds as well as wealthy private investors, sought for alternatives to earn above current market yields. Snowball floaters are very attractive in these environments because they usually promise high coupons. However, the risk an investor faces is substantial for two reasons: First, the complex coupon payment structure depends considerably on the interest rates. If it turns out that interest rates fall over time, coupons accumulate quickly, very much like a snowball, but they might also quickly melt down to zero when interest rates rise. Second, callable snowball floaters are endowed with an option, giving the issuer the right to buy back the floater at par at specific points in time (e.g., semi-annually). Hence, the investor is exposed to prepayment risk due to his short position in the call-back option.

In fact, investors who bought snowball floaters denominated in Euro during the low interest rate environment of, e.g. 2004–2005, faced considerable losses because it turned out that short-term

interest rates steadily grew in the market upturn. Hence, coupon payments immediately vanished. Furthermore, in a market of rising interest rates, investments in snowball floaters became more and more unattractive, which lead prices fall. Selling them off at a large discount was no alternative for many investors. Finally, with the onset of the sub-prime crisis central banks lowered key interest rates to stimulate the economy. This pushed up prices of snowball floaters again, but unfortunately, some issuers then exercised their right to buy back the floater at par and, hence, investors missed the opportunity to generate business in the low interest rate environment.

The actual prepayment risk as well as the risk of the complex coupon payments is often unpredictable for non-professional investors. Actually, some investors who faced considerable losses already brought issuers of snowball floaters to trial, suing the issuers for not having been properly and fully clarified about the embedded risks. In fact, alone the pricing of a snowball floater is demanding. Various different approaches are suggested in the academic literature, each having their pros and cons. Moreover, the price is very sensitive on changes in the level of interest rates, their volatilities and correlation.

With this paper we address three issues: First, by focusing on a specific callable snowball floater we analyze the sensitivity of the price and the value of the embedded option to changes in the market environment. This reveals the main driving factors and provides important insights about the functionality and the risk of callable snowball floaters. In doing so, we estimate the price of the embedded option on a quarterly basis over the floater's lifetime between October 10, 2005 and April 10, 2008. Moreover, we conduct a scenario analysis where we use the market data of the issue date as the base case and then shift or twist the forward rate curve up and down or scale the forward rate volatilities. The second objective of this study is to highlight the importance of decomposing complex instruments with embedded options, such as callable snowball floaters into two parts, i.e. the non-callable instrument and the Bermudan option, and price each part separately. We demonstrate that if we directly solve the optimal stopping problem on the callable instrument, we would drastically underestimate its price.

The third issue we focus on is to compare different valuation algorithms. We use a popular and wide-spread valuation algorithm proposed in the standard academic literature to price Bermudan options, namely least squares Monte Carlo simulation. This algorithm is applied to estimate conditional expectations in a Monte Carlo simulation framework, which then allows to judge if either immediate exercise or continuation is more valuable. In particular, to value the embedded option we focus our analysis on a specific least squares Monte Carlo simulation algorithm introduced by Longstaff and Schwartz (2001), which we refer to as the *LSM approach*. In contrary to a *standard approach* introduced by Tsitsiklis and Van Roy (1999), where discounted option values of the subsequent period are regressed on basis functions evaluated on the current state vector, Longstaff and Schwartz (2001) suggest to use the ex post realized payoffs of only in-the-money option scenarios to estimate the continuation value. They point out, that their approach is, on the one hand, computationally more efficient and, on the other hand, yields downward biased option values which can then be used to control accuracy when selecting the optimal type and number of basis functions. However, even if their approach is very promising it still needs further empirical testing, particularly for instruments that are based on interest rates (see, e.g., Brigo and Mercurio

(2006)). Hence, the third objective is to compare the approach of Longstaff and Schwartz (2001) with the standard approach.

As an objective criterion to evaluate the different approaches, we use an elaborate nested Monte Carlo or random-tree simulation to obtain the callable snowball floater price and the value of the embedded option. This is the most accurate estimate we get, but admittedly, computationally also the most expensive one. Therefore, we consider the nested Monte Carlo estimate as the benchmark. In addition, we also provide an easy to compute lower bound of the floater value, by conditioning on the complete knowledge of the simulated paths. This perfect foresight of interest rates allows choosing the best exercise strategy among all strategies and, hence, yields the largest option value and the lowest achievable price for the callable snowball floater, respectively.

We find in our empirical analysis that the LSM algorithm yields price estimates that are substantially downward biased, while estimates obtained from the standard approach are much more accurate. We also find that there is no optimal type of basis function that can always be recommended to price interest rate instruments. Instead, we suggest using the LSM algorithm to decide which basis function to use and then rely on the standard approach to price the instrument. Moreover, we demonstrate the importance to decompose callable instruments into a non-callable part and the embedded option and price each part separately. One can lower the standard errors of the callable floater when the embedded option is valued based on the same simulated paths as the non-callable instrument because both values are positively correlated.

The paper is structured as follows: In Section 2 we introduce a generic valuation framework for pricing callable snowball floaters while relying on a Libor market model. In Section 3 we discuss two algorithms for solving the dynamic programming problem: nested Monte Carlo simulation and least squares Monte Carlo simulation. Section 4 presents the methodology and the data used for pricing the callable snowball floater. Section 5 reports the findings of the empirical analysis, where we conduct a time-series, scenario and basis function analysis. Section 6 concludes and the appendix contains the figures and tables.

## 2 Valuation Framework

In this section we discuss the framework to value a callable snowball floater. We assume a finite time interval  $[T_0, T_m]$ , where  $T_0$  is the current time and  $T_m$  the maturity date of the callable snowball floater. Moreover, we consider a complete probability space  $(\Omega, \mathcal{F}, \mathbb{Q})$ , where  $\Omega$  is the set of all possible realizations of the economy,  $\mathcal{F}$  is the sigma field of distinguishable events of the economy at time  $T$ , and  $\mathbb{Q}$  the associated probability measure. The probability space is equipped with a discrete time filtration  $(\mathcal{F}_{T_j})_{j=0, \dots, m}$ .

A snowball floater is a path-dependent instrument with variable coupon payments  $K(T_i)$ ,  $i = 1, \dots, m$  that are contingent on the prevailing interest rates in arrears, and recursively also on the previous coupon payments.<sup>1</sup> A typical snowball floater is endowed with an option that allows the issuer to buy back the floater at par on a fixed set of exercise opportunities or stopping times

---

<sup>1</sup>For a description of the exact coupon payment structure of the specific snowball floater under consideration, see Section 4.1.

$\mathcal{T}_{k,m}$  with values in  $\{T_k, \dots, T_m\}$ . Usually, the option has a lock-out period  $T_k$ , meaning that the issuer cannot call back the floater before  $T_k$ . Since the option can be exercised only at discrete points in time, it is called a Bermudan option. We assume that the set of stopping times  $\mathcal{T}_{k,m}$  is a subset of the coupon payment dates  $T_i$ , which in reality is most common case.

We price the snowball floater under the Libor market model, which we now briefly recap: Let us denote by  $D_i(t)$  the price at time  $t$  of a discount bond with maturity  $t_i$  and by  $L_i(t)$  the forward Libor rate, which is the interest rate fixed at time  $t$  for the interval  $[t_i, t_i + \delta_i]$ :

$$L_i(t) = \frac{D_i(t) - D_{i+1}(t)}{\delta_i D_{i+1}(t)}, \quad (1)$$

where  $\delta_i = t_{i+1} - t_i$  denotes the accrual period measured in years. Under the spot measure  $\mathbb{Q}$  the dynamics of the forward Libor rate are given by

$$\frac{dL_n(t)}{L_n(t)} = \sum_{i=\eta(t)}^n \frac{\delta_i L_i(t) \sigma_n(t)^\top \sigma_i(t)}{1 + \delta_i L_i(t)} dt + \sigma_n(t)^\top dW(t), \quad 0 \leq t \leq t_n, \quad n = 1, \dots, m \quad (2)$$

where  $\eta(t)$  denotes the next following reset date at time  $t$  and  $W$  is a  $d$ -dimensional standard Brownian Motion.<sup>2</sup> The spot measure  $\mathbb{Q}$  is associated with a simply compounded bank account  $B(t)$  as the numeraire asset which is defined as

$$B(t) = B_{\eta(t)}(t) \prod_{i=0}^{\eta(t)-1} [1 + \delta_i L_i(T_i)], \quad (3)$$

where  $\eta(t)$  gives the index of the next tenor date at time  $t$ . The volatility  $\sigma_n$  depends on  $t$  but may also depend on the current Libor rates  $L_i(t)$ . However, in our setup the volatility  $\sigma_n$  is assumed to be deterministic and piecewise constant.

The callable snowball floater can either be directly priced using the dynamic programming principle on the floater payoff or by first decomposing the floater into two parts, i.e. a corresponding non-callable floater and the embedded option, where each part then is valued separately. We discuss each approach in detail in the following sub-sections and also provide an easy to compute lower bound of the callable floater price.

## 2.1 Direct Valuation of the Callable Floater

Following the direct valuation approach the value of the callable snowball floater is given by minimizing the total discounted cash flow to the investor over all admissible stopping times  $\tau$ . Hence, we introduce the payoff process  $(Z(T_j))_{k, \dots, m}$  defined as the sum of all previously paid coupons  $K(T_i)$  plus a notional payment of 1 on exercise, measured in time-zero dollars, i.e.

$$Z(T_j) = \sum_{i=1}^j \frac{K(T_i)}{B(T_i)} + \frac{1}{B(T_j)}. \quad (4)$$

---

<sup>2</sup>We simulate the Libor market model under the spot instead of the forward measure, because the bias due to the discretization error in the drift is more distributed among the different Libor rates (see, e.g., Brigo and Mercurio (2006)). Moreover, Glasserman and Zhao (2000) find that simulating under the spot measure usually results in smaller variance.

The payoff  $Z(T_j)$  is paid by the issuer to the investor if the bank exercises the call-back option at time  $T_j$ . Let us define another process  $H(T_j)$  as the largest submartingale dominated by  $Z$ . Then the pricing problem can be formulated as:

$$H(T_j) = \begin{cases} \inf_{\tau \in \mathcal{T}_{j,m}} \mathbb{E} [Z(\tau) | \mathcal{F}_{T_j}] & k \leq j \leq m \\ \mathbb{E} [H(T_k) | \mathcal{F}_{T_j}] & j < k \end{cases} \quad (5)$$

This is an optimal stopping problem which can be solved using the backward dynamic programming principle where  $H(T_j)$  is the value of the callable snowball floater at  $T_j$  assuming the floater has not been called previously. The process  $H(T_j)$  for  $T_k \leq T_j < T_m$  is defined by a terminal condition and backward induction:

$$\begin{cases} H(T_m) = Z(T_m) \\ H(T_j) = \min \{ Z(T_j), \mathbb{E} [H(T_{j+1}) | \mathcal{F}_{T_j}] \}, j = k, \dots, m-1. \end{cases} \quad (6)$$

Note that, by construction,  $H(T_m)$  equals the total cash flow paid to the investor until expiration, i.e., the sum of the coupon payments up to maturity time  $T_m$  including the notional value of 1. For exercise dates  $T_j$ ,  $T_k \leq T_j < T_m$ , the value of the callable snowball floater is the minimum of (a) the sum of the coupon payments and the notional when exercising immediately and (b) the expected cash flow to be paid when continuing. Since the option cannot be exercised before  $T_k$  the dynamic programming stops at  $T = T_k$ . Hence, the value of the callable snowball floater at time  $T_0$  is given by

$$V^{dir}(T_0) = \mathbb{E}[H(T_k) | \mathcal{F}_{T_0}]. \quad (7)$$

We may also write this pricing problem in terms of stopping rules:

$$\tau_j = \min \{ T_n \geq T_j | H(T_n) = Z(T_n) \}, \quad (8)$$

and the callable snowball floater value is then given by

$$V^{dir}(T_0) = \mathbb{E} [Z(\tau_k) | \mathcal{F}_{T_0}]. \quad (9)$$

## 2.2 Valuation by Decomposing the Callable Floater

Instead of solving the dynamic programming problem given in (6) we now first decompose the callable snowball floater value into the difference of a non-callable snowball floater  $V^{NC}(T_0)$  and the value of the Bermudan call-back option  $C(T_0)$ . This is a standard procedure for pricing instruments with embedded options. It turns out that, in doing so, we empirically get more accurate results (see Section 5.3). We first define the price of the non-callable floater at time  $T_j$  as the expected value of the remaining coupon payments  $K(T_i)$  plus the notional payment of 1 at maturity  $T_m$  measured in time-zero dollars, i.e.

$$V^{NC}(T_j) = \mathbb{E} \left[ \sum_{i=j+1}^m \frac{K(T_i)}{B(T_i)} + \frac{1}{B(T_m)} | \mathcal{F}_{T_j} \right], \quad (10)$$

where by convention, we set  $\sum_{i=j+1}^m \cdot = 0$  for  $i > m$ , hence  $V^{NC}(T_m) = 1/B(T_m)$ . For any stopping time  $\tau_k$  we may alternatively write for  $V^{dir}(T_0)$

$$\begin{aligned}
\mathbb{E} \left[ \sum_{i=1}^{\xi(\tau_k)} \frac{K(T_i)}{B(T_i)} + \frac{1}{B(\tau_k)} \middle| \mathcal{F}_{T_0} \right] &= \mathbb{E} \left[ \sum_{i=1}^m \frac{K(T_i)}{B(T_i)} + \frac{1}{B(T_m)} \middle| \mathcal{F}_{T_0} \right] \\
&\quad - \mathbb{E} \left[ \mathbb{E} \left[ \sum_{i=\xi(\tau_k)+1}^m \frac{K(T_i)}{B(T_i)} + \frac{1}{B(T_m)} \middle| \mathcal{F}_{\tau_k} \right] - \frac{1}{B(\tau_k)} \middle| \mathcal{F}_{T_0} \right] \\
&= V^{NC}(T_0) - \mathbb{E} \left[ V^{NC}(\tau_k) - \frac{1}{B(\tau_k)} \middle| \mathcal{F}_{T_0} \right] \\
&= V^{NC}(T_0) - C(T_0), \tag{11}
\end{aligned}$$

where  $\xi(\tau)$  gives the index of the stopping time  $\tau$ . Thus, the option value  $C(T_0)$  is defined as the expected difference of the non-callable snowball floater at time  $\tau_k$  and the notional value of 1 in time-zero dollars. To compute the option value we define a virtual payoff function  $Z(T_j)$ , i.e.

$$Z(T_j) = V^{NC}(T_j) - \frac{1}{B(T_j)}. \tag{12}$$

It is virtual in the sense that an option holder on exercise would receive the remaining coupons of a non-callable snowball floater and pays par. We introduce the Snell envelope  $(U(T_j))_{j=k, \dots, m}$  for the corresponding payoff process  $(Z(T_j))_{j=k, \dots, m}$ . The option value is then given by the following optimal stopping problem:

$$U(T_j) = \begin{cases} \sup_{\tau \in \mathcal{T}_{j,m}} \mathbb{E}[Z(\tau) | \mathcal{F}_{T_j}] & k \leq j \leq m \\ \mathbb{E}[U(T_k) | \mathcal{F}_{T_j}] & j < k \end{cases} \tag{13}$$

which, again, can be solved using the backward dynamic programming scheme. Note, that the Snell envelope at the terminal time  $T_m$  is 0, by construction, i.e. the option value is worthless at maturity  $T_m$ . At intermediate times  $T_k \leq T_j < T_m$  the price of the callable snowball floater is the maximum of the continuation value  $\Gamma(T_j) = \mathbb{E}[U(T_{j+1}) | \mathcal{F}_{T_j}]$  and the exercise value  $Z(T_j)$ :

$$\begin{cases} U(T_m) = Z(T_m) = 0 \\ U(T_j) = \max\{Z(T_j), \Gamma(T_j)\}, j = k, \dots, m-1 \end{cases} \tag{14}$$

The option value at time  $T_0$  is then given either by the Snell envelope  $U(T_k)$  or in terms of stopping rules, i.e.

$$C(T_0) = \mathbb{E}[U(T_k) | \mathcal{F}_{T_0}] = \mathbb{E}[Z(\tau_k) | \mathcal{F}_{T_0}]. \tag{15}$$

Hence, the value of the callable snowball floater is defined as the difference between the value of the non-callable snowball floater  $V^{NC}(T_0)$  and the call-back option  $C(T_0)$ :

$$V(T_0) = V^{NC}(T_0) - C(T_0) \tag{16}$$

### 2.3 Lower Bound

We next find an easy to compute lower bound  $V^{low}(T_0)$  for  $V(T_0)$  by maximizing the Bermudan call-back option or minimizing the cash flow of the callable floater to the investor, respectively, over all paths conditional on the complete knowledge of the entire path, i.e.

$$V^{low}(T_0) = \mathbb{E} \left[ \min_{k \leq i \leq m} \left\{ \sum_{j=1}^i \frac{K(T_j)}{B(T_j)} + \frac{1}{B(T_i)} \right\} | \mathcal{F}_{T_0} \right]. \quad (17)$$

The perfect foresight of the evolution of the interest rates allows to choose the best exercise strategy among all strategies and, hence, yields the lowest achievable price for the callable snowball floater and the highest option value, i.e.  $C^{up}(T_0) = V^{NC}(T_0) - V^{low}(T_0)$ .

## 3 Valuation Algorithms

While in the last section we provided a framework to value a generic callable snowball floater, we now focus on two valuation algorithms to implement the pricing framework. Since a callable snowball floater is not just a path-dependent, but also an early-exercise instrument, we need to solve an optimal stopping problem. This involves the estimation of the continuation value such that we can decide if either immediate exercise or continuation is more valuable. This value can be estimated by a number of approaches suggested in the literature, each having their pros and cons. In this study we focus on two algorithms: First, nested Monte Carlo or random-tree simulation and, second, least squares Monte Carlo simulation. We discuss each of these approaches in more detail in the following two sub-sections.

### 3.1 Nested Monte Carlo Simulation

Nested Monte Carlo simulation is a pure simulation approach to obtain estimates of conditional expectations. The procedure works as follows: Starting from an initial state  $X_0$  which corresponds to, e.g., the forward Libor rates  $L_i(T_0)$ ,  $i = 0, \dots, m$ , we simulate  $b$  independent subsequent states  $X_1(1), \dots, X_1(b)$ , and then, again, for each of the subsequent states  $X_1(i)$  we draw  $b$  further paths  $X_2(i1), \dots, X_2(ib)$ , given the most recent information available at time  $T_1$ .<sup>3</sup> The estimated continuation value is given by the average of the discounted value functions plus the coupon payment of the subsequent states. Hence, we need to branch our tree at each exercise date. For a given number of branches  $b$  at each node the demand of computational power grows exponentially as the number of exercise dates increases.<sup>4</sup> Obviously, such an approach is computationally very expensive and time-consuming and, hence, in practice often simply not applicable.<sup>5</sup> However, for

<sup>3</sup>Here we use the notation  $X_j(i_1 \dots i_j)$  to refer to the state vector at time  $T_j$ , where  $i_1 \dots i_j$  denotes the sequence of paths in the random tree to node  $i_j$  at time step  $T_j$ , i.e. for a branching parameter  $b = 3$  we have  $i \in \{1, 2, 3\}$ .

<sup>4</sup>For example, given a maturity of 8 years with semi-annual exercise opportunities and a lock-out period of one year, we have  $k = 15$  exercise dates. Assuming a branching parameter of  $b = 2$ , i.e. for each state we simulate two subsequent states, we need to simulate  $b^k = 32768$  paths in total. For  $b = 3$  we already need to simulate 14.3 million paths.

<sup>5</sup>Given that our specific snowball floater has, when issued,  $k = 15$  exercise opportunities, we apply a branching parameter of  $b = 2$ . However, for lower time-to-maturities we adjust the branching parameter accordingly, such

this study we nevertheless rely on a nested Monte Carlo simulation to value our snowball floater because the obtained value is the most accurate achievable price. We consider the estimate of the nested Monte Carlo simulation our benchmark against other valuation algorithms such as the least squares Monte Carlo approach.

Instead of estimating the value of the callable snowball floater based on a direct approach, we use the more accurate approach and value the embedded option  $C(T_0)$  and the non-callable snowball floater  $V^{NC}(T_0)$  separately, as discussed in Section 2. To solve the dynamic programming problem given in (14) we need to know the conditional expectation of the non-callable snowball floater  $V^{NC}(T_j)$  and the continuation value  $\Gamma(T_{j+1})$  for each time step. We get both implicitly from the nested Monte Carlo simulation. The estimate of the conditional non-callable snowball floater value at a particular node  $i_j$  is approximated by

$$\widehat{V}_j^{NC}(i_1 \dots i_j) = \frac{1}{b} \sum_{i=1}^b \widehat{V}_{j+1}^{NC}(i_1 \dots i_j i) + \frac{K_{j+1}(i_1 \dots i_j i)}{B(i_1 \dots i_j i)}. \quad (18)$$

At a final node  $i_m$ , the value of the non-callable floater equals the notional value and, hence, the option value is zero. Before expiration an estimate of the continuation value  $\Gamma_j$  is given by averaging the option values of the subsequent step, i.e.

$$\widehat{\Gamma}_j(i_1 \dots i_j) = \frac{1}{b} \sum_{i=1}^b \widehat{U}(T_{j+1})(i_1 \dots i_j i). \quad (19)$$

Finally, the value of the callable floater  $V^{nmc}(T_0)$  is then given by the difference of the non-callable floater  $\widehat{V}^{NC}(T_0)$  and the Bermudan option  $\widehat{C}^{nmc}(T_0)$ . The estimate of the embedded option value obtained by a nested Monte Carlo simulation is consistent but biased high because the same information is used in deciding whether to exercise immediately as for estimating the continuation value. To put it differently, the estimation of the continuation value is based on information from subsequent nodes of the state variable, i.e. future information is used to decide if exercise is preferable or not (for a formal proof see Broadie and Glasserman (1997)). We produce another consistent estimator, one that is biased low but also converges to the true value. In general, to remove the high bias we need to estimate the value of continuation and the value on which we decide to exercise on different samples. There are various ways to do this, but we follow the method described in Broadie and Glasserman (1997). They use all but one of the following paths for the exercise decision and the remaining path to estimate the continuation value. Again, we refer to Broadie and Glasserman (1997) for a formal proof, verifying that this approach yields an estimate biased low.

Given the downward biased estimate  $\widehat{C}_L^{nmc}$  and the upward biased estimate  $\widehat{C}_H^{nmc}$  of the embedded call-back option at time  $T_0$ , we know for  $n$  simulation runs and branching parameter  $b$  that:

$$\mathbb{E} \left[ \widehat{C}_L^{nmc}(b, n) \right] \leq C \leq \mathbb{E} \left[ \widehat{C}_H^{nmc}(b, n) \right] \quad (20)$$

---

that the computational burden is approximately the same. For example, for  $k = 8$  we use a branching parameter of  $b = 4$  to further increase accuracy.

This allows us to compute a  $(1 - \alpha)$  confidence interval for the call-back option value, i.e.

$$\left( \widehat{C}_L^{nmc}(b, n) - z_{\alpha/2} \frac{\sigma_L^{nmc}}{\sqrt{n}}, \widehat{C}_H^{nmc}(b, n) + z_{\alpha/2} \frac{\sigma_H^{nmc}}{\sqrt{n}} \right), \quad (21)$$

where  $z_{\alpha/2}$  denotes the  $\alpha/2$ -quantile of the standard normal distribution. The standard deviations  $\sigma_L^{nmc}$  and  $\sigma_H^{nmc}$  can be estimated from  $n$  independent runs of the nested Monte Carlo simulation. The true value  $C$  is within the confidence interval with a probability of at least  $1 - \alpha$ . Note, that as the branching parameter  $b \rightarrow \infty$  and the number of paths  $n \rightarrow \infty$  the interval collapses to the point  $C(T_0)$ . We finally use the mid-point as a point estimator for the embedded option price, i.e.

$$\widehat{C}^{nmc}(T_0) = 1/2 \left( \widehat{C}_L^{nmc} + \widehat{C}_H^{nmc} \right). \quad (22)$$

Despite the fact that the computation of a mid-point is somewhat arbitrary, Glasserman (2004) shows that it is highly effective.<sup>6</sup>

### 3.2 Least Squares Monte Carlo Simulation

The least squares Monte Carlo technique is a prominent approach, which is more appealing in practice than nested Monte Carlo simulation. Here, the continuation value  $\Gamma(T_j)$  is estimated by regressing the value function of the subsequent period on a vector of basis functions  $\Psi$ . Clément et al. (2002) point out that the least squares Monte Carlo approach involves two approximations: First, the conditional expectation  $\Gamma(T_j) = \mathbb{E}[U(T_{j+1})|\mathcal{F}_{T_j}]$  is replaced by an orthogonal projection  $P_j$  on a vector of  $d$  basis functions  $\Psi^d$  evaluated on the current state vector  $X(T_j)$ , i.e.

$$P_j(U(T_{j+1})) = \alpha_j \Psi^d(X(T_j)), \quad (23)$$

where the  $d$  dimensional parameter  $\alpha_j$  is given by

$$\alpha_j = (A_j)^{-1} \mathbb{E} [\Psi^d(X(T_j))U(T_{j+1})], \quad (24)$$

and the  $d \times d$  matrix  $A_j$  is defined as<sup>7</sup>

$$A_j = \mathbb{E} [\Psi^d(X(T_j))\Psi^d(X(T_j))^\top]. \quad (25)$$

We could use different basis functions for different exercise dates. However, we stick on the same type of basis function when solving the dynamic programming problem. The second approximation involves Monte Carlo simulation and least squares regression to estimate the continuation value  $\Gamma(T_j)$ . Hence, an estimate for  $\alpha_j$  is obtained by minimizing the squared differences between the value function  $\widehat{U}(T_{j+1})$  and the linear combination of the  $d$  basis functions  $\Psi^d(X(T_j))$ , i.e.

<sup>6</sup>In order to get sufficient accuracy we simulate 1000 random trees. Furthermore, to decrease standard errors we use antithetic sampling for branching parameters  $b > 2$ , i.e. to compute the low estimates we average over the antithetic pairs. Note, that for  $b = 2$  we cannot compute low biased estimates using antithetic sampling.

<sup>7</sup>Due to numerical issues, we might experience singularities for  $A$ . Hence, we calculate the Moore-Penrose generalized inverse of the matrix  $A$ .

$$\hat{\alpha}_j = \arg \min_{\alpha_j \in \mathbb{R}^d} \sum_{i=1}^n \left( \hat{U}^{(i)}(T_{j+1}) - \alpha_j \cdot \Psi^d(X^{(i)}(T_j)) \right)^2. \quad (26)$$

This is a rather general description of using least squares Monte Carlo simulation to estimate Bermudan options which was first applied by Tsitsiklis and Van Roy (1999). We denote this method the *standard algorithm* to distinguish it from a more specific approach introduced by Longstaff and Schwartz (2001). The difference is that they regress the ex post realized payoffs  $Z(\tau_{j+1})$  on basis functions evaluated on the state variables. They point out that their algorithm results in option values that are less than or equal to the value implied by the optimal stopping rule, which provides an objective criterion for convergence. We call their approach the *LSM algorithm*. Moreover, Longstaff and Schwartz (2001) recommend to use only in-the-money option scenarios ( $Z(T_j) > 0$ ) for regression to improve computational efficiency. Hence, we estimate  $\hat{\alpha}^{lsm}$  by minimizing the following problem:

$$\hat{\alpha}_j^{lsm} = \arg \min_{\alpha_j^{lsm} \in \mathbb{R}^d} \sum_{i=1}^n \left( \hat{Z}^{(i)}(\tau_{j+1}) - \alpha_j^{lsm} \cdot \Psi^d(X^{(i)}(T_j)) \right)^2. \quad (27)$$

We use the least squares Monte Carlo technique once to value the callable snowball floater directly by solving the optimal stopping problem given in (6), and once where we first decompose the floater into a non-callable floater and the embedded option. Valuing the embedded option corresponds to the optimal stopping problem given in (14). The latter approach requires the knowledge of the conditional expected values of the non-callable snowball floater  $V^{NC}(T_j)$  which can also be obtained by least squares regression. We compute the projection of the non-callable snowball floater value  $V^{NC}(T_j)$  on a set of basis functions  $\Psi^d(X(T_j))$ , i.e.

$$P_j(V^{NC}(T_j)) = P_j(V^{NC}(T_{j+1}) + K(T_{j+1})) = \beta_j \Psi^d(X(T_j)), \quad (28)$$

where the  $d$  dimensional parameter  $\beta_j$  is given by

$$\hat{\beta}_j = \arg \min_{\beta_j \in \mathbb{R}^d} \sum_{i=1}^n \left( \hat{V}^{NC,(i)}(T_{j+1}) + K^{(i)}(T_{j+1}) - \beta_j \cdot \Psi^d(X^{(i)}(T_j)) \right)^2. \quad (29)$$

## 4 Methodology

### 4.1 The Callable-Snowball Floater Under Consideration

We conduct the analyses in this study on a specific callable snowball floater, which was issued by Bank Austria Creditanstalt AG in October 10, 2005.<sup>8</sup> The instrument has a maturity of  $T_m = 8$  years and semi-annual coupon payments (i.e., each October 10th and April 10th) which depend on the prevailing 6-month EURIBOR in arrears and recursively on previous coupon payments by the following relation:

$$K(T_i) = \delta_i \cdot (2K(T_{i-1}) + r(T_i) - L_i(T_i))^+, \quad 2 < i \leq 16, \quad (30)$$

<sup>8</sup>ISIN: AT0000248687.

where  $K(T_{i-1})$  is the previously paid coupon,  $L_i(T_i)$  the 6-month EURIBOR valid for the interval  $[T_i, T_i + \delta_i]$  where we assume that  $\delta_i = 0.5, \forall i$ . The fixed interest rate premia  $r(T_i)$  at the coupon payment dates  $T_i$  are given in Table 1. Note, that the investor receives fixed coupon payments of 5.25% p.a. for the first two coupon payment dates  $i = \{1, 2\}$ . Furthermore, the floater is endowed with an option giving the issuer the right to buy back the floater at par from which on no further coupon payments are made. The option can be exercised only on the coupon payment dates. However, there is a lock-out period of one year, meaning that the first exercise date is not before one year. This guarantees that the investor, at least initially, receives two rather large coupon payments.

To price the callable snowball floater we use a Libor market (or log normal forward-Libor) model where we calibrate the model on observable cap prices with maturities between 3 and 8 years. In doing so, we need market discount factors which we compute based on the 6- and 12-month EURIBOR and by bootstrapping the swap rates on the 6-month EURIBOR with a maturity spectrum between 2 and 8 years. We obtain all these data from Datastream. Cap prices allow us to estimate the forward rate volatilities. To keep things simple we do not account for forward rate correlations. The accuracy of the algorithms, however, is not affected by the, admittedly, more realistic assumption of correlated forward Libor rates. We estimate the forward rate volatilities by a bootstrapping approach where we assume a piecewise constant volatility structure. This simplification does also not affect the conclusions drawn from the empirical results.

Based on a simple time-series analysis, we first examine the main driving forces of the floater price based on the price estimate obtained from a nested Monte Carlo simulation. In particular, we show how coupon payments, the value of the call-back option and the forward rates affect the callable snowball floater price. In doing so, we gain important insights about the instrument's functionality and the risk an investor is exposed to. We price the floater on a quarterly basis over its lifetime, starting on October 10, 2005 and ending on April 10, 2009, when the floater finally was called back by the issuer. A scenario analysis permits us to gain further important insights. In this analysis we use the market values of the issue date (i.e. October 10, 2005) as the base case. We then move the forward curve up and down by a parallel shift of 25 and 50 bp, respectively. In a second step we twist the forward curve up and down at the short end by, again, 25 and 50 bp, where we set the pivot point equal to half the initial time-to-maturity. In a third step we scale the volatility of all forward Libor rates by a factor of 150% and 50%. On the basis of these scenarios we examine the impact of a change in the market environment on the price of a callable snowball floater and the embedded option.

## 4.2 The Choice of Appropriate Basis Functions

To price path-dependent instruments with early-exercise features least squares Monte Carlo simulation is widely used in practice. This requires choosing a specific type and number of basis functions to estimate the continuation value. Recent studies (e.g., Clément et al. (2002)) show that least squares Monte Carlo simulation yields unbiased estimates when the number of basis functions and paths approaches infinity. Clearly, for practical purposes, one is restricted on a finite set of basis functions and a limited number of paths. Moreover, the accuracy of the estimates

does not just depend on the number of basis functions but also on the specific choice. Since it is still not clear which type of basis function serves best to approximate conditional expectations for pricing interest rate derivatives with embedded options, we analyze various different types of basis functions given in Table 2.

We start the analysis using a simple linear combination of the level and the slope of the interest rate curve, as some studies suggest that such a restricted set of basis functions already provides sufficient accuracy for pricing exotic interest rate derivatives. We approximate the average interest rate level by the swap rate  $S_m$  of a swap with a maturity  $T_m$ , which corresponds to the remaining maturity of the floater. The slope of the interest rate curve is approximated by the difference between the swap rate  $S_m$  and the spot Libor rate  $L_i(T_i)$ . To distinguish it from other types of basis functions we denote this set with  $\Psi_1^p(S_m, S_m - L_i)$ , where we introduce the notation that the super-index indicates the type of function and the sub-index its order.

In a next step we extend the number of basis functions using a linear combination of the simulated forward Libor rates or discount factors including a constant. These sets are denoted by  $\Psi_1^p(L_{i,\dots,m})$  and  $\Psi_1^p(D_{i,\dots,m})$ . Hence, at any time  $T_i$  we regress on  $m - i + 2$  basis functions, where  $T_m$  refers to the maturity, i.e. one time step before maturity we have three basis functions at hand, because  $m = 16$  and  $i = 15$ . Moving stepwise forward in time, the number of basis functions increases by one. Right after the lock-out period we regress on 16 basis functions.

We also analyze other more sophisticated types of basis functions using polynomials of order two, including quadratic polynomials  $\Psi_2^p$ , weighted Laguerre polynomials  $\Psi_2^{wl}$ , and Fourier series  $\Psi_2^f$ . Again, these values are evaluated on the simulated forward Libor rates or discount factors. Recent numerical tests (e.g., Longstaff and Schwartz (2001)) indicate that a Fourier series or a simple power transformation already provides sufficient accurate estimates. We compare the efficiency of all these basis functions by computing the option price based on the LSM algorithm, because it provides a useful convergence criterion which gives guidance in determining the type and the required number of basis functions to obtain a desired level of accuracy. Longstaff and Schwartz (2001) show that their algorithm results in low biased estimates of the option price. The criterion works as follows: Simply increase the number of basis functions until the estimate no longer increases. We analyze the efficiency of these basis functions empirically by quarterly estimating the option price over the floater's lifetime.

### 4.3 An Analysis of Valuation Algorithms

A snowball floater can either be valued directly or by first decomposing the instrument into a non-callable snowball floater and the embedded option and then pricing each part separately. The former approach relies on solving the optimal stopping problem given in (6), whereas for the latter we need to solve the problem given in (14) and, additionally, estimate the conditional value of the non-callable instrument. We apply both approaches to demonstrate that solving the optimal stopping problem directly on the callable floater results in estimates that are biased low or option values that are biased high, respectively. To demonstrate the magnitude in bias we also compute an upper bound of the embedded option price by assuming the complete knowledge of the paths of the forward Libor rates. The complete knowledge of the paths allows choosing the best exercise

strategy among all strategies and hence yields the largest possible option value.

For valuing the snowball floater we focus on a specific implementation of the least squares Monte Carlo simulation algorithm introduced in Longstaff and Schwartz (2001). Longstaff and Schwartz (2001) varies somewhat from standard procedures usually applied in dynamic programming. First, they regress the ex post *realized* payoff from continuation on functions of the values of the state variables, instead of simply regressing the discounted option values of the subsequent period. They argue that their algorithm yields downward biased estimates of the option value given a finite choice of basis function. Although this property is appealing when we need guidance in determining the number of basis functions to obtain a desired level of accuracy, it might still result in considerable bias. This downward bias results from the fact that the stopping rule implied by the LSM algorithm does not necessarily be the optimal stopping rule that maximizes the option value. Note, that this is not necessarily the case for estimates obtained from the standard approach which can also be upward biased as a consequence of Jensen's inequality. In estimating the optimal exercise decision, i.e. the optimal stopping rule, one unfairly uses information about the future evolution of the state variable.

The second characteristic of the algorithm by Longstaff and Schwartz (2001) is that they suggest using solely in-the-money option paths for the least squares regression. They point out that when focusing only on in-the-money scenarios they restrict the region over which the conditional expectation must be estimated, thereby considerably reducing the needed number of basis functions. Moreover, they also find that the estimates have usually lower standard errors. We denote the algorithm of Longstaff and Schwartz (2001) the *LSM approach* and refer to the algorithm that uses the option values of the subsequent period for regression as the *standard approach*.

It is the main purpose of this paper to analyze the approach of Longstaff and Schwartz (2001) by comparing the estimate with results obtained from the standard approach. We also analyze the effect of using only in-the-money option scenarios for regression rather than all scenarios. Hence, we compute an option value  $\hat{C}^{lsm1}$  using the original LSM algorithm, which is based on in-the-money scenarios and another estimate  $\hat{C}^{lsm2}$  that considers all scenarios for regression. The option price obtained from a nested Monte Carlo simulation serves as a benchmark which allows us to compare the various least squares Monte Carlo algorithms. Given a downward and an upward biased option value resulting from the nested Monte Carlo simulation we compute a 95% confidence interval such that we can judge the statistical precision of the other estimates. Hence, we statistically test if the various estimators of the embedded option value are within the option bounds at a 95% confidence level. We compute the option values on a quarterly basis over the floater's lifetime and for a set of various scenarios introduced in Section 4.1. As the scenario base case we use the market data of the issue date on October 10, 2005.

## 5 Empirical Results

### 5.1 The Inherent Risk of a Callable Snowball Floater

A snowball floater is a risky instrument due to the additional risk of stochastic rather than fixed coupons and the potential prepayment risk. Hence, the price depends on the expected future

coupon payments and the value of the embedded Bermudan option. In the following, we discuss the functionality and the risk factors by means of a time-series and scenario analysis.

Table 3 shows the quarterly estimated values of the callable snowball floater  $\widehat{V}^{nmc}$ , the corresponding non-callable snowball floater  $\widehat{V}^{NC}$  and the embedded call-back option  $\widehat{C}^{nmc}$  over the floater's lifetime, starting on October 10, 2005 and ending on April 10, 2009, when the floater finally was called back by the issuer. The estimates  $\widehat{V}^{nmc}$  and  $\widehat{C}^{nmc}$  are based on an elaborate nested Monte Carlo simulation with 1000 paths. The bank issued the floater at 100.15% in October 10, 2005 which is slightly lower than the estimated price of 100.89%.<sup>9</sup> Its value then steadily decreased because market participants expected interest rates to rise during the market upturn. Around the third quarter of 2007, when the floater value was at its lowest (77.40%), it gained again in value until the second quarter of 2009, where the issuer finally exercised the option and called back the floater at par. The table also reveals that the non-callable snowball floater  $\widehat{V}^{NC}$  shows a very similar pattern, albeit the price is somewhat larger. The difference amounts to the value of the embedded call-back option. Initially, the value of the option was 9.17% but then almost got worthless in the middle of the lifetime because of the very low coupon payments the issuer had to pay. From the third quarter of 2007 on, its value increased until the second quarter of 2009, where the option value reaches its maximum at 19.53%.

The relation between option price, coupon size and the 6-month EURIBOR is shown in Figure 1. In the first year the investor receives rather large fixed coupons of 5.25% p.a. However, after another year the coupon immediately melts down to zero due to the rise in interest rates. For the same reason, the embedded option value markedly decreases, because the incentive for the issuer to buy back the floater is low in a high interest rate environment. In the middle of sub-prime crisis, when Lehman Brothers defaulted in September 2008, central banks dramatically lowered key interest rates to stimulate economic growth, which led the 6-month EURIBOR fall from 5.43% on October 10, 2008 to 1.62% on April 10, 2009 and, hence, made investments in snowball floaters more attractive.<sup>10</sup> We observe that in April 2009 investors received a positive coupon payment again and also the embedded option markedly gained in value. This indicates that the market anticipated a lower interest rate level. However, the issuer of the snowball floater exercised the option in April 2009 and called back the floater at par.

We also conduct a scenario analysis to get further insights into the functionality and the risk of the callable snowball floater. As the base case we use the market data at the issue date of October 10, 2005. In the first analysis we shift the forward rate curve parallel up and down by 25 and 50 bp, respectively. The results are reported in Table 4a. As expected, the value of the non-callable as well as the callable snowball floater decreases when the forward rate curve is shifted upwards, because, first, expected coupon payments decrease and, second, these coupons are discounted at higher interest rates. We also find that for higher interest rate levels the value of the callable floater is only slightly lower compared to the non-callable floater, indicating a low option value (2.46% for a +50 bp shift). When the forward rate curve is shifted downwards, the value of the non-callable

<sup>9</sup>The discount is due to the fact that we do not account for correlated forward rates. The assumption of independent forward rate dynamics increases the price of the callable snowball floater.

<sup>10</sup>The European Central Bank lowered key interest rates from 3.25% to 0.25% between October 2008 and April 2009.

snowball floater considerably increases, because investors expect to receive larger future coupon payments due to the floater's payoff structure. Very similar, albeit far less pronounced, grows the value of the callable snowball floater. Hence, the value of the embedded Bermudan option increases substantially (28.88% for a  $-50$  bp shift). This clearly highlights the downside risk of the callable snowball floater and its limited upside potential to the investor. When interest rates are large the coupon quickly diminishes which makes the floater worth less and leaves the investor with a loss if he decides to sell off the instrument. On the other hand, when interest rates turn out to be low, coupons quickly accumulate which is unfavorable for the issuer. Since the issuer must pay larger coupons to the investor, it is more attractive for the issuer to terminate the contract by exercising the call-back option. Hence, the option gains in value which means that the investor faces larger prepayment risk.

The impact of a twist in the forward rate curve on the snowball floater price is reported in Table 4b. As the pivot point we choose the 4-year forward rate, which corresponds to half the time-to-maturity of the floater. We twist the curve up and down by 25 and 50 bp, respectively, at the short end and observe that a down-twist of the curve increases the non-callable and the callable snowball floater value. Expected short term coupon payments increase, while long-term payments decrease, which eventually makes the floater more valuable. An up-twist of the forward rate curve has a similar effect in the opposite direction, although slightly less pronounced. A twisting of the forward rate curve obviously does not much change the value of the embedded option, which roughly remains at a price level of 9.3%, since the callable and the non-callable floater are affected similarly.

Finally, Table 4c shows the effect on the price when the volatility of the forward rates is scaled by 150% and 50%. A higher forward rate volatility increases the price of the non-callable snowball floater. The downside risk for the investor is limited because coupon payments are floored at zero. Hence, high volatile forward rates increase the value of the non-callable floater due to the asymmetric payoff structure. In contrast, the price of the callable floater changes only marginally as indicated by the benchmark price, because the option is affected similarly as the non-callable snowball floater by a change in the forward rate volatilities. Therefore, the callable snowball floater is more sensitive to a change in the level of the forward rates and is much less affected by changes in the forward rate volatilities.

## 5.2 Analysis of Various Basis Functions

An important issue when least squares Monte Carlo is used in dynamic programming is the specific choice and the required number of basis functions to obtain a desired level of accuracy. In this section we conduct an analysis using nine different sets of basis functions to value the callable snowball floater. The estimation is based on the LSM algorithm, because Longstaff and Schwartz (2001) provides a convergence criterion. This criterion gives guidance in choosing an appropriate type and the required number of basis functions. Since their algorithm results in low biased estimates of the option price the criterion works as follows: Simply increase the number of basis functions until the estimate no longer increases. The criterion can also be used to choose between different types of basis functions. We analyze this convergence criterion empirically on our specific

callable snowball floater.

The results of the basis function analysis are reported in Table 5, which shows the quarterly embedded option values between October 10, 2005 and April 10, 2009. The first column shows the most restrictive set in terms of the number of basis functions, i.e. a linear combination of two proxies: an estimate for the level and the slope of the interest rate curve. We find that this set yields almost always the lowest option value and which is also considerably lower than the benchmark price. Hence, it is not recommended to price interest rate instruments based solely on the level and the slope of the interest rate curve as is often suggested in the literature. We also see that a linear combination of simulated forward Libor rates or discount factors including a constant is still not optimal on average compared to the other sets of basis functions. A further increase in the number of basis functions by considering also second order terms in the functions considerably improves the accuracy of the estimates. The results indicate that most of the time a quadratic polynomial or a Fourier series yields the largest option values. A finding that was also made in Longstaff and Schwartz (2001). However, it is not clear if these functions should be evaluated preferably on the forward Libor rates or discount factors, even though there is weak evidence that forward Libor rates are slightly more appropriate. Interestingly, either the quadratic polynomial or the Fourier series always outperforms the weighted Laguerre polynomial of order two.

This analysis makes clear that increasing the number of basis functions can significantly improve the accuracy of the estimates. We find that the level and the slope of the interest rate curve is the worst set and yields mostly the lowest option values. Using quadratic polynomials or Fourier series of order two were the best choice in our analysis. Extending the polynomials to even higher order might further improve accuracy but will also lower computational efficiency. We also find that there is no specific type of basis function in our set that outperforms, i.e. results always in the largest option value. Obviously, the best choice depends on the market environment, i.e. on the forward rate curve and the forward rate volatilities. Hence, it is important to use the LSM algorithm as a performance or convergence criterion.

### 5.3 Analysis of Least Squares Monte Carlo Algorithms

In this sub-section we focus our analysis on two different least squares Monte Carlo valuation algorithms: the LSM algorithm where the ex post realized payoffs are regressed on basis functions and the standard algorithm, which uses the estimated value functions of the subsequent period for regression. To improve the efficiency of the least squares Monte Carlo algorithm Longstaff and Schwartz (2001) suggest using only on in-the-money option scenarios. We show the impact of this rather ad-hoc approach by comparing it with the results when all scenarios are considered for regression. We start the analysis by first demonstrating the importance of decomposing a callable instrument into a corresponding non-callable part and the embedded option and pricing each part individually. At each point in time, we use the optimal choice of basis functions found in Section 5.2 that maximizes the embedded option value. For comparing the various different estimation approaches we compute a benchmark price for the option based on a nested Monte Carlo simulation for which we also provide a 95% confidence interval.

The time-series of the quarterly estimated call-back option values between October 10, 2005

and April 10, 2009 are given in Table 6. We observe that the option value obtained from a direct estimation approach is, except for one case, well above the benchmark price. This is due to the fact, that the regression for estimating the conditional expectation for continuation is not order preserving and indicates the importance of separating the embedded option from the floater and estimating each part separately. Note, that the latter approach is always preferable regardless which estimation algorithm is used. We see that due to the large bias the estimate is mostly significantly larger than the option price upper bound given a 95% confidence level. However, We find that the price upper bound is only a strict upper bound for the least squares Monte Carlo estimates based on the decomposition approach.

Comparing these estimates with each other, we observe that the standard estimator lies closer to the benchmark price and, moreover, is in many cases not significantly outside the price bounds of the nested Monte Carlo estimator given a 95% confidence level, indicating that the standard estimator is more accurate. Interestingly, whenever the difference of  $\hat{C}^{std}$  to the benchmark price is statistically significant we find that the estimates are lower than the benchmark price. This indicates that not just the LSM but also the standard estimator is downward biased in our time-series samples. The relative difference of the standard estimator to the benchmark price in our time-series is on average 8% vs. 29% for the LSM estimate (based on  $\hat{C}^{lsm1}$ ). However, the difference might be substantially as in, e.g., the fourth quarter of 2008, where it amounts to more than 30% also for the standard estimate. It is unclear to us what causes such a large deviance. Comparing the two LSM estimates with each other we find that the original LSM estimate ( $\hat{C}^{lsm1}$ ), which considers only in-the-money option scenarios for regression is always larger. Hence, the option value is not just closer to the true option price but the estimation approach is computationally also more efficient, since fewer observations are used for regression. A  $t$ -test indicates that for lower option prices, as in the middle of the floater's lifetime, the LSM estimates are significantly different from each other. This makes perfectly sense because more out-of-the-money scenarios are observed on these specific dates and, hence, both approaches are based on rather different samples.

Table 7 reports the various estimates of the embedded Bermudan option for a set of different scenarios, where the base case is the issue date, i.e. October 10, 2005. Regarding the accuracy of our estimates we find that the standard errors are larger when the forward rate curve is shifted downwards (see Table 7a). Lower interest rates result in higher coupons and, hence, more volatile coupon payments. Table 7b provides the option values when the forward rate curve is twisted. This has only a minor effect on the standard errors of the estimates. Finally, Table 7c reports the option values when the forward rate volatilities are scaled by 150% or 50%. Clearly, more volatile forward rates come along with higher standard errors for all estimators. According to Longstaff and Schwartz (2001) estimates based on all simulated paths, and hence, not conditioning on events of in-the-money options should result in higher standard errors. However, we do not make this finding based on our empirical analysis. In the scenario analysis the standard errors of both estimates are virtually identical, whereas in the time-series we even observe slightly lower standard errors for the estimate that considers all paths. Comparing the standard errors of the LSM and the standard estimators with each other, we observe slightly lower standard errors for  $\hat{C}^{std}$ . Lower

standard errors would be in favor for the standard estimator. However, considering the standard errors for the estimate of the corresponding callable snowball floater, we find lower standard errors for the LSM estimates (see Table 8). The variance of the non-callable floater and the embedded option do not simply add up, but instead are positively correlated. The reason is that  $\widehat{V}^{NC}$  and the call option values are not independent, because they are based on the same simulated paths.

## 6 Conclusion

Callable snowball floaters are both path-dependent and early-exercise interest rate instruments. Such products are usually priced using least squares Monte Carlo simulation, where the conditional expected values of continuation are obtained by regression. In this paper we focus on a specific least squares Monte Carlo (LSM) algorithm introduced by Longstaff and Schwartz (2001). The difference to the standard approach discussed in Tsitsiklis and Van Roy (1999) is that in the LSM approach conditional expectations are estimated by regressing ex post realized payoffs on the current vector of state variables of in-the-money option scenarios, instead of regressing the option values of all scenarios from the subsequent period. Longstaff and Schwartz (2001) point out that their algorithm provides a useful convergence criterion because the obtained option values are less than or equal to the values implied by the optimal stopping problem. This would suggest increasing the number of basis functions and scenarios until the obtained estimate no longer increases.

The main objective of this paper was to compare the LSM approach with the standard approach by applying both approaches on a specific callable snowball floater. In doing so, we obtained estimates from an elaborate nested Monte Carlo simulation, which we used as a benchmark. We also calculated a 95% confidence interval for the nested Monte Carlo simulation estimate. Moreover, we computed a simple upper price bound for the embedded option value, conditional on the knowledge of the simulated paths.

Our main findings of the empirical analysis are that the LSM estimate of the embedded option value might be substantially downward biased whereas the standard estimate is much more accurate and closer to the benchmark price. We also find that there is no optimal type of basis function which can always be recommended when pricing complex interest rate instruments such as callable snowball floaters. Instead, we recommend choosing the most appropriate basis function based on the LSM algorithm, such that the embedded option value is maximized. Given the optimal basis function, the price of the instrument should then be obtained using the standard approach, which turns out to be less biased. Moreover, we demonstrated the necessity of decomposing callable instruments into a non-callable instrument and the embedded option and pricing each part individually. We showed that a direct estimation approach leads to option estimates that are considerable upward biased, mostly even larger than the obtained option price upper bound. We also find that for the least squares Monte Carlo estimation approach the standard errors of the callable instrument can be reduced if the embedded option is valued based on the same simulated paths as the non-callable instrument. These are important issues to consider for pricing interest rate instruments, particularly for callable snowball floaters.

## A Figures and Tables

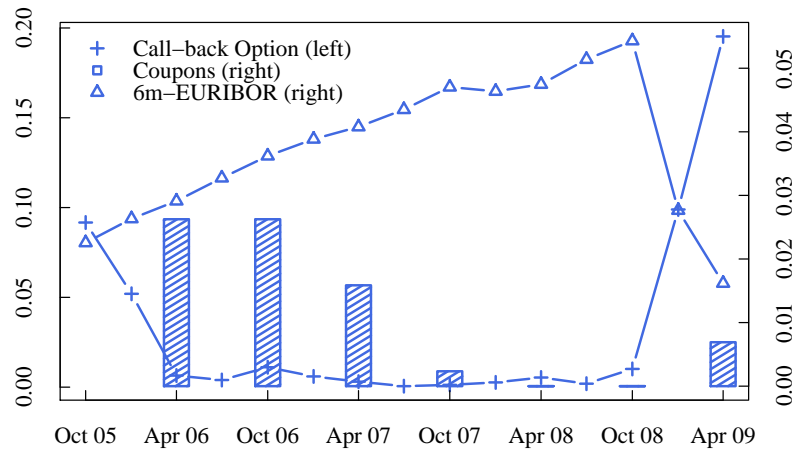


Figure 1: This figure shows the relation between the quarterly estimated embedded option value of a callable snowball floater, the semi-annual coupon payments and the quarterly 6-month EURIBOR between October 10, 2005 and April 10, 2008. The option value is given on the left, the coupon payments and the 6-month EURIBOR on the right axis.

$i$	3, 4	5, 6	7, 8	9, 10	11, 12	13, 14	15, 16
$r_i$ (% p.a.)	2.00	2.50	3.00	3.50	4.00	4.50	5.00

Table 1: This table reports the fixed interest rate premia  $r(T_i)$  at the coupon payment dates  $T_i$ ,  $i = 3, \dots, 16$  of the callable snowball floater under consideration. For  $i < 3$  the issuer receives fixed coupon payments of 5.25% p.a.

Type	Basis Functions
$\Psi_1^p(S_m, S_m - L_i)$	1, $S_m$ , $S_m - L_i$
$\Psi_1^p(L_{i,\dots,m})$	1, $L_{i,\dots,m}$
$\Psi_2^p(L_{i,\dots,m})$	1, $L_{i,\dots,m}$ , $L_{i,\dots,m}^2$
$\Psi_2^{wl}(L_{i,\dots,m})$	$\exp(-L_{i,\dots,m}/2)$ , $\exp(-L_{i,\dots,m}/2)(1 - L_{i,\dots,m})$ , $\exp(-L_{i,\dots,m}/2)(1 - 2L_{i,\dots,m} + L_{i,\dots,m}^2/2)$
$\Psi_2^f(L_{i,\dots,m})$	1, $\cos(L_{i,\dots,m})$ , $\sin(L_{i,\dots,m})$ , $\cos(2L_{i,\dots,m})$ , $\sin(2L_{i,\dots,m})$
$\Psi_1^p(D_{i,\dots,m})$	1, $D_{i,\dots,m}$
$\Psi_2^p(D_{i,\dots,m})$	1, $D_{i,\dots,m}$ , $D_{i,\dots,m}^2$
$\Psi_2^{wl}(D_{i,\dots,m})$	$\exp(-D_{i,\dots,m}/2)$ , $\exp(-D_{i,\dots,m}/2)(1 - D_{i,\dots,m})$ , $\exp(-D_{i,\dots,m}/2)(1 - 2D_{i,\dots,m} + D_{i,\dots,m}^2/2)$
$\Psi_2^f(D_{i,\dots,m})$	1, $\cos(D_{i,\dots,m})$ , $\sin(D_{i,\dots,m})$ , $\cos(2D_{i,\dots,m})$ , $\sin(2D_{i,\dots,m})$

Table 2: Various sets of basis functions that we use for least squares Monte Carlo simulation. The set of basis function  $\Psi_1^p$  denotes a linear combination,  $\Psi_2^p$  a quadratic polynomial,  $\Psi_2^{wl}$  a weighted Laguerre polynomial, and  $\Psi_2^f$  a Fourier series of order 2. The basis functions are evaluated on the level and the slope of the interest rate curve, where  $S_m$  denotes the swap rate with maturity  $T_m$ , the current forward Libor rates  $L_{i,\dots,m}$  or discount factors  $D_{i,\dots,m}$  available at time  $T_i$ .

	$V^{NC}$	$V^{nmc}$	$C^{nmc}$
Q4-2005	110.141	100.887	9.168
Q1-2006	102.685	97.309	5.198
Q2-2006	84.573	86.616	0.646
Q3-2006	82.052	84.305	0.398
Q4-2006	84.326	85.838	1.109
Q1-2007	81.763	83.736	0.597
Q2-2007	79.315	80.650	0.305
Q3-2007	75.827	77.399	0.055
Q4-2007	78.494	78.677	0.128
Q1-2008	81.258	81.214	0.264
Q2-2008	83.996	83.477	0.537
Q3-2008	79.779	79.586	0.194
Q4-2008	86.568	85.554	1.014
Q1-2009	107.287	97.319	9.894
Q2-2009	120.105	101.087	19.529

Table 3: This table reports the quarterly estimates of the non-callable snowball floater  $V^{NC}$ , the callable snowball floater  $V^{nmc}$  and the embedded option  $C^{nmc}$  between October 10, 2005 and April 10, 2009. The callable snowball floater and the embedded option value are estimated based on a nested Monte Carlo simulation with 1000 paths, the non-callable snowball floater is estimated based on standard Monte Carlo simulation with 10000 paths. All numbers are given in percentage points.

(a) Shifting the forward rate curve.			
	$V^{NC}$	$V^{nmc}$	$C^{nmc}$
+50 bp	94.746	92.009	2.459
+25 bp	101.249	96.617	4.496
$\pm 0$ bp	110.141	100.887	9.168
-25 bp	121.022	103.454	17.199
-50 bp	133.684	104.531	28.875

(b) Twisting the forward rate curve.			
	$V^{NC}$	$V^{nmc}$	$C^{nmc}$
+50 bp	108.994	99.291	9.308
+25 bp	109.380	100.145	9.296
$\pm 0$ bp	110.141	100.887	9.168
-25 bp	110.914	101.644	9.328
-50 bp	111.752	102.251	9.518

(c) Scaling of the forward rate volatilities.			
	$V^{NC}$	$V^{nmc}$	$C^{nmc}$
150%	114.166	100.862	12.939
100%	110.141	100.887	9.168
50%	107.766	101.540	6.097

Table 4: This table reports the estimates of the non-callable snowball floater  $V^{NC}$ , the callable snowball floater  $V^{nmc}$  and the embedded option  $C^{nmc}$  for various scenarios. The base case of the scenario analysis is given by the market data in October 10, 2005. Panel (a) reports the results when the forward rate curve is parallel shifted up and down by 25 and 50 bp, respectively. Panel (b) shows the results when the forward rate curve is twisted up and down by 25 and 50 bp at the short end, using the 4-year forward rate as the pivot point. In Panel (c) we report the results when the forward rate volatilities are either scaled up by 150% or down by 50%. The callable snowball floater and the embedded option value are estimated based on a nested Monte Carlo simulation with 1000 paths, the non-callable snowball floater is estimated based on standard Monte Carlo simulation with 10000 paths. All numbers are given in percentage points.

	$\Psi_1^p(S_m, S_m - L_{i_t})$	$\Psi_1^p(L_{i_t, \dots, m})$	$\Psi_2^p(L_{i_t, \dots, m})$	$\Psi_2^{wt}(L_{i_t, \dots, m})$	$\Psi_2^f(L_{i_t, \dots, m})$	$\Psi_1^p(D_{i_t, \dots, m})$	$\Psi_2^p(D_{i_t, \dots, m})$	$\Psi_2^{wt}(D_{i_t, \dots, m})$	$\Psi_2^f(D_{i_t, \dots, m})$
Q4-2005	7.917 (0.033)	8.329 (0.032)	8.098 (0.033)	8.212 (0.033)	8.480 (0.033)	8.298 (0.033)	8.285 (0.034)	8.341 (0.033)	8.283 (0.035)
Q1-2006	4.729 (0.027)	4.970 (0.027)	5.100 (0.028)	4.949 (0.027)	4.941 (0.028)	5.007 (0.028)	4.891 (0.028)	4.929 (0.026)	5.102 (0.030)
Q2-2006	0.342 (0.005)	0.362 (0.006)	0.425 (0.007)	0.360 (0.006)	0.453 (0.007)	0.362 (0.006)	0.386 (0.006)	0.350 (0.006)	0.438 (0.008)
Q3-2006	0.153 (0.003)	0.167 (0.003)	0.238 (0.005)	0.184 (0.004)	0.242 (0.005)	0.197 (0.004)	0.182 (0.004)	0.164 (0.003)	0.252 (0.005)
Q4-2006	0.775 (0.009)	0.835 (0.010)	0.912 (0.011)	0.814 (0.009)	0.862 (0.011)	0.851 (0.010)	0.791 (0.009)	0.828 (0.009)	0.907 (0.011)
Q1-2007	0.336 (0.005)	0.363 (0.006)	0.444 (0.007)	0.367 (0.006)	0.437 (0.007)	0.390 (0.006)	0.397 (0.006)	0.356 (0.005)	0.434 (0.007)
Q2-2007	0.151 (0.003)	0.156 (0.003)	0.193 (0.004)	0.152 (0.003)	0.190 (0.004)	0.147 (0.003)	0.156 (0.003)	0.151 (0.003)	0.160 (0.003)
Q3-2007	0.006 (0.000)	0.007 (0.000)	0.022 (0.001)	0.008 (0.001)	0.022 (0.001)	0.008 (0.001)	0.008 (0.001)	0.006 (0.000)	0.009 (0.001)
Q4-2007	0.033 (0.001)	0.034 (0.001)	0.067 (0.002)	0.039 (0.001)	0.067 (0.002)	0.037 (0.001)	0.038 (0.001)	0.036 (0.001)	0.037 (0.001)
Q1-2008	0.090 (0.002)	0.104 (0.003)	0.150 (0.004)	0.105 (0.003)	0.150 (0.004)	0.106 (0.003)	0.101 (0.003)	0.098 (0.003)	0.124 (0.003)
Q2-2008	0.268 (0.005)	0.285 (0.005)	0.344 (0.006)	0.294 (0.005)	0.318 (0.006)	0.286 (0.005)	0.312 (0.005)	0.292 (0.005)	0.300 (0.005)
Q3-2008	0.058 (0.002)	0.069 (0.002)	0.101 (0.003)	0.063 (0.002)	0.100 (0.003)	0.061 (0.002)	0.073 (0.002)	0.061 (0.002)	0.112 (0.003)
Q4-2008	0.379 (0.006)	0.391 (0.006)	0.514 (0.008)	0.405 (0.006)	0.479 (0.008)	0.423 (0.007)	0.455 (0.007)	0.388 (0.006)	0.540 (0.009)
Q1-2009	8.996 (0.028)	9.235 (0.028)	9.385 (0.030)	9.334 (0.029)	9.368 (0.030)	9.301 (0.030)	9.225 (0.030)	9.371 (0.029)	9.298 (0.030)
Q2-2009	19.496 (0.024)	19.553 (0.024)	19.569 (0.023)	19.503 (0.024)	19.638 (0.024)	19.494 (0.023)	19.388 (0.024)	19.387 (0.024)	19.508 (0.024)

Table 5: This table reports the quarterly estimated embedded option values for various types of basis functions based on the LSM algorithm introduced in Longstaff and Schwartz (2001) between October 10, 2005 and April 10, 2008. We use a set of nine different types of basis functions: A linear combination of the level and the slope of the interest rate curve  $\Psi_1^p(S_m, S_m - L_{i_t})$ , where  $S_m$  refers to the swap rate with maturity  $T_m$ . Other basis functions are linear transformations  $\Psi_1^p$ , quadratic polynomials  $\Psi_2^p$ , weighted Laguerre polynomials  $\Psi_2^{wt}$ , and Fourier series  $\Psi_2^f$  evaluated either on the simulated forward Libor rates  $L_{i_t, \dots, m}$  or discount factors  $D_{i_t, \dots, m}$  at time  $T_i$ . Standard errors are reported in parenthesis and all numbers are given in percentage points.

	$\widehat{C}^{nmc}$	$\widehat{C}^{dir}$	$\widehat{C}^{std}$	$\widehat{C}^{lsm1}$	$\widehat{C}^{lsm2}$	$\widehat{C}^{up}$
Q4-2005	[7.586; 9.168; 10.742]	11.992 (0.137)	8.741* (0.021)	8.480* (0.033)	8.438* (0.034)	11.186 (0.109)
Q1-2006	[3.924; 5.198; 6.446]	8.115 (0.107)	5.175* (0.010)	5.102* (0.030)	5.023* (0.030)	6.740 (0.071)
Q2-2006	[0.549; 0.646; 0.737]	1.171 (0.048)	0.656* (0.001)	0.453 (0.007)	0.409 (0.007)	0.756 (0.015)
Q3-2006	[0.346; 0.398; 0.448]	0.645 (0.041)	0.355* (0.000)	0.252 (0.005)	0.228 (0.005)	0.452* (0.011)
Q4-2006	[0.934; 1.109; 1.276]	1.646 (0.049)	1.134* (0.002)	0.912 (0.011)	0.847 (0.010)	1.309 (0.020)
Q1-2007	[0.520; 0.597; 0.672]	0.885 (0.040)	0.601* (0.001)	0.444 (0.007)	0.403 (0.006)	0.681* (0.012)
Q2-2007	[0.288; 0.305; 0.325]	0.445 (0.031)	0.286 (0.001)	0.193 (0.004)	0.175 (0.004)	0.315* (0.007)
Q3-2007	[0.053; 0.055; 0.058]	0.075* (0.024)	0.042 (0.000)	0.022 (0.001)	0.020 (0.001)	0.060* (0.003)
Q4-2007	[0.121; 0.128; 0.135]	0.192 (0.026)	0.113 (0.000)	0.067 (0.002)	0.060 (0.002)	0.137* (0.004)
Q1-2008	[0.249; 0.264; 0.281]	0.431 (0.034)	0.233 (0.000)	0.150 (0.004)	0.135 (0.003)	0.295 (0.008)
Q2-2008	[0.501; 0.537; 0.575]	0.843 (0.039)	0.488 (0.001)	0.344 (0.006)	0.315 (0.006)	0.594 (0.012)
Q3-2008	[0.184; 0.194; 0.205]	0.266 (0.033)	0.167 (0.000)	0.112 (0.003)	0.102 (0.003)	0.212 (0.006)
Q4-2008	[0.993; 1.014; 1.035]	1.512 (0.061)	0.682 (0.002)	0.540 (0.009)	0.496 (0.008)	1.155 (0.024)
Q1-2009	[9.785; 9.894; 10.004]	9.711* (0.095)	9.418 (0.013)	9.385 (0.030)	9.380 (0.030)	10.397 (0.078)
Q2-2009	[19.049; 19.529; 20.049]	19.706* (0.094)	19.638* (0.023)	19.638* (0.024)	19.638* (0.024)	19.892* (0.091)

Table 6: This table reports estimates for the embedded option value on a quarterly basis between October 10, 2005 and April 10, 2009. The estimate  $\widehat{C}^{nmc}$  is obtained by a nested Monte Carlo simulation and serves as a benchmark price. For this estimate we report a 95% confidence interval. The estimate  $\widehat{C}^{up}$  is an upper bound of the option price, which we obtain by conditioning on the complete knowledge of the paths. All other estimates are based on least squares Monte Carlo simulation. The estimate  $\widehat{C}^{dir}$  is computed by directly solving the dynamic programming problem on the callable snowball floater. The estimates  $\widehat{C}^{std}$ ,  $\widehat{C}^{lsm1}$ , and  $\widehat{C}^{lsm2}$  are all obtained by first decomposing the floater into a non-callable part and the embedded option and estimating each part individually. The difference between these estimates is that  $\widehat{C}^{std}$  uses option values of the subsequent period for regression, whereas the  $\widehat{C}^{lsm1}$  and  $\widehat{C}^{lsm2}$  are based on a specific algorithm introduced in Longstaff and Schwartz (2001). The former uses only in-the-money scenarios for regression, the latter all scenarios. Estimates which are not significantly outside of the lower and the upper bound of the benchmark at a 95% confidence level are marked with an asterisk. Standard errors are reported in parenthesis and all numbers are given in percentage points.

(a) Shifting the forward rate curve.

	$\widehat{C}^{nmc}$	$\widehat{C}^{dir}$	$\widehat{C}^{std}$	$\widehat{C}^{lsm1}$	$\widehat{C}^{lsm2}$	$\widehat{C}^{up}$
+50 bp	[1.855; 2.459; 3.038]	4.546 (0.093)	2.375* (0.005)	1.936* (0.018)	1.797 (0.018)	3.256 (0.050)
+25 bp	[3.418; 4.496; 5.550]	7.402 (0.111)	4.301* (0.009)	3.731* (0.024)	3.548* (0.024)	5.981 (0.074)
±0 bp	[7.586; 9.168; 10.742]	11.992 (0.137)	8.741* (0.021)	8.480* (0.033)	8.438* (0.034)	11.186 (0.109)
-25 bp	[15.445; 17.199; 18.963]	18.536* (0.154)	17.818* (0.030)	17.881* (0.032)	17.879* (0.032)	19.513 (0.138)
-50 bp	[27.167; 28.875; 30.595]	30.344* (0.159)	30.157* (0.030)	30.157* (0.030)	30.157* (0.030)	30.732 (0.152)

(b) Twisting the forward rate curve.

	$\widehat{C}^{nmc}$	$\widehat{C}^{dir}$	$\widehat{C}^{std}$	$\widehat{C}^{lsm1}$	$\widehat{C}^{lsm2}$	$\widehat{C}^{up}$
+50 bp	[7.496; 9.308; 11.112]	12.889 (0.130)	9.775* (0.017)	9.458* (0.034)	9.392* (0.034)	11.652 (0.100)
+25 bp	[7.595; 9.296; 10.991]	12.288 (0.129)	9.075* (0.018)	8.750* (0.034)	8.685* (0.034)	11.224 (0.101)
±0 bp	[7.586; 9.168; 10.742]	11.992 (0.137)	8.741* (0.021)	8.480* (0.033)	8.438* (0.034)	11.186 (0.109)
-25 bp	[7.865; 9.328; 10.785]	11.566 (0.138)	8.577* (0.024)	8.346* (0.034)	8.286* (0.034)	11.234 (0.112)
-50 bp	[8.134; 9.518; 10.895]	11.178 (0.142)	8.725* (0.027)	8.604* (0.033)	8.549* (0.034)	11.455 (0.118)

(c) Scaling of the forward rate volatilities.

	$\widehat{C}^{nmc}$	$\widehat{C}^{dir}$	$\widehat{C}^{std}$	$\widehat{C}^{lsm1}$	$\widehat{C}^{lsm2}$	$\widehat{C}^{up}$
150%	[10.717; 12.939; 15.158]	15.955 (0.178)	12.921* (0.031)	12.591* (0.048)	12.509* (0.050)	15.792 (0.149)
100%	[7.586; 9.168; 10.742]	11.992 (0.137)	8.741* (0.021)	8.480* (0.033)	8.438* (0.034)	11.186 (0.109)
50%	[5.226; 6.097; 6.967]	8.225 (0.082)	5.654* (0.012)	5.575* (0.018)	5.563* (0.019)	7.257 (0.062)

Table 7: This table reports estimates for the embedded option value on a quarterly basis for various scenarios. The estimate  $\widehat{C}^{nmc}$  is obtained by a nested Monte Carlo simulation and serves as a benchmark price. For this estimate we report a 95% confidence interval. The estimate  $\widehat{C}^{up}$  is an upper bound of the option price. All other estimates are based on least squares Monte Carlo simulation. The estimate  $\widehat{C}^{dir}$  is computed by directly solving the dynamic programming problem. The base case of the scenario analysis is given by the market data in October 2005. Panel (a) reports the results when the forward rate curve is parallel shifted up and down by 25 and 50 bp, respectively. Panel (b) shows the results when the forward rate curve is twisted up and down by 25 and 50 bp at the short end, using the 4-year forward rate as the pivot point. In Panel (c) we report the results when the forward rate volatilities are either scaled up by 150% or down by 50%. Estimates which are not significantly outside of the lower and the upper bound of the benchmark at a 95% confidence level are marked with an asterisk. Standard errors are reported in parenthesis and all numbers are given in percentage points.

	$\widehat{V}^{NC}$	$\widehat{V}^{nmc}$	$\widehat{V}^{dir}$	$\widehat{V}^{std}$	$\widehat{V}^{lsm1}$	$\widehat{V}^{lsm2}$	$\widehat{V}^{low}$
Q4-2005	110.141 (0.142)	100.887 (0.062)	98.149 (0.011)	101.400 (0.134)	101.661 (0.126)	101.703 (0.125)	98.955 (0.047)
Q1-2006	102.685 (0.110)	97.309 (0.067)	94.570 (0.009)	97.510 (0.109)	97.583 (0.098)	97.662 (0.098)	95.945 (0.053)
Q2-2006	84.573 (0.050)	86.616 (0.047)	83.402 (0.007)	83.917 (0.050)	84.121 (0.047)	84.164 (0.048)	83.818 (0.041)
Q3-2006	82.052 (0.041)	84.305 (0.034)	81.407 (0.005)	81.697 (0.041)	81.799 (0.040)	81.824 (0.040)	81.600 (0.035)
Q4-2006	84.326 (0.051)	85.838 (0.044)	82.680 (0.007)	83.192 (0.051)	83.414 (0.046)	83.479 (0.047)	83.018 (0.037)
Q1-2007	81.763 (0.040)	83.736 (0.030)	80.878 (0.005)	81.161 (0.040)	81.319 (0.038)	81.360 (0.038)	81.081 (0.032)
Q2-2007	79.315 (0.032)	80.650 (0.016)	78.870 (0.006)	79.029 (0.032)	79.122 (0.031)	79.141 (0.031)	79.000 (0.028)
Q3-2007	75.827 (0.024)	77.399 (0.012)	75.752 (0.004)	75.785 (0.024)	75.806 (0.024)	75.808 (0.024)	75.767 (0.023)
Q4-2007	78.494 (0.028)	78.677 (0.011)	78.302 (0.006)	78.382 (0.028)	78.427 (0.027)	78.434 (0.027)	78.358 (0.025)
Q1-2008	81.258 (0.035)	81.214 (0.011)	80.826 (0.005)	81.024 (0.035)	81.108 (0.034)	81.123 (0.034)	80.962 (0.031)
Q2-2008	83.996 (0.041)	83.477 (0.017)	83.153 (0.007)	83.508 (0.041)	83.652 (0.039)	83.681 (0.039)	83.402 (0.033)
Q3-2008	79.779 (0.034)	79.586 (0.013)	79.512 (0.006)	79.612 (0.034)	79.667 (0.033)	79.676 (0.033)	79.567 (0.031)
Q4-2008	86.568 (0.064)	85.554 (0.006)	85.056 (0.011)	85.886 (0.064)	86.028 (0.061)	86.072 (0.061)	85.413 (0.048)
Q1-2009	107.287 (0.097)	97.319 (0.009)	97.576 (0.006)	97.869 (0.094)	97.902 (0.082)	97.907 (0.082)	96.890 (0.028)
Q2-2009	120.105 (0.095)	101.087 (0.016)	100.399 (0.003)	100.467 (0.086)	100.466 (0.085)	100.466 (0.085)	100.212 (0.010)

Table 8: This table reports an estimate of the non-callable snowball floater value  $\widehat{V}^{NC}$  and various estimates for the callable snowball floater on a quarterly basis between October 10, 2005 and April 10, 2009. The estimate  $\widehat{V}^{nmc}$  is obtained by a nested Monte Carlo simulation and serves as a benchmark price. The estimate  $\widehat{V}^{low}$  is a lower bound of the floater, which we obtain by conditioning on the complete knowledge of the paths. All other estimates are based on least squares Monte Carlo simulation. The estimate  $\widehat{V}^{dir}$  is computed by directly solving the dynamic programming problem on the callable snowball floater. The estimates  $\widehat{V}^{std}$ ,  $\widehat{V}^{lsm1}$ , and  $\widehat{V}^{lsm2}$  are all obtained by first decomposing the floater into a non-callable part and the embedded option and estimating each part individually. The difference between these estimates is that  $\widehat{V}^{std}$  uses all option values of the subsequent period for regression, whereas the  $\widehat{V}^{lsm1}$  and  $\widehat{V}^{lsm2}$  are based on the specific algorithm introduced in Longstaff and Schwartz (2001). The former uses only in-the-money scenarios for regression, the latter all scenarios. Standard errors are reported in parenthesis and all numbers are given in percentage points.

## References

- Andersen, L. and M. Broadie (2004). Primal-dual simulation algorithm for pricing multidimensional american options. *Management Science* 50 (9), 1222–1234.
- Bally, V., G. Pagès, and J. Printems (2005). A quantization tree method for pricing and hedging multidimensional american options. *Mathematical Finance* 15 (1), 119–168.
- Belomestny, D., C. Bender, and J. Schoenmakers (2008). True upper bounds for bermudan products via non-nested monte carlo. *Mathematical Finance* 19 (1), 53–71.
- Bossaerts, P. (1989). Simulation estimators of optimal early exercise. Technical report, Carnegie-Mellon University.
- Boyle, P., M. Broadie, and P. Glasserman (1997). Monte carlo methods for security pricing. *Journal of Economic Dynamics and Control* 21 (8-9), 1267–1321.
- Brigo, D. and F. Mercurio (2006). *Interest Rate Models - Theory and Practice* (2nd ed.). Springer Finance.
- Broadie, M. and J. Detemple (1996). American option valuation: New bounds, approximations, and a comparison of existing methods. *The Review of Financial Studies* 9 (4), 1211–1250.
- Broadie, M. and P. Glasserman (1997). Pricing american-style securities using simulation. *Journal of Economic Dynamics and Control* 21, 1323–1352.
- Broadie, M., P. Glasserman, and G. Jain (1997). Enhanced monte carlo estimates for american option prices. *Journal of Derivatives* 5, 25–44.
- Carriere, J. F. (1996). Valuation of early-exercise price of options using simulations and nonparametric regression. *Insurance: Mathematics and Economics* 19, 19–30.
- Clément, E., D. Lamberton, and P. Protter (2002). An analysis of a least squares regression method for american option pricing. *Finance and Stochastics* 6, 449–471.
- Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. Springer-Verlag New York, Inc.
- Glasserman, P. and X. Zhao (2000). Arbitrage-free discretization of lognormal forward libor and swap rate models. *Finance and Stochastics* 4, 35–68.
- Haugh, M. B. and L. Kogan (2004). Pricing american options: A duality approach. *Operations Research* 52 (2), 258–270.
- Jamshidian, F. (2003). Minimax optimality of bermudan and american claims and their monte carlo upper bound approximation. *Stochastics* 79, 27–60.
- Kolodko, A. and J. Schoenmakers (2006). Iterative construction of the optimal bermudan stopping time. *Finance and Stochastics* 10, 27–49.

- Longstaff, F. A. and E. S. Schwartz (2001). Valuing american options by simulation: A simple least-squares approach. *Review of Financial Studies* 14, 113–147.
- Rogers, L. (2002). Monte carlo valuation of american options. *Mathematical Finance* 12 (3), 271–286.
- Tilley, J. A. (1993). Valuing american options in a path simulation model. *Transactions of the Society of Actuaries* 45, 83–104.
- Tsitsiklis, J. N. and B. Van Roy (1999). Optimal stopping of markov processes: Hilbert spaces theory, approximation algorithms, and an application to pricing high-dimensional financial derivatives. *IEEE Transactions of the Society of Actuaries on Automatic Control* 44, 1840–1851.
- Tsitsiklis, J. N. and B. Van Roy (2001). Regression methods for pricing complex america-style options. *IEEE Transactions on Neural Networks* 12, 694–703.