# Efficient Control Variates and Strategies for Bermudan Swaptions in a LIBOR Market Model

**Malene Shin Jensen and Mikkel Svenstrup**

**Malene Shin Jensen** is a postdoctoral fellow in the School of Economics and Management at the University of Aarhus in Denmark.
msjensen@econ.au.dk

**Mikkel Svenstrup** is a chief analyst at Scanrate Financial Systems in Denmark.
mikkel.svenstrup@scanrate.dk

*Bermudan options can be valued efficiently using a primal-dual simulation algorithm. We apply trigger strategies, replacing a maximum over several option values by a carefully chosen single option. We study reducing variance by applying the method of control variates, sampling the controls at the time of exercise of a Bermudan swaption.*

*Application of these efficiency improvements in combination with an optimal simulation path allocation rule significantly enhances numerical efficiency. The approach reduces the bias of upper- and lower-bound estimates, improves precision, and shortens estimation time.*

The history of numerical methods for pricing American-style contingent claims is long and rich. The earliest approaches include low-dimensional lattice methods such as trees and finite-difference methods, methods that work well only for American-style options on a single underlying asset or state variable. Boyle [1988] and Boyle, Evnine, and Gibbs [1989] propose multidimensional generalizations of the binomial tree method and extensions of finite-difference methods to higher dimensions, including the alternating direction implicit method. Even so, lattice methods suffer from the curse of dimensionality, implying that high-dimensional problems are often infeasible in practice.

Various Monte Carlo simulation approaches have been proposed. Early examples include Bossaerts [1988] and Tilley [1993]. Broadie and Glasserman [1997] make a key contribution in proposing a simulated tree algorithm for estimation of asymptotically unbiased upper and lower bounds on the option value. Generalization is difficult, however, and the method is most suitable for options with only a few exercise decisions.

For most applications, the size of the pricing problem quickly increases with the number of exercise opportunities; a more recent approach has been to map the true high-dimensional early exercise boundary on realizations of a much smaller number of representative proxy variables, generating lower-bound estimates on the true option price. Popular methods include least squares Monte Carlo (Longstaff and Schwartz [2001]), non-parametric methods (Andersen [2000]), and low-dimensional lattice methods (Longstaff, Santa-Clara, and Schwartz [2001]).

Haugh and Kogan [2004] and Andersen and Broadie [2004] outline a primal-dual simulation algorithm for estimation of upper and lower bounds on the value of an American-style option by solving the primal optimization and dual minimization problem. Any primal solution algorithm providing an exercise strategy, and hence a lower-bound estimate, is an input to the algorithm. This is a highly time-consuming algorithm, however, as it requires nested simulations.

We examine how to value Bermudan options efficiently using the primal-dual simulation algorithm. For implementation, we consider the valuation of Bermudan swaptions in a two-factor LIBOR market model with deterministic volatility. Although our example is contract-specific, many of the ideas presented are easily extended to more general settings.

Our primary contributions are as follows:

- We study reduction in variance from applying a range of control variates to the primal-dual algorithm. Following Rasmussen [2002], we sample the controls at the time of exercise of the Bermudan swaption. We expand this technique to handle dividend-paying control assets. The results show significant reduction in standard deviation.
- To reduce estimation time, we simplify the exercise strategies proposed by Andersen [2000]. For Bermudan swaptions, the structure of the contracts implies that the most valuable core European swaption at each exercise date will almost always be the next to mature, so we propose including only this particular swaption in the exercise decision. We find the proposed modified strategies highly preferable, as they reduce both the bias of the price estimates and estimation time.
- Jensen and Svenstrup [2004] address the variance-bias trade-off between reducing the bias or the variance of the duality gap estimator of the primal-dual problem, and propose a general allocation rule for choosing the optimal number of simulation paths. We apply this allocation rule to the estimation of duality gaps in combination with control variates and the enhanced strategies. The results show significant numerical improvements on both lower and upper bounds.

## I. THE PRIMAL-DUAL PROBLEM

Our presentation requires understanding of the theory underlying the computation of lower and upper bounds on the value of a Bermudan option. We consider a complete financial market, where all uncertainty is described by a standard filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{0\leq t\leq T}, \mathbb{P})$ satisfying the usual conditions, and assuming a $\mathbb{P}$-equivalent probability measure $\mathbb{Q}$ under which asset prices discounted with the numeraire asset $B(t)$ are martingales. $E_t^{\mathbb{Q}}(\cdot)$ denotes the expectation under $\mathbb{Q}$, conditional on the information available until time $t$.

A Bermudan option is an American-style option with only a discrete set of exercise dates $\mathcal{T} = \{t_s, t_{s+1}, \ldots, t_x\}$. We can write the value $V$ at time 0 of a Bermudan option with payoff $h$ paid upon exercise as an optimal stopping problem:

$$\frac{V(0)}{B(0)} = \sup_{\tau\in\Gamma(0)} E_0^{\mathbb{Q}}\left(\frac{h(\tau)}{B(\tau)}\right) \tag{1}$$

where $\Gamma(t)$ is the set of optional stopping times, taking values in $\mathcal{T} \cap [t, T]$. Bermudan options are optimally exercised when the intrinsic value $h$ exceeds the continuation value of the option $V^+$.

This implies that the optimal stopping time $\tau^*$ for the Bermudan option is given as:

$$\tau^* = \inf_{t_k\in\mathcal{T}}\left(h(t_k) \geq V^+(t_k)\right) \tag{2}$$

This criterion is not easily evaluated because $V^+$ is not readily available in a forward simulation algorithm. In this case, the challenge of the valuation problem is to find a stopping time close to the optimal one.

From Equation (1) it follows that any optional stopping time $\tau$ belonging to $\Gamma(0)$ will generate a lower bound on the true option value $V(0)$. Hence, for any stopping time $\tau$, we define the value of the *lower bound* at time $t$ as:

$$L(t) = B(t)E_t^{\mathbb{Q}}\left(\frac{h(\tau_t)}{B(\tau_t)}\right) \tag{3}$$

where $\tau_t$ denotes some stopping time from time $t$ onward (we also use the term, *exercise strategy*, as a synonym for stopping time).

The tightness of the lower bound depends entirely on the stopping time. The closer $\tau$ is to $\tau^*$, the tighter the lower bound. The lower bound can be computed by simulation or lattice methods. We estimate the lower bound using Monte Carlo simulation and refer to the simulation paths used as the *outer paths*.

Following Haugh and Kogan [2004] and Andersen and Broadie [2004], we refer to the right-hand side of (1) as the *primal problem* and define the *dual problem* in (1) as the right hand side of:

$$\frac{V(0)}{B(0)} \leq \inf_{\pi\in\mathcal{H}}\left(\pi(0) + E_0^{\mathbb{Q}}\left[\max_{t_k\in\mathcal{T}}\left(\frac{h(t_k)}{B(t_k)} - \pi(t_k)\right)\right]\right) \tag{4}$$

where $\pi$ denotes a martingale process, and the infimum is taken over the space $\mathcal{H}$ of all adapted martingales.[1]

By definition of the infimum operator, any adapted martingale will define an *upper bound* on the true option value whose tightness depends critically on the choice of $\pi$. Andersen and Broadie [2004] propose $\pi$ as the martingale part of the deflated lower-bound price process generated by a given exercise strategy and constructed as follows:

$$\pi(t_s) = L(t_s)/B(t_s)$$

and

$$\pi(t_k) = \pi(t_{k-1}) + \frac{L(t_k)}{B(t_k)} - \frac{L(t_{k-1})}{B(t_{k-1})} - I(t_{k-1})E^{Q}_{t_{k-1}}\left[\frac{L(t_k)}{B(t_k)} - \frac{L(t_{k-1})}{B(t_{k-1})}\right] \quad (5)$$

for $k = s + 1, s + 2, \ldots, x$. $I(t_k)$ denotes the indicator function for whether the exercise strategy signals exercise at time $t_k$. For this choice of $\pi$, we can write the upper bound as the lower bound plus a term resembling the value of a lookback option.[2]

We refer to this term as the *duality gap* $D(\cdot)$. Replacing terms in inequality (4), we have that

$$\frac{V(0)}{B(0)} \leq \frac{L(0)}{B(0)} + E^{Q}_{0}\left(\max_{t_k \in \mathcal{T}}\left(\frac{h(t_k)}{B(t_k)} - \pi(t_k)\right)\right) = \frac{L(0)}{B(0)} + D(0) \quad (6)$$

The essence of Equation (6) is that an upper bound for the true option value is the lower bound plus the duality gap. We can therefore use the duality gap as a worst case error of the option price estimate. The duality gap also depends on the exercise strategy used to estimate the lower bound, as the martingale process depends on the lower bound.

In effect, the tightness of the bounds on the true option value depends on the optimality of the exercise strategy. The duality gap is equal to zero under the optimal exercise strategy. This observation is not very constructive, though, since the whole valuation problem consists of finding a strategy close to the optimal one.

A primal-dual simulation algorithm for computing upper-bound estimates on the true option price can be derived from (6). It follows that to estimate the upper bound, we need to estimate both the lower bound and the duality gap. Hence, for each outer simulation path, we need to sample a value for the lower bound and for the duality gap. Sampling a value for the latter requires knowledge of the martingale process $\pi(t_k)$ for all $t_k \in \mathcal{T}$ for each outer path.

From Equation (5), it follows that estimation of $\pi$ requires knowledge of the lower bound along each path, but Equation (3) gives the lower bound as a conditional expectation. Hence, to estimate $\pi(t_k)$ we need to estimate either

$$E^{Q}_{t_k}\left(\frac{h(\tau_{t_k})}{B(\tau_{t_k})}\right)$$

or

$$E^{Q}_{t_{k-1}}\left(\frac{h(\tau_{t_k})}{B(\tau_{t_k})}\right)$$

depending on the exercise decision at time $t_k$. To do so, we simulate a set of *nested paths* or *inner paths* starting from the current state of the world (i.e., the outer paths) to estimate the conditional expectations.[3]

Hence, for each outer path we need to simulate a set of inner paths for each possible exercise date $t_k$ until exercise. Naturally, this makes estimation of upper bounds much more time-consuming than estimation of lower bounds, especially for long-maturity options.

## II. BERMUDAN SWAPTIONS IN LIBOR MARKET MODELS

The primal-dual simulation algorithm is applicable to many high-dimensional pricing problems. Applications in the literature, such as Svenstrup [2002], Andersen and Broadie [2004], and Jensen and Svenstrup [2004], have primarily dealt with pricing callable LIBOR exotics, particularly Bermudan swaptions. Because of its usefulness for hedging callable bonds, the Bermudan option is one of the most liquid American-style fixed-income instruments. It is also challenging to price and hedge and thus ideal for illustrative purposes.

We follow the literature and apply the algorithm to price Bermudan swaptions. The focus is on numerical efficiency, as improvements of the standard Monte Carlo simulation are essential in practical applications, with the complexity of the algorithm and the vast amount of estimation time spent on nested simulation paths. We think we are the only ones to target numerical efficiency of the algorithm in Jensen and Svenstrup [2004].

We use the LIBOR market model introduced by Brace, Gatarek, and Musiela [1997], Jamshidian [1997], and Miltersen, Sandmann, and Sondermann [1997]. This is one of the most commonly used models for pricing callable LIBOR exotics.

For completeness, we describe the model and explain Bermudan swaptions. We present the exercise strategies used for valuation of Bermudan swaptions and the control variates used in the Monte Carlo simulation. Our motivation is to introduce ways to improve the numerical efficiency of the pricing algorithm.

### LIBOR Market Models

The fundamental state variables of the forward LIBOR market model are the discretely compounded forward LIBOR $F_k(\cdot)$ defined on an increasing tenor structure $0 = t_0 < t_1 \ldots < t_K$. The dynamics of the standard lognormal LIBOR market model under the spot LIBOR measure $\mathbb{Q}$, using a savings account as the numeraire, are given by:

$$dF_k(t) = F_k(t)\lambda_k(t)\left(\sum_{j=\eta(t)}^{k} \frac{\delta_j F_j(t)}{1+\delta_j F_j(t)}\lambda_j(t)' dt + dW_t^{\mathbb{Q}}\right)$$

for $k = 1, 2, \ldots, K$. $\eta(t) = \{k\colon t_{k-1} < t \leq t_k\}$ is the left-continuous mapping function returning the index of the next tenor date; $\delta_k$ is the distance between tenor dates; $\lambda_k$: $\mathbb{R}^+ \rightarrow \mathbb{R}^{1\times m}$ is a bounded deterministic volatility function; and $W^{\mathbb{Q}}$ is an $m$-dimensional Brownian motion under $\mathbb{Q}$. $\lambda_k(t)$ is basically the time $t$ loading of the $k$-th forward rate on the underlying stochastic process.

The dimension of the model is given by the number of tenors $K$, which may easily be high. For example, for an option with quarterly payments and a 30-year time horizon, a total of 120 state variables are needed.

### Bermudan Swaptions

The standard Bermudan (payer) swaption is an option to enter into a (payer) swap contract with coupon $\kappa$ and final swap maturity $t_e$ on a discrete set of exercise dates $\mathcal{T} = \{t_s, t_{s+1}, \ldots, t_x\}$. $t_s$ is the first exercise date, often denoted the *lockout date*, and $t_x$ is the last exercise date. We consider the fixed-maturity case, implying that all underlying swaps have maturity $t_e$, and assume that $t_x = t_{e-1}$. Hence, the holder of the option can at time $t_k \in \mathcal{T}$ decide to enter into an $e - k$ period swap contract with start date $t_k$ and fixed rate $\kappa$.

If the option is exercised at time $t_k$, the holder of the option will receive the payoff from the swap:

$$BSwap_{s,x,e}(t_k, \kappa) = \left(Swap(t_k, t_k, t_e, \kappa)\right)^+$$

The payoff is exactly that of a European swaption $ESwap_{k,e}(.,\kappa)$ with exercise date $t_k$, $e - k$ periods, and strike $\kappa$. We call this set of European swaptions the *core* European swaptions, and we denote a Bermudan swaption with lock-out date $t_s$ and final swap maturity $t_e$ as a $t_e$ no-call $t_s$.

Looking back at Equation (1), the value of the Bermudan swaption at time $t$ is given as the solution to an optimal stopping problem:

$$BSwap_{s,x,e}(t, \kappa) = \sup_{\tau\in\Gamma(t)} E_t^{\mathbb{Q}}\left[\frac{B(t)}{B(\tau)}\left(Swap(\tau, \tau, t_e, \kappa)\right)^+\right] \quad (7)$$

### Exercise Strategies

An exercise strategy is basically the only element that has to be specified to run the primal-dual algorithm. Ideally, we would want to find the optimal exercise strategy given by (2). Yet the optimal exercise strategy $\tau^*$ that solves the valuation problem given by (7) is not readily available. We therefore have to find (or rather specify) an exercise strategy that is close to optimal.

The possibilities include exercise strategies from lower-dimensional models, which can be solved by using lattice methods, least squares Monte Carlo by Longstaff and Schwartz [2001], neural networks, or trigger strategies as explored in Andersen [2000]. We choose the latter because of its simplicity and speed.

In general, the optimal exercise decision of the Bermudan swaption will depend not only on the current yield curve but also on the entire dynamics of the yield curve—but to include that much information in the exercise decision is computationally infeasible in a Monte Carlo setting. The idea behind trigger strategies is to reduce the dimensionality of the exercise decision. For Bermudan swaptions, we assume the exercise decision at a given time depends only on the values of all still-alive core European swaptions.

We set out five exercise strategies that we use throughout for Bermudan swaptions. The proposed ideas do not apply solely to the pricing of Bermudan swaptions in a LIBOR market model setting, but extend to many other settings as well.

Let:

$$M(t_k) = \max_{j=k+1,\ldots,e} ESwap_{j,e}(t_k, \kappa)$$

denote the most valuable of the still-alive core European swaptions at time $t_k$. For strategy $l$, $I^l(\cdot)$ denotes the early exercise indicator function, which equals 1 if exercise is optimal, and 0 if not. Let $H_l(\cdot)$ denote a time-dependent *critical barrier* (*bar*) for strategy $l$. The strategies are as follows:

Strategy Bar:

$$I^B(t_k) = \begin{cases} 1 & \text{if } ESwap_{k,e}(t_k, \kappa) > H_B(t_k) \\ 0 & \text{otherwise} \end{cases}$$

Strategy Bar & Max:

$$I^{B\&M}(t_k) = \begin{cases} 1 & \text{if } ESwap_{k,e}(t_k, \kappa) > H_{B\&M}(t_k) \\ & \text{and } ESwap_{k,e}(t_k, \kappa) > M(t_k) \\ 0 & \text{otherwise} \end{cases}$$

Strategy Bar + Max:

$$I^{B+M}(t_k) = \begin{cases} 1 & \text{if } ESwap_{k,e}(t_k, \kappa) > H_{B+M}(t_k) + M(t_k) \\ 0 & \text{otherwise} \end{cases}$$

Strategy Bar & First:

$$I^{B\&F}(t_k) = \begin{cases} 1 & \text{if } ESwap_{k,e}(t_k,\kappa) > H_{B\&F}(t_k) \\ & \text{and } ESwap_{k,e}(t_k,\kappa) > ESwap_{k+1,e}(t_k,\kappa) \\ 0 & \text{otherwise} \end{cases}$$

Strategy Bar + First:

$$I^{B+F}(t_k) = \begin{cases} 1 & \text{if } ESwap_{k,e}(t_k,\kappa) > H_{B+F}(t_k) + ESwap_{k+1,e}(t_k,\kappa) \\ 0 & \text{otherwise} \end{cases}$$

One can think of the critical barrier as a predetermined function specifying the time of exercise. Estimation of the critical barrier requires a set of presimulated Monte Carlo paths and a series of one-dimensional optimizations, one for each exercise date. We first choose a strategy $l$ and simulate a set of Monte Carlo paths. Then, for each path and all exercise dates $t_k$, we estimate and store the discounted intrinsic values of the underlying swaps and the information needed to evaluate $I_l(t_k)$ for a given value of $H_l(t_k)$. Finally, for each $t_k$, we search for the value of the critical barrier $H_l(t_k)$ that maximizes the value of the Bermudan swaption using the stored information.

The Bermudan swaption is valued in a two-step procedure. First, the critical barrier is estimated using one set of paths, and, second, lower and upper bounds are estimated using a new set of paths. It is necessary to estimate the critical barrier only once for each strategy, type of contract, and set of model parameters.

Strategies Bar, Bar & Max, and Bar + Max are proposed by Andersen [2000]. In strategy Bar, the Bermudan swaption is exercised when the option payoff exceeds the critical barrier. The exercise decision in strategies Bar & Max and Bar + Max depends additionally on the maximum value of the remaining European swaptions. The rationale is that the holder of the Bermudan swaption can always sell it for the value of the most valuable core European swaption.

The latter two strategies are much more time-consuming than strategy Bar, especially for contracts with many exercise dates, due to the additional European swaption valuations. For example, finding the maximum value of the remaining European swaptions for a Bermudan swaption with 40 exercise dates requires as many as $40(40+1)/2 = 820$ European swaption valuations for each outer simulation path.

To enhance numerical efficiency, we propose strategies Bar & First and Bar + First as a simplification of Bar & Max and Bar + Max. We propose replacing the maximum of the remaining European swaptions with the first-to-mature European swaption. The rationale is that the most valuable core European swaption is almost always the first to mature, as its underlying swap is the longest. Alternatively, we could use a small subset of the still-alive core European swaptions, but we have found this does not improve the pricing bounds significantly and only increases calculation time.

### Control Variates

To improve the standard Monte Carlo estimator, we apply the method of control variates, a common variance reduction technique. Its efficiency depends crucially on the degree of correlation between the sampled option payoff and control value. Typically, a single variable is selected as a control variate, but we will use several. An appendix provides a brief introduction to multiple control variates.

When control variates are applied to the pricing of Bermudan and American options, it is not immediately clear how or when to sample the controls. Rasmussen [2002] suggests sampling the controls at the time of exercise rather than at expiration of the option, and shows that this makes the control and the option significantly more correlated. We follow this approach, and sample the controls at the time of exercise. We also extend the sampling technique to handle dividend-paying controls.

We first show how to sample dividend-paying controls and then describe control variates suitable for application to Bermudan swaptions. Although the controls are intended to apply to Bermudan swaptions, many of the ideas, such as sampling dividend-paying control assets at the time of exercise and using hedge portfolios as controls, can be extended to other applications as well.

***Sampling Dividend-Paying Controls.*** We extend the Rasmussen [2002] sampling method to include dividend-paying control assets. Let $X$ denote the control asset with dividend process $\{\gamma(t_k)\}_{1\le k\le K}$, and let $\tau \in \Gamma(t)$ be a stopping time specifying when to exercise. The value of the remaining discounted dividends is a martingale under $\mathbb{Q}$, and we have the relation:

$$\frac{X(t)}{B(t)} = E_t^{\mathbb{Q}}\left[\sum_{k=\eta(t)}^{K}\frac{\gamma(t_k)}{B(t_k)}\right]$$

Applying the *law of iterated expectations* to the right-hand side, and using the martingale property of the dividend process, we have that:

$$\begin{aligned}\frac{X(t)}{B(t)} &= E_t^{\mathbb{Q}}\left[E_\tau^{\mathbb{Q}}\left(\sum_{k=\eta(t)}^{K}\frac{\gamma(t_k)}{B(t_k)}\right)\right]\\ &= E_t^{\mathbb{Q}}\left[\sum_{k=1}^{\eta(\tau)}\frac{\gamma(t_k)}{B(t_k)} + E_\tau^{\mathbb{Q}}\left(\sum_{k=\eta(\tau)+1}^{K}\frac{\gamma(t_k)}{B(t_k)}\right)\right]\\ &= E_t^{\mathbb{Q}}\left[\sum_{k=1}^{\eta(\tau)}\frac{\gamma(t_k)}{B(t_k)} + \frac{X(\tau)}{B(\tau)}\right] \qquad (8)\end{aligned}$$

The expectation in Equation (8) consists of the sum of two terms: the discounted value of the control asset at time $\tau$, and the sum of discounted future dividends. The intuition behind this is that, as the initial value of the control asset includes all future dividends received from the asset, we have to include all dividends reinvested in the numeraire asset in the sampled control variate $Y$. Therefore, we sample the value for the controls:

$$Y(\tau) = \sum_{k=1}^{\eta(\tau)}\frac{\gamma(t_k)}{B(t_k)} + \frac{X(\tau)}{B(\tau)} \qquad (9)$$

We apply this sampling rule to the control variates below.

***Caps and Caplets.*** Caps and caplets have closed-form solutions in a standard LIBOR market model and may therefore be suitable as controls. A caplet $Caplet(\cdot, t_k, \kappa)$ is basically a call option on a forward LIBOR $F_k$ with strike $\kappa$, expiration date $t_k$, and payoff at time $t_{k+1}$. A cap is a portfolio of caplets with different expiration dates, so the value is

$$Cap(t, t_e, \kappa) = \textstyle\sum_{k=s}^{e-1} Caplet(t, t_k, \kappa)$$

Let $Y_k^{Caplet}$ denote the sampled value of the caplets. Following (9), we have:

$$Y_k^{Caplet}(\tau) = \frac{Caplet(\tau\wedge t_k, t_k, \kappa)}{B(\tau\wedge t_k)},\ k \in \mathcal{T}$$

where $\wedge$ denotes the minimum operator.[4]

Similarly, with a cap control $Y^{Cap}$, we sample:

$$Y^{Cap}(\tau) = \sum_{k=s+1}^{\eta(\tau)}\frac{(F_{k-1}(t_{k-1}) - \kappa)^+}{B(t_k)} + \frac{Cap(\tau, t_e, \kappa)}{B(\tau)}$$

By using a series of caplets or a single cap as controls, we essentially sample the same set of caplets, although we can choose which caplets to use and how many. Theoretically, including all $x - s$ caplets produces the greatest variance reduction, but in practice, problems with multicollinearity and numerical efficiency arise because using all caplets implies estimation of a variance-minimizing parameter vector $\beta^* \in \mathbb{R}^{x-s}$ (see the appendix for definition). When we use a cap, we reduce the dimension of $\beta^*$ to one.

***Swaps and Zero-Coupon Bonds.*** Standard fixed-income securities, such as zero-coupon bonds maturing before the Bermudan swaption expires, can be used as controls and are easy to price.[5] We test a series of zero-coupon bonds $P(\cdot, t_k)$ with maturities $t_k$ equal to the set of exercise dates. Let $Y_k^{Z\text{-}Bond}$ denote the sampled value of the zero-coupon bonds. Following (9), we have:

$$Y_k^{Z\text{-}Bond}(\tau) = \frac{P(\tau \wedge t_k, t_k)}{B(\tau \wedge t_k)} \text{for } k \in \mathcal{T}$$

We also test a single swap as a control variate. The swap is basically a portfolio of zero-coupon bonds and is therefore easy to price. We apply the longest swap underlying the Bermudan swaption and sample the swap control $Y^{Swap}$ as follows:

$$Y^{Swap}(\tau) = \sum_{k=s}^{\eta(\tau)} P(t_k, t_{k+1}) \frac{\delta_k\left(F_k(t_k) - \kappa\right)}{B(t_k)} + \frac{Swap(\tau, \tau + \delta_k, t_e, \kappa)}{B(\tau)}$$

As with the cap and the caplets, the essential difference between the two control types is that $\beta^*$ is multidimensional when using the zero-coupon bonds as a control and only one-dimensional in the case of the swap.

***Approximate Delta-Hedge.*** Discounted payoffs from the core European swaptions seem to be well suited as controls, as these are $\mathbb{Q}$-martingales, but only approximate closed-form solutions are available in the standard LIBOR market model. Rather than use these approximations, we propose creating a self-financing portfolio that replicates as closely as possible the core European swaptions by using the hedge ratios from the approximate solutions.

As the value of any self-financing portfolio discounted with the pricing numeraire is a martingale, we can use the payoff from a portfolio that approximately replicates the option that we are trying to value as a control. Hence, a hedge portfolio can be used as a general type of control variate.

In the Bermudan swaption application, we form a portfolio delta-hedging the first of the remaining core European swaptions to expire, using the delta implied by the approximate swaption formula given in Andersen and Andreasen [2000]. We set:

$$\Delta_{s,e}(t) = \frac{\partial ESwap_{s,e}(t,\kappa)}{\partial S_{s,e}(t)} = \Phi(d_+) \sum_{k=s+1}^{e} \delta_{k-1} P(t, t_k)$$

where $S_{s,e}(\cdot)$ denotes the *forward swap rate* defined by:

$$S_{s,e}(t) = \frac{P(t,t_s) - P(t,t_e)}{\sum_{k=s+1}^{e} \delta_{k-1} P(t,t_k)}$$

$\Phi$ denotes the standard normal cumulative density function, and

$$d_+(t,t_s) = \frac{\log \frac{S_{s,e}(t)}{\kappa} + \frac{1}{2}\xi^2(t,t_s)}{\xi(t,t_s)}$$

where $\xi(t,t_s)$ is the approximated integrated volatility and $\kappa$ is the strike. We resort to Andersen and Andreasen [2000] for details on how to construct $\xi$.

When setting up the hedge portfolio, we buy at time $t_k$:

$$\Delta_{s,e}(t_k) S_{s,e}(t_k) \text{ for } t_k < t_s$$
$$\Delta_{k+1,e}(t_k) S_{k+1,e}(t_k) \text{ for } t_s \le t_k \le t_e$$

and finance it with a short position in a bank account. At the next tenor date, we liquidate the swap position and enter into a new one, putting the profits or losses in the bank account. This is repeated until exercise or maturity.

The drawback of this approach is that the hedge portfolio must be updated along each path, which is computationally expensive. Furthermore, as the hedging errors are reduced with the number of resettings, and we reset only at the tenor dates, a high correlation between the payoff and the Bermudan swaption cannot be expected. A more precise hedging strategy might be to take into account the dynamics of the swap rate and try to hedge each driving factor of the yield curve. Under such a hedging strategy, improvements in precision should be measured against longer estimation times.

## III. NUMERICAL RESULTS

LIBOR market models cannot be simulated exactly, but discretization biases are manageable and negligible compared to the variance as shown in Andersen and Andreasen [2000]. We simulate the LIBOR market model using a standard log Euler scheme under the spot LIBOR measure $\mathbb{Q}$ on an equally spaced tenor space.

We first introduce model parameters and specifications of the Bermudan swaptions. Then we benchmark the model. Second, we test for the effect of application of control variates and the choice of exercise strategies on the lower bounds. Finally, we apply the efficiency improvements to the upper-bound estimations.

### Benchmarking

To compare our results with those of Andersen and Broadie [2004], we use one of their test scenarios: a stan-

## EXHIBIT 1

**Benchmark Scenario**

| $t_s$ | $t_e$ | *Strike* | *Lower (Std)* | *Time Low* | $\hat{D}_0$ *(Std)* | *Time* $\hat{D}_0$ |
|---|---|---|---|---|---|---|
| 1 | 3 | 0.08 | 339.14 (0.24) | 00:00:02 | 0.32 (0.05) | 00:00:34 |
| 1 | 3 | 0.10 | 124.71 (0.34) | 00:00:02 | 0.47 (0.06) | 00:00:56 |
| 1 | 3 | 0.12 | 36.26 (0.24) | 00:00:03 | 0.27 (0.05) | 00:01:11 |
| 1 | 6 | 0.08 | 750.82 (0.56) | 00:00:04 | 3.26 (0.27) | 00:04:34 |
| 1 | 6 | 0.10 | 317.81 (0.68) | 00:00:09 | 5.19 (0.36) | 00:08:18 |
| 1 | 6 | 0.12 | 127.79 (0.60) | 00:00:11 | 2.90 (0.28) | 00:10:44 |
| 1 | 11 | 0.08 | 1248.12 (1.13) | 00:00:13 | 18.89 (1.23) | 00:30:06 |
| 1 | 11 | 0.10 | 621.30 (1.14) | 00:00:26 | 20.16 (1.06) | 00:51:01 |
| 1 | 11 | 0.12 | 328.07 (1.18) | 00:00:34 | 14.17 (1.02) | 01:04:21 |

*$t_s$ denotes the lock-out period of the Bermudan swaption, and $t_e$ is the final swap maturity. Strike denotes the coupon of the underlying swap. The initial yield curve is flat at 0.1, so the various strikes represent ITM, ATM, and OTM options. Lower (Std) is the lower bound point estimate and its standard error, estimated by 50,000 antithetic paths and strategy Bar. $\hat{D}_0$ (Std) is the point estimate of the duality gap and its standard deviation based on 750 antithetic outer paths and 300 antithetic inner paths and strategy Bar. Time Low and Time $\hat{D}_0$ denote the computation time for the lower-bound estimates and duality gap estimates in hours, minutes, and seconds. Critical barriers are estimated using 50,000 antithetic paths.*

dard two-factor lognormal LIBOR market model with a quarterly tenor structure. The initial yield curve is flat at 10%, and the volatility structure is given by:

$$\lambda_k(t) = (0.15, 0.15 - \sqrt{0.009 \cdot (t_k - t)})', \; k = 1, 2, \ldots, K$$

To benchmark the model implementation, we include for completeness Exhibit 1, which replicates parts of Table 4 in Andersen and Broadie [2004]. The exhibit shows lower-bound estimates based on 50,000 antithetic simulation paths and duality gap estimates based on 750 antithetic outer paths and 300 antithetic inner paths for a set of Bermudan swaptions of varying maturities and strikes. Both lower- and upper-bound estimates are indistinguishable from those in Andersen and Broadie [2004].

Estimation time is also given. It is clear that estimation of the duality gap is much more time-consuming than estimation of the lower bound, particularly for long-maturity options.[6] To ascertain the suboptimality of the applied exercise strategies, however, estimation of the duality gap is necessary.

### Testing Control Variates

Comparing the efficiency of different variance reduction techniques is always difficult as they depend on the implementation, but it is still relevant to compare particular implementations. To evaluate the efficiency of different variance reduction techniques, we must compare not only the variance of the estimators but also the actual estimation time. For example, we typically have to choose between two estimators, one with a lower variance, but more time-consuming.

Following Glynn and Whitt [1992], we compare two unbiased estimators with variances $\sigma_1^2$ and $\sigma_2^2$ requiring expected work per replication of $EWork_1$ and $EWork_2$, respectively, by comparing the product of the variance scaled with the workload $\sigma_i^2\, EWork_i$. We denote this product the *efficiency factor* $EF_i$. The estimator with the lowest efficiency factor is preferred.

We want to explore the relative performance of alternative choices of control variates. Exhibit 2 lists efficiency test results for all the control variates. To test the numerical efficiency of multiple caplet controls, we include two caplet setups: seven caplets, and all possible caplets.[7] The control variates are applied to estimation of lower bounds for 6 no-call 1 and 11 no-call 1 Bermudan swaptions of various strikes using 25,000 antithetic paths and strategy Bar.

To compare performance improvements, we provide efficiency factors and standard deviations of each control variate simulation relative to those of the standard antithetic variate simulation; that is, $EF_{AV}/EF_{CV}$ and $Std_{AV}/Std_{CV}$. High values of the efficiency factors are preferred, and relative values strictly higher than 1 indicate improved performance.

There are several interesting observations to be made in Exhibit 2. First, antithetic variates work best for in-the-money (ITM) options, since the payoff function is close to linear. Second, application of control variates improves the estimation results significantly. For instance, application of the seven caplets as controls reduces the standard deviation by a factor of 3.25 with only a marginal increase in estimation time for the 11 no-call 1 ITM Bermudan swaption. Obtaining the same reduction in standard deviation by using only antithetic variates requires increasing the number of simulation paths by a factor of $3.25^2$, i.e., a total of more than 250,000 antithetic paths.

The reduction in standard deviation is slightly better for the 6 no-call 1 than for the 11 no-call 1 in most cases, but the efficiency factors clearly show that application of

## EXHIBIT 2

**Efficiency of Control Variates on Lower Bounds**

| | | 6 *no-call* 1 | | | | 11 *no-call* 1 | | | |
|---|---|---|---|---|---|---|---|---|---|
| *Strike* | *Control Variate* | *Std* | *Rel. Std* | *EF* | *Rel. EF* | *Std* | *Rel. Std* | *EF* | *Rel. EF* |
| 0.08 | Antithetic | 0.7928 | 1.0 | 0.0609 | 1.0 | 1.6048 | 1.0 | 0.6799 | 1.0 |
| 0.08 | Cap | 0.1752 | 4.5 | 0.0040 | 15.2 | 0.5121 | 3.1 | 0.0843 | 8.1 |
| 0.08 | Caplets | 0.1554 | 5.1 | 0.0028 | 21.8 | 0.4939 | 3.2 | 0.0681 | 10.0 |
| 0.08 | Caplets all | 0.1437 | 5.5 | 0.0031 | 19.6 | 0.4635 | 3.5 | 0.0782 | 8.7 |
| 0.08 | Delta hedge | 0.3796 | 2.1 | 0.0352 | 1.7 | 0.8451 | 1.9 | 0.4185 | 1.6 |
| 0.08 | Swap | 0.3702 | 2.1 | 0.0188 | 3.2 | 0.9120 | 1.8 | 0.3068 | 2.2 |
| 0.08 | Z-Bond | 0.3695 | 2.1 | 0.0153 | 4.0 | 0.9095 | 1.8 | 0.2372 | 2.9 |
| 0.10 | Antithetic | 0.9605 | 1.0 | 0.1715 | 1.0 | 1.6049 | 1.0 | 1.3940 | 1.0 |
| 0.10 | Cap | 0.2061 | 4.7 | 0.0086 | 19.9 | 0.5319 | 3.0 | 0.1642 | 8.5 |
| 0.10 | Caplets | 0.2315 | 4.1 | 0.0106 | 16.2 | 0.5812 | 2.8 | 0.1873 | 7.4 |
| 0.10 | Caplets all | 0.2020 | 4.8 | 0.0090 | 19.1 | 0.5137 | 3.1 | 0.1637 | 8.5 |
| 0.10 | Delta hedge | 0.5205 | 1.8 | 0.1099 | 1.6 | 0.9105 | 1.8 | 0.9003 | 1.5 |
| 0.10 | Swap | 0.7065 | 1.4 | 0.1096 | 1.6 | 1.3898 | 1.2 | 1.2543 | 1.1 |
| 0.10 | Z-Bond | 0.7040 | 1.4 | 0.0963 | 1.8 | 1.3948 | 1.2 | 1.0949 | 1.3 |
| 0.12 | Antithetic | 0.8544 | 1.0 | 0.1723 | 1.0 | 1.6621 | 1.0 | 1.8878 | 1.0 |
| 0.12 | Cap | 0.1895 | 4.5 | 0.0087 | 19.8 | 0.4754 | 3.5 | 0.1593 | 11.9 |
| 0.12 | Caplets | 0.1868 | 4.6 | 0.0085 | 20.3 | 0.4450 | 3.7 | 0.1376 | 13.7 |
| 0.12 | Caplets all | 0.1693 | 5.0 | 0.0075 | 23.0 | 0.3986 | 4.2 | 0.1194 | 15.8 |
| 0.12 | Delta hedge | 0.4025 | 2.1 | 0.0818 | 2.1 | 0.7574 | 2.2 | 0.7999 | 2.4 |
| 0.12 | Swap | 0.6252 | 1.4 | 0.1055 | 1.6 | 1.3518 | 1.2 | 1.4502 | 1.3 |
| 0.12 | Z-Bond | 0.6231 | 1.4 | 0.0952 | 1.8 | 1.3587 | 1.2 | 1.3036 | 1.4 |

*Std denotes the standard deviation (in bp) of the lower-bound estimate. EF denotes the efficiency factor, i.e., the variance in bp scaled with the average time per path, measuring the numerical efficiency of the various controls. Low values are preferred to high. For comparison, we list for each control the efficiency ratio and standard deviation relative to those of the antithetic variate application.*

control variates is more efficient for the 6 no-call 1 Bermudan swaptions. What is especially interesting is that the cap and the caplets outperform the other control variates for all strikes and maturities. The series of seven caplets works slightly better for ITM options and the series of all caplets for out-of-the-money (OTM) options. Furthermore, the more an option is in the money, the less efficient it is to use all caplets than to use only seven.

Differences in efficiency between the cap and the caplets are negligible compared to the other controls. Note that it is easier to use the cap as a control than the caplets, as no decision on how many caplets to use is required. Furthermore, using too many or too few caplets may cause the caplet control to be less efficient than the single cap.

A comparison of the other types of control variates shows that the series of zero-coupon bonds is more efficient than the single swap in all cases. Furthermore, in most cases the hedge portfolio results in more reduction in variance than the zero-coupon bonds. The efficiency factors, however, show that zero-coupon bonds are numerically more efficient than the hedge portfolio for the ITM options and partially so for the at-the-money (ATM) options.

As there may be additional gains in combining various types of controls, we analyze the use of multiple controls. We run a set of simulations estimating lower bounds for the 11 no-call 1 Bermudan swaption, using the same setting and applying all possible combinations of the controls. We sample the controls as before, except that for each path we sample, say, a cap and a swap rather than just the cap.

The 15 most efficient multiple control variates are listed in Exhibit 3. Interestingly enough, the caplet control still performs reasonably well compared to the multiple control variates, especially for OTM options. Combining a series of caplets or a single cap with a series of zero-coupon bonds is still more efficient.

As expected, using the zero-coupon bond in combination with the cap and caplets is more efficient than using the swap in combination with the cap and caplets. The triple combination of swap, zero-coupons, and cap/caplets also performs well, but not as well as the zero-coupons in combination with the cap/caplet only. Finally, we see the delta hedge portfolio is far more time-consuming than other controls.

### Testing Strategies

The motivation for testing strategies is to reduce estimation time and the bias of the estimator, while control variates are applied to reduce the variance of the estimator. Both approaches are used to improve numerical efficiency, but in two different ways.

We use the approximate European swaption formula of Andersen and Andreasen [2000] for valuation of the core European swaptions. Andersen [2000] demonstrates that strategy Bar works well for most Bermudan swaptions across several scenarios and that only for short

## EXHIBIT 3

**Efficiency of Combinations of Control Variates on Lower Bounds**

| *Rank* | *Strike* = 0.08 *EF* | *Control* | *Strike* = 0.10 *EF* | *Control* | *Strike* = 0.12 *EF* | *Control* |
|---|---|---|---|---|---|---|
| 1 | 0.0526 | ZB+Cap | 0.1403 | ZB+CLA | 0.1172 | ZB+CLA |
| 2 | 0.0530 | ZB+CL | 0.1493 | ZB+Cap | 0.1194 | CLA |
| 3 | 0.0620 | ZB+CLA | 0.1625 | SW+CLA | 0.1252 | ZB+Cap |
| 4 | 0.0654 | SW+Cap | 0.1636 | ZB+CL | 0.1298 | ZB+CL |
| 5 | 0.0661 | SW+CL | 0.1637 | CLA | 0.1352 | SW+CLA |
| 6 | 0.0681 | CL | 0.1642 | Cap | 0.1376 | CL |
| 7 | 0.0697 | SW+CLA | 0.1701 | SW+Cap | 0.1428 | SW+Cap |
| 8 | 0.0747 | ZB+SW+CL | 0.1816 | ZB+SW+CLA | 0.1446 | SW+CL |
| 9 | 0.0762 | ZB+SW+Cap | 0.1824 | ZB+SW+Cap | 0.1476 | ZB+SW+CLA |
| 10 | 0.0782 | CLA | 0.1857 | SW+CL | 0.1498 | ZB+SW+Cap |
| 11 | 0.0840 | ZB+SW+CLA | 0.1873 | CL | 0.1565 | ZB+SW+CL |
| 12 | 0.0843 | Cap | 0.2023 | ZB+SW+CL | 0.1593 | Cap |
| 13 | 0.1000 | ZB+Cap+DH | 0.2228 | ZB+CLA+DH | 0.2027 | ZB+CLA+DH |
| 14 | 0.1019 | ZB+CL+DH | 0.2372 | SW+CLA+DH | 0.2032 | CLA+DH |
| 15 | 0.1061 | ZB+CLA+DH | 0.2420 | ZB+Cap+DH | 0.2157 | ZB+CL+DH |

*Combinations of controls are ranked according to performance; only the 15 best multiple controls are listed. Control indicates the combination of control variates. ZB = Zero-Coupon, SW = Swap, CL = Subsample of Caplets, CLA = Caplets All, and DH = Delta Hedge.*

options on long swaps in a multifactor model do the strategies Bar & Max and Bar + Max add significant value. As our tests show this too, we concentrate on the 11 no-call 1 Bermudan swaption. Computations are based on 50,000 antithetic paths. To reduce the variance of the estimator, we apply the cap as a control variate.[8]

Lower-bound estimates and estimation times are shown in Exhibit 4. We benchmark strategies Bar & Max, Bar + Max, Bar & First, and Bar + First against the simple strategy Bar, computing the pickup value as the difference between the lower bound estimated by strategy Bar and the bounds estimated by the more advanced strategies. As all estimates are lower bounds on the true option price, high pickup values are preferred.

Several interesting observations can be drawn from Exhibit 4. First of all, strategies Bar & Max, Bar + Max, Bar & First, and Bar + First add value over strategy Bar. Second, strategy Bar + Max adds more value than strategy Bar & Max, but is more time-consuming. Estimation time increases because all remaining core European swaptions must be valued at every exercise decision of strategy Bar + Max, while valuation of the remaining core European swaptions can be skipped as soon as one is more valuable than the intrinsic value in strategy Bar & Max. Naturally, this is particularly visible for OTM options because there are more expected exercise decisions.

An intuitive explanation as to why strategy Bar + Max adds more value than strategy Bar & Max is that the critical barrier or strategy Bar + Max can be interpreted as the value of the deferred exercise premium at the exercise boundary, remembering that we are searching for the barrier $H_{B\&M}(t_i)$ that would make the intrinsic value equal to the continuation value, $X(t_i) = H_{B\&M}(t_i) + M(t_i)$. The barriers of strategies Bar and Bar & Max approximate the whole sum $H(t_i) + M(t_i)$ rather than just parts of it, and since the early/deferred exercise value goes to zero when the option is deep in or out of the money, the sum might be more difficult to approximate.

Overall, this suggests that in the case of the LIBOR market model, where good approximations for the core European options are available, the parameterization of the exercise strategies should be formulated in terms of early/deferred exercise values rather than simply tests of the European option value.

Finally, it is particularly interesting that strategies Bar & First and Bar + First yield lower-bound estimates that are only marginally lower than those of strategies Bar & Max and Bar + Max, and at significantly reduced computational effort.

Strategies Bar & First and Bar + First enhance strategies Bar & Max and Bar + Max in two ways:

- The European swaption valuations needed for each exercise decision are reduced from several to one.
- Short-maturity swaptions are much cheaper to evaluate than long-maturity swaptions since less volatility has to be integrated. The particular European swaption is therefore priced quickly.

The efficiency gains of the modified strategies Bar & First and Bar + First are indisputable. Strategy Bar + First is only about 10%-30% slower than strategy Bar, but 10 to 30 times faster than strategy Bar + Max. Time spent in presimulation of the critical barriers is not included in the timing of the lower bounds; including it would make the performance improvement even more obvious. Even so, the conclusion is quite clear. Strategy Bar + First is highly preferable as the pickup value is indistinguishable

# EXHIBIT 4

**Performance of Strategies on Lower Bounds**

| *Strike* | *Strategy* | *Low bps* | *Pickup bps* | *Time min : sec* | *Relative Time* |
|---|---|---|---|---|---|
| 0.08 | Bar | 1247.62 (0.36) | | 00:16 | 1.0 |
| 0.08 | Bar&Max | 1251.72 (0.37) | 4.1 | 05:39 | 21.2 |
| 0.08 | Bar+Max | 1256.01 (0.34) | 8.4 | 09:39 | 36.2 |
| 0.08 | Bar&First | 1251.63 (0.36) | 4.0 | 00:22 | 1.4 |
| 0.08 | Bar+First | 1255.32 (0.34) | 7.7 | 00:20 | 1.3 |
| 0.10 | Bar | 621.88 (0.38) | | 00:28 | 1.0 |
| 0.10 | Bar&Max | 628.59 (0.35) | 6.7 | 04:40 | 10.0 |
| 0.10 | Bar+Max | 634.95 (0.31) | 13.1 | 11:17 | 24.2 |
| 0.10 | Bar&First | 628.89 (0.35) | 7.0 | 00:33 | 1.2 |
| 0.10 | Bar+First | 634.62 (0.31) | 12.7 | 00:33 | 1.2 |
| 0.12 | Bar | 329.41 (0.33) | | 00:35 | 1.0 |
| 0.12 | Bar&Max | 334.29 (0.31) | 4.9 | 02:18 | 3.9 |
| 0.12 | Bar+Max | 338.80 (0.30) | 9.4 | 06:43 | 11.5 |
| 0.12 | Bar&First | 334.15 (0.33) | 4.7 | 00:39 | 1.1 |
| 0.12 | Bar+First | 338.55 (0.30) | 9.1 | 00:38 | 1.1 |

*Strategies Bar, Bar & Max, and Bar + Max are strategies proposed by Andersen [2000]; strategies Bar & First and Bar + First represent our simplifying strategies. The pickup from using the advanced strategies instead of the simple strategy Bar is included as well as the estimation time relative to that of strategy Bar.*

# EXHIBIT 5

**Performance of Strategies on Duality Gaps—11 No-Call 1 Bermudan Swaption for Various Strikes**

| *Strike* | *Strategy* | *Duality bps* | *Time hour : min : sec* | *Relative Time* |
|---|---|---|---|---|
| 0.08 | Bar | 20.77 (0.81) | 00:08:35 | 1.0 |
| 0.08 | Bar&Max | 17.14 (0.64) | 01:57:38 | 13.7 |
| 0.08 | Bar+Max | 13.6 (0.48) | 03:51:15 | 26.9 |
| 0.08 | Bar&First | 16.72 (0.64) | 00:10:55 | 1.3 |
| 0.08 | Bar+First | 14.19 (0.50) | 00:11:10 | 1.3 |
| 0.10 | Bar | 25.26 (0.85) | 00:14:28 | 1.0 |
| 0.10 | Bar&Max | 19.31 (0.59) | 01:36:45 | 6.7 |
| 0.10 | Bar+Max | 12.55 (0.40) | 04:12:37 | 17.5 |
| 0.10 | Bar&First | 17.73 (0.58) | 00:16:50 | 1.2 |
| 0.10 | Bar+First | 13.02 (0.42) | 00:17:05 | 1.2 |
| 0.12 | Bar | 15.44 (0.63) | 00:18:18 | 1.0 |
| 0.12 | Bar&Max | 10.61 (0.43) | 01:01:12 | 3.3 |
| 0.12 | Bar+Max | 6.00 (0.27) | 02:51:01 | 9.3 |
| 0.12 | Bar&First | 10.89 (0.46) | 00:20:18 | 1.1 |
| 0.12 | Bar+First | 6.79 (0.30) | 00:20:22 | 1.1 |

*Numbers are computed using 1,600 antithetic outer paths and 40 antithetic inner paths in the primal-dual algorithm. Critical barriers are estimated using 50,000 antithetic paths.*

from that of strategy Bar + Max and the additional computational cost relative to Bar alone is negligible.

## Upper-Bound Calculations

It takes much longer to estimate the upper bound than the lower bound. It is therefore especially interesting to apply the efficiency improvement techniques to estimation of the duality gap. As the nested simulations are computationally expensive, it is also important not to use too many inner paths compared to outer paths.

In Jensen and Svenstrup [2004] we address the variance-bias trade-off between reducing the bias or the variance of the duality gap estimator, propose a general optimal allocation rule for the relation between the number of inner and outer paths, and present a general procedure for estimation of this allocation rule. We find $n(m) \propto m^2$, where $n$ and $m$ are the number of outer and inner paths, respectively. In other words, when one doubles the number of outer paths, one should increase the number of inner paths by a factor of only $\sqrt{2}$. As a result, we use 1,600 antithetic outer paths and 40 antithetic inner paths rather than the 750 outer and 300 inner paths we use for the benchmark case in Exhibit 1.

First, we test the impact of the exercise strategies on the duality gap, focusing on the 11 no-call 1 Bermudan swaption. Duality gaps and estimation times are reported in Exhibit 5. Low values of the duality gap are preferred, as the gap can be interpreted as a measure of the suboptimality of the exercise strategy.

Again strategy Bar + Max clearly outperforms strategies Bar and Bar & Max. The duality gap estimates of strategies Bar & First and Bar + First are not significantly different from those of Bar & Max and Bar + Max, although the strategies are much more computationally efficient. Strategy Bar + First is 8.4 to 20.7 times faster than strategy Bar + Max, depending on the strike.

The conclusion is very clear. Strategy Bar + First is highly preferable since it generates duality gap estimates that are indistinguishable from those of the advanced strategy Bar + Max and almost as quickly as the simple strategy Bar.

Second, we test the single cap as a control variate on the duality gap estimate.[9] Because we expect minor variations across maturity, we compute duality gaps for the 3, 6, and 11 no-call 1 Bermudan swaption with various strikes. Duality gap estimates and estimation times are presented in Exhibit 6 for strategy Bar + First.

The results very clearly show considerably reduced duality gaps at a reasonable computational cost across all maturities and strikes. In fact, when we combine the cap control with strategy Bar + First and the improved path

## EXHIBIT 6
**Effect of Control Variate on Duality Gaps**

| | | *None* | | *Cap* | | *Reduce* | *Increase* |
|---|---|---|---|---|---|---|---|
| $t_e$ | *Strike* | $\hat{D}_0$ (*Std*) | *Time* | $\hat{D}_0^{Cap}$ (*Std*) | *Time* | $\hat{D}_0$ | *Time* |
| 3 | 0.08 | 0.17 (0.02) | 00:15 | 0.09 (0.01) | 00:20 | 47% | 33% |
| 3 | 0.10 | 0.36 (0.03) | 00:20 | 0.20 (0.02) | 00:25 | 44% | 25% |
| 3 | 0.12 | 0.20 (0.02) | 00:21 | 0.12 (0.02) | 00:25 | 40% | 19% |
| 6 | 0.08 | 2.14 (0.12) | 01:47 | 0.94 (0.06) | 02:27 | 56% | 37% |
| 6 | 0.10 | 2.61 (0.15) | 02:44 | 1.27 (0.08) | 03:17 | 51% | 20% |
| 6 | 0.12 | 1.69 (0.13) | 03:15 | 1.31 (0.11) | 03:42 | 23% | 14% |
| 11 | 0.08 | 14.18 (0.53) | 10:38 | 7.57 (0.29) | 13:18 | 47% | 25% |
| 11 | 0.10 | 12.82 (0.42) | 16:20 | 6.71 (0.24) | 18:35 | 48% | 14% |
| 11 | 0.12 | 5.99 (0.28) | 19:33 | 3.86 (0.18) | 21:20 | 36% | 9% |

allocation rule in Jensen and Svenstrup [2004], we get a much smaller duality gap estimate that is much more precise and much faster to estimate than the benchmark duality gap in Exhibit 1.

## IV. SUMMARY

We have addressed numerically efficient valuation of Bermudan options in the primal-dual simulation algorithm of Andersen and Broadie [2004]. Implementation has focused on valuation of Bermudan swaptions using a two-factor LIBOR market model with deterministic volatility. We test the efficiency gains from applying control variates on lower-bound estimates, using a range of carefully chosen control variates. Following Rasmussen [2002], we sample all control variates at the time of exercise of the Bermudan swaption, and we expand this sampling technique to handle dividend-paying controls.

Using control variates reduces the variance of the estimators significantly. In particular, for the 11 no-call 1 Bermudan swaption, the standard deviations of the lower-bound estimates are reduced by a factor of four in the best case. While the most efficient control variate is clearly a series of caplets or a single cap, the cap is preferred because use of caplets requires choosing which and how many caplets to use. Combining the caplets or the cap with a series of zero-coupon bonds improves the variance reduction slightly.

Second, we have proposed a general simplification of the exercise strategies presented in Andersen [2000], particularly replacing the maximum taken over several option values by one single carefully chosen option. Application to Bermudan swaptions shows that the simplified strategies are highly preferable, as the lower-bound estimates are indistinguishable from the estimates in the advanced strategies, and estimation time is considerably shorter. For the 11 no-call 1 Bermudan swaption, simplification of the strategies produces 10 to 30 times faster lower-bound estimations, depending on the moneyness of the option.

Finally, we demonstrate the effect of the enhanced exercise strategies and the method of control variates on the duality gap. The estimation results clearly show significant improvements in both cases. When we combine these two efficiency improvements with the optimal simulation path allocation rule, we are able to reduce bias in lower- and upper-bound estimates, improve precision, and shorten estimation time significantly.

## APPENDIX

### Control Variates

The most straightforward implementation of the method of control variates is to replace the evaluation of an unknown expectation with evaluation of the difference between the unknown quantity and another expectation whose value is known. Suppose that at time $t$ we know the expectation $E_t^{\mathbb{Q}}[Y]$ of an $M$-dimensional stochastic vector $Y$. Assuming that we can sample $I$ realizations of a scalar variable $Z$ exactly, an unbiased estimator $\bar{Z}^{CV}$ of $E_t^{\mathbb{Q}}[Z]$ is given by:

$$\bar{Z}_t^{CV}(\beta) = \frac{1}{I}\sum_{i=1}^{I}\left(Z^i - \beta'\left(Y^i - E_t^{\mathbb{Q}}[Y]\right)\right)$$

with variance

$$Var\left(\bar{Z}_t^{CV}(\beta)\right) = \frac{1}{I}\left(\sigma_Z^2 - 2\beta'\Sigma_{YZ} + \beta'\Sigma_Y\beta\right) \tag{A-1}$$

for some appropriately chosen parameter vector $\beta \in \mathbb{R}^M$. Here $\sigma_Z^2$ is the variance of $Z$, $\Sigma_Y$ is the covariance matrix of the controls, and $\Sigma_{YZ}$ is the vector of covariances between $Z$ and $Y$.

The variance-minimizing choice of $\beta$ is given by

$$\beta^* = \Sigma_Y^{-1}\Sigma_{YZ} \tag{A-2}$$

Inserting (A-2) into (A-1), we have at optimality:

$$Var\left(\bar{Z}_t^{CV}(\beta^*)\right) = \frac{1}{I}\left(1 - R^2\right)\sigma_Z^2$$

where

$$R^2 = \frac{\Sigma_{YZ}' \Sigma_Y^{-1} \Sigma_{YZ}}{\sigma_Z^2}$$

Thus, the variance reduction is determined by the coefficient of multiple correlation $R$ between $Z$ and the control variates $Y$. Because $R^2 \in (0,1)$, variance is always reduced, but as use of controls requires additional estimation time, efficiency is not necessarily improved. In effect, good control variates should have closed-form solutions, be highly correlated with $Z$, and be relatively easy to compute.

## ENDNOTES

The authors acknowledge comments and insights from Leif B.G. Andersen, Ren Raw Chen, Bent Jesper Christensen, Tom Engsted, and Nicki S. Rasmussen; and seminar participants at the Stockholm School of Economics, the 2003 FMA Annual Meeting Doctoral Student Consortium, the Sixth International Conference on Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing, and the 2004 Second International Conference on Monte Carlo and Probabilistic Methods for Partial Differential Equations. This research was supported by ScanRate Financial Systems A/S and the Danish Social Science Research Councils.

[1]We require that $\sup|\pi(t)| < \infty$ and $\pi(0) = 0$ for all $\pi \in \mathcal{H}$

[2]A lookback option is an option whose payoff is determined not only by the settlement price but also by the maximum or minimum price of the underlying asset over the life of the option.

[3]Monte Carlo simulation is only one way to estimate the conditional expectations. For particular applications, analytical solutions or other numerical methods might be faster or more accurate.

[4]For upper-bound calculations we will use at-the-money caplets. Thus, when calculating $E_t(\cdot)$ we fix the strike at time $t$ such that the caplet is ATM.

[5]The LIBOR market model is in the Heath, Jarrow, and Morton [1992] class of models where the initial yield curve is known.

[6]All timed estimations are performed on a Pentium M 1.5 GHz laptop.

[7]Results of unreported studies show that a subsample of seven caplets equally spaced over the exercise dates of the Bermudan swaption is adequate.

[8]We expect a limited interaction between the effect of the controls and the strategy, as the exercise decisions of the strategies are close to optimal and therefore similar. Hence, we argue that applying the cap as a control does not change the outcome of the analysis significantly. The alternative is to add to the number of sample paths considerably and hence also the estimation time.

[9]When we test multiple control variates for estimation of the duality gap, we experience problems in determining the variance-minimizing parameter due to the small number of inner loops.

## REFERENCES

Andersen, L. "A Simple Approach to the Pricing of Bermudan Swaptions in the Multifactor LIBOR Market Model." *Journal of Computational Finance*, 3 (2000), pp. 5-32.

Andersen, L., and J. Andreasen. "Volatility Skews and Extensions of the LIBOR Market Model." *Applied Mathematical Finance*, 7 (2000), pp. 1-32.

Andersen, L., and M. Broadie. "A Primal-Dual Simulation Algorithm for Pricing Multi-Dimensional American Options." *Management Science*, 50 (2004), pp. 1222-1234.

Bossaerts, P. "Simulation Estimators of Optimal Early Exercise." Working paper, Carnegie Mellon University, 1988.

Boyle, P. "A Lattice Framework for Option Pricing with Two State Variables." *Journal of Financial and Quantitative Analysis*, 23 (1988), pp. 1-12.

Boyle, P., J. Evnine, and S. Gibbs. "Numerical Evaluation of Multivariate Contingent Claims." *The Review of Financial Studies*, 2 (1989), pp. 241-250.

Brace, A., D. Gatarek, and M. Musiela. "The Market Model of Interest Rate Dynamics." *Mathematical Finance*, 7 (1997), pp. 127-154.

Broadie, M., and P. Glasserman. "Pricing American-Style Securities Using Simulation." *Journal of Economic Dynamics and Control*, 21 (1997), pp. 1323-1352.

Glynn, P.W., and W. Whitt. "The Asymptotic Efficiency of Simulation Estimators." *Operations Research*, 40 (1992), pp. 505-520.

Haugh, M.B., and L. Kogan. "Pricing American Options: A Duality Approach." *Operations Research*, 52 (2004), pp. 258-270.

Heath, D., R. Jarrow, and A. Morton. "Bond Pricing and the Term Structure of Interest Rates: A New Methodology." *Econometrica*, 60 (1992), pp. 77-105.

Jamshidian, F. "LIBOR and Swap Market Models and Measures." *Finance and Stochastics*, 1 (1997), pp. 293-330.

Jensen, M.S., and M. Svenstrup. "Optimal Allocation of Simulation Paths in the Primal-Dual Algorithm." Working paper, Centre for Analytical Finance, 2004. Available at www.caf.dk.

Longstaff, F.A., P. Santa-Clara, and E. S. Schwartz. "Throwing Away a Billion Dollars: The Cost of Suboptimal Exercise Strategies in the Swaptions Market." *Journal of Financial Economics*, 62 (2001), pp. 39-66.

Longstaff, F.A., and E.S. Schwartz. "Valuing American Options by Simulation: A Simple Least Squares Approach." *The Review of Financial Studies*, 14 (2001), pp. 113-147.

Miltersen, K.R., K. Sandmann, and D. Sondermann. "Closed Form Solutions for Term Structure Derivatives with Log-normal Interest Rates." *Journal of Finance*, 52 (1997), pp. 409-430.

Rasmussen, N.S. "Efficient Control Variates for Monte Carlo Valuation of American Options." Working paper, SSRN, 2002. Available at http://ssrn.com/abstract=325260.

Svenstrup, M. "On the Suboptimality of Single-Factor Exercise Strategies for Bermudan Swaptions." Working paper, SSRN, 2002. Forthcoming, *Journal of Financial Economics*. Available at http://ssrn.com/abstract=388141.

Tilley, J. "Valuing American Options in a Path Simulation Model." *Transactions of the Society of Actuaries*, 45 (1993), pp. 83-104.

*To order reprints of this article, please contact Ajani Malik at amalik@iijournals.com or 212-224-3205.*