

MONTE CARLO BOUNDS FOR CALLABLE PRODUCTS WITH NON-ANALYTIC BREAK COSTS

MARK S. JOSHI

ABSTRACT. The pricing of callable derivative products with complicated pay-offs is studied. A new method for finding upper bounds by Monte Carlo simulation is introduced, this relies on modelling the callable product directly. The method has a wide range of applicability and is shown to be effective for Asian tail products.

1. INTRODUCTION

The valuation of derivative securities which feature early exercise by Monte Carlo simulation is a classic hard problem. The problem essentially being that the decision to exercise relies on knowing the value of the unexercised product which is not available in a process which evolves through time. Backwards methods such as trees and PDEs are therefore better adapted to such problems. However, where the problem is high-dimensional, such as the LIBOR market model, or involves path-dependence trees and PDEs are also problematic, and it becomes necessary to adapt the Monte Carlo simulation technique.

In recent years, one method of finding lower bounds by Monte Carlo simulation has become very popular: the Longstaff–Schwartz method of least-squares regression, [8]. This has been found to be effective in a wide range of cases. However, to be truly sure of effectiveness, it is necessary to have an upper bound, and see that the bounds are close together. Rogers, [11], suggested using a hedging and duality approach which involved optimizing over a set of hedging portfolios. Anderson and Broadie, [2], adapted this method by choosing the product with a sub-optimal exercise strategy as hedge. This method proved effective. For more background, see [3].

However, both of the Longstaff–Schwartz and Anderson–Broadie methods are designed for the case of a product which effectively pays a cash sum on exercise. They can therefore be applied to the right to exercise into products whose value is analytically computable. For example, a

Date: January 6, 2006.

Key words and phrases. Monte Carlo simulation, Bermudan option, American option, upper bounds.

Bermudan swaption is the right to enter into a swap at various times. The value of a swap is immediately computable, and so the method applies. In fact, most callable products actually involve the right to terminate a stream of cashflows rather than the right to enter into one. Application of these methods therefore relies on the decomposition of the product into the analagous non-callable product and the right to enter into the opposite product. For example, a breakable swap is a vanilla swap plus a Bermudan option on the opposite swap.

Here we focus on products where the exercise value is not analytically computable. For example, a CMS steepener is a swap that pays a floating rate against the capped and floored difference of two CMS rates, (e.g. the 2 and 10 year rates,) the break value of this is not analytically evaluable.

Another example, is an Asian tail callable bond. An investor puts his principal into a bond that pays the growth of an index over a fixed period, e.g. 3 years, and has his principal guaranteed. The growth is assessed by taking the average of the index during the final year with bimonthly monitoring. He also has the right to redeem the bond at a pre-agreed rate, e.g., 0.98, at the end of year one and year two. The Asian tail means that the value of the opposite contract is not immediately evaluable at either of the two break dates. Here we will focus on the Asian tail bond, since it is sufficiently complex to illustrate our methods, whilst being sufficiently simple to avoid obscuring the method with irrelevant details. The extension to more complex cases such as the LIBOR market model is straightforward.

For lower bounds, Piterbarg, [10], suggested using a regression approach to estimate the exercise value, but his method does not appear applicable to upper bounds. We examine the lower bound approach of Amin, [1], and show how related ideas can be developed to apply to upper bounds. Both on the upper and lower methods rely on working with the callable product directly rather than decomposing it. This has the virtue that all cash-flow values that need to be computed are model-independent being simply the cash-flows generated by the product's pay-off.

We therefore find the exercise strategy by comparing the value of the regression of the continuation cash-flows to the value obtained on early exercise. For the upper bound, we work out the hedging error along a path by netting the cash-flows from product and hedge. A key part of our approach is the observation that Rogers' approach is really a statement about the equality of buyer's and seller's prices of callable derivatives.

The structure of this paper is as follows. We establish notation and the general set-up in Section 2. We review Amin's adaptation of Longstaff-Schwartz in Section 3. We recall the upper bound methods of Rogers and Anderson–Broadie in Section 4. We also reinterpret Rogers' result, and this allows us to develop the upper bound method. In Section 5, we present numerical results for an Asian tail bond in a jump-diffusion model, and see that tight bounds are obtained. We conclude in Section 6.

I am grateful to Vladimir Piterbarg for helpful comments on an earlier version of this paper.

2. THE PROBLEM

We have a stock price, S_t , following some process, and a product D which pays at a set of increasing times, t_i , $i = 0, \dots, k$, cash-flows which depend on the stock price path up to that time. We also have a set of call times, c_j , $j = 0, \dots, l$, at each of which the investor has the right to terminate the product and receive a rebate R_j which may depend on the stock price. For simplicity and definiteness, we assume that the sets $\{t_i\}$ and $\{c_j\}$ are distinct. One can always move a time forward or backwards by a very small quantity to ensure this distinctness. We will only deal with the case where both sets of times are finite, which is the only practical case in a discrete simulation in a computer in any case. Thus American options would need to be approximated by a Bermudan option with many call dates.

We give a couple of examples. In the first, $k = 0$, and $l = N$ large. Suppose the derivative D pays $(K - S)_+$ at time $t_0 = T$, and at each time c_j the rebate is $(K - S)_+$, we then have a Bermudan put option.

A second example, is $k = 0$ and $t_0 = 3$, and $l = 1$ with $c_0 = 1$ and $c_1 = 2$. The cash-flow paid at time 3 is

$$\max(A/S_0, 1.0) = \max(A/S_0 - 1.0, 0) + 1.0,$$

where

$$A = \frac{1}{6} \sum_{j=1}^6 S_{2+j/6}.$$

Both rebates are equal to 0.98. This is the Asian tail bond.

We will find a lower bound by developing an exercise strategy. This will be a rule for breaking the contract which depends only on information gathered up to that time. Once the breaking strategy has been chosen, the pricing is then that of a path-dependent derivative without optionality and is straight-forward.

3. LOWER BOUNDS FOR CALLABLE PRODUCTS

The essence of the Longstaff–Schwartz method is to use least-squares regression to estimate the value of the unexercised product. This estimate is then compared to the exercise value, and exercise occurs according to which is greater. For Amin’s adaptation, the approach is similar, one regresses the value of the cash-flows one would receive on non-exercise to gain an approximation function. This function is then compared to the rebate value and the exercise decision is made on the basis of which is bigger.

As with Longstaff–Schwartz one proceeds in a backwards fashion. We work with discounted cash-flows. Our numeraire would be the continuously compounding money-market account in a stock model but could equally well be a zero-coupon bond in the LIBOR market model. Thus for each path, we need to know the value of the ratio of cash-flows to numeraire between each pair of break dates and after the last one. Our only change to the algorithm is that our exercise opportunity gives us a choice between the rebate and the sum of cash-flows up to the next exercise plus the value of the ratio on the next exercise date.

Our modified algorithm is therefore as follows:

- First generate the paths. For each path store the sum of the discounted value of the cash-flows that occur between each exercise opportunity and the next, and after the last one.
- For the last exercise opportunity, regress the basis functions against the discounted cash-flows.
- Decide whether the contract would continue at the last exercise opportunity for each path by seeing whether the regression function is greater than the rebate. Store the value at the last time frame to be either the break value or the continuation value on that path according to the result of the exercise decision.
- Go back to the previous time frame, taking the continuation value for each path to be the sum of the discounted cash-flows up to the next time frame and the value at the next time frame, and repeat the same regression.
- Work all the way back repeating.

One advantage of this approach is that coding of the pay-off of the contract is wholly independent of the model.

Note that there is an implicit assumption in this method that the cash-flows between two exercise dates (or after the last one) only depend upon the value of the stock price between those two dates. Whilst the method could be applied without this assumption, the state-dependence

arising from the previous stock price path is not taken into account and we would expect poor performance.

Once one has the regression functions, one runs a second Monte Carlo simulation to obtain an unbiased estimate of the lower bound.

4. UPPER BOUNDS

We first recall and reinterpret Rogers' method. The price of an American derivative, C , is usually written as

$$C_0 = Z_0 \sup_{\tau} \mathbb{E}(Z_{\tau}^{-1} C_{\tau}), \quad (4.1)$$

where C is the pay-off and Z is the numeraire. The expectation is taken in the pricing measure associated with the numeraire Z , and the supremum is taken over the set of stopping times. The restriction to the set of stopping times corresponds to the fact that the holder of the option cannot see the future. We assume that the product has a finite time horizon and non-exercise corresponds to exercise at the final time horizon receiving zero. This ensures that the stopping times are finite. If we introduce a hedge, H , of zero initial value, a hedge being a self-financing trading strategy, then its discounted value will be a martingale and it will continue to be a martingale if stopped at a bounded stopping time by the Optional Sampling theorem, [6]. We therefore have for any hedge, H , letting $M_t = H_t/Z_t$,

$$C_0 = Z_0 \sup_{\tau} \mathbb{E}(Z_{\tau}^{-1} C_{\tau} - M_{\tau}). \quad (4.2)$$

If we now pass to the set of all random times, instead of all stopping times, we can only increase the value. There is also a clear best strategy, exercise with maximal foresight, that is take the maximum along the path. We therefore have

$$C_0 \leq Z_0 \inf_M \mathbb{E}(\max(Z_t^{-1} C_t - M_t)), \quad (4.3)$$

where the maximum is taken over all exercise times, and the infimum is taken over all martingales of initial value zero. Rogers' result, proved using the Doob decomposition of the Snell envelope, is that equality holds.

To understand intuitively why it holds, we can interpret the right hand side as the seller's price. The seller of an option cannot choose the exercise strategy but has to cover its pay-off regardless of when the buyer exercises. In particular, he cannot be sure that the buyer will use the optimal strategy. Indeed the buyer could choose the exercise date randomly. This means that the buyer has a non-zero probability of exercising at the maximum. Since in derivatives pricing we wish to

cover the cost no matter what happens, the seller has to be able to cope with a buyer exercising at the maximum. However, the seller can hedge. Equality in equation (4.4), therefore states that by investing the buyers price the seller can hedge in such a way as to cover the pay-off even if the buyer is exercising at the maximum. In other words, the buyer's and seller's prices are the same. A version of this result is in [7].

The key to the Anderson–Broadie method and to our method lies in identifying the optimal hedge. We assume a complete market. Suppose we have sold the derivative C and wish to cover its pay-off. We dynamically replicate (or purchase which is equivalent in a complete market) C with the optimal exercise strategy. At each exercise date, there are four possibilities, according to the buyer's exercise strategy and the optimal one.

In the two cases, where they agree the seller is certainly hedged. If the buyer exercises and the optimal strategy says not to, then the derivative with optimal strategy is worth more than the exercise value and the seller is ahead when he dissolves his portfolio. If the buyer does not exercise and the seller's strategy is to exercise, then the seller exercises and then continues to dynamically replicate. This will result in a profit for the seller as the continuation value will necessarily will be less than the exercise value for the optimal strategy. The excess profit is put into numeraire bonds (or any positive instrument).

This means that by adopting the trading strategy of buy the instrument to be hedged, exercise according to optimal strategy, and at points of exercise, exercise and buy the remaining product, with excess cash going into numeraire bonds, the seller can hedge against exercise with perfect foresight.

We therefore take this optimal hedge and evaluate the expression

$$Z_0 \mathbb{E}(\max(Z_t^{-1}C_t - M_t)), \quad (4.4)$$

to get the seller's price. Unfortunately, we do not know the optimal exercise strategy so this is not possible. However, we can carry out the same trading strategy based on any exercise rule. We therefore take the rule developed for the lower bound and apply the same hedging approach. We will no longer have a perfect hedge, and there will be some slippage but a good approximation to the optimal strategy will lead to an effective upper bound. In particular, the accumulated cash from exercising and rebuying may be negative. Additionally, for this method, we need the value of the hedge at each exercise time and this is evaluated by Monte Carlo. This introduces extra error which causes a further upwards bias, as the errors are inside the maximum, see [2].

We can apply this same technique to a callable bond directly. Thus we first develop an exercise strategy as in Section 3 to obtain the lower bound. We then run a second simulation to estimate the upper bound. We do this by simulating holding the callable bond as a hedge. As before, at points of exercise we rebuy the remaining product, hopefully making a profit but occasionally a loss as our strategy is not optimal.

We need to compute the maximum of the product minus the hedge along the path's exercise times. For this the product at an exercise time must be interpreted as the accumulated cash-flows invested in numeraire bonds plus the rebate obtained by exercise at that time. The hedge will similarly incorporate all cash-flows generated by the product, as well as any cash-flows received from exercise and repurchase. The hedging argument above goes over verbatim to this case and so we know that if the optimal strategy were used, the upper bound would equal the lower bound, except for Monte Carlo. We can expect a good bound for a good strategy.

Note that the seller's price must always be least as much as the buyer's price by the principle of no arbitrage: if not adopt a strategy giving a price higher than the seller's price; replicate the negative of the product with this strategy; hedge using the seller's price; and pocket the difference. We have shown that the hedging is perfect using the optimal strategy so buyer's and seller's prices agree.

We restricted ourselves to hedges of zero initial value. The hedge will therefore be the product minus the initial value of the product invested in numeraire bonds. To obtain the upper bound, we therefore have to add the estimated hedging error to the lower bound. Our estimation of the upper bound therefore has error coming both from the standard error in the lower bound and the standard error in the upper bound.

Along each path, the value will be the lower bound plus the maximum along that path. Up to the point where the lower strategy first exercises, (or the end if it never does,) the cash-flows from the product and hedge will agree, so at that point the net value be zero. This means that on every path the upper bound value will be at least as high as the lower bound, and therefore after averaging over many paths, it will still be at least as high.

Note that our method has made no assumptions about the product other than it can be exercised at most once. In particular, if we had a lower bound strategy, it would also apply to products where the cash-flows are fully path-dependent.

We summarize the upper bound algorithm for each path:

- (1) Set the maximum for the path to minus infinity.

- (2) Start at beginning.
- (3) Accumulate cash flows for the product and for the hedge up to next exercise time.
- (4) Compute the unexercise value of the hedge by Monte Carlo.
- (5) See whether the lower strategy says to exercise.
- (6) If strategy says exercise, value at this point is difference of accumulated cash-flows up to this point. Update total cashflows for hedge by difference of price of hedge and exercise value.
- (7) If strategy says do not exercise, value is exercise value plus accumulated cashflows for product minus accumulated cashflows for hedge and the unexercise value of the hedge.
- (8) Update maximum.
- (9) If end not reached return to 3.

Note we could reinterpret a breakable product as being a Bermudan option which pays dividends. Our approach can then be viewed as an extension of the Anderson–Broadie method which takes account of dividends.

5. NUMERICAL RESULTS

We consider the Asian tail bond introduced in Section 3. We price it in a jump-diffusion model, [9]. (See [5] for a discussion of the jump-diffusion model.)

Our parameters for the jump-diffusion model with log-normal jumps are

$$r = 5\%, \quad (\text{interest rate,}) \tag{5.1}$$

$$d = 2\%, \quad (\text{dividend rate,}) \tag{5.2}$$

$$\sigma = 0.1, \quad (\text{volatility,}) \tag{5.3}$$

$$\mu_J = 0.9, \quad (\text{jump mean,}) \tag{5.4}$$

$$\sigma_J = 0.1, \quad (\text{jump sigma.}) \tag{5.5}$$

We use a variety of values of jump intensity, λ .

For our regression functions, we use cubic polynomials in the stock price that is linear combinations of $1, S, S^2, S^3$. We use ten thousand paths to generate the exercise strategy.

In order to obtain meaningful probabilistic bounds, we use the Mersenne twister random number generator with a different seed for each of the regression paths, the lower bound price, the upper bound outer simulation and the generation of the hedge prices within the simulation. Faster convergence could be obtained by a number of means, including

Intensity	0	0.1	0.2	0.4	0.8	1.6	3.2
lower	0.9735	0.9826	0.9916	1.0083	1.0363	1.0772	1.1352
lower s.e.	0.0001	0.0001	0.0001	0.0002	0.0002	0.0003	0.0005
upper	0.0002	0.0004	0.0004	0.0005	0.0004	0.0006	0.0014
upper s.e.	3E-05	7E-05	6-05	8E-05	7E-05	0.0001	0.0001
upper paths	1887	1703	1660	1729	1649	1796	1653
lower time	9.844	10.453	11.656	11.156	12.031	10.953	11.093
total s.e.	0.0001	0.0001	0.0001	0.0002	0.0002	0.0003	0.0005
2 s.e. upper	0.9739	0.9833	0.9923	1.0092	1.0372	1.0784	1.1376

TABLE 1. Numerical results for a range of jump intensities.

the use of Sobol numbers, but we avoid these in order to retain the probabilistic interpretation of the standard error.

We use one million paths for the lower bound estimation. For the upper bound, we use 256 paths for the inner estimation of the price, and use a variable number of paths restricting to ten seconds, for the outer simulation. Timings taken on a 2GHz Pentium IV. We present results in table 1 for a variety of jump intensities. We present the unbiased estimate of the lower bound (lower), the standard error of the lower bound (lower s.e.), the upper bound minus the lower bound (upper), the standard error of this difference (upper s.e.), the number of paths for the upper bound (upper paths), the time taken for lower bound (lower time), the two standard errors together (total s.e.), and the upper bound with standard errors added.

We see that for low jump intensities the upper and lower bounds are very close with only a few basis points between them. For high intensities the value increases but it is still reasonably right. Note that the upper bounds have much smaller standard errors than the lower bounds, this is because the upper bound is a hedged simulation and so there is much less variation between paths.

We have presented results for the jump-diffusion case to emphasize that the techniques presented here make no assumptions on the process in the pricing measure beyond that we can price within it by Monte Carlo simulation. In particular, we do not need analytic prices nor approximations thereof for the product nor its pay-off. Note we have not presented comparisons with existing Monte Carlo methods for this product since no such methods exist.

6. CONCLUSION

We have presented a new approach to upper bounds for callable products with Monte Carlo simulations. This approach requires few

assumptions on product and process, allowing a much wider range of applicability than previous approaches. We have seen for one particular case that would formerly been hard that the bounds are tight.

REFERENCES

- [1] A. Amin, Multi-factor cross currency LIBOR market model: implementation, calibration and examples, preprint, available from <http://www.geocities.com/anant2999/>
- [2] L. Andersen, M. Broadie, A primal-dual simulation algorithm for pricing multi-dimensional American options, *Management Science*, 2004, Vol. 50, No. 9, pp. 1222-1234.
- [3] P. Glasserman, *Monte Carlo Methods in Financial Engineering*, Springer Verlag, 2003.
- [4] M. Haugh, L. Kogan, Pricing American Options: A Duality Approach, MIT Sloan Working Paper No. 4340-01
- [5] M.S. Joshi, *The concepts and practice mathematical finance*, Cambridge University Press 2003
- [6] I. Karatzas, E. Shreve, *Brownian Motion and Stochastic Calculus*, Second edition, Springer Verlag, 1997
- [7] I. Karatzas, E. Shreve, *Methods of Mathematical Finance*, Springer Verlag, 1998
- [8] F. Longstaff, E. Schwartz, Valuing American options by simulation: a least-squares approach. *Review of Financial Studies*, 14:113–147, 1998.
- [9] R. Merton, Option pricing when underlying stock returns are discontinuous, *J. Financial Economics* **3**, 125–144, 1976
- [10] V. Piterbarg, A practitioner's guide to pricing and hedging callable LIBOR exotics in forward LIBOR models, *Journal of Computational Finance*, Vol. 8, No. 2, 65–119, Winter 2004/05
- [11] L.C.G. Rogers: Monte Carlo valuation of American options, *Mathematical Finance*, *Mathematical Finance*, Vol. 12, pp. 271-286, 2002

CENTRE FOR ACTUARIAL STUDIES, DEPARTMENT OF ECONOMICS, UNIVERSITY OF MELBOURNE, VICTORIA 3010, AUSTRALIA

E-mail address: mark@markjoshi.com