Bermudan Option Pricing with Monte-Carlo Methods

Raphaël Douady*

January 30, 2002

Abstract

We explain, compare and improve two algorithms to compute American or Bermudan options by Monte-Carlo. The first one is based on threshold optimisation in the exercise strategy (Andersen 1999). The notion of "fuzzy threshold" is introduced to ease optimisation. The second one uses a linear regression to get an estimate of the option price at intermediary dates and determine the exercise strategy (Carriere 1997, Longstaff-Schwartz 1999). We thoroughly study the convergence of these two approaches, including a mixture of both.

1. Introduction

American and Bermudan option pricing in a Monte-Carlo framework is, in theory, impossible because it requires the value of the option at intermediary dates — in order to decide whether to exercise or to keep it — an information that is usually not provided. Actually, rather than the value of the option, one needs an *optimal exercise strategy,* that is, for each trajectory, an "optimal" date when to exercise the option.

In this situation, the word "optimal" means that it maximizes the price of the option — which is computed as usual by averaging the pay-out, including early exercise — among all "acceptable" strategies (probabilists use the word *adapted*).

* C.N.R.S. and Ecole Normale Supérieure, C.M.L.A., 61 av. du Pdt Wilson, 94235 Cachan Cedex, France. E-mail: rdouady@cmla.ens-cachan.fr

A strategy is acceptable if the decision to exercise or not only depends on the available information and not on the future of the trajectory. For instance, it is not acceptable to exercise the option at the maximum of its exercise value along the trajectory, because at a given date t, one doesn't know whether the trajectory is at its maximum. It is easy to check that pricing a standard American put according to this principle would lead to much too high a value.

Dynamic programming theory works backwards in time, starting from the maturity of the option. Let $T_1, ...T_n$ be the exercise dates (where T_n is the option maturity). One first determines the optimal exercise strategy at date T_{n-1}. This strategy should only depend on the state of the market at this date, therefore there exists a criterion $\Phi_{n-1}(x_1, ..., x_m)$, where the x_i are the market variables at date T_{n-1}, such that exercise should occur whenever $\Phi_{n-1} \geq 0$ and the option be kept until maturity in the other case. The zero-set[1] of Φ_{n-1} is called the *exercise boundary.* The optimal criterion is not unique because only its sign matters. Any other criterion that always has the same sign (hence zero-set) would fit. Once the optimal Φ_{n-1} is found, one determines the optimal exercise strategy at T_{n-2} by assuming that, is the case the option is kept, then exercise occurs at T_{n-1} according to the already computed criterion Φ_{n-1}, and so forth.

It is important to notice that, at any date T_k, the optimal exercise criterion Φ_k does not depend on whether the option can be exercised before T_k or not. Only exercise *afterwards* matters. A possible exercise criterion is simply the difference between the exercise value and the option price when ignoring immediate exercise possibility.

In this note, we first describe a general technique to compute optimal exercise criterions. This technique requires a multi-dimensional nonlinear optimizer. We then mention a simplified version of it, described in Andersen [1, 1999], which only needs one-dimensional optimization.[2] The notion of "fuzzy threshold" is introduced in sect. 3.2, in order to ease the optimization. The next section is devoted to an algorithm of Carriere (see [2, 1996]), enhanced by Longstaff and Schwartz (see [3, 1998]), based on an approximation of the option price. Finally, we develop our current approach, which uses Longstaff-Schwartz as a starting point, then applies the first technique to improve accuracy.

As, generally speaking, this type of algorithm is rather time consuming, we shall insist on the gradation of complexity with respect to accuracy, in order to let the user chose between a "quick and dirty" price and a more accurate one, but

[1] The set of all $(x_1, ..., x_m)$ such that $\Phi_{n-1}(x_1, ..., x_m) = 0$.

[2] This approach was first mentioned to me in an oral communication by B. Dupire in 1996.

longer to compute.

A general overview of the literature on the topic can be found in Dupire [4, 1998], where several seminal articles are fully reproduced.

2. General Approach to Optimal Exercise

2.1. Theory

Let $\mathbf{X}(t) = (X_1(t), ..., X_m(t))$ be the random process followed by market variables $x_1, ..., x_m$ and $r(t)$ be the short rate process (possibly random). The actualization factor $A(t, T)$ is defined by:

$$A(t, T) = \exp\left(-\int_t^T r(s)\, ds\right)$$

This actualization factor is the *actual discounting* on the period $[t, T]$, which is only known at the end date T, and should be distinguished from the discount factor $DF(t, T)$, which applies the corresponding rate *as known at date t*. In fact, in the risk-neutral probability, one has:

$$\mathbf{E}_t\left[A(t, T)\right] = DF(t, T)$$

The *exercise value* $V_k(x_1, ..., x_m)$ is the amount provided by the option if exercised at date T_k in the case $\mathbf{X}(T_k) = (x_1, ..., x_k)$. If $k = n$, this is the option pay-out. In this note, we assume that V_k is a known function of market variables $x_1, ..., x_n$. In particular, compound options, which, when exercised, lead to another option of a lesser "compounding degree" are excluded (see remark 3).[3]

If the option bears no other cash flows than when exercised, its value at date t is:

$$P(t) = \mathbf{E}_t\left[\sum_{T_k \geq t} A(t, T_k)\, V_k(\mathbf{X}(T_k))\, \mathbf{1}_{\{\tau = T_k\}}\right]$$

where τ is the optimal exercise date, characterized as the first date T_k at which the criterion Φ_k is nonnegative, or T_n if they are all negative (τ is of course trajectory

[3] An example of such options is a cap one can enter at times $T_1, ..., T_n$ with a strike equal to the current Libor, plus a fixed spread. A more complex option, that is, with a higher "compounding degree", is a cap with n settlements out of which only $m < n$ can be exercised, with discretionary choice.

dependent). In general, one should add the discounted value of extra cash flows, conditioned by the existence of the option at the corresponding date:

$$P(t) = \mathbf{E}_t \left[\sum_{T_k \geq t} A(t, T_k) \, \left(V_k(\mathbf{X}(T_k)) \, \mathbf{1}_{\{\tau = T_k\}} + F_k \, \mathbf{1}_{\{\tau > T_k\}} \right) \right]$$

where F_k is the sum of cash flows that occur at dates $T \in (T_k, T_{k+1}]$ discounted to T_k (i.e. multiplied by $A(T_k, T)$), assuming the option hasn't been exercised at T_k or before. These cash flows are assumed not to depend on market variables at all or, at most, on the knowledge of market variables at date T_k.

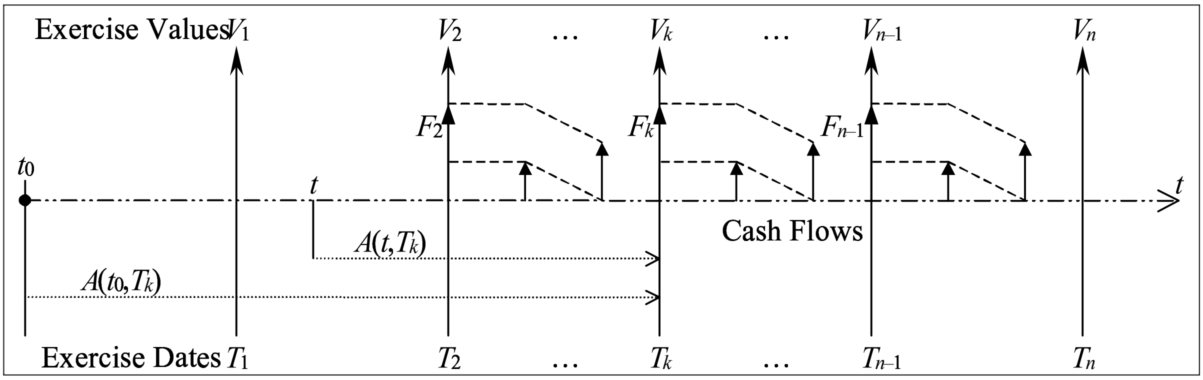

Fig1

Remark 1. *In practice, if they do depend on market variables at date $T > T_k$ but do not bear a strong convexity — such as non capped or non floored rates — then they can be replaced by their discounted expectation (or forward value) at T_k. The same remark applies when the actual discounting $A(T_k, T)$ depends on market variables posterior to T_k.*

In order to ease notations in an already complicated setting, intermediary cash flows F_k will be omitted in the sequel. We shall reintroduce them in the last section where the algorithm is described (see sect. 6).

In this context, the exercise date τ is optimal in the sense that it maximizes the price $P(t)$ among all acceptable exercise strategies, where *acceptable* means that exercise decision is taken under the only knowledge of available data at the current date.

Let the *holding value* $H_k(x_1, \dots, x_m)$ be defined by:

$$H_k(x_1, \dots, x_m) = \mathbf{E}_{T_k} \left[\sum_{\ell=k+1}^{n} A(T_k, T_\ell) \, V_\ell(\mathbf{X}(T_\ell)) \, \mathbf{1}_{\{\tau = T_\ell\}} \mid \mathbf{X}(T_k) = (x_1, \dots, x_m) \right]$$

that is, the value of the option at date T_k and market variables $\mathbf{X}(T_k) = (x_1, \ldots, x_m)$, when assuming no immediate exercise. One has:

$$P(T_k \,|\, \mathbf{X}(T_k) = (x_1, \ldots, x_m)) = \max(V_k(x_1, ..., x_m), H_k(x_1, ..., x_m))$$

and a possible criterion Φ_k is the difference $V_k - H_k$. The whole difficulty of the Monte-Carlo framework resides in the fact that H_k is not known (except when $k = n - 1$).

Remark 2. *Thanks to the Bayesian rule, if exercise is forbidden in the window* $[t_0, t]$ *then:*

$$P(t_0) = \mathbf{E}_{t_0}\left[A(t_0, t)\, P(t)\right]$$

therefore, $P(t_0)$ *is maximal whenever* $P(t)$ *is, and the optimal exercise strategy after date* t *does not depend on whether the option can be exercised before* t *or not.*

This remark justifies a backward induction on k to compute the exercise criterions Φ_k. Fix a date T_k and assume that exercise criterions $\Phi_{k+1}, ..., \Phi_{n-1}$ have already been computed. Then define the optimal exercise date $\tau \geq T_{k+1}$ as given by criterions Φ_ℓ, $k+1 \leq \ell \leq n-1$. According to the Bayesian rule, the criterion Φ_k will be a function of market variables $(X_1(T_k), ..., X_m(T_k))$ which maximizes $P(T_k)$, but also $P(t_0)$ when ignoring possible exercise at dates $T_1, ..., T_{k-1}$. This criterion determines when $\tau = T_k$ and when $\tau \geq T_{k+1}$. The backward induction process then may go on.

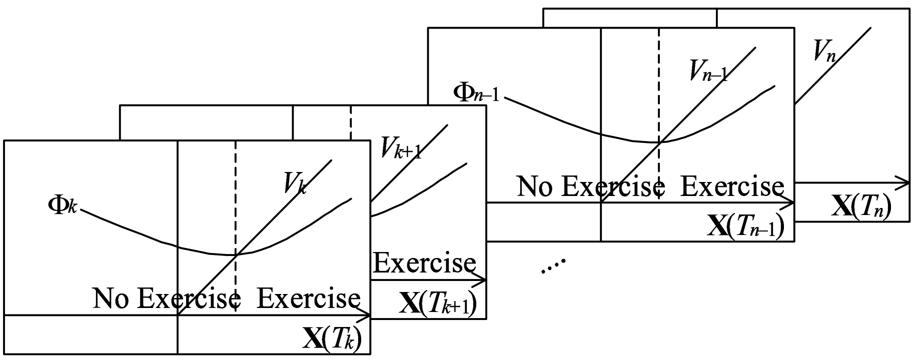

Fig2

Remark 3. *Compound options are handled by performing several times the backward induction. Every induction processing is used to provide the exercise value of the next processing. From the algorithmic point of view, induction* *processing of different compounding levels can be performed all at once. This simplifies the procedure but does not reduce the amount of computation.*

Remark 4. *American options are treated as Bermudan ones with narrowly sampled exercise dates.*

2.2. Monte-Carlo Sampling

In a Monte-Carlo framework, expectations are replaced by averages over the paths. Let N be the number of paths, which we denote by $\mathbf{x}_j(t) = (x_{1j}(t), ..., x_{mj}(t))$, $j = 1, ..., N$. Let also $A_j(t, t')$ be the actualization factor to apply over the period $[t, t']$ in path j (the short term interest rate r may differ from one path to another). For a given exercise strategy, that can be characterized by a sequence of criterions $\Phi = (\Phi_1, \dots, \Phi_n)$, one can deduce, for each path j, an exercise date $\tau_\Phi(j)$. This is simply the first date T_k at which $\Phi_k \geq 0$ on the path. The corresponding option value is:

$$P_\Phi(t_0) = \frac{1}{N} \sum_{j=1}^{N} A_j(t_0, \tau_\Phi(j))\, V_{\tau_\Phi(j)}(\mathbf{x}_j(\tau_\Phi(j)))$$

The temptation is to maximize $P_\Phi(t_0)$ over all possible sequences of criterions $(\Phi_1, ..., \Phi_n)$. Then the algorithm goes as follows:

1. For a given criterion Φ_{n-1} at date T_{n-1} , define $\tau_{n-1}(j)$ to be either T_{n-1} if $\Phi_{n-1}(\mathbf{x}_j(T_{n-1})) \geq 0$ or T_n otherwise. Then set:

$$\begin{aligned} P_{n-1,\Phi_{n-1}}(t_0) &= \frac{1}{N} \sum_{j=1}^{N} A_j(t_0, \tau_{n-1}(j))\, V_{\tau_{n-1}(j)}(\mathbf{x}_j(\tau_{n-1}(j))) \\ &= \frac{1}{N} \sum_{\Phi_{n-1} \geq 0} A_j(t_0, T_{n-1})\, V_{n-1}(\mathbf{x}_j(T_{n-1})) \\ &\quad + \frac{1}{N} \sum_{\Phi_{n-1} < 0} A_j(t_0, T_n)\, V_n(\mathbf{x}_j(T_n)) \end{aligned}$$

2. Parameterize Φ_{n-1} as a function of $\mathbf{x}_j(T_{n-1})$ and find the parameters that maximize $P_{n-1,\Phi_{n-1}}(t_0)$.

3. Once Φ_{n-1} is known, take the exercise strategy at T_{n-1} for granted and resume the procedure with Φ_{n-2} , etc. until date T_1 .

There are two major problems. The most obvious is complexity. If we deal with 10 market variables and we want Φ_k to be a 4^{th} degree polynomial in these variables, we speak of several thousands of coefficients to optimize, and the price $P_{n-1,\Phi_{n-1}}(t_0)$ depends in a highly non quadratic manner in these coefficients. Such an optimization problem, that must be solved at each possible exercise date, is out of scope.

The second one is more subtle. If criterions Φ_k are too accurate, then as we optimize the actual pay-out of each single trajectory, we are actually cheating by using information from the future, supposedly not known. Therefore there appears a kind of race between the degree of flexibility of criterions and the amount of Monte-Carlo paths. Intuitively, the number of paths must remain large with respect to that of parameters, inducing a lower bound to complexity. Both questions will be assessed in the sequel.

A third difficulty is the non smoothness of the function to be maximized. We shall also provide a solution to this problem (see sect. 3.2 on "fuzzy threshold").

3. The Exercise Value Threshold Method

3.1. Optimal Threshold

The following technique is described in Andersen [1, 1999]. In order to bound complexity, as well as avoid using future information in a forbidden way, criterions Φ_k only depend on the exercise value $V_k(\mathbf{x}_j(T_k))$. Hence, the exercise criterion is, at each date T_k , determined by a threshold θ_k : exercise occurs as soon as $V_k \geq \theta_k$. The value of the threshold is that which maximizes the function:

$$\psi(\theta) = P_{k,\theta}(t_0)$$

where $P_{k,\theta}(t_0)$ is the value of the option at t_0 when exercise at T_k is given by θ and exercises afterwards are given by the optimal $\theta_{k+1}, ..., \theta_{n-1}$.

There are many efficient algorithms the maximize a function of just one variables. We warn the reader that, in the present situation, this function ψ has bad features. First it is not smooth, because changing the threshold by a small amount can make a trajectory suddenly exercised at T_k , hence possibly change the pay-out of the precise trajectory by a big amount, albeit dumped by the $1/N$ factor. The second problem is asymptotic values when $\theta \to 0$ or $\theta \to +\infty$. In the first case, the option is always exercised, hence it is similar to an European option maturing at T_k . In the second case, the option is never exercised at T_k and its price is $P_{k+1,\theta_{k+1}}(t_0)$. These two values, and especially the second one, are often very close to the maximum. It is not totally absurd to simply scan the real axis (within reasonable bounds) and pick the maximum value.

The following graph shows the shape of function ψ in the case of a standard European put on a stock (one year, 40% volatility, 10% interest rate), with possible exercise at mid-maturity. One can observe the irregularities mentioned above. The left hand side asymptotic value is the price of a 6 months European option, the right hand side one is that of a 1 year European option.

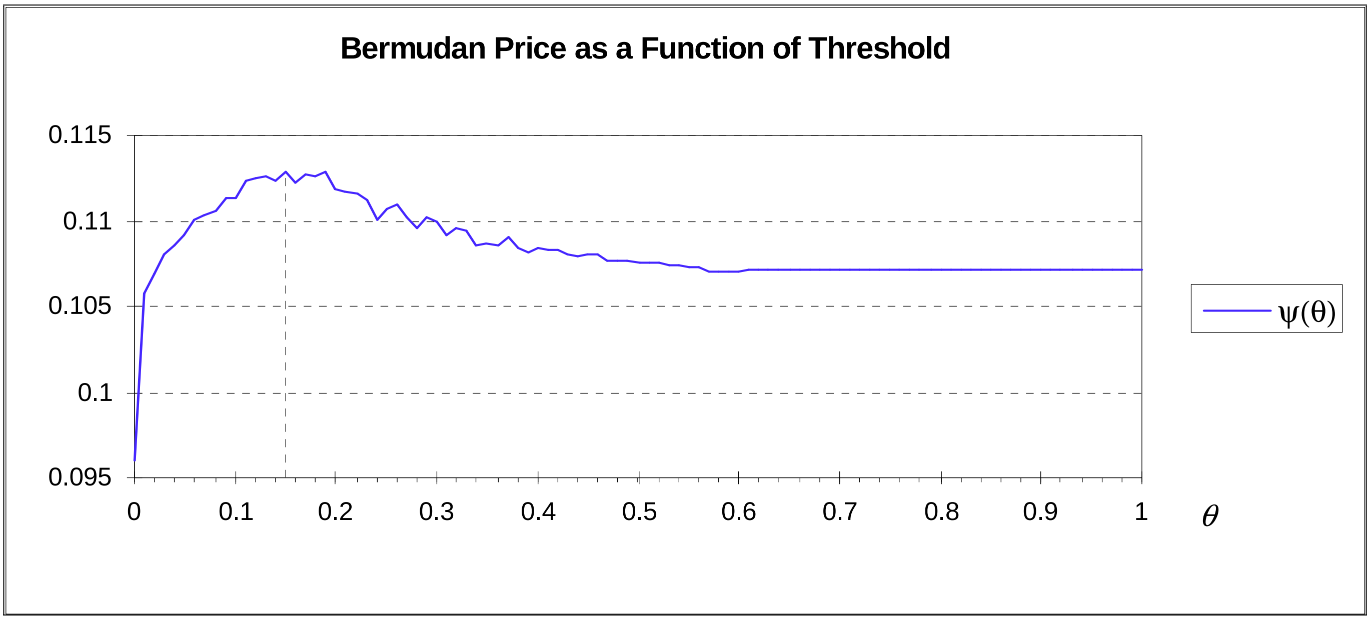

The limitation of this approach is that, when the number of paths tends to infinity, it converges to a lower value than the correct one. Indeed, the exercise strategy is sub-optimal, because it uses only partial information. When the underlying is only one-dimensional (as in the example above), then it does converge to the true price, as information is totally contained in the exercise value. However, for basket options, stochastic volatility models, interest rate options, etc. one can expect a negative bias in the Monte-Carlo price. If, mostly, one underlying price or rate matters, as it is the case with Bermudan swaptions, then it seems that the bias is not larger than the Monte-Carlo noise. The graph in sect. 5 shows a 10% downward bias in the early exercise premium of a "best of" option on a pair of uncorrelated stocks.

3.2. Fuzzy Threshold

Theoretically, there exists for each exercise date — and for each date in the case of an American option — an optimal exercise boundary. Beneath the boundary, the option should be kept and beyond, it should be exercised. By no mean should the decision involve any type of randomness. However, in the Monte-Carlo framework, the criterion ψ to be optimized is not smooth. The solution is to ignore the "all or nothing" rule and introduce "exercise probabilities". At a given exercise date T_k, each path $\mathbf{x}_j$ is exercised with probability p_{kj} and kept with probability $1 - p_{kj}$. Then the "portion of trajectory" that is not exercised eventually follows the exercise strategy defined by the backward induction.

In practice, assume that probabilities $p_{\ell j}$ have been determined for all $\ell > k$ and define:

$$\bar{p}_{\ell j} = (1 - p_{k+1j}) \dots (1 - p_{\ell-1j})\, p_{\ell j}$$

The "holding value" of trajectory $\mathbf{x}_j$ at date T_k is:

$$\hat{H}_{kj} = \sum_{\ell=k+1}^{n} \bar{p}_{\ell j}\, A_j(T_k, T_\ell)\, V_{\ell j}$$

where

$$V_{\ell j} = V_{T_\ell}(\mathbf{x}_j(T_\ell))$$

The quotes on "holding value" and the hat on $\hat{H}_{kj}$ come from the fact that the expectation has been omitted. The value $\hat{H}_{kj}$ summarizes the true future history of trajectory $\mathbf{x}_j$. In particular, the following induction formula holds:

$$\hat{H}_{k-1j} = A(T_{k-1}, T_k) \left(p_{kj}\, V_{kj} + (1 - p_{kj})\, \hat{H}_{kj} \right)$$

The "fuzzy threshold" technique consists of setting exercise probabilities as follows:

$$p_{kj} = \eta(V_{kj} - \theta)$$

where $\eta : \mathbb{R} \to [0, 1]$ is smooth, increasing and satisfies:

$$\lim_{s \to -\infty} \eta(s) = 0 \qquad \lim_{s \to +\infty} \eta(s) = 1$$

For example:

$$\eta_\alpha(s) = \frac{1}{1 + e^{-\alpha s}}$$

where the parameter α measures the "fuzziness" of the threshold. A non fuzzy threshold corresponds to $\alpha \to +\infty$.

The new "fuzzy criterion" ψ_α is defined by:

$$\psi_\alpha(\theta) = \frac{1}{N} \sum_{j=1}^{N} A(t_0, T_k) \left(p_{kj}(\theta)\, V_{kj} + (1 - p_{kj}(\theta))\, \hat{H}_{kj} \right)$$

where

$$p_{kj}(\theta) = \eta_\alpha(V_{kj} - \theta)$$

This is a smooth function of the "threshold" θ because each single probability $p_{kj}(\theta)$ depends smoothly on θ. It tends to the non fuzzy — and non smooth — criterion ψ when $\alpha \to +\infty$.

The following graph shows functions ψ and ψ_α in the same setting as previously with different values of α.

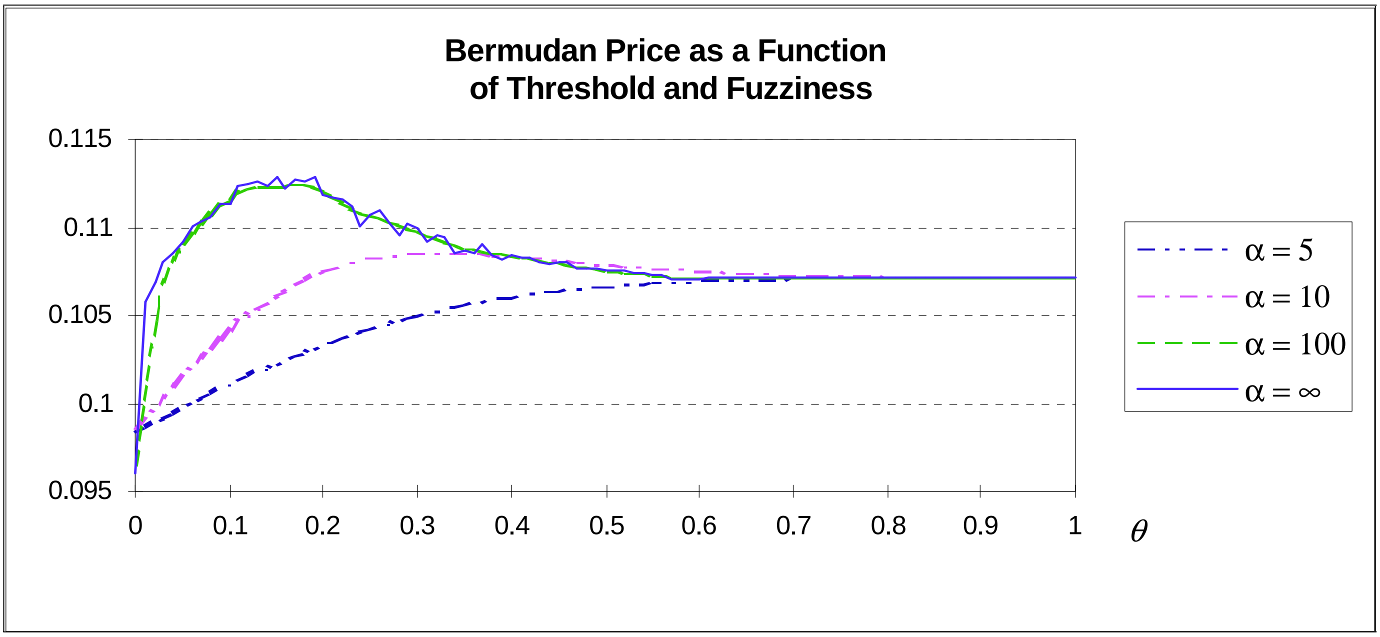

One sees that, if α is too small, then either ψ_α does not have a maximum, or it is far away from the correct value. However, an appropriate parameter α smoothes the criterion without incuring bias in the result.

4. The Carriere-Longstaff-Schwartz Approach

4.1. Linear Regression

J.F. Carriere's approach [2, 1996], improved by F. Longstaff and E. Schwartz [3, 1998], goes back to one of the definitions of criterions Φ_k as the difference between the exercise value V_k and the holding value, that is, the option "net present value" if not exercised, denoted H_k :

$$\Phi_k(\mathbf{X}(T_k)) = V_k(\mathbf{X}(T_k)) - H_k(\mathbf{X}(T_k))$$

As mentioned earlier, the problem is that H_k is not known. Nevertheless, like any option price, it should depend smoothly on variables $X_1(T_k), ..., X_m(T_k)$, thus it can be approximated by a polynomial of these variables, or any type of expansion in "basic functions", which we shall call $f_1, ..., f_q$:

$$H_k(\mathbf{X}(T_k)) \simeq \lambda_0 + \sum_{i=1}^{q} \lambda_i \, f_i(\mathbf{X}(T_k))$$

The number q of basic functions can potentially be much larger than the number m of variables (if these are all polynomials of degree at most d, then it is of the order of $m^d/d!$). As no exercise can occur between T_k and T_{k+1} (in the American case, we assume T_{k+1} to be sufficiently close for this approximation), one has:

$$H_k(\mathbf{X}(T_k)) = \mathbf{E}_{T_k}\left[\sum_{\ell=k+1}^{n} A(T_k, T_\ell)\, V_\ell(\mathbf{X}(T_\ell))\, \mathbf{1}_{\{\tau = T_\ell\}}\right]$$

Therefore, the difference:

$$\mathbf{E}_{T_k}\left[\sum_{\ell=k+1}^{n} A(T_k, T_\ell)\, V_\ell(\mathbf{X}(T_\ell))\, \mathbf{1}_{\{\tau = T_\ell\}}\right] - \sum_{i=1}^{q} \lambda_i \, f_i(\mathbf{X}(T_k))$$

should be made as small as possible.

In the Monte-Carlo setting, where expectations are replaced by averages, this will be achieved by a linear regression. As in the previous section, we define $\tau(j) = \tau_{k+1}(j)$ to be the exercise time of the option, as known at T_{k+1}. The actual discounted pay-off of each trajectory j is:

$$\hat{H}_k(j) = A(T_k, \tau(j))\, V_{\tau(j)}(\mathbf{x}_j(\tau(j)))$$

The difference with its expectation at T_k should be independent of variables $X_1(T_k), ..., X_m(T_k)$, hence of any function of these, in particular basic ones $f_1, ..., f_q$. One performs thus a linear regression of the array $\hat{H}_k(j)$ with respect to arrays $f_i(\mathbf{x}_j(T_k))$, $i = 1, ..., q$:

$$\hat{H}_k(j) = \lambda_0 + \sum_{i=1}^{q} \lambda_i f_i(\mathbf{x}_j(T_k)) \; + Z(j)$$

where the remainder $Z(j)$ is decorrelated from $f_i(\mathbf{x}_j(T_k))$, $i = 1, ..., q$. Then set:

$$\Phi_k(\mathbf{x}) = V_k(\mathbf{x}) - \lambda_0 - \sum_{i=1}^{q} \lambda_i f_i(\mathbf{x})$$

A possible improvement of this method is, after the first run, to perform a second linear regression that is weighted to augment the importance of trajectories close to the exercise boundary that has been found by the first regression.

The convergence behaviour of the algorithm is as follows. For a given set of basic functions, when the number of paths tends to infinity, the price converges to a sub-optimal price, because the exercise strategy is forced to depend linearly on basic functions of rates (albeit less sub-optimal than in the exercise value threshold method). On the other hand, for a "low" number of paths, we face the second problem mentioned in sect. 2.2 and the option is, on the contrary, overpriced. As the number of path increases, one can see the price going from much too high to slightly too low, the words "much" and "slightly" being strengthen when the number of basic functions increases. Eventually, when the number of basic functions itself tends to infinity, but the number of Monte-Carlo paths tends even faster to infinity (probably at least like the square of the number of basic functions), then the limit is equal to the true Bermudan premium. This is shown in the following picture, which graphs the early exercise premium of an American option on two stocks with pay-off $\max(K - \max(S_1, S_2), 0)$.

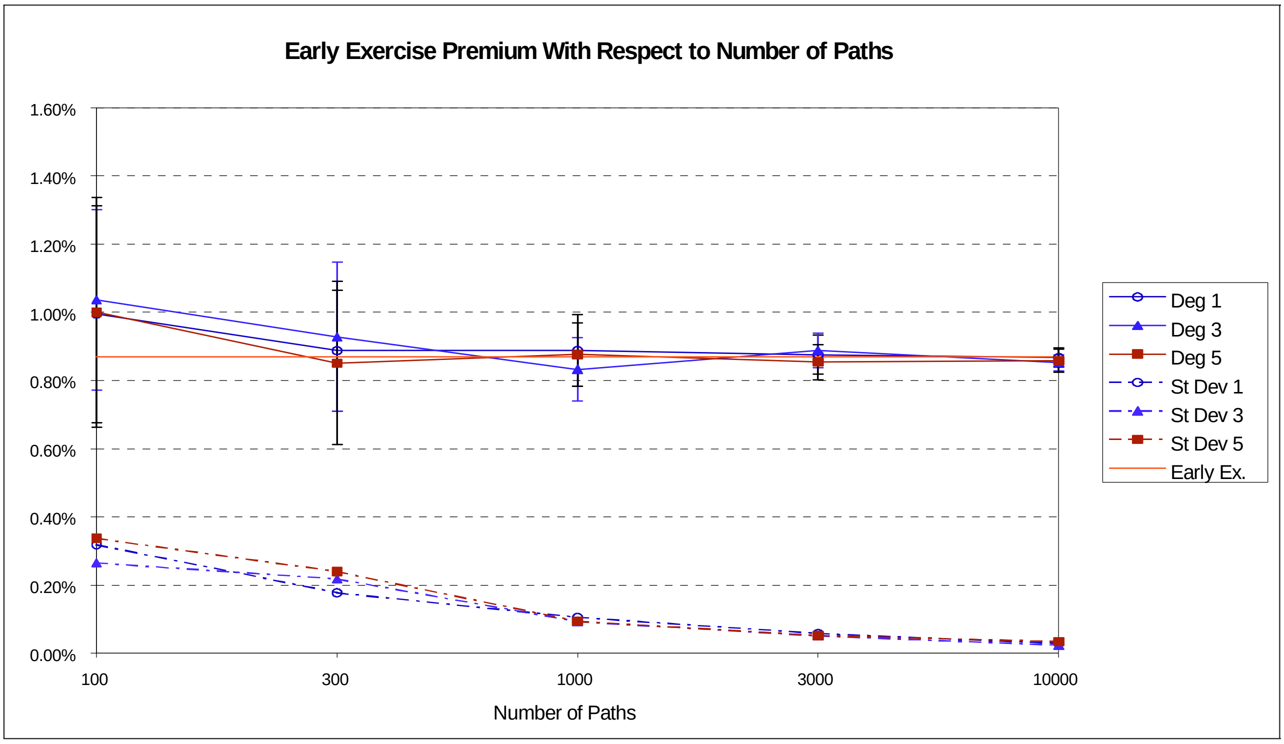

Early exercise premium in percentage of notional (assumed equal for the two stocks). $\mathrm{Vol}(S_1) = 20\%$, $\mathrm{Vol}(S_2) = 25\%$, $\rho = 0$, $r = 5\%$, 10 exercise dates. The "American" price is an estimation, because no benchmark is available.

Remark 5. *Let $\tilde{P}_k(j) = \lambda_0 + \sum_{i=1}^{q} \lambda_i f_i(\mathbf{x}_j(T_k))$. This is what the model assumes to be the option price at date T_k for trajectory $\mathbf{x}_j$. One could be tempted to replace the regression of $\hat{H}_k(j)$ by that of $A(T_k, T_{k+1})\tilde{P}_{k+1}(j)$, which stands for the discounted pay-off if the option is sold back on the market at T_{k+1}. This change (which looks like a simplification but, in fact is not) introduces important biases and should be avoided. What happens is that $\tilde{P}_k(j)$ is close to the option price when $\mathbf{x}_j(T_k)$ is close to its expectation, but it may be in average far away when it is close to the exercise boundary, where the comparison with V_k matters.*

4.2. Choice of Basic Functions

It is often a good idea to include, when possible, V_k as one of the basic functions. If this is the only one, then we are in the threshold situation, although the linear regression does not necessarily provide the optimal threshold. Generally speaking, let $f_1 = V_k$ and assume that $\lambda_1 \neq 1$. Then Φ_k can be replaced by:

$$\bar{\Phi}_k(\mathbf{x}) = V_k(\mathbf{x}) - \bar{\lambda}_0 - \sum_{i=2}^{q} \bar{\lambda}_i f_i(\mathbf{x})$$

where $\bar{\lambda}_i = \lambda_i \,/\, (1 - \lambda_1)$, $i = 0, 2, ..., q$. In general, the option value if not exercised makes a narrow angle with the exercise value along the exercise boundary (in the limit case of American options, theorems by Shiryaev and by McKean, known as "pasting conditions", state that the two surfaces are tangent along the exercise boundary). This usually implies a value of λ_1 close to 1. The criterion Φ_k is almost degenerate near 0, but $\bar{\Phi}_k$ isn't because of the large coefficient $1 \,/\, (1 - \lambda_1)$.

The choice of other basic functions is, in practice, less sensitive a matter than one could expect. Simple monomials such as $x_1^a x_2^b ...$ provide a very good set of basic functions, on top of the constant λ_0 and the exercise value V_k. A higher degree provides greater flexibility for, obviously, a higher complexity. In order to master the number of basic functions, while allowing flexibility, the following rule can be applied: select, for each index i a maximum degree for the variable x_i, say d_i. A monomial $x_1^a x_2^b ...$ is selected as a basic function if its total degree $a + b + ...$ is not greater than each degree d_i for which the variable x_i appears in the monomial (asking only that the exponent of x_i be not greater than d_i would lead to too many basic functions). For example, if $d_1 = 1$ and $d_2 = 2$ then monomials x_1, x_2 and x_2^2 are accepted but not $x_1 x_2$, which has total degree $2 > d_1$.

5. Combination of the two methods

The Longstaff-Schwartz technique is inaccurate for a low number of paths but, for a high number, it is slightly biased downwards because the exercise strategy is not optimal. It is possible to combine the two above techniques to shrink down this bias. Select some of the coefficients λ_i, $i \in J$, where J is a subset of $\{0, ..., q\}$ and, at each step, modify the values provided by Longstaff-Schwartz by amounts ε_i, $i \in J$ chosen to maximize the discounted expectation at t_0. More precisely, consider one of the exercise date T_k and let $\boldsymbol{\varepsilon} = (\varepsilon_i)_{i \in J}$. Define $\psi_k(\boldsymbol{\varepsilon})$ as the option value at t_0 if exercise at $T_1, ..., T_{k-1}$ is not allowed, exercise at $T_{k+1}, ..., T_n$ is given by formerly determined values of λ_i and exercise at T_k is given by criterion $\Phi_{k,\varepsilon}$ where λ_i is replaced by $\lambda_i + \varepsilon_i$, whenever $i \in J$. Standard optimization techniques provide the optimal vector $\boldsymbol{\varepsilon}_k = \arg\max \psi_k$ or, equivalently, the optimal set of coefficients $(\lambda_0, ..., \lambda_n)$, in the sense that only indices in J are optimized.

Remark 6. *The "fuzzy threshold" technique can be applied in this context to obtain a smooth function ψ_k which, in general, is not.*

The next graph shows the valuation of an option on two stocks S_1 and S_2 with payoff $\max(K_1 - S_1, K_2 - S_2, 0)$ and 10 possible exercise dates $T_1, ..., T_{10}$. In fact, only the early exercise premium (i.e. the difference between the European and the Bermudan option) is displayed. Three series of valuation are shown. The first one is the Longstaff-Schwartz method as explained in the previous section, referred to as LS. The second, referred to as Optim, optimizes all coefficients λ_i at every step. In the third one, referred to as Paral for "parallel shift", only λ_0 is optimized, but $\lambda_1, ..., \lambda_n$ keep the values provided by the linear regression. Needless to say that the third technique is much faster and robust than the second one — one-dimensional optimization can be made extremely efficient, especially when the shape of the function is known — although it seems to provide very similar results. Basic functions are $S_1^a S_2^b$, $a + b \leq d$ where the degree d is an input. With the above notations, this would correspond to $d_1 = d_2 = d$. For each series of simulation, the standard deviation of the result is displayed (obtained by repeating the simulation several times).

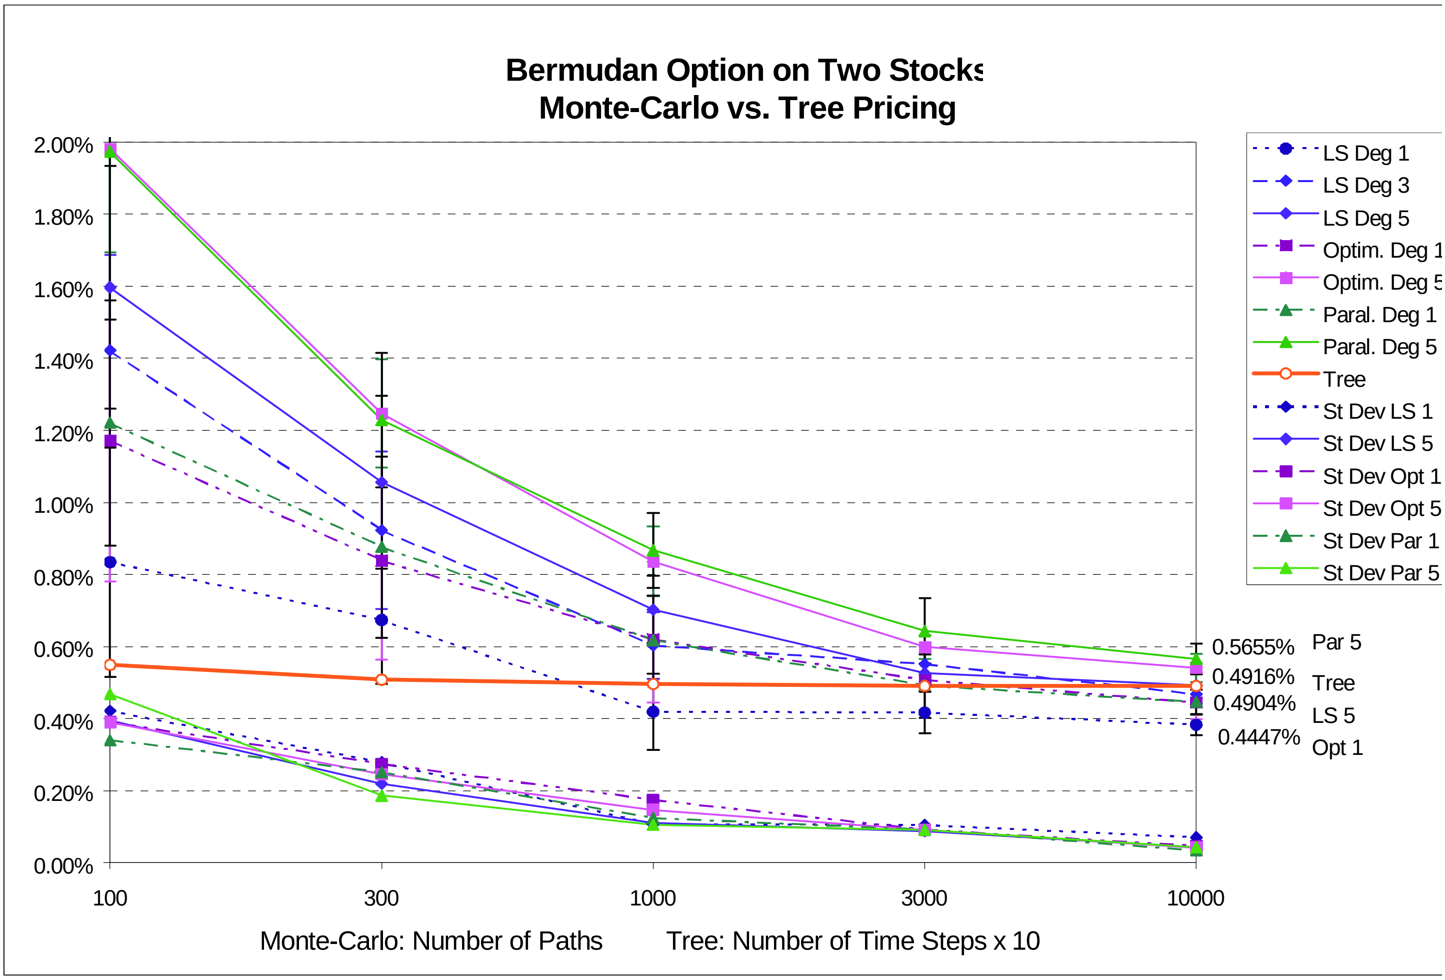

One sees that, for each degree d, there is a critical number of paths $N(d)$ for which the estimation is unbiased. For a given option that we may have to compute several times, but with not a high accuracy, one can first estimate $N(d)$ by a thorough simulation and, for future evaluation, only use basic functions up to degree d and simulate only $N(d)$ paths.

6. Practical implementation

The algorithm explicited in this section combines both "regression" and "optimal threshold" approach, and uses the "fuzzy threshold" technique.

1. Let $t_1, ..., t_N$ be the diffusion sample dates and $T_1, ..., T_n$ be the exercise dates. At each t_i some cash flow may occur, with an amount depending on the sample path. We first add, for each *exercise date* T_k all cash flows occurring at t_i such that $T_k \leq t_i < T_{k+1}$, discounted from t_i to T_k and we call $F(j,k)$ the array of such cash flows, where j is the path index.

2. Load the array $A(k,j)$ of the actual discounting factors from t_0 to T_k in path j.

3. For each T_k and each path j, we compute $V(j,k)$, the exercise value if the option is exercised at date T_k in the path j.

4. Start from the one but last date T_{n-1} and load arrays of basic rates $R_i(j)$ and compute basic functions:

$$f_q(j) = f_q(R_1(j), ..., R_\ell(j)) \ , \quad q = 1, ..., m$$

   up to the desired degree. Set when possible:

$$f_0(j) = V(j, n-1)$$

5. Regress the "holding value":

$$H(j, n-1) = F(j, n-1) + V(j,n)\, A(j,n) \,/\, A(j, n-1)$$

   over arrays $f_q(j)$, $q = 0, ..., m$ (including the exercise value, denoted f_0).

6. Let $\lambda_0, ..., \lambda_m$ be the regression coefficients and μ be the constant and define the *conditional expectation:*

$$E(j, n-1) = \mu + \sum \lambda_q f_q(j, n-1)$$

7. Define the *exercise probability:*

$$p(j, n-1) = \eta_\alpha(V(j, n-1) - E(j, n-1))$$

   where α is the fuzziness parameter of the threshold.

8. Define the *Bermudan value:*

$$P(j, n-1) = p(j, n-1)\, V(j, n-1) + (1 - p(j, n-1))\, H(j, n-1)$$

9. Optimize over desired $\mu, \lambda_1, ..., \lambda_h$ the price at t_0 :

$$P_0(n-1) = \frac{1}{N} \sum A(j, n-1)\, P(j, n-1)$$

10. Repeat procedure for $k = n-2, ..., 0$.

Remark 7. *For compound options, the exercise value $V(j,k)$ should be replaced by the price $P'(j,k)$ of the option with one level less of compounding (i.e. the option one gets when exercising the compound option).*

References

[1] L. Andersen *A Simple Approach to the Pricing of Bermudan Swaptions in the Multi-factor Libor Market Model* Jour. of Comput. Finance, Winter 1999, **3(2)**, p. 5-32.

[2] J.F. Carriere *Valuation of the Early-Exercise Price for Derivative Securities using Simulations and Splines* Insurance: Math. and Econ., 1996, **19**, p. 19-30.

[3] F. Longstaff, E. Schwartz *Valuing American Options by Simulation: A Simple Least-Square Approach* working paper, 1998, Anderson School at UCLA, forthcoming in Review of Fin. Studies.

[4] B. Dupire *Monte-Carlo: Methodologies and Applications for Pricing and Risk Management* Risk Books, London, 1998.

Other articles on Monte-Carlo pricing of American options:

[5] J. Barraquand, D. Martineau *Numerical Valuation of High Dimensional Multi-variate American Securities* J. of Fin. and Quant. Analysis, 1995, **30**, p. 383-405 (also in [4, 1998]).

[6] M. Broadie, P. Glasserman *Pricing American Style Securities using Simulations* J. of Econ. Dyn. and Control, 1997, **21**, p. 1323-52.

[7] M. Broadie, P. Glasserman *A Stochastic Mesh Method for Pricing High-dimensional American Options* working paper, 1997, Columbia University.

[8] D. Grant, G. Vora, D. Weeks *Simulations and the Early-exercise Option Problem* J. of Fin. Engin., 1996, **5(3)**, p. 211-27

[9] S. Raymar, M. Zwecher *A Monte-Carlo Valuation of American Call Options on the Maximum of Several Stocks* J. of Derivatives, 1997, **5(1)**, p. 7-23.

[10] J.A. Tilley *Valuing American Options in a Path Simulation Model* Trans. of the Soc. of Actuaries, 1993, **45**, p. 83-104 (also in [4, 1998]).