

EQF 013/024: Exercise boundary optimization methods

21st July 2008

1 Introduction

One of the most important characteristics of the Monte Carlo simulation method is its intrinsically forward-looking nature. Whilst this feature enables us to take into account inherently path-dependent payoff structures of any derivative contract with comparative ease, it makes it difficult to accommodate the inclusion of any early exercise rights into the contract. The two most commonly used methods to handle products with early exercise opportunities within a Monte Carlo simulation framework are *regression-based* techniques (which are covered in EQF 013 / 006) and *exercise boundary optimization* approaches which are discussed in this chapter.

2 Optimal stopping time as exercise boundary optimization

For the sake of generality, we assume that we are dealing with the valuation of a financial product Π_0 with embedded exercise optionality and discrete cashflows that span m time horizons $t_1 < \dots < t_m$, with the current time being $t \equiv t_0$. The product is considered to have two underlying contracts A and B of contingent cashflows which both, individually, contain no exercise optionality. The product Π_0 , initially, pays the same cashflows as product A , but, in addition, permits the exercise option holder to switch into B at one of the time horizons t_j . This formulation is fairly generic and encompasses practically all callable structures, including what is sometimes referred to as “options on options” or “higher order options” since those can be reformulated as a sequence of payable cashflows to continue until a final contingent payoff is attained, or to opt out at any of the intermediate exercise times.

Assuming a finite-dimensional Markovian representation of our usual probability space given a state vector \mathbf{x} , the contingency of the cashflows a_j in product A means that $a_j = a_j(\mathbf{x}(t_j))$, and likewise for product B . Risk-neutral valuation of the exercisable financial product Π_0 requires that we identify the optimal exercise strategy represented by the *optimal stopping time*. Valuation of Π_0 in the filtration \mathcal{F}_{t_0} , denoted by $\mathcal{V}_{t_0}[\Pi_0]$, is given by

$$\mathcal{V}_{t_0}[\Pi_0] = N(\mathbf{x}(t_0)) \cdot \sup_{\tau} \mathbb{E}_{t_0}^{\mathcal{M}(N)} \left[\sum_{j=1}^{\tau-1} \frac{a_j(\mathbf{x}(t_j))}{N(\mathbf{x}(t_j))} + \sum_{j=\tau}^m \frac{b_j(\mathbf{x}(t_j))}{N(\mathbf{x}(t_j))} \right] \quad (2.1)$$

wherein the discrete random variable $\tau \in 1, \dots, m$ is the stopping time index, $N(\cdot)$ is the chosen numéraire, and $\mathbb{E}_{t_0}^{\mathcal{M}(N)}[f]$ is the expectation of any given f under the measure induced by the numéraire in the filtration \mathcal{F}_{t_0} . Define Π_k in complete analogy to the financial product Π_0 except that it can only be exercised on or after t_k . A crucial observation in the following is that, if we had knowledge of the optimal stopping time process

$$\tau_k^* := \arg \sup_{\tau | \tau \geq k} \mathbb{E}_{t_0}^{\mathcal{M}(N)} \left[\sum_{j=1}^{\tau-1} \frac{a_j(\mathbf{x}(t_j))}{N(\mathbf{x}(t_j))} + \sum_{j=\tau}^m \frac{b_j(\mathbf{x}(t_j))}{N(\mathbf{x}(t_j))} \right], \quad (2.2)$$

we could value Π_k by means of the simple expectation

$$\mathcal{V}_{t_0}[\Pi_k] = N(\mathbf{x}(t_0)) \cdot \mathbb{E}_{t_0}^{\mathcal{M}(N)} \left[\sum_{j=1}^{k-1} \frac{a_j(\mathbf{x}(t_j))}{N(\mathbf{x}(t_j))} + \sum_{j=k}^{\tau_k^*-1} \frac{a_j(\mathbf{x}(t_j))}{N(\mathbf{x}(t_j))} + \sum_{j=\tau_k^*}^m \frac{b_j(\mathbf{x}(t_j))}{N(\mathbf{x}(t_j))} \right]. \quad (2.3)$$

*OTC Analytics

†Bank of America Securities

By virtue of the assumed Markovian representation in the state vector $\mathbf{x}(t)$, it is possible to rephrase the valuation problem based on an indicator function $\mathcal{I}_k(\mathbf{x}(t_k))$ which is one when exercise is optimal and zero otherwise. This gives us the recursive formulation

$$\mathcal{V}_{t_0}[\Pi_k] = N(\mathbf{x}(t_0)) \cdot \mathbb{E}_{t_0}^{\mathcal{M}(N)} \left[\mathcal{I}_k(\mathbf{x}(t_k)) \cdot \left(\sum_{j=1}^{k-1} \frac{a_j(\mathbf{x}(t_j))}{N(\mathbf{x}(t_j))} + \sum_{j=k}^m \frac{b_j(\mathbf{x}(t_j))}{N(\mathbf{x}(t_j))} \right) + \tilde{\mathcal{I}}_k(\mathbf{x}(t_k)) \cdot \frac{\mathcal{V}_{t_k}[\Pi_{k+1}]}{N(\mathbf{x}(t_k))} \right] \quad (2.4)$$

wherein $\tilde{\mathcal{I}}_k = 1 - \mathcal{I}_k$. If we define the product $\hat{\Pi}_k$ as Π_k minus all cashflows occurring before t_k such that the first cashflow of $\hat{\Pi}_k$ is b_k or a_k at t_k , depending on whether exercise was invoked at t_k or not, we have

$$\mathcal{V}_{t_0}[\hat{\Pi}_k] = N(\mathbf{x}(t_0)) \cdot \mathbb{E}_{t_0}^{\mathcal{M}(N)} \left[\tilde{\mathcal{I}}_k(\mathbf{x}(t_k)) \cdot \frac{a_k(\mathbf{x}(t_k)) + \mathcal{V}_{t_k}[\hat{\Pi}_{k+1}]}{N(\mathbf{x}(t_k))} + \mathcal{I}_k(\mathbf{x}(t_k)) \cdot \sum_{j=k}^m \frac{b_j(\mathbf{x}(t_j))}{N(\mathbf{x}(t_j))} \right]. \quad (2.5)$$

In this form, it seems that we need to know the conditional value $\mathcal{V}_{t_k}[\hat{\Pi}_{k+1}]$ in order to value Π_0 (which is by construction equal to $\hat{\Pi}_0$ since no cashflows occur before t_1). However, since this conditional expectation appears *inside an expectation over its conditioning filtration*, by virtue of the *tower law*, we can replace it by the sequence of (numéraire-deflated) contingent cashflow values:

$$\mathcal{V}_{t_0}[\hat{\Pi}_k] = N(\mathbf{x}(t_0)) \cdot \mathbb{E}_{t_0}^{\mathcal{M}(N)} \left[\tilde{\mathcal{I}}_k(\mathbf{x}(t_k)) \cdot \left(\sum_{j=k}^{\tau_{k+1}^* - 1} \frac{a_j(\mathbf{x}(t_j))}{N(\mathbf{x}(t_j))} + \sum_{j=\tau_{k+1}^*}^m \frac{b_j(\mathbf{x}(t_j))}{N(\mathbf{x}(t_j))} \right) + \mathcal{I}_k(\mathbf{x}(t_k)) \cdot \sum_{j=k}^m \frac{b_j(\mathbf{x}(t_j))}{N(\mathbf{x}(t_j))} \right]. \quad (2.6)$$

In this form, we can view

$$c_k = \sum_{j=k}^{\tau_{k+1}^* - 1} \frac{a_j(\mathbf{x}(t_j))}{N(\mathbf{x}(t_j))} + \sum_{j=\tau_{k+1}^*}^m \frac{b_j(\mathbf{x}(t_j))}{N(\mathbf{x}(t_j))} \quad (2.7)$$

as the numéraire-deflated *continuation value*, and

$$q_k = \sum_{j=k}^m \frac{b_j(\mathbf{x}(t_j))}{N(\mathbf{x}(t_j))} \quad (2.8)$$

as the numéraire-deflated *cancellation value*, to arrive at

$$\mathcal{V}_{t_0}[\hat{\Pi}_k] = N(\mathbf{x}(t_0)) \cdot \mathbb{E}_{t_0}^{\mathcal{M}(N)} \left[\tilde{\mathcal{I}}_k(\mathbf{x}(t_k)) \cdot c_k + \mathcal{I}_k(\mathbf{x}(t_k)) \cdot q_k \right]. \quad (2.9)$$

The only trouble now is that we do not know, a priori, the optimal exercise indicator function $\mathcal{I}_k(\cdot)$. We can confidently say, though, that given any trial exercise decision function $\mathcal{E}_k(\mathbf{x}; \boldsymbol{\lambda}_k)$ (which we choose to indicate exercise if its value is positive), with parameter vector $\boldsymbol{\lambda}_k$, we have

$$\mathcal{V}_{t_0}[\hat{\Pi}_k] \geq N(\mathbf{x}(t_0)) \cdot \max_{\boldsymbol{\lambda}} \mathbb{E}_{t_0}^{\mathcal{M}(N)} \left[\mathbf{1}_{\{\mathcal{E}_k(\mathbf{x}(t_k); \boldsymbol{\lambda}_k) \leq 0\}} \cdot c_k + \mathbf{1}_{\{\mathcal{E}_k(\mathbf{x}(t_k); \boldsymbol{\lambda}_k) > 0\}} \cdot q_k \right]. \quad (2.10)$$

The key idea of *exercise boundary optimization methods* for the valuation of exercisable financial products in a Monte Carlo simulation framework is to use a suitably chosen exercise decision function $\mathcal{E}_k(\mathbf{x}; \boldsymbol{\lambda}_k)$, and to find the value for $\boldsymbol{\lambda}_k$ that maximizes the Monte Carlo estimator objective function

$$F_k := \frac{1}{n} \sum_{i=1}^n \left(\mathbf{1}_{\{\mathcal{E}_k(\mathbf{x}(t_k); \boldsymbol{\lambda}_k) \leq 0\}} \cdot c_{ik} + \mathbf{1}_{\{\mathcal{E}_k(\mathbf{x}(t_k); \boldsymbol{\lambda}_k) > 0\}} \cdot q_{ik} \right), \quad (2.11)$$

with c_{ik} and q_{ik} representing the continuation and cancellation values of the i -th path at time horizon t_k for a set of n simulated discrete evolutions of the state vector in the chosen measure. Incidentally, it becomes apparent in this formulation that the exercise boundary optimization method allows readily for the switch product B to be, in principle, a different one at each exercise time horizon since the switch payments enter only in the form of the term q_{ik} as the sum of all numéraire-deflated cashflows that are payable if exercise is made at time t_k . Once the maximizing value $\boldsymbol{\lambda}_k$ has been established, the procedure continues at t_{k-1} , in analogy to a backwards induction method. An important observation in this context is that throughout the entire backwards iteration the same set of simulated paths and associated values can be used. This is a key point for the efficiency of this method. An initial simulation only needs to store the required state variables $\mathbf{x}(t_k)$ and associated continuation and cancellation values for each exercise time horizon. All subsequent optimization calculations can then be done over this precomputed *training set*. For typical training set sizes in the range of 8191–65535 paths, the evaluation of the objective function (2.11) is thus very fast indeed.

3 Choices of exercise boundary specification

The choice of exercise decision function $\mathcal{E}(x; \lambda)$ is crucial for the performance of exercise boundary optimization methods. This applies to both the actual choice of the functional form and with it the number of free parameters, as well as the effective financial variables that are considered the primary arguments of $\mathcal{E}(\cdot)$.

3.1 Assessment by related financial contracts

An intuitive approach to the choice of a reduced set of state variables is to monitor related financial contracts. Andersen [And00] and Piterbarg [Pit03] describe how swap rate levels and European swaptions values, even if only attainable in an approximate analytical fashion in any one given model, can be used to capture most of the callability value for Bermudan swaptions. They also show that for this family of financial contracts, a one-parameter choice for $\mathcal{E}_k(x; \lambda)$ is often sufficient. This particularly holds when the used model is itself driven by a single Brownian motion, even if there is no one-dimensional Markovian representation as is the case for Libor market models. In this case, the exercise decision function can be as simple as

$$SR(x) - \lambda,$$

with $SR(x)$ denoting the coterminal swap rate, for a payer's Bermudan swaption. More complicated decision rules for Bermudan swaptions can be found in [And00].

3.2 Financial coordinate transformations

It is not always intuitive and easy to find a related financial contract that can be used as an exercise indicator and whose value is attainable (semi-)analytically. Jäckel [Jäc02] discusses the exercise boundary optimization method in more general terms and suggests, when necessary, the use of tree methods and non-linear transformations of x for the assessment of the suitability of any particular functional form and choice of variables prior to its use in an exercise boundary optimization context. The useful observation is made that a functional form that works well for a given financial contract's exercise domain delineation with one model, even if the contract is simplified to a significant extent, practically always also works very well with other models for the fully fledged product version. This makes it possible, for instance, to visualize exercise domains computed with a non-recombining tree implementation of a two-factor Libor market model for a short contract life such as 6-non-call-2, and to apply the same functional form with a fully factorized model for very long-dated contracts.

3.3 A useful generic functional form

In practice, the functional form

$$\mathcal{E}_{\text{hyperbolic}}(x, y; a, b, c, d, g) = a - y + c(x - b) + g\sqrt{(c(x - b))^2 + d^2} \quad (3.1)$$

suits many practical applications where two financial variables are required to unlock most of the callability value.

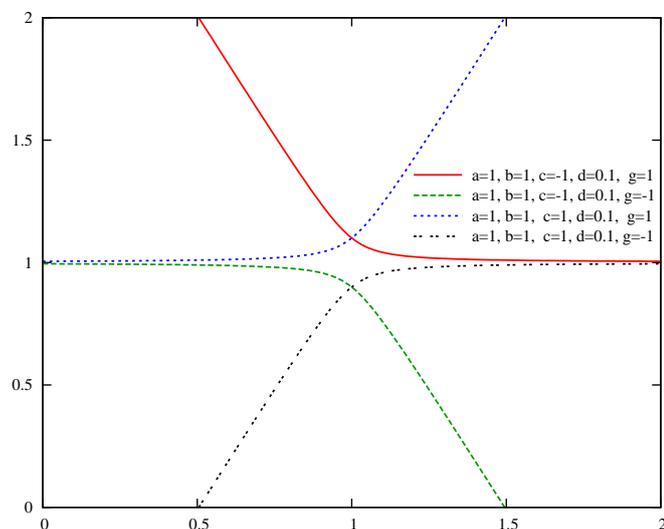


Figure 1: The hyperbolic exercise boundary given by equation (3.1).

The typical shapes of its zero-level contour line (which is the exercise boundary) are shown in figure 1. As an example for the use of this functional form, consider the callability of a payer's Bermudan swaption in a multi-factor model, and consider x to represent the front Libor and y as the coterminal swap rate. In the limit of both x and y very large, it is clearly beneficial to exercise as soon as possible, whence \mathcal{E} ought to be large in this limit. In contrast, when both are very low, we should not exercise, and \mathcal{E} must be negative in this limit. When x is small, and y is moderate but larger than the fixed rate, exercise should not be done now but is likely to become optimal at a later stage. When y is very small, exercise should be avoided, even if x is large. This simple analysis already suggests that $-\mathcal{E}_{\text{hyperbolic}}$ with $a > 0$, $b > 0$, $c < 0$, $g > 0$ might be a good choice, and empirical tests support that this indeed works very well.

4 The full algorithm

The key stages of the exercise boundary optimization method are:-

- Step 1.* Decide on a functional form for the exercise decision functions \mathcal{E} for all exercise time horizons. Note that different functions may need to be chosen for different time horizons if the product exhibits strong inhomogeneity in its features over time. Also note that the exercise domain may not be singly connected whence the implicit formulation of the exercise domain in the form of $\mathcal{E}(x) > 0$ is generally preferable over explicit specifications of the boundary. A simple example for this is a multi-callable best-of option paying $(\max(S_1, S_2) - K)_+$.
- Step 2.* Generation of the n -path training set. The only values that need to be stored are each path's continuation values, cancellation values, and exercise decision function argument values for each exercise decision horizon. Note that for complex models, with contemporary computer's typical memory capacities, this reduction in storage requirements is typically necessary in order to be able to store all data in memory. Also note that the reduction of required memory typically leads to significant speed-up since cache memory access speed and main memory access speed differ considerably.
- Step 3.* In reverse chronological order, optimize the discretely sampled objective function (2.11) for each exercise time horizon in turn. Note that the objective function, at a high resolution level, appears to be piecewise constant in its parameters whence an optimization method ought to be used that can cope with the fact that the function appears to change only at scales compatible with the granularity of the Monte Carlo sampling. One of the simplest methods that allows for a scale change during the optimization is the *Downhill-Simplex* algorithm [PTVF92]. For the case that λ is one-dimensional, Golden Section Search [PTVF92] or outright sorting also works well.
- Step 4.* Using the exercise strategy now defined by the fully specified exercise decision functions established in stage 3, re-evaluate the callable financial contract by an independent N -path Monte Carlo simulation with $N \gg n$. In practice, $N \sim 4 \cdot n$ has been found to be a good ratio when low discrepancy numbers are used throughout.

The final result, by virtue of the inequality (2.10), is of course biased low since the valuation based on an optimized (implicit) functional approximation can only be as good as the exercise domain boundary is represented in the approximation. It can be shown readily, though, that for a small difference ε (defined in any suitable way) between the truly optimal exercise boundary and the one used in the numerical approximation, the difference between the numerically obtained value and the truly optimal value scales like the second order in ε , i.e. like $\mathcal{O}(\varepsilon^2)$ whence small differences tend to have negligible influence on the calculation. Another mitigating factor with respect to the exercise boundary representation is that any mismatches only contribute proportional to the probability of actually reaching this part of state space, i.e. the exercise boundary only need be matched well were probability densities are high. It is these features of second-order-only exercise domain matching error propagation and the fact that the boundary only needs to be represented accurately in the centre of the state vector distribution that makes this method so effective in practice.

References

- [And00] L. Andersen. A simple approach to the pricing of Bermudan swaptions in the multifactor LIBOR market model. *Journal of Computational Finance*, 3(2):5–32, Winter 1999/2000.
- [Jäc02] P. Jäckel. *Monte Carlo methods in finance*. John Wiley and Sons, February 2002.
- [Pit03] V. Piterbarg. Computing deltas of callable Libor exotics in forward Libor models. *Journal of Computational Finance*, 7(3), 2003. ssrn.com/abstract=396180.
- [PTVF92] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C*. Cambridge University Press, 1992. www.library.cornell.edu/nr/cbookcpdf.html.