Monte Carlo Methods for Financial Instruments with American Exercises

January 22-24, 2007

Springer
Berlin Heidelberg New York
Hong Kong London
Milan Paris Tokyo

v

René CARMONA
Bendheim Center for Finance
Department of Operations Research & Financial Engineering
Princeton University
Princeton, N.J. 08544

Monte Carlo Methods for Financial Instruments with American Exercises

Short Course, Stieltjes Winter School, 1/22-24/2007

VERY PRELIMINARY DRAFT

Contents

Preface

In the last few years, academicians have touted practitioners with many recipes to price more efficiently American options, especially when the latter are written on large baskets. Most of these methods are based on Monte Carlo simulation principles. The purpose of these notes is to review the most talked about methods, to expose the mathematical underpinnings, with the hope to debunk some of exaggerated claims which have been made on their practicality and the universality.

As some of these methods go under the name of Monte Carlo regression methods, we felt compelled to devote a chapter to a short primer on the basics of Monte Carlo computations, and an appendix to nonparametric regression. However, these components of the course are not central to the subject, and they can be skipped in a first lecture: they are included for the sake of completeness.

Disclaimer. The present document is still incomplete. These notes are more of a work in progress than a finished product. They still need quite a lot of TLC. They are distributed as a guide to the lectures. My apology for the many missing parts, the unfinished sections and the remaining typos.

René Carmona
Princeton, N.J.

February 27, 2006

1

Optimal Stopping and Snell Envelop for Markov Chains

This first chapter presents the basics of American options. Our emphasis is on numerical methods for computing prices and identifying optimal exercise strategies. As a consequence, we mostly limit ourselves to the discrete time case spending little time explaining how to reduce a continuous time model to this set-up. So for all practical purposes, we study Bermudan options rather than American options.

We review the fundamental mathematical results of the theory of optimal stopping in discrete time. We rephrase the problem in terms of Snell envelops. We concentrate on the particular case of Markovian dynamics for the instruments underlying the options, and we illustrate the important role played by the Bellman dynamic programming principle.

1.1 Optimal Stopping and American Options

Because of our emphasis on numerical implementations, we will exclusively work in a discrete time setting. So in order to avoid technical difficulties, we only consider the optimal stopping problem for discrete time stochastic processes or sequences. Moreover, when we talk about American option, we actually mean what is known as a Bermudan option, namely an option which can only be exercised on a finite set of discrete dates $t = 0 < t_1 < \cdots < t_N$, this whether or not the underlying instrument or index on which the option is written evolves in continuous or discrete time. As a result, the theory which we develop and use will not depend upon what is happening in between two successive possible exercise times t_i and t_{i+1}, only the transition from time t_i to time t_{i+1} matters.

We first fix an integer N. This integer has the interpretation of the horizon of the problem, and the date of maturity when we consider an American option. N could possibly be infinite if we wanted to consider perpetual options. Throughout this chapter we use the following notation. $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space on which all the random variables are defined, and $\{\mathcal{F}_n\}_n$ is a filtration, the σ-field $\mathcal{F}_n$ giving the information available at any given time. We shall use the notation $\mathbb{E}_n\{\ \cdot\ \}$ and alternatively $\mathbb{E}\{\ \cdot\ |\mathcal{F}_n\}$ for the conditional expectation with respect to $\mathcal{F}_n$. $\mathcal{S}$ denotes the set of $\mathcal{F}_n$-stopping times, and $\mathcal{S}_{\alpha,\beta} = \{\tau \in \mathcal{S}; \alpha \leq \tau \leq \beta\}$ the set of stopping times taking values in the interval $[\alpha, \beta]$. We denote by $\{\xi_n\}_n$ an adapted stochastic process with the interpretation of a reward, or pay-off in the case of American options: should the process be stopped at time n or should the owner of the option decide to exercise her right at time n, she will be rewarded by the amount ξ_n. The following technical assumption:

$$\mathbb{E}\{\sup_{0\le n\le N}\xi_n^+\}<\infty \tag{1.1}$$

will be needed for some of the results stated below. Notice that this assumption reduces to the integrability of the positive part ξ_n^+ of each ξ_n when N is finite. We shall often assume that ξ_n is non-negative. This is not a restriction when we deal with option pay-offs. As usual, we use the notation x^+ for the positive part $\max(x,0)$ of a real number x.

1.1.1 Optimal Stopping Problem

In the present context, the optimal stopping problem can be stated as the optimization problem

$$\sup_{\tau\in\mathcal{S}_{0,N}}\mathbb{E}\{\xi_\tau\},$$

the goal being to

1. find the value of the supremum;
2. prove that the supremum is actually attained;
3. characterize the optimal τ's (i.e. the stopping times at which the supremum is attained).

1.1.2 The Snell Envelop

Mathematicians have an uncanny ability to find the solution of a problem by just giving a new name to its solution! So in this spirit we introduce

Definition 1. *The stochastic process $\{\hat{\xi}_n\}_n$ defined by*

$$\hat{\xi}_n=\sup_{\tau\in\mathcal{S}_{n,N}}\mathbb{E}_n\{\xi_\tau\} \tag{1.2}$$

is called the Snell envelop of the process $\{\xi_n\}_n$.

Notice that we used a conditional expectation, so the Snell envelop is a stochastic process which is adapted to the filtration $\{\mathcal{F}_n\}_n$. According to this definition, $\hat{\xi}_n$ is the solution of the initial problem if we assume that the process was not stopped before time n, and if we use all the information available at time n. In particular, $\hat{\xi}_0$ is the value of our original problem since

$$\hat{\xi}_0=\sup_{\tau\in\mathcal{S}_{0,T}}\mathbb{E}_0\{\xi_\tau\}.$$

The following bullet points summarize the properties of the Snell envelop which we use throughout.

- $\hat{\xi}_N=\xi_N$ when N is finite;
- Bellman **dynamic programming principle**:

$$\hat{\xi}_n=\max\{\xi_n,\mathbb{E}_n\{\hat{\xi}_{n+1}\}\}\qquad n<N. \tag{1.3}$$

  Notice that when the horizon N is finite, these first two bullet points give a computational algorithm for the solution of the optimal stopping problem as a function of the terminal reward ξ_N provided we can compute the conditional expectations. Indeed, starting with $\hat{\xi}_N=\xi_N$, one can repeatedly apply the dynamic programming principle to compute $\hat{\xi}_n$ for n decreasing from $n=N-1$ to $n=0$, giving the value $\hat{\xi}_0$ of the optimum.

- $\{\hat{\xi}_n\}_n$ is the smallest super-martingale majorizing $\{\xi_n\}_n$
- τ optimal if and only if
  - $\hat{\xi}_\tau = \xi_\tau$ and
  - $\{\xi_{n\wedge\tau}\}_n$ martingale
- $\tau_0 = \inf\{n; \hat{\xi}_n = \xi_n\}$ is the smallest optimal stopping time

1.1.3 The Markovian Case

We first consider a functional analytic result not involving directly probability or stochastic processes.

Réduite of a Function

We assume that P is a sub-Markovian transition kernel on a measurable space $(E, \mathcal{E})$. In the case of American options, E can be thought of as the space of all possible values of the underlying instrument. For each nonnegative measurable function $\varphi : E \hookrightarrow \mathbb{R}_+$ on E, the function φ^* defined as $\varphi^* = \lim_{k\to\infty} \varphi_k$ with

$$\varphi_0 = \varphi, \qquad \varphi_{k+1} = \max\{\varphi, P\varphi_k\}$$

is the smallest super-harmonic function majorizing φ. A non-negative measurable function f on E is said to be super-harmonic (for the kernel P) if $Pf \leq f$. φ^* is called the *réduite* of φ for the kernel P. The limit appearing in the definition of φ^* exists because the sequence $\{\varphi_k\}_k$ is non-decreasing as we can easily see by induction. Indeed, $\varphi_0 \leq \varphi_1$ by the above definitions of φ_0 and φ_1. Now, if we assume that $\varphi_k \leq \varphi_{k+1}$ we have

$$\begin{aligned}\varphi_{k+2} &= \max\{\varphi, P\varphi_{k+1}\} \\ &\geq \max\{\varphi, P\varphi_k\} \\ &= \varphi_{k+1}.\end{aligned}$$

where we used the induction hypothesis $\varphi_k \leq \varphi_{k+1}$, the monotonicity of P, and the definition of φ_{k+1}. The fact that φ^* is super-harmonic is a straightforward consequence of the monotone convergence theorem. Indeed:

$$P\varphi^* = P(\lim_{k\to\infty} \varphi_k = \lim_{k\to\infty} P\varphi_k = \varphi^*.$$

Notice that (using once more monotone convergence)

$$\boxed{\varphi^* = \max\{\varphi, P\varphi^*\} \qquad (1.4)}$$

which is a system of coupled equations (as many as elements of the set E) which we could try to solve for the values of the reduite. Indeed, this functional equation says

$$\varphi^*(x) = \max\{\varphi(x), \int_E \varphi^*(y)P(x, dy)\}, \qquad x \in E$$

Functions of a Markov Chain

We now consider a stochastic process $\{X_n\}_n$ which, in the case of the pricing problem for American options, will represent the underlying indexes or securities on which the options are written, and we assume that this process is a Markov chain with transition probability given by a transition kernel which we denote by P. This means that the kernel P gives the conditional expectations of functions of the Markov chain in the sense that

$$\mathbb{E}\{g(X_{n+1})|\mathcal{F}_n\} = \mathbb{E}\{g(X_{n+1})|X_n\} = [Pg](X_n)$$

for any function g on the state space for which the above expectations/integrals make sense. The main result is that if the reward ξ_n at time n is given by a function $\varphi(X_n)$ of the Markov chain, then the Snell envelop is given by the reduite of that function computed along the chain. In other words:

$$\hat{\xi}_n = \varphi^*(X_n). \tag{1.5}$$

Among the many consequences of this theoretical result, the following has far reaching numerical implications. The smallest optimal stopping time is the hitting time of the set

$$D = \{x \in E;\, \varphi(x) = \varphi^*(x)\} \tag{1.6}$$

which is called the exercise region. Indeed, we saw earlier that the smallest optimal stopping time τ^* was the first time the Snell envelop process coincides with the reward process. Given the representation (1.5) of the Snell envelop, this implies

$$\tau^* = \inf\{n;\, X_n \in D\}.$$

Remarks

The above results are stated in the time-homogeneous case when the transition kernel does not change with time.

$\diamond$ However, the results above also hold when the transition kernel P depends upon time, in other words, when the transition from time n to time $n+1$ is given by a kernel P_n which can change with n;

$\diamond$ Moreover, if the transition kernel is allowed to change with time, there is no reason why the state space should not also be allowed to change with time. When appropriate, we shall use the notation E_k for the set of possible states at time k.

As we shall see in the next section, these generalizations provide a convenient way to recast the binomial tree model in the Markov chain framework considered in this subsection.

1.1.4 American / Bermudan Options

We now concentrate on the case of an American option, and we assume that the pay-off to the owner of the option is $\varphi(X_n)$ if the owner chooses to exercise her right at time n. Here φ is a non-negative function which can be computed for all the values of the underlying instruments X_n at time n. Typical examples include

- $\varphi(x) = (x-K)^+$ in the case of a call option with strike K on a single stock or index with value x;

- $\varphi(x) = (K - x)^+$ in the case of a put option with strike K on a single stock or index with value x;
- $\varphi(x) = (\sum_i w_i x^{(i)} - K)^+$ (resp. $\varphi(x) = (K - \sum_i w_i x^{(i)})^+$) in the case of a call (resp. put) basket option with strike K and weights w_i on a basket of stocks or indexes with components $x^{(i)}$;

As the components $x^{(i)}$ of the vector x from which the basket is constructed are allowed to depend upon each other, one could think of the $x^{(i)}$ as values of the same underlying instrument at different time, including in this way Asian option in the modelling framework of basket options.

As before, we assume that $\{X_n\}_n$ is a Markov chain. For the sake of generality, we do not assume that this chain is homogeneous, and we denote by P_n the transition kernel from time n to time $n+1$.

Discounting

As explained in our discussion of the dynamic programming principle, when considering exercising the option at time n, the owner of the option has to compare the immediate reward $\varphi(X_n)$ to the expected future rewards should she choose not to exercise immediately. For this comparison to be meaningful, she needs to express the values of the future cash flows in "time n currency". Hence the need for discounting the future cash flows to compute their present value equivalents. We denote by $D(n_1, n_2)$ the discounting factor giving the time n_1 value of a unit cash flow taking place at time $n_2 > n_1$. In the discrete time setting we have;

$$D(n_1, n_2) = \prod_{n=n_1}^{n_2-1} \frac{1}{1 + r_n} \tag{1.7}$$

where r_n stands for the yield over the period from time n to time $n + 1$. This discounting factor is multiplicative in the sense that

$$D(n_1, n_3) = D(n_1, n_2)D(n_2, n_3)$$

whenever $n_1 < n_2 < n_3$. We explain in formula (1.30) how to compute this discount factor from the exponential term used in the continuous time when we discretize a model from continuous time finance. We will allow for *stochastic interest rates* in the sense that $\{r_n\}_n$ could be a random sequence as long as it is adapted, the typical case being when r_n is a function of the X_m for $m \leq n$.

The value at time $n = 0$ of the American option is given by the solution of the optimal stopping problem:

$$\sup_{\tau \in \mathcal{S}_{[0,N]}} \mathbb{E}\{D(0, \tau)\varphi(X_\tau)\}. \tag{1.8}$$

Despite its clear intuitive meaning, the fact that this optimal value is the right price requires a justification, and we shall not discuss the no-arbitrage argument leading to this result.

In order to recast this optimal stopping problem in the Snell envelop framework discussed above, we introduce the discounted transition kernels $\tilde{P}_n$

$$[\tilde{P}_n f](x) = \mathbb{E}\{D(n, n + 1)f(X_{n+1})|X_n = x) \tag{1.9}$$

which can be written in the simpler form

$$[\tilde{P}_n f](x) = D(n, n+1)[P_n f](x) \tag{1.10}$$

when the interest rates (and hence the discounting factors) are deterministic. Even though $\tilde{P}_n$ is not exactly the transition kernel of the Markov chain $\{X_n\}_n$, the results stated earlier in the case of a function of time-homogeneous Markov chain can be adapted to the present situation. As in the homogeneous case described above, the Snell envelop

$$\hat{\xi}_n = \sup_{\tau \in \mathcal{S}_{[n,N]}} \mathbb{E}\{D(n,\tau)\varphi(X_\tau)|\mathcal{F}_n\} \tag{1.11}$$

can still be written as a function of the underlying Markov chain (recall our discussion of the réduite), but in the present case, this function is time dependent, namely:

$$\hat{\xi}_n = v_n(X_n) \tag{1.12}$$

for a sequence $\{v_n\}_n$ of functions of the state variable x. $\hat{\xi}_n$ gives the value of the option at time n if it has not yet been exercised by that time. As we concentrate on the finite horizon case, this time dependent value function is given by the dynamic programming principle

$$\begin{cases} v_N(x) = \varphi(x) \\ v_n(x) = \max\{\varphi(x), [\tilde{P}_n v_{n+1}](x)\}, \qquad n = 0, 1, \cdots, N-1. \end{cases} \tag{1.13}$$

The smallest optimal time τ_n^* if the option was not exercised before time n is given by the first time after n for which the Snell envelop coincides with the pay-off

$$\tau_n^* = \inf\{n' \geq n;\ v_{n'}(X_{n'}) = \varphi(X_{n'})\}. \tag{1.14}$$

and in particular, the Snell envelop can be expressed as a conditional expectation

$$\hat{\xi}_n = v_n(X_n) = \mathbb{E}_n\{D(n, \tau_n^*)\varphi(X_{\tau_n^*})\} \tag{1.15}$$

instead of a supremum of such conditional expectations.

As usual, we drop some of the subscripts n when $n = 0$ so that the value of the option is given by

$$\hat{\xi}_0 = v_0(X_0) = \sup_{\tau \in \mathcal{S}} \mathbb{E}\{D(0,\tau)\varphi(X_\tau)\}$$

and the smallest optimal exercise time τ^* is given by

$$\tau^* = \inf\{n \geq 0;\ v_n(X_n) = \varphi(X_n)\}.$$

In particular

$$\hat{\xi}_0 = v_0(X_0) = \mathbb{E}\{D(0,\tau^*)\varphi(X_{\tau^*})\} \tag{1.16}$$

Continuation Function

In many instances it will be convenient to rewrite the dynamic programming principle in terms of the value of not exercising immediately. We formalize this intuitive idea by defining the function q_n for $n = 0, 1, \cdots, N-1$ by

$$q_n = \tilde{P} v_{n+1} \tag{1.17}$$

for each $x \in E_n$, then $q_n(x)$ represents the current expectation, conditioned by the fact that $X_n = x$ at time n, of future rewards assuming that future exercise will be decided optimallly, if there is no exercise at time n. In other words q_n gives the **continuation value** at time n and the dynamic programming principle can be rewritten as

$$v_n(x) = \max\{\varphi(x), q_n(x)\}.$$

Obviously, the full backward induction can be rewritten in term of the continuation value functions:

$$\begin{cases} q_{N-1} = \tilde{P}_{N-1}\varphi \\ q_n = \tilde{P}_n \max\{\varphi, q_{n+1}\}, \qquad n = 0, 1, \cdots, N-2 \end{cases} \tag{1.18}$$

It is interesting to compare (1.18) with (1.13). A major difference lies in the relative placements of the maximum and the action of the discounted transition operator $\tilde{P}_n$. The latter enters linearly in (1.18) and this will make computations easier.

Still Another Formula

For each integer $n \geq 0$ we define the non-linear operator Q_n by

$$[Q_n f](x) = \max\{\varphi(x), [\tilde{P}_n f](x)\} \tag{1.19}$$

whenever the right hand side makes sense (e.g. when $f \geq 0$). With this notation, the backward induction providing the value function at time $n = 0$ can be re-written as the following equality:

$$\boxed{v_0(x) = [Q_0 Q_1 \cdots Q_{N-1}\varphi](x) \qquad (1.20)}$$

This form of the dynamic programming principle is quite useful when the state space E or the state spaces E_n are finite. Indeed, a function f on E_n can be identified to a vector comprising all the values $f(x)$ (so the length of this vector is equal to the cardinality of E_n). Since the transition probability operator $\tilde{P}_n$ is linear, it can be viewed as a matrix, and computing the function $\tilde{P}_n f$ can be done by multiplying the matrix corresponding to $\tilde{P}_n$ by the vector of the entries of f. The computation of the vector comprising the values of the function $Q_n f$ can then be completed by computing entry by entry, the maximum of each entry $\varphi(x)$ with the corresponding entry $[\tilde{P}_n f](x)$.
We will use this remark in our implementation of the quantization method introduced in Chapter 3.

These results can be written in a simpler form when the Markov chain $\{X_n\}_n$ and the discount factors are time-homogeneous. In this case, formula (**??**) reduces to the more compact expression

$$\boxed{v_0(x) = [Q^N\varphi](x) \qquad (1.21)}$$

where the exponent N means that we apply the operator Q N times to the pay-off function φ, and where the non-linear operator Q is defined by:

$$[Qf](x) = \max\{\varphi(x), [\tilde{P}f](x)\} \tag{1.22}$$

for all non-negative measurable function f, or for any function f for which $\mathbb{E}\{|f(X_1)|\} < \infty$.

American Call and Put Options in the Black-Scholes Model

This subsection is a little excursion in the world of continuous finance., but before we even consider specific models for the underlying risky assets, we derive a general fact about the exercise of American call and put options which is useful both in discrete and continuous time models as long as the maturity is finite.

Let us assume momentarily that the discounted reward process $\tilde{\xi}_t = D(0,t)\varphi(X_t)$ is a *submartingale*. Contrary to the case of super-martingales, expectations and conditional expectations increase with time for such processes. So in order to find a stopping time maximizing $\mathbb{E}\{\tilde{\xi}_\tau\} = \mathbb{E}\{D(0,\tau)\varphi(X_\tau)\}$ it always pays to wait. Consequently, for such processes, the optimal exercise time is the maturity date T. In other words, the fact that the exercise is of the American type is of no benefit to the holder of the option if the discounted reward process is a submartingale: there is **no early exercise premium** in this case.

This innocent looking remark has important applications in several cases including the Black-Scholes model of call and put options on a geometric Brownian motion $\{X_t\}_t$. In this case, and in the absence of dividend payments the risk neutral dynamics of the underlier are given by

$$dX_t = X_t[rdt + \sigma dW_t],$$

we assume that the short interest rate is deterministic and constant over the life of the option. We use the discount $D(0,t) = e^{-rt}$. We also assume that the maturity T and the strike K are positive numbers. A simple application of Ito's stochastic calculus gives

- in the case of an American call option, i.e. when the pay-off function is $\varphi(x) = (x-K)^+$, the discounted reward process $\tilde{\xi}_t$ is a submartingale if $r \geq 0$.
- in the case of an American put option, i.e. when the pay-off function is $\varphi(x) = (K-x)^+$, the discounted reward process $\tilde{\xi}_t$ is a submartingale if $r \leq 0$.

Hence, unless the underlying stock is paying dividends (and the drift r is replaced by $r-q$ if q is the rate of continuous dividend payment) the problem of pricing and exercising an American call option is trivial in the sense that it is not any different from the case of a European call. This is why when considering American options, we consider put options most of the time.

1.2 MULTIPLICATIVE BINOMIAL TREES

Tree models are often proposed as simple models capturing stylized facts of real markets. They are popular because of their conceptual simplicity and the fact that they are easy to implement. Moreover, they are also used as discrete approximations of more sophisticated continuous time models. The following notation will be used for discrete time models as well as for continuous time models.

- N is the number of periods, or the number of intervals in the time discretization
- Δt is the length of the single period;
- $T = N\Delta t$ is the maturity of the option;
- K is the strike of the option;
- X_0 is the value at time $t = n = 0$ of the asset underlying the option;
- r is the (annual) short interest rate.

Time is measured in years, and the discount factor used in computing present values is $D(n, n+1) = e^{-r\Delta t}$ if the tree is used as a discretization of a continuous time model, and $D(n, n+1) = 1/(1 + r\Delta t)$ if the tree is directly considered as a model of a discrete multi-period economy.

In the general setting of a multi-period discrete time model, the assumption is that, from one period to the next, the value of the underlying asset, say X, is *multiplied* by a random quantity, say ξ, which represents the return on the asset over the period in question. The premise of a multiplicative tree is that this quantity can only take finitely many values. We implicitly assume that the random quantities corresponding to different periods are independent. The tree is binomial when the random variable ξ can only take two values. So over a single period, the value of the underlying asset can go from $X_0 = x_0$ to $X_1 = X_0\xi$ where the random variable ξ can only take two values $\xi_d < \xi_u$. A schematic of the evolution of the value of the underlying asset is given on the left part of the following diagram.

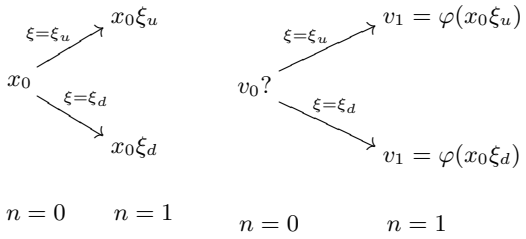

The right part of the above diagram illustrates the analysis of an American option with pay-off function φ. If the option is exercised at time $n = 1$, the reward is given by the cash flow given by the pay-off, and the value at time $n = 0$ of exercising this call option at time $n = 1$ is given by the expected discounted pay-off. This is the same thing as valuing at time $n = 1$ the European option with maturity $n = 1$ and pay-off $\varphi(X_1)$. A standard arbitrage argument implies that this value is given by the expected discounted pay-off at time $n = 1$:

$$v_0^{(e)} = \mathbb{E}\{D(0,1)v_1\} = D(0,1)[p_u v_u + (1 - p_u)v_d] \tag{1.23}$$

provided we define the risk neutral probability p_u by:

$$p_u = \frac{e^{r\Delta t} - \xi_d}{\xi_u - \xi_d}. \tag{1.24}$$

The optionality of the American exericise is to give the holder the right to exercise the option at time $n = 0$, in which case, she would get $\varphi(x)$. A time $n = 0$, the holder of the option can choose to exercise the option immediately, or if it is more advantageous, to wait for time $n = 1$ to exercise her right. Hence, in accordance with the dynamic programming principle, her value at time $n = 0$ is

$$\begin{aligned} v_0^{(e)} = \max\{\varphi(x_0), v_0^{(e)}\} &= \max\{\varphi(x_0), \mathbb{E}\{D(0,1)v_1\}\} \\ &= \max\{\varphi(x_0), D(0,1)p_u v_u + (1 - p_u)v_d\}. \end{aligned}$$

1.2.1 **Pricing by Backtracking in Multiperiod Models**

We now consider a binomial tree as a model for a N period horizon. The price dynamics for the risky asset underlying the option are given by the following recombining tree:

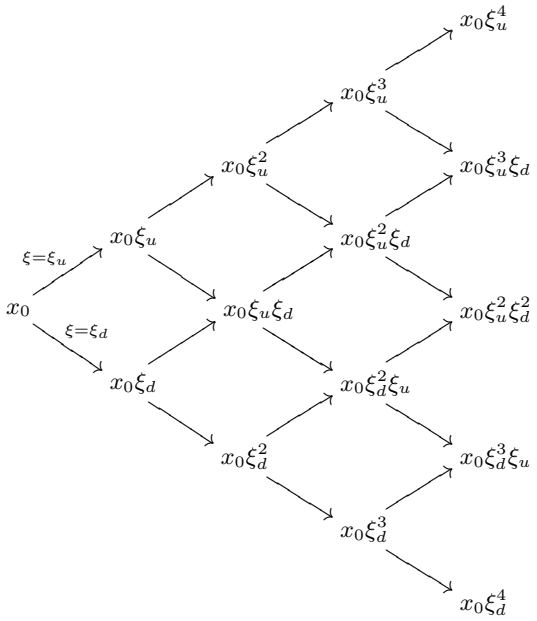

For the sake of consistency, we recast this price tree in the Markov chain framework introduced earlier in the chapter. For each $n = 0, 1, \cdots, N$ the underlying risky asset takes exactly $n + 1$ values, and if we denote by $x_n[i]$ the i-th value, the latter can be characterized by i down-moves and $n - i$ up-moves for the underlying asset in the first n periods. So

$$x_n[i] = x_0\xi_d^i\xi_u^{n-i}, \qquad i = 0, 1, \cdots, n,\ n = 0, 1, \cdots, N.$$

According to the notation of Section 2.3, the state space at time n is given by

$$E_n = \{x_n[i];\ i = 0, 1, \cdots, n\}$$

The transition kernel P_n is easily defined since from $X_n = x_n[i]$ the risky asset can only go to up or down, and with probabilities p_u and $p_d = 1 - p_u$ respectively. So

$$P_n(x, y) = \begin{cases} p_u \text{ if } y = x\xi_u \text{ and } x \in E_n \\ p_d \text{ if } y = x\xi_d \text{ and } x \in E_n \\ 0 \text{ otherwise} \end{cases}$$

Notice that the returns ξ_u and ξ_d can be made time dependent (i.e. functions of n) without affecting the pricing procedure as long as the risk neutral probabilities p_u (which automatically become time dependent) are adjusted to satisfy (1.23). Notice that discounted prices are martingales under this probability structure. Indeed:

$$X_n = D(n, n+1)\mathbb{E}\{X_{n+1}|\mathcal{F}_n\}, \qquad n = 0, 1, \cdots, N-1.$$

The dynamic programming principle (1.13) can be used to compute the value of the option by filling the tree from right to left. Indeed, if the option has not been exercised yet, its value at maturity

$n = N$ (4 in the diagram below) is given by the pay-off of the option, since exercise can only take place at that time. In other words, we can compute all the values $v_N(i)$ for $i = 0, 1, \cdots, N$ as

$$v_N(i) = \varphi(x_N[i]) = \varphi(x_0 \xi_d^i \xi_u^{n-i}), \qquad i = 0, 1, \cdots, N.$$

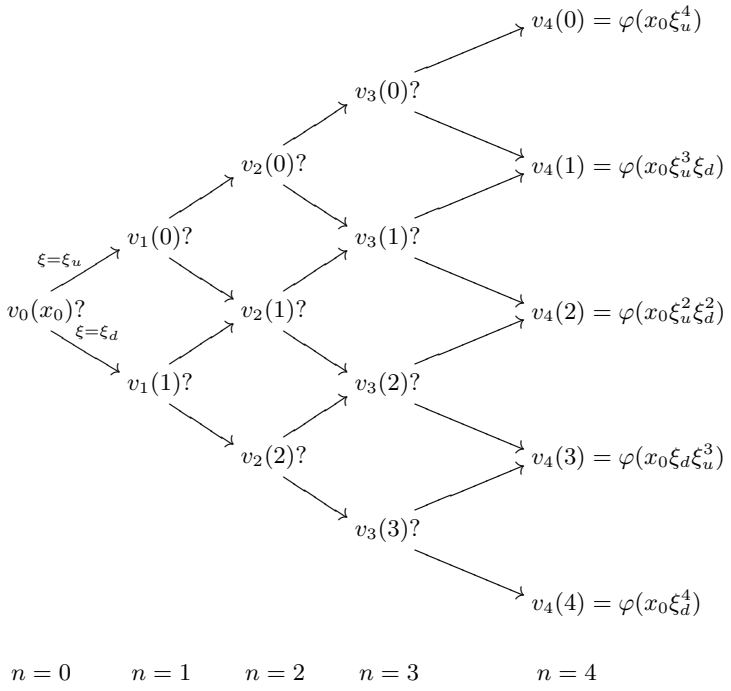

Next, we backtrack one step in time and compute the value of the option at time $n = N - 1$. Following the prescription of the dynamic programming equation (1.3), we compute for each node, the maximum of the reward for immediate exercise, i.e. $\varphi(x_n[i])$ and the expected reward $[\tilde{P}_n v_{n+1}(\,\cdot\,)](x_n[i])$ we would get if instead of exercising immediately we were to wait one unit of time, and then act optimally from that point in time.

What we just did for $n = N - 1$ can be done for $n = N - 2, \cdots, n = 1$ and finally $n = 0$ to fill in the tree from right to left by implementing the dynamic programming principle at each node. Proceeding in this way, we fill in the entire tree all the way to the root. According to the dynamic programming principle, this last value is the (risk neutral) price of the option. This algorithm is encapsulated by the `pseudo code`

```
for i varying from 0 to N v(N,i) = phi(x_N[1])
for n varying from N-1 down to 0
    for i varying from 0 to n
        v_n(i)=max{phi(x_n[i]),D(n,n+1)*
                [p_u * v(n+1,i)+(1-p_u) * v(n+1,i+1)]}
```

We can now move down the tree (i.e. from left to right) to identify the smallest exercise strategy: along each path, the option is exercised as soon as the value function $v_n(i)$ computed above is greater than the reward $\varphi(x_n[i])$ for immediate exercise.

1.3 BINOMIAL TREES STILL, BUT FROM A DIFFERENT PERSPECTIVE

In this section we change our perspective on trees. Instead of viewing them as stand alone models for an economy and bona fide supports for the risk neutral stochastic dynamics of the underlying asset, we assume that these dynamics are in fact given by the geometric Brownian motion model of continuous time finance proposed by Samuelson, and we view the tree as an intuitively appealing approximation tool from which one constructs numerical schemes to compute prices, hedges, ...

The over-arching model is:

$$dX_t = X_t[\mu dt + \sigma dW_t] \tag{1.25}$$

where μ is the annual rate of growth of the underlying asset, σ is its volatility, and dW_t represents the infinitesimal increment of a standard Wiener process also called process of Brownian motion. We are interested in pricing and hedging contingent claims, and for that reason, we need to work with the risk neutral dynamics of the asset. Since we assume that the latter does not pay dividends, this forces us to assume that the rate of growth is the short term interest rate, i.e. $\mu = r$. As before, we restrict ourselves to the particular case of a European call option with maturity T and strike K for the sake of definiteness. Its price is given by the risk neutral expectation of the discounted pay-off of the option, hence:

$$C_{T,K} = \mathbb{E}\{e^{-rT}(X_T - K)^+\}. \tag{1.26}$$

Notice that the stochastic differential equation (1.25) can be solved explicitly. The solution is given by the geometric Brownian motion:

$$X_t = X_0 e^{(r-\sigma^2/2)t+\sigma W_t} \tag{1.27}$$

and this shows that X_t has a log-normal distribution. The unexpected term $-\sigma^2 t/2$ appearing in the exponential is usually called the Itô correction. Using the log-normal distribution, we can compute explicitly the expectation in (1.26), leading to the famous Black-Scholes formula which we recalled and programmed in Chapter **??**.

Our goal in this section is to construct multiplicative binomial tree models providing approximations to the desired price, e.g. (1.26). Once a time horizon T is chosen, an approximating binomial tree can be constructed by choosing N typically large (recall that the length of one time period is then given by $\Delta t = T/N$), two numbers ξ_u and ξ_d for the distribution of the random return ξ, and a probability p.

Remark 1. At this stage of our discussion of the approximation strategy, we do not require that the probability p be given by the risk neutral condition (1.24). In other words, we do not assume that $p = p_u$ given by this formula. This means that the tree may not be exactly risk neutral, and consequently, computing prices by expectations of prices over the next period (i.e. filling in the tree from right to left as we did earlier) may not be justified rigorously, and such a computation procedure may introduce significant errors if the risk neutral condition (1.24) is far from satisfied.

We review two standard tree construction procedures, both relying on choices of the parameters ξ_u, ξ_d and p intended to match the first two statistical moments of some of the random variables in the continuous time model (1.25) with those of their analogs in the tree in question.

1.3.1 The CRR Prescriptions

Formula (1.27) states that $X_{\Delta t}$ is of the form:

$$X_{\Delta t} = X_0 e^{(r-\sigma^2/2)\Delta t} e^{\sigma W_{\Delta t}}$$

Cox, Ross and Rubinstein use a binomial tree where

$$X_{\Delta t} = X_0 e^{(r-\sigma^2/2)\Delta t} e^{\epsilon}$$

(hence $\xi = e^{(r-\sigma^2/2)\Delta t} e^{\epsilon}$) where the random variable ϵ takes only two values in such a way to have the same mean and the same variance as $\sigma W_{\Delta t}$. In other words, ϵ has mean 0 and variance $\sigma^2 \Delta t$. These two constraints give two equations for the three variables ξ_u, ξ_d and p, which is not enough to determine them uniquely. So they choose arbitrarily $p = .5$, which is in some sense, a symmetry condition on the probability distribution of ϵ. Solving explicitly one finds:

$$p = \frac{1}{2}, \qquad \xi_u = e^{(r-\sigma^2/2)\Delta t + \sigma\sqrt{\Delta t}}, \qquad \xi_d = e^{(r-\sigma^2/2)\Delta t - \sigma\sqrt{\Delta t}}. \tag{1.28}$$

In line with the above remark, we notice that the corresponding tree is not risk exactly neutral, but simple computations with the Taylor expansions of these quantities when Δt is small show that the risk neutral condition (1.24) is satisfied in the limit $\Delta t \to 0$.

One of the main criticism of the CRR approximation is its slow convergence toward the value of the Black-Scholes formula.

1.3.2 The JR Prescriptions

In order to improve on the poor convergence properties of the CRR prices, Jarrow, and Rudd propose to use a binomial tree where

$$p = \frac{1}{2} + \frac{r - \sigma^2/2}{2\sigma}\sqrt{\Delta t}, \qquad \xi_u = e^{\sigma\sqrt{\Delta t}}, \qquad \xi_d = e^{-\sigma\sqrt{\Delta t}}. \tag{1.29}$$

As in the case of the CRR trees, the JR trees are only risk neutral in the limit $\Delta t \to 0$.

1.4 EXTENSIONS

This final section is a patchwork of topics which are connected to the problems discussed in this chapter, but which we do not want to cover in full because of their technical nature.

1.4.1 Dividend Payments

If the underlying asset is a stock paying dividends and if the payments of these dividends can be modelled with a a continuous flow at a (continuous) rete $\delta > 0$, then the Black-Scholes theory shows that the pricing problem is equivalent to pricing the derivative in a model without dividend payments and where the effective discounting rate r giving the drift in the equation for the dynamics of the underlying asset is replaced by the effective rate $r - \delta$. The construction above can be adapted without any significant change to handle this case.

On the other hand, the problem is more difficult when owning the stock entitles the investor to cash payments.

1.4.2 **Multivariate Derivatives**

Higher Dimensional Trees

The concept of binomial tree can be extended to the pricing of options on several underlying assets. This is for example the case of spread options or more general basket options. We consider the case of spreads for the case of definiteness. Basket options are never priced with trees because even a reasonably small number of components in the basket is enough to make the size of the tree computationally prohibitive.

Let us consider for example an option with strike K and maturity T on the difference $X_t = X_t^{(2)} - X_t^{(1)}$ of two assets. A first hunch could be to model directly the dynamics of the underlying index X_t and to build a binomial tree for this model, This has been done with additive trees, with a reasonable success for some sets of parameters. But because of the limitations of these models, it makes more sense to model separately the underlying assets $X_t^{(1)}$ and $X_t^{(2)}$, construct a binomial tree for each of them, and then combine these two separate trees into a single tree capturing the dependencies of the two assets.

Let us assume for example that we want to build a tree with N time periods. At time $t = n\Delta t$, i.e. at the end of the n-th period, the first asset $X^{(1)}$ can move up or down and move to two different values, and for each of them, the second asset $X^{(2)}$ can also have two values. This makes a total of 4 nodes to which the couple $(X^{(1)}_{n\Delta t}, X^{(2)}_{n\Delta t})$ can move during the period. Hence four probabilities to choose. In order to choose these probabilities, not only should we use the separate dynamics of the two underlying assets, but also the way they *correlate*.

Such a construction is reasonable if one is interested in the analysis of qualitative properties of economic models. However, it is unsatisfactory if one needs precise approximations of continuous models. Indeed, because of poor rates of convergence, the size of the tree becomes rapidly unmanageable.

Discretized Black-Scholes Model

In this subsection we consider the case of an American option written on a basket of d underlying risky assets. Here d can be as large as 50 or 100. However, and despite many claims to the contrary, numerical methods available today work well only for dimensions much smaller if we do not rely on the Monte Carlo methods discussed in Chapter 3. For the sake of these lectures, we shall only consider values of d between 2 and 6.

We assume for the sake of simplicity that the dynamics of the underlying assets $X^{(i)}$ are given by a non-degenerated d-dimensional geometric Brownian motion. As we shall see later in Chapter 4, in this particular case, the time discretization leads to exact simulation avoiding errors produced by a scheme such as Euler's. So we assume that there exists a d-dimensional Wiener process $\{W_t^{(i)}\}_{i=1,\cdots,d,\, t\geq 0}$ such that:

$$X_t^{(i)}(x) = x_0^{(i)} \exp\left[\left(r - \frac{1}{2}\sum_{j=1}^{d}\sigma_{ij}^2\right)t + \sum_{j=1}^{d}\sigma_{ij}W_t^{(j)}\right]$$

where the process start from $X_0^{(i)} = x_0^{(i)}$, $i = 1, \cdots, d$, and where $\Sigma = [\sigma_{ij}]_{i,j=1,\cdots,d}$ is a positive definite $d\times d$ matrix, and if we denote by $t_0 = 0 < t_1 < \cdots < t_N = T$ the fixing dates which we use

to discretize the continuous time model. These are the dates at which the options which we consider can be exercised. So as explained earlier, the option is Bermudan instead of being American.

With these notation, the transition kernel P_n governing the transition from time t_n to t_{n+1} is given by:

$$[P_n f](x) = \mathbb{E}\{f(X_{t_{n+1}-t_n}(x))\}, \qquad x \in \mathbb{R}^d$$

for $n = 0, 1, \cdots, N-1$. Notice also that

$$r_n = e^{r(t_{n+1}-t_n)} - 1, \qquad n = 0, 1, \cdots, N-1. \tag{1.30}$$

and that P_n is independent of n (i.e. the Markov chain is homogeneous) whenever the interest rate is constant and the fixings are regularly spaced.

Notice that, because of the availability of exact simulation in this case, we do not need to consider intermediate times between two successive fixing dates, as we would have needed if we were to rely on a discretization scheme (such as Euler's or Milshtein's) for a stochastic differential equation. See next chapter for a discussion of this issue.

NOTES & COMPLEMENTS

A detailed discussion (including proofs) of the Snell envelop of discrete time stochastic processes can be found in Neveu's textbook [**?**]. Applications to the analysis of American options are given in the masterful little textbook of Lamberton and Lapeyre [**?**] where we can also find the no-arbitrage argument showing that the value of the American option is given by the value function of the optimal stopping problem.

The original paper of Cox, Ross and Rubinstein [**?**] is still one of the best places to learn about the intricacies of binary trees as a model for option pricing. Our presentation of the financial material of this chapter was influenced by the first chapter of Clewlow and Strickland's book [**?**]. As emphasized in the text, the binomial tree can be viewed as a financial model of its own, or alternatively, it can also be viewed as a numerical scheme set up to compute approximate values for prices and hedges of a continuous time finance model such as the Samuelson model of the Black-Scholes theory for example. In doing so, we may relax the risk neutrality restriction, but we need to address the accuracy of the approximation and the rate of convergence when convergence toward the limiting continuous model does occur. This is addressed in a recent paper by J. Walsh [**?**] where the author cleverly embeds all the different approximating trees in the same continuous model using Skorohod embedding arguments. As shown by W. Schachermayer in a recent intriguing note [**?**], this convergence is very subtle due to the meaning of risk neutrality, and minor changes can lead to quite unexpected limiting behavior.

The reader interested in spread options is referred to the recent survey by Carmona and Durrleman [**?**] where she will find among other things, a detailed discussion of the various numerical methods used to price these instruments.

2

Monte Carlo Primer

This chapter is intended as a crash course on Monte Carlo computations. We review the standard methods of random simulations (often called Monte Carlo simulations) and we emphasize the case of time discretization of continuous time stochastic differential equation. The second part is devoted to the numerical computations of expectations by the so-called Monte Carlo method. Our goal is not to give an in depth presentations of these subjects. We merely review the results which we use in the following chapter on Monte Carlo methods for the pricing of American options.

2.1 RANDOM GENERATION

Most financial engineering problems lead to the computation of expectations of the form

$$\mathbb{E}\{\psi(X_T)\} \qquad \text{or} \qquad \mathbb{E}\{\Psi(X_{[0,T]})\}.$$

In many cases, X_T is a random vector giving the values of economic factors and underlying securities at time T and ψ give the present value of a specific cash flow or pay-off depending on these factors and underlying instruments. These random pay-offs often depend not only on the values of the factors at time T, but on the entire past history, say $X_{[0,T]}$, of these factors. In these cases, we need to evaluate the second of the above expectations. Notice that computations of probabilities, such as those probabilities $\mathbb{P}\{L \leq x\}$ involved in the computations of values at risk, (VaR for short) fall under the present framework as the probability of an event is nothing more than the expectation of the indicator random variable which is 1 when the event occurs and 0 otherwise. The goal of this chapter is to present Monte Carlo methods to compute approximations to the values of these expectations and probabilities.

In this first section, we address the most fundamental issues of random number generation. We use bullet points to highlight the main topics which we intend to illustrate with examples.

- *Random Number Generators.*
- *Random Simulations.*
- *Variance Reduction Techniques.*

A reasonable treatment of these topics would require more time and space, and our cavalier discussion does not do justice to the importance of the concepts. We encourage the reader to consult the

Notes & Complements at the end of the chapter for references to systematic presentations of the issues skimmed over in this chapter. Each bullet point is now explained and illustrated with simple introductory examples. We use concepts and tools introduced in the analysis and mitigation of credit risk to illustrate the computational methods introduced in this chapter.

2.1.1 **Random Number Generators and Random Samples**

All computing environments comprise a random number generator. Repeated calls to this generator return a sequence of numbers, say $u_1, u_2, \ldots, u_n$ which form a sample of random numbers uniformly distributed over the unit interval $[0, 1]$. For us, a "sample" means a set of realizations of a finite number of random variables which have the same distribution and which are independent. In the present situation, the u_j's are realizations of independent random variables $U_1, U_2, \ldots, U_n$ with the uniform distribution $U(0, 1)$ over the unit interval $[0, 1]$. Also, we shall try to adhere to the convention of using lower cases for realizations of random variables, and upper cases for the actual random variables.

Random simulations and Monte Carlo computations require non-uniform random samples. The easiest way to construct such samples with a uniform random number generator is to use the method of the cumulative distribution (cdf for short). The latter is based on the following elementary result from probability theory.

> If F^{-1} is the inverse of a cdf F, and if $U \sim U(0, 1)$ is a uniform random variable on the unit interval $[0, 1]$, then the random variable $X = F^{-1}(U)$ has cdf F.

This theoretical result is used in the following practical way: in order to construct a sample $x_1, x_2, \ldots, x_n$ from a distribution with cdf F, we first generate a sample $u_1, u_2, \ldots, u_n$ from the uniform distribution on the unit interval $[0, 1]$, and we compute:

$$x_1 = F^{-1}(u_1), x_2 = F^{-1}(u_2), \ldots, x_n = F^{-1}(u_n).$$

2.1.2 **Simulation of Finite Random Variables**

The first example of non-uniform random generator which we discuss is simple enough to be understood without requiring the above theoretical result. It concerns the simulation of random variables taking only finitely many values. Indeed, in this case, the cdf method takes a very intuitive form which can be explained at a very intuitive level.

Let us assume for example that the random variable X can only take finitely many values $x^{(1)}, \cdots, x^{(k)}$ with probabilities $p_1, \cdots, p_k$ respectively

$$\mathbb{P}\{X = x^{(j)}\} = p_j, \quad j = 1, \cdots, k \quad \text{with} \quad p_j \geq 0 \quad \text{and} \quad p_1 + \cdots + p_k = 1.$$

Let us partition the unit interval $[0, 1)$ in the k disjoint intervals $I_1 = [0, p_1)$, $I_1 = [p_1, p_1 + p_2)$, $\cdots$, $I_k = [p_1 + \cdots + p_{k-1}, 1)$. Since

$$\mathbb{P}\{U \in I_j\} = p_j, \qquad j = 1, \cdots, k.$$

if $U \sim U(0, 1)$ is a random variable with uniform distribution on the interval $[0, 1]$, in order to generate a sample of size n from the distribution of X, we can generate a sample $u_1, \cdots, u_n$ from the uniform distribution on $[0, 1]$ and set

$$x_i = x^{(j)} \text{ whenever } u_i \in I_j, \qquad i = 1, \cdots, n.$$

This procedure is simple, intuitive, and fully general. It is very efficient for small values of k. This is not to say that it is used in all the cases of random simulations of finite random variables. Indeed, this method can be very difficult to implement when k is large and the partial sums $p_1 + \cdots + p_j$ are difficult to compute. Indeed, finding out which of the interval I_j does a given random number u_i belong may require many costly evaluations and comparisons in too large a number to make this method reasonable. This is why alternative algorithms have been developed in many cases, including but not limited to binomial and multinomial distributions for large values of the parameters. As we will need it later in the construction of non-central χ^2 random generators, we present an example of a random generator for the Poisson distribution.

Poisson Random Number Generator

Recall that for any $\lambda > 0$, the Poisson distribution $P(\lambda)$ with parameter λ is the distribution of an integer valued random variable satisfying

$$\mathbb{P}\{N = j\} = e^{-\lambda}\frac{\lambda^j}{j!}, \qquad j = 0, 1, 2, \cdots \tag{2.1}$$

For the Poisson distribution $P(\lambda)$, we will use a random number generator based on the inverse cdf method. To be specific, we start with a uniform random number $U \sim U(0, 1)$ and we compute $F^{-1}(U)$ as the smallest integer $n$satisfying $F(n) \le U$, and in order to compute the cdf, we use the fact that $\mathbb{P}\{N = j + 1\} = \lambda\mathbb{P}\{N = j\}/(j + 1)$. We summarize the algorithm in the following box.

- Set $F = p = e^{-\lambda}$, $j = 0$ and generate $U \sim U(0, 1)$
- Do
  - Add 1 to j
  - Replace p by $p\lambda/j$ and add p to F
- while $U > F$
- Return j

Credit Rating Migrations and (Finite State) Markov Chains

Motivated by future applications to credit default simulation, in this subsection, we consider the simulation of Markov chains with finitely many states. A stochastic process $\{X_n\}_{n=0,1,\cdots}$ is a (finite state) Markov chain if the random variables X_n can only take finitely many values, say $x^{(1)}, \cdots, x^{(k)}$, and if

$$\mathbb{P}\{X_{n+1} = x^{(j)} | X_n = x^{(i)}, X_{n-1} = x^{(i_1)}, \cdots, X_1 = x^{(i_{n-1})}\} = \mathbb{P}\{X_{n+1} = x^{(j)} | X_n = x^{(i)}\}. \tag{2.2}$$

It is usual to introduce a special notation for this conditional probability, say

$$\mathbb{P}\{X_{n+1} = x^{(j)} | X_n = x^{(i)}\} = P(i, j), \qquad i, i = 1, 2, \cdots, k. \tag{2.3}$$

P is a $k \times k$ matrix called the transition matrix of the Markov chain. Its entries are non-negative numbers and on each row, they sum up to 1.

Markov chain simulation is concerned with the generation of large numbers, say M, of independent samples of length N, say $x_{i,0}, x_{i,1}, x_{i,2}, \cdots, x_{i,N}$ for $i = 1, 2, \cdots, M$ having the same joint distribution as the dependent random variables $X_0, X_1, X_2, \cdots, X_N$ forming the Markov chain. The Markov property makes this quite easy. For each i between 1 and M, one first generates a sample $x_{i,0}$ from the distribution of the random variable X_0. Then, one generates a sample $x_{i,1}$ from the conditional distribution of X_1 given that $X_0 = x_{i,0}$. This conditional probability is given by the row of the matrix P corresponding to which $x^{(j)}$, $x_{i,0}$ is equal to. Then, one generates a sample $x_{i,2}$ from the conditional distribution of X_2 given that $X_1 = x_{i,1}$. Again, this conditional probability is given by the corresponding row of the matrix P. So it is straightforward, though tedious to articulate, to generate the desired Markov chain samples, tarting with the distribution of X_0 (which is a finite random variable) and using the rows of the transition matrix as dictated by the successive sample values of the chain appearing in the simulation.

We already saw several examples of finite state Markov chains in our discussion of trees in Chapter 1, and the above discussion will come handy when we try to price American options in these models. For the sole purpose of illustration, we discuss a concrete example borrowed from the theory of corporate bonds and credit ratings. Rating agencies such as S&P and Moody's rate debt issues. However, one often talks about the rating of a company. Since companies have many loans, it is not clear how to infer a rating for a company from the ratings of its debt issues. A typical convention is to assign to the company the rating of its senior unsecured bond. A senior unsecured debt is an unsecured debt with the highest payment priority level in the event of default.

	Aaa	Aa	A	Baa	Ba	B	C	D
Aaa	93.66	5.83	0.40	0.08	0.03	0.00	0.00	0.00
Aa	0.66	91.72	6.94	0.49	0.06	0.09	0.02	0.01
A	0.07	2.25	91.76	5.19	0.49	0.20	0.01	0.04
Baa	0.03	0.25	4.83	89.26	4.44	0.81	0.16	0.22
Ba	0.03	0.07	0.44	6.67	83.31	7.47	1.05	0.98
B	0.00	0.10	0.33	0.46	5.77	84.19	3.87	5.30
C	0.16	0.00	0.31	0.93	2.00	10.74	63.96	21.94
D	0.00	0.00	0.00	0.00	0.00	0.00	0.00	100.00

One-year ratings migration probabilities based upon bond rating data from 1981-2000. Data is adjusted for rating withdrawals. Source: Standard & Poor's.

Using the above matrix as a transition matrix P, we can simulate at will scenarios for the time evolution of the rating of a corporate bond if we know its initial rating.

We now discuss the examples of the exponential and Gaussian distributions which we motivate with credit risk modelling issues.

Arrival Model

Exponential distributions plays a crucial role in modelling *arrival processes.* We introduce notation specific to the theory of these processes with an emphasis on the modelling of credit events. However, these arrival times can also be used in modelling crashes, jumps, ... as the latter became popular in the financial engineering circles.

In the late nineties spectacular defaults on debts issued by corporations, counties, municipalities and sovereigns have created havoc in the financial market. The market of credit derivatives has emerged as a way to mitigate credit risk. We use this important new area of financial engineering as a source of examples for our discussion of random simulation. For the sake of illustration, let us consider a single bond issued by a corporation. Associated to this *loan* we introduce a random variable τ taking value in $[0, \infty)$ which we think of as the time of default on the associated payments (interests, coupons, principal). Default usually means a missed or delayed payment, a filing for bankruptcy (Chapter 7 or 11 in the US) or any form of debt restructuring to diminish financial obligations. Note that the random variable τ can take the value ∞, which corresponds to the company never defaulting. Obviously,we shall assume that $\mathbb{P}\{\tau = 0\} = 0$. The cdf is denoted by F_τ, i.e.

$$F_\tau(t) = \mathbb{P}\{\tau \le t\}, \qquad t > 0.$$

Even though the cdf characterizes completely the distribution of τ, we shall use most often the survival function G_τ defined by:

$$G_\tau(t) = 1 - F_\tau(t) = \mathbb{P}\{\tau > t\}, \qquad t > 0.$$

The function:

$$\Lambda(t) = -\log G_\tau(t) = -\log[1 - F_\tau(t)]$$

is called the hazard function, and whenever τ has a density, say f_τ, the function

$$\lambda_\tau(t) = \frac{f_\tau(t)}{G_\tau(t)}$$

is called the intensity, or hazard rate. Notice that with these definitions, we have

$$\lambda(t) = -\frac{d}{dt}\log G_\tau(t) = -\frac{F'_\tau(t)}{G_\tau(t)} \tag{2.4}$$

and consequently that

$$G_\tau(t) = e^{-\int_0^t \lambda_\tau(u)du}$$

which bears a striking resemblance with the formula giving the price of a zero coupon bond as a function of the instantaneous forward rate.

First Example: the Exponential Distribution

One of the most popular models for default arrival is to choose for τ the first time of arrival of a point process of arrivals, and the simplest one of them is obviously the Poisson process. In this case, the random variable τ has an exponential distribution $E(\lambda)$ with parameter the rate λ of the Poisson process in question. The practical significance of this parameter is clear in this case because it gives the inverse of the mean time of arrival in the sense that

$$\mathbb{E}\{\tau\} = \frac{1}{\lambda}.$$

Notice also that in this case

$$F_\tau(t) = 1 - e^{-\lambda t} \quad \text{hence} \quad G_\tau(t) = e^{-\lambda t} \quad \text{and} \quad \lambda_\tau(t) = \lambda.$$

Here is our first example of a non-uniform random generator. We assume, that the time of default of a given corporate bond is given by the first time of arrival of a Poisson process. Since pricing this corporate bond only involves the time of the first arrival of the Poisson process (and possibly the recovery rate which we ignore at this stage), and since such a distribution is well known, the problem of random simulation or Monte Carlo computation is very simple. Since the inverse cdf is the function $F_\tau^{-1}(u) = -\frac{1}{\lambda}\log(1-u)$, and since a random variable U is uniformly distributed on $[0,1]$ if and only if $1-U$ is, we can generate a sample from the distribution $E(\lambda)$ in the following way: for given a sample size n, we generate a sample $u_1, u_2, \ldots, u_n$ from the uniform distribution on the unit interval $[0,1]$, and we compute:

$$\tau_1 = -\frac{1}{\lambda}\log u_1,\ x_2 = -\frac{1}{\lambda}\log u_2, \ldots, x_n = -\frac{1}{\lambda}\log u_n.$$

2.1.3 **Random Samples from the Gaussian and Student Distributions**

Gaussian Distributions

The Gaussian distribution plays a very important role in most of the models for mathematical finance. The well known example of Samuelson's model in which prices are modelled as log-normal random variables (i.e. exponentials of Gaussian random variables) in which the famous Black-Scholes formula was derived is a case in point. Hence the need for efficient random number generators producing Gaussian samples.

We use the notation φ and Φ to denote the density and the cdf of the standard Gaussian distribution $N(0,1)$ respectively. In other words

$$\varphi(x) = \frac{1}{\sqrt{2\pi}}e^{-x^2/2}, \quad \Phi(x) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{x} e^{-z^2/2}dx, \quad x \in \mathbb{R}. \tag{2.5}$$

The appeal of the cumulative distribution function method is due to its great generality. However, its implementation can be inefficient when the c.d.f. or its inverse cannot be computed explicitly and one has to rely on numerical inversion methods. This is the case for the Gaussian distribution. Because of the crucial importance of the normal distribution, specific random generation algorithms have been designed to provide code for generators of Gaussian samples. The most popular of these algorithms is the Box-Mueller method which constructs Gaussian variate two-by-two, using the characterization of the bivariate normal distribution in polar coordinates.

Student or t Distributions

The Student distribution with ν degrees of freedom is the distribution of the ratio $\xi/\sqrt{\chi/\nu}$ where ξ and χ are independent random variables, ξ having the standard Gaussian distribution $N(0,1)$ and χ having the χ_ν chi-square distribution with ν degrees of freedom. Because of the properties of the Gamma function, the parameter ν does not need to be an integer. However, when ν is an integer, a simple random sample generator can be set up for t-variates. Indeed, if $X_1, X_2, \cdots, X_\nu$ and $X_{\nu+1}$ are independent $N(0,1)$ random variables, then $\xi = X_{\nu+1}$ and $\chi = X_1^2 + X_2^2 + \cdots + X_\nu^2$ can be used to produce a random variable $X = \xi/\sqrt{\chi/\nu}$ whose distribution is by definition, the t-distribution with ν degrees of freedom.

The density of such a distribution is given below in formula (2.6). This formula defines a probability density even when the parameter ν is not an integer. In fact, we shall need random generators

for t-distributions with non-integer parameter ν. See our discussion of the CIR model later in the chapter.

Notice that a random variable X with a t-distribution can be written in the form $X = \sigma\xi$ if we set $\sigma = \sqrt{\nu/\chi}$ which shows that the distribution is a (continuous) mixture of mean-zero Gaussian random variables, the mixing being done via a random variance. In financial jargon X appears as some form of *stochastic volatility* return. In any case, this is the trade-mark of heavy tails as we are about to see.

Simple calculus shows that the t-distribution has a density which is given by the formula

$$f_\nu(x) = \frac{\Gamma((\nu+1)/2)}{\sqrt{\nu\pi}\Gamma(\nu/2)} \frac{1}{(1+x^2/\nu)^{(1+\nu)/2}} \tag{2.6}$$

where the notation Γ is used for the Gamma function

$$\Gamma(x) = \int_0^\infty s^{x-1}e^{-s}ds. \tag{2.7}$$

The Student distribution is a cornerstone in the distribution theory developed for the statistical analysis of normal samples. Our interest in this distribution has little to do with the original reasons for its introduction. One of our reasons is the fact that this distribution has a natural generalization to the multivariate case which we will use extensively in our discussion of copulas. However, one of very attractive feature of this distribution is the fact that it has heavy tails. We shall make use of this property later in our discussion of risk control by Monte Carlo methods. Indeed, the density (2.26) decays polynomially at both ends of the real axis, and if we measure the size of the upper tail of a distribution by the survival function we see that

$$G_\nu(x) = \mathbb{P}\{X > x\} \approx \frac{\text{cst}}{x^\nu} \qquad \text{when } x \nearrow +\infty$$

with a similar behavior for the lower tail when $x \searrow -\infty$. This shows that the t-distribution belongs to the family of Generalized Pareto Distributions (GPD for short).

Gamma Distributions

The Gamma distribution with scale parameter $a > 0$ and shape parameter $b > 0$ is the distribution $\Gamma(a, b)$ with density $f_{a,b}$ defined by

$$f_{a,b}(x) = \frac{1}{\Gamma(a)b^a}x^{a-1}e^{-x/b}, \qquad x > 0 \tag{2.8}$$

and $f_{a,b}(x) = 0$ when $x \leq 0$. Recall formula (2.7) for the definition of the Gamma function Γ. It has mean ab and variance ab^2. The Gamma distribution family includes the family of exponential distributions which are obtained by setting the shape parameter a to 1. It also include the family of χ^2 (Chi-square) distributions. Indeed, if d is an integer, the χ^2_d distribution with d degrees of freedom is nothing but the Gamma distribution with shape $a = d/2$ and scale $b = 2$. Using this equality

$$\chi^2_d = \Gamma\left(\frac{d}{2}, 2\right),$$

one can extend to definition of the χ^2 distribution to non-integer number of degrees of freedom! We will need this extension for the simulation of the square root diffusion process of the CIR model.

We now discuss random number generators for Gamma distributions. Many different algorithms have been proposed. The encyclopedic volume [**?**] by De Vroye reviews most of them in detail. For the purpose of this book, we shall trust the choice made by Glasserman in [**?**]. Since the density $f_{a,b}(x)$ is bounded for $a > 1$ and unbounded for $a < 1$, not surprisingly, we are led to treat these two alternative separately.

- if $a \leq 1$
  - Set $\beta = (a + e)/e$
  - Repeat
    - Generate U_1 and U_2 independent $U(0, 1)$ and set $Y = \beta U_1$
    - If $Y \leq 1$
      - Set $Z = Y^{1/a}$
      - If $U_2 < e^{-Z}$ `accept` otherwise set $Z = -\log((\beta - Y)/a)$
      - if $U_2 \leq Z^{a-1}$ `accept`

    until `accept`
  - Return bZ
- if $a > 1$
  - Set $\alpha = a - 1$, $\beta = (a - (1/(6a)))/\alpha$, $m = 2/\alpha$, $d = m + 2$
  - Repeat
    - Generate U_1 and U_2 independent $U(0, 1)$ and set $V = \beta U_2/U_1$
    - if $mU_1 - d + V + (1/V) \leq 0$ `accept` else if $m \log U_1 - \log V + V - 1 \leq 0$ `accept`

    until `accept`
  - Return αbV

2.2 NUMERICS FOR STOCHASTIC DIFFERENTIAL EQUATIONS

Our earlier discussion of exotic and path-dependent options emphasized the need for the simulation of entire sample paths of a continuous time process. We tackle this problem when the continuous process is the solution of a stochastic differential equation, as it is often the case in the mathematical models of quantitative finance. We already encountered several times the stochastic differential equation of the geometric Brownian motion

$$dS_t = \mu S_t dt + \sigma S_t dW_t. \tag{2.9}$$

It is a particular case of more general stochastic differential equations of the form

$$dX_t = b(t, X_t)dt + \sigma(t, X_t)dW_t \tag{2.10}$$

as we can see by choosing $b(t, x) = \mu x$ and $\sigma(t, x) = \sigma x$. It was proposed as a reasonable model for the value of a firm, and hence, its simulation will help us understand Merton's model for default. For the purposes of illustration, we also consider a few of the other commonly used stochastic differential equations in the world of financial mathematics. For example, the Vasicek model used for interest rates. It reads

$$dr_t = -\lambda(r_t - \overline{r})dt + \sigma dW_t \tag{2.11}$$

and is nothing but the classical Ornstein-Uhlenbeck process of mathematical physics. In financial applications it is often used as a model for the short interest rate because of its mean reversion

feature, and because it leads to explicit closed form formulae for many derivative prices. The Vasicek model fits in the general framework of (2.10) by setting:

$$b(t,x) = -\lambda(x - \overline{r}) \qquad \text{and} \qquad \sigma(t,x) = \sigma.$$

Unfortunately, at any time $t > 0$ the solution r_t of such an equation has a positive probability to be negative: not that good for an interest rate! Using an equilibrium argument, Cox, Ingersoll and Ross suggested to use an alternative model which retains mean reversion and which remains positive at all time (as long as it initial value is positive.) This model is known as the CIR model. It reads

$$dr_t = -\lambda(r_t - \overline{r})dt + \sigma\sqrt{r_t}dW_t. \tag{2.12}$$

Notice that in this case we have:

$$b(t,x) = -\lambda(x - \overline{r}) \qquad \text{and} \qquad \sigma(t,x) = \sigma\sqrt{x}.$$

So far, we only considered scalar equations where X_t and W_t are real valued. Popular models involve multivariate stochastic differential equations for vector valued state X_t and driving Wiener process W_t. Let us consider for example Heston's model for stochastic volatility.

$$\begin{aligned} dS_t &= \mu S_t dt + \sigma_t S_t dW_t^{(S)} \\ d\sigma_t &= -\lambda(\sigma_t - \overline{\sigma})dt + \sigma_\sigma\sqrt{\sigma_t}dW_t^{(\sigma)} \end{aligned} \tag{2.13}$$

which fits in the framework of (2.10) if we set:

$$X_t = \begin{bmatrix} S_t \\ \sigma_t \end{bmatrix},\ b(t,X) = \begin{bmatrix} \mu S \\ -\lambda(\sigma - \overline{r}) \end{bmatrix},\ \sigma(t,X_t) = \begin{bmatrix} \sigma S & 0 \\ 0 & \sigma_\sigma\sqrt{\sigma} \end{bmatrix} \text{ and } W_t = \begin{bmatrix} W_t^{(S)} \\ W_t^{(\sigma)} \end{bmatrix}.$$

Notice that in all the examples above, the drift and diffusion coefficients $b(t,x)$ and $\sigma(t,x)$ are independent of t. For this reason, we restrict our presentation to time independent coefficients in what follows, even though most of the results have analogs in the case of time dependent coefficients.

In order to understand the point of view chosen in the discussion below, it is important to keep in mind that at each given time t, X_t comprises the values of economic factors and financial instruments of interest, and that depending upon the specifics of the model used for the dynamics of X_t and depending upon the acutal problem at hand, it might be necessary to construct Monte Carlo samples from the distribution of X_t at a specific time $t = T$, or of the entire history $X_{[0,T]}$ of X up until time T, or in other instances, the values $(X_{t_0}, X_{t_1}, X_{t_2}, \cdots, X_{t_N})$ of X at some fixing dates $t_0, T-1, t_2, \cdots, t_N$. The discussion below should be relevant to these different situations.

2.2.1 **Discretization Schemes for SDE's**

As before we set $T > 0$ for the time horizon, we choose an integer N for the number of time steps, and $\Delta t = T/N$ for the mesh of the regular subdivision $t_0 < t_1 < \cdots < t_N$ where $t_i = i\Delta t$ for $i = 0, 1, \cdots, N$. The goal of this section is to identify simply implementable random simulation procedures which can lead to sample realizations of the solution $\{X_t\}_t$ of (2.10) at the times t_i of the subdivision, or at least to sample realizations of reasonable approximations of the solution.

Explicit Solutions

In some very rare cases, the solution of (2.10) is given by an explicit formula. In those cases, random samples can usually be generated using the explicit form of the solution. For example, in the geometric Brownian motion case of (2.9) one has:

$$S_{(i+1)\Delta t} = S_{i\Delta t} \exp\left[(\mu - \sigma^2/2)\Delta t + \sigma[W_{(i+1)\Delta t} - W_{i\Delta t}]\right], \tag{2.14}$$

for $i = 0, 1, 2 < \cdots, N-1$, and the *exact* values of the solution can be computed at the points of the subdivision from the knowledge of the initial value S_0 and the increments $W_{(i+1)\Delta t} - W_{i\Delta t}$ of the Wiener process. Since the latter are independent random variables with the distribution $N(0, \Delta t)$, random samples for the vector $\mathbf{S} = \{S_{i\Delta t}\}_{i=0,1,\cdots,N}$ can be simulated by generating samples $\{\epsilon_i\}_{i=1,\cdots,N}$ of independent variates with the standard Gaussian distribution $N(0,1)$, and replacing the increments $W_{(i+1)\delta t} - W_{i\delta t}$ by $\sqrt{\Delta t}\epsilon_{i+1}$ for $i = 0, 1, \cdots, N-1$.

Remark. The geometric Brownian motion stochastic differential equation is not the only one to have an explicit solution in closed form. Indeed, (2.11) admits the solution

$$r_t = r_s e^{-\lambda(t-s)} + \int_s^t e^{-\lambda(t-u)} dW_u. \tag{2.15}$$

which is Gaussian as long as r_s is, and well suited for random simulations since the stochastic integrals over disjoint intervals $[s,t)$ are independent and have the Gaussian distribution $N(0, [1 - e^{-2\lambda(t-s)}]/(2\lambda))$. Moreover, and even though this fact is much less obvious, we shall see later that the CIR model is also amenable to exact simulation.

If one is only interested in random samples of the vector $\mathbf{X} = \{X_{i\Delta t}\}_{i=0,1,\cdots,N}$ an exact solution is possible in some (rare) cases where one does not have such a nice explicit solution as (2.14) or (2.15). Indeed the random variables $X_{i\Delta t}$ can be simulated inductively if one has random number generators for the conditional distributions $\mathbb{P}\{X_{i\Delta t}|X_{(i-1)\Delta t}, \cdots, X_{\Delta t}, X_0\}$. Indeed, given the initial condition X_0, one can generate samples from $X_{\Delta t}$ using a random number generator for the conditional distribution $\mathbb{P}\{X_{\Delta t}|X_0\}$. Next, for each value of the couple $(X_0, X_{\Delta t})$ so generated, one can use a random generator for the conditional distribution $\mathbb{P}\{X_{2\Delta t}|X_{\Delta t}, X_0\}$ to generate samples for the third component of $(X_0, X_{\Delta t}, X_{2\Delta t})$, etc.

Unfortunately, this situation does not occur very often, even if in most of the situations of interest the conditional distributions are simplified by the equality

$$\mathbb{P}\{X_{i\Delta t}|X_{(i-1)\Delta t}, \cdots, X_{\Delta t}, X_0\} = \mathbb{P}\{X_{i\Delta t}|X_{(i-1)\Delta t}\}$$

which follow from the Markov property of the solutions of the stochastic differential equations which we are considering in this chapter.

Euler's Scheme

The situations described above are not generic, and exact simulation is not available in general. In order to find reasonable approximations, we rewrite the stochastic differential equation (2.10) over an arbitrary interval $[s,t)$ in integral form

$$X_t = X_s + \int_s^t b(X_u)du + \int_s^t \sigma(X_u)dW_u \tag{2.16}$$

and we use Taylor expansions to approximate the integrands by simpler functions. Our first attempt is based on the approximation by a constant given by a Taylor expansion of order one. In this case, equality (2.16) becomes the approximate equality

$$X_t \approx X_s + b(X_s)(t-s) + \sigma(X_s)[W_t - W_s]. \tag{2.17}$$

This approximation suggest to define the vector $\hat{\mathbf{X}} = \{\hat{X}_i\}_{i=0,1,\cdots,N}$ inductively by $\hat{X}_0 = X_0$ and

$$\hat{X}_{i+1} = \hat{X}_i + b(\hat{X}_i)\Delta t + \sigma(\hat{X}_i)\sqrt{\Delta t}\epsilon_{i+1}, \qquad i = 0, 1, \cdots, N-1 \tag{2.18}$$

for a vector $\{\epsilon_i\}_{i=1,\cdots,N}$ of independent identically distributed (iid for short) standard Gaussian random variables. This construction of a random vector as a proxy for the solution on the subdivision is called the Euler scheme.

2.2.2 **Controlling the Errors**

One of the most important problems facing financial engineers is the computation of the expected values of functions of the entire path of underlying assets and indexes, i.e. expectations of the form $\mathbb{E}\{f(X_{[0,T]})\}$ where we use the notation $X_{[0,T]}$ to emphasize the fact that we consider the whole entire path. The idea is to approximate this function by an appropriate function say $\tilde{f}$ of the random vector produced by the Euler scheme. For example, in the case of an Asian option, ignoring the discounting factor,

$$f(X_{[0,T]}) = \left(\int_0^T X_t dt - K\right)^+$$

can be approximated by

$$\tilde{f}((\hat{X}_0, \hat{X}_1, \cdots, \hat{X}_N)) = \left(\Delta t \sum_{i=1}^N \hat{X}_i - K\right)^+.$$

Similarly, in the case of a down and out call,

$$f(X_{[0,T]}) = (X_T - K)^+ \mathbf{1}_{\{\inf\{X_t;\, 0\le t\le T\} > H\}}$$

can be approximated by

$$\tilde{f}((\hat{X}_0, \hat{X}_1, \cdots, \hat{X}_N)) = (\hat{X}_N - K)^+ \mathbf{1}_{\{\inf\{\hat{X}_i;\, 0\le i\le N\} > H\}}.$$

We discuss later the problems created by the fact that the average and the barrier crossing involve only the fixing times $t_i = i\Delta t$ in Subsection 2.2.4 below. Pay-offs of European options are simple to manipulate because they only involve the path at maturity. In the case of a plain vanilla European call,

$$f(X_{[0,T]}) = (X_T - K)^+$$

can be approximated by

$$\tilde{f}((\hat{X}_0, \hat{X}_1, \cdots, \hat{X}_N)) = (\hat{X}_N - K)^+.$$

A discretization scheme is expected to provide approximations of the solution of (2.10) in the sense that

$$\lim_{\Delta t \to 0} \mathbb{E}\{\|X_T - \hat{X}_N\|^2\} = 0$$

or even that

$$\lim_{\Delta t \to 0} \mathbb{E}\{\sup_{0 \le t \le T} \|X_t - \hat{X}_{[Nt]\Delta t}\|^2\} = 0.$$

Even though this second condition seems to be more restrictive than the first one, in many proofs, the same argument seems to imply both. In any case, we say that a scheme is of strong order α if

$$\mathbb{E}\{\|X_T - \hat{X}_N\|^2\}^{1/2} \le c(\Delta t)^\alpha$$

for some constant $c > 0$. The Euler scheme presented above is of strong order $1/2$.

2.2.3 **Other Discretization Schemes**

One of the strange consequences of Ito's stochastic integration theory is that, despite the fact that both integrands were expanded to the same order, the approximation of the first integral ended up being of order Δt while the approximation of the stochastic integral is only of order $\sqrt{\Delta t}$. A remedy could be to expand the integrand $\sigma(X_s)$ in the stochastic integral to higher order. If we do that, after re-arranging the terms we get a discretization scheme given by the formula

$$\hat{X}_{i+1} = \hat{X}_i + b(\hat{X}_i)\Delta t + \sigma(\hat{X}_i)\sqrt{\Delta t}\epsilon_{i+1} + \frac{1}{2}\sigma'(\hat{X}_i)\sigma(\hat{X}_i)\Delta t(\epsilon_{i+1}^2 - 1). \tag{2.19}$$

This discretization procedure is called Milstein's scheme. Notice that it requires the differentiability of the diffusion coefficient. As expected, it provides a better approximation than the Euler scheme since it is of strong order 1. Notice that this improvement comes at a cost since in a typical Monte Carlo computation, the extra term in (2.19) needs to be computed a very large number of times, one time for each step of the subdivision, and one time for each Monte Carlo sample.

2.2.4 **Warning: Random Simulation can be a Touchy Business**

We consider several examples to illustrate the care which needs to be taken in using random simulations based on discretization schemes when dealing with subtle properties of continuous time models.

Time of Default as a First Hitting Time

Let us assume that the dynamics of the value of a firm are given by a stochastic differential equation like (2.10), and that default on its debt is believed to occur at the (random) time τ defined as

$$\tau = \inf\{t > 0;\ X_t < L(t)\}$$

the first time the value of the firm drops below a level $L(t)$ which could possibly depend upon time. There are two obvious sources of error in replacing the computation of the infimum of the times at which the continuous time process X is below the level L. First, a discretization scheme such as Euler or Milstein can only provide samples of X at pre-assigned times t_i, and the time τ can be wrongly estimated if X is not below L at these specific times. The second source of error is the fact that the *discrete* dynamics implemented by the discretization scheme is only an approximation of the theoretical continuous dynamics given by the stochastic differential equation defining the model.

These two shortcomings have to be kept in mind when dealing with discrete scheme. We address them both in the case of the CIR model.

Despite its high level of abstraction, continuous time finance is very powerful in providing sophisticated models and tools of analysis to resolve the challenges posed by complex instruments. Fortunately, in a practical situation, the indenture of a contract specifies a set of fixing dates at which the values of the instruments underlying the contract are monitored, and discrete time considerations can be substituted for the continuous time models. Running averages are often computed as plain averages instead of integrals, barrier crossing is checked on a finite set of dates. American options can only be exercised on certain dates (and are called Bermudan options)

Sample Paths of the CIR model

This last example is presumably the most disturbing. The CIR model (2.12) for the short interest rate was first derived from an equilibrium argument. However, its popularity is mostly due to the fact that it offers an alternative to the Vacicek model (2.11) resolving the issue of its major shortcoming. Indeed, it is well known that

$$r_t > 0 \qquad \text{for all } \ t > 0$$

whenever $r_0 > 0$ and $\sigma^2 \leq 2\overline{r}/\lambda$. There is an intuitive reason for that (a mathematical proof would be more involved). Notice that whenever the value of r_t gets dangerously close to zero, the mean reverting drift will pull r_t toward $\overline{r}$. However, this is not enough to keep r_t from becoming negative. Indeed, the Vasicek's model has the same mean reverting drift, and nevertheless at any given time, it can be negative with positive probability. Things are different with the CIR model. Indeed, the instantaneous standard deviation of the a solution of the square root equation (2.12) is proportional to the square root of its value, so the closer the solution gets to 0, the smaller the standard deviation of the Gaussian kick produced by the dW_t term, and if $\sigma^2 \leq 2\overline{r}/\lambda$, then the random kick will not be larger enough to overcome the effect of the mean reverting drift, and it will not be able to force r_t to cross to the negative side.

The important thing to keep in mind when computing with this model, is that the positivity property is a feature of the continuous time model. It is because r_t approaches 0 in a continuous fashion that the decay of the volatility can take place at the right pace and control the size of the random kicks of the noise term dW_t. Things are different in the discrete case. If one considers a plain implementation of the Euler scheme and one generates a real sequence from the recursive definition

$$X_{i+1} = X_i - \lambda(X_i - \overline{r})\Delta t + \sigma\sqrt{X_i}\epsilon_{i+1}$$

for an i.i.d. sequence $\{\epsilon_i\}_i$ of standard Gaussian random variables, and $X_0 = r_0 > 0$, then at any time i, the probability that the next value becomes negative is strictly positive. Indeed, conditioned on the knowledge of X_i, X_{i+1} is a Gaussian random variable with mean $X_i - \lambda(X_i - \overline{r})\Delta t$ and variance $\sigma^2 X_i$, and consequently, it has a positive probability to be negative. Several tricks are used to handle occurrences of negative values, the most frequent one being to replace any negative value trying to appear by a small positive number chosen a priori.

Remark 1. A word of caution is needed if one wants to recast the square root stochastic differential equation (2.12) in the framework of the general theory of stochastic differential equations. Indeed, the square root nature of the diffusion coefficient is so singular that the CIR model does not satisfy the usual assumptions under which existence and uniqueness of a solution is guaranteed. Indeed, the coefficient or the CIR stochastic differential equation are not Lipschitz. However, it was proved

a long time ago by Feller that not only the square root stochastic differential equation (2.12) has a unique strong solution, but this solution remains positive at all times if its initial value is positive, i.e. if $r_0 > 0$ and if $\sigma^2 \leq 2\overline{r}/\lambda$.

Remark 2. Once a grid $t_0 < t_1 < \cdots < t_N$ of time points has been fixed, exact simulation is possible for the transition probability of the Markov process solving the stochastic differential equation (2.12) can be determined explicitly. It is given by a non-central χ^2 (Chi-square) distribution. See below for details.

Exact Simulation for the CIR Process

As before, we choose a discrete grid $0 = t_0 < t_1 < \cdots < t_N$ of time points, and we try to generate n samples of the random vector $(r_{t_1}, \cdots, r_{t_N})$ given that $r_0 = r$ is given. Despite the above horror stories, the situation is not as dire as the above discussion could lead to believe. Indeed, like the Vasicek model, the CIR model is amenable to exact simulation. As we are about the demonstrate, the formulae are not as simple, and the simulation is more involved. The main theoretical result was proven by Feller. It states that if $s < t$, conditioned on the knowledge of r_s, the distribution of r_t is a multiple of a non-central χ^2 distribution. Specifically:

$$r_t | r_s \sim \frac{1 - e^{-\lambda(t-s)}}{4\lambda} \sigma^2 \chi^2_d(\alpha) \tag{2.20}$$

where the number of degrees of freedom d is given by

$$d = \frac{4\lambda\overline{r}}{\sigma^2}$$

and the non-centrality parameter α is given by

$$\alpha = \frac{4\lambda e^{-\lambda(t-s)}}{\sigma^2(1 - e^{-\lambda(t-s)})} r_s.$$

We can now put all the elements reviewed above together to describe the exact simulation algorithm for CIR samples over a grid $t_1 < t_2, \cdots < t_N$ starting from r_0.

- For $i = 1, 2, \cdots, n$ do
  - For $j = 1, 2, \cdots, N$ do
    - Set $c_j = \sigma^2(1 - e^{-\lambda(t_j - t_{j-1})})/(4\lambda)$
    - Set $\alpha_j = 4\lambda e^{-\lambda(t_j - t_{j-1})} r_{i,j-1}/(\sigma^2(1 - e^{-\lambda(t_j - t_{j-1})}))$
    - Generate $\chi \sim \chi^2_d(\alpha_j)$
    - Set $r_{i,j} = c_j \chi$

Obviously, this exact simulation algorithm is possible only if one can simulate samples from non-central χ^2 distributions. This is done on the basis of the following two remarks.

- If ν is a Poisson random variable with mean $\alpha/2$, and if conditioned on the value of ν, χ is a χ^2 random variable with $d + 2\nu$ degrees of freedom, i.e. if $\chi|\nu \sim \chi^2_{d+2\nu}$, then the unconditional distribution of χ is the χ^2 distribution with d degrees of freedom and non-centrality α, i.e. $\chi \sim \chi^2_d(\alpha)$. This first remark reduces the problem of the generation of non-central χ^2 random variables to the generation of regular χ^2 distributions.
- The remaining problem is to simulate χ^2_d random variates when the number of degrees of freedom is not necessarily an integer. Indeed, when d is an integer, χ^2_d is merely the distribution of the sum of the squares of d independent $N(0,1)$ random variables. In general, one uses the fact that χ^2_d is a Gamma distribution with shape parameter $a = d/2$ and scale parameter $b = 2$.

2.3 MONTE CARLO COMPUTATIONS

Monte Carlo computations have become the tool of choice in quantitative finance. They are used for the simulation analysis based on scenarios generated from a stochastic model (typically given by a system of stochastic differential equations), but most often they are used to compute numerically probabilities and expectations which cannot be easily evaluated by other methods, e.g. quadratures, solving partial differential equations, etc. Most Monte Carlo computations are rooted in the following result known as law of large numbers. If $\mu = \mathbb{E}\{Z\}$ is the expected value of an integrable random variable Z, then for any sequence $\{Z_n\}_n$ of independent random variables having the same distribution as Z, it holds almost surely that:

$$\mathbb{E}\{Z\} = \lim_{n\to\infty} \frac{1}{n}\sum_{j=1}^{n} Z_j \tag{2.21}$$

In most cases, Z appears as a deterministic function, say f, of random factors $X_0, X_1, \ldots, X_N$ which can be grouped in a vector $\mathbf{X} = (X_0, X_1, \ldots, X_N)$ so that $Z = \psi(\mathbf{X})$. We shall see several instances of this structure later on. In any case, if we have a way to generate random samples for the distribution of Z, for any such sample $z_1, \ldots, z_n$, the sample average

$$\frac{z_1 + \cdots + z_n}{n}$$

gives an approximation of the expectation $\mathbb{E}\{Z\}$, and the larger the sample, the better the approximation.

Remark. If the distribution of Z has a density, say $f(z)$, then:

$$\mathbb{E}\{Z\} = \int z f(z) dz$$

which explains why Monte Carlo computations are often used for the numerical computation of integrals outside of any probabilistic context.

No approximation can be reasonably used if one does not have a sense of the size of the error incurred. In the present situation, the error is

$$\begin{aligned}\mathbb{E}\{Z\} - \frac{1}{n}\sum_{j=1}^{n} Z_j &= \frac{Z_1 + \cdots + Z_n - n\mathbb{E}\{Z\}}{n} \\ &= \frac{1}{\sqrt{n}} \frac{Z_1 + \cdots + Z_n - n\mathbb{E}\{Z\}}{\sqrt{n}}\end{aligned}$$

The central limit theorem tells us that (at least when the common distribution of the Z's has a second moment)

$$\mathbb{E}\{Z\} - \frac{1}{n}\sum_{j=1}^{n} Z_j \approx \frac{1}{\sqrt{n}}\sigma_Z \xi \tag{2.22}$$

with $\xi \sim N(0,1)$. Here σ_Z denotes the common variance of the Z_j's. So as announced, the error goes to zero when the sample size n increases without bound. However, we can derive much more

information from this result. Indeed, we learn that the rate of convergence is of the order $n^{-1/2}$. So the size of the error can be controlled by the choice of a large sample size. Furthermore, even though the size of ξ cannot be predicted, after all this is a random variable, we know that it will be typically in the range $[-2, +2]$, and the only remaining way to control the error is to lower the variance of Z. We shall come back to this idea later in the chapter.

A First Example: Pricing a European Option

Financial derivatives are based on well defined cash flows. The payments are made on specific dates in the future, and the pay-offs are uncertain when viewed from time $t = 0$. Their prices are most often given by risk neutral expected values of discounted pay-offs, and Monte Carlo computations offer the easiest way to get an approximation of these prices. Let us give a simple example.

For the sake of illustration, we price a European call at the money, with maturity one year on an underlying stock valued at US $ 50. We also assume that the annualized volatility is $\sigma = 30\%$ and that the yearly continuously compounded interest rate is $r = 3.5\%$. The price C of this option can be derived from the value of an expectation since

$$C = e^{-.035}\mathbb{E}\{(S_1 - 50)^+\}$$

where S_1 is a log-normal random variable with parameters $\log 50 + .035 - .3^2/2$ for the mean and $.3^2$ for the variance. So, assuming that we have a random number generator for the standard Gaussian distribution, we can price this option without appealing to the Black-Scholes formula. Indeed we have

$$Z = (S_1 - 50)^+ \qquad \text{with} \qquad S_1 = 50e^{.035-.3^2/2+.3\epsilon}$$

where ϵ is a mean-zero variance-one Gaussian random variable. So for any value of the integer n we can generate a sample $\epsilon_1, \cdots, \epsilon_n$ of size n of standard Gaussian variates, compute the corresponding sample

$$z_1 = (50e^{.035-.3^2/2+.3\epsilon_1} - 50)^+, \cdots\cdots, z_n = (50e^{.035-.3^2/2+.3\epsilon_n} - 50)^+$$

of Z's, and finally compute the average

$$\frac{z_1 + \cdots + z_n}{n}$$

which gives the desired approximation of the call option price. Obviously, the above cannot be a very convincing example. Black-Scholes formula offers a more efficient way to compute the value of the option. Moreover, the value obtained by evaluating this formula is exact (as long as the log-normal model holds). This example is merely a first instance of Monte Carlo computation of a derivative price.

2.3.1 Monte Carlo Computation from A Discretization: Overall Error

We now consider the typical case of the use of the approximation provided by the Euler scheme for the purpose of pricing an option by a Monte Carlo method. For the sake of simplicity we discuss only the case of an European contingent claim with pay-off $f(X_T)$, and we ignore the discounting factor. Because of the European nature of the exercise, we can take $\tilde{f} = f$.

We need to take into account the errors produced separately by the two steps of the procedure

- replacing the correct pay-off $f(X_T)$ by the approximation $f(\hat{X}_N)$ provided by the Euler scheme;
- replacing the expectation $\mathbb{E}\{f(\hat{X}_N)\}$ by the Monte Carlo average $(f(\hat{X}_N^{(1)}) + \cdots + f(\hat{X}_N^{(n)}))/n$;

So, computing the root mean square error (RMSE for short) we get

$$\begin{aligned}\text{RMSE} &= \mathbb{E}\left\{\mathbb{E}\{f(X_T)\} - \frac{1}{n}(f(\hat{X}_N^{(1)}) + \cdots + f(\hat{X}_N^{(n)}))\right\}^{1/2} \\ &= |\mathbb{E}\{f(X_T)\} - \mathbb{E}\{f(\hat{X}_N)\}| \\ &\qquad + \mathbb{E}\left\{|\mathbb{E}\{f(\hat{X}_N)\} - \frac{f(\hat{X}_N^{(1)}) + \cdots + f(\hat{X}_N^{(n)})}{n}\right\}^{1/2} \\ &\leq c_1\sqrt{\Delta t} + \frac{c_2}{\sqrt{n}}.\end{aligned}$$

The computing time is proportional to the number n of Monte Carlo simulations. Moreover, this computing time is also proportional to the number N of time steps in the discretization scheme. Consequently, it is inversely proportional to the mesh Δt. Hence we see that the computing time c is of the form $c = c_3 n/\Delta t$ for some constant c_3. Using this *computer time budget constraint* we can rewrite the estimate for RMSE as

$$\text{RMSE} \leq c_1\sqrt{\Delta t} + \frac{c_2\sqrt{c_3}}{\sqrt{c\Delta t}}$$

and minimizing the right hand side with respect to $\sqrt{\Delta t}$ we find that the critical value for the latter is a multiple of $c^{-1/4}$, and that the minimum upper bound for the RMSE is

$$\text{RMSE} \leq c' c^{-1/4}$$

which shows that to improve the RMSE by a factor of k we need to increase the computing time by a factor of k^4!

2.4 More Monte Carlo Computations

2.4.1 Stochastic Volatility Models

Stochastic volatility models comprise an equation for the dynamics of an underlying stock or index of the standard form

$$dS_t = S_t[\mu dt + \sigma(t) dW_t^{(S)} \tag{2.23}$$

but the volatility $\sigma(t)$ appearing in such an equation is not assumed to be a deterministic constant. Instead, it is assumed to have its own stochastic dynamics. To be more specific, $\sigma(t)$ is assumed to be of the form $\sigma(t) = \sigma(Y_t)$ for some function $y \hookrightarrow \sigma(y)$ of a random factor Y_t whose dynamics are given by the solution of a stochastic differential equation of the form

$$dY_t = p(t, Y_t)dt + q(t, Y_t)dW_t^{(Y)} \tag{2.24}$$

Hull and White proposed an Ornstein-Uhlenbeck for Y_t in which case the function σ has to take only positive values. On an intuitive level, stochastic differential equations (2.24) for which the solution

remains positive are very attractive proposal because on can use the volatility itself as a factor by choosing $\sigma(y) = y$. A natural proposal is to use the geometric Brownian motion stochastic differential equation (2.9) in lieu of (2.24). One of the shortcomings of this model is that it does not have the mean reversion property which is empirically observed. Another popular proposal is the Heston model for which the dynamics (2.24) of Y_t are given by the square-root diffusion (2.12) which is typically used as a model for the short interest rate in the CIR model. The Wiener processes $W_t^{(S)}$ and $W_t^{(Y)}$ appearing in equations (2.23) and (2.24) can be assumed to be independent for the convenience of computations. However, in order to give an account of the leverage effect (observed negative empirical correlation between the returns and the changes in volatility) these Wiener processes are often assumed to be negatively correlated.

Random simulations of stochastic volatility models can be done with the tools developed up to now. The major difference is obviously the dimension of the model: we need to generate samples for both the volatility σ_t (or the factor Y_t) and the underlying asset S_t. Euler schemes can be used for both. The only point in need of attention is the possible correlation between the two Wiener processes. This problem is considered in full generality in the subsection. For the time being, it is enough to know that the innovations $\epsilon_{i+1}^{(S)}$ and $\epsilon_{i+1}^{(Y)}$ can be obtained from two independent standard Gaussian random variates $\epsilon_{i+1}^{(1)}$ and $\epsilon_{i+1}^{(2)}$ from the formula

$$\epsilon_{i+1}^{(S)} = \epsilon_{i+1}^{(1)}, \quad \text{and} \quad \epsilon_{i+1}^{(Y)} = \rho\epsilon_{i+1}^{(1)} + \sqrt{1-\rho^2}\epsilon_{i+1}^{(2)}.$$

2.4.2 **Multivariate Distributions**

We prepare the ground for our discussion of portfolio risk management by discussing multivariate distributions.

The Multivariate Gaussian Distribution

We say that a m-variate random vector $\mathbf{X} = (X_1, \cdots, X_m)$ has the m-variate Gaussian distribution with mean $\mu = (\mu_1, \cdots, \mu_m)$ and variance/covariance matrix Σ if any linear combination $\alpha_1 X_1 + \cdots + \alpha_m X_m$ is a univariate Gaussian random variable with mean

$$\alpha \cdot \mu = \alpha_1 + \cdots + \alpha_m \mu_m$$

and variance

$$\alpha^t \Sigma \alpha = \sum_{i,j=1}^{m} \alpha_i \alpha_j \Sigma_{i,j} \tag{2.25}$$

where we use the notation α for the vector $\alpha = (\alpha_1, \cdots, \alpha_m)$.

Here $m \geq 1$ is an integer, and Σ is a $m \times m$ symmetric matrix. This distribution is denoted by $N_m(\mu, \Sigma)$. This distribution has a density (the joint density of the X_i's) when the matrix Σ is of full rank (i.e. non-degenerate). However, we shall no need and/or use the form of this density. We shall only need a way to generate samples of random vectors with a given multivariate Gaussian distribution. A general random generator procedure can be deduced from the following properties of the Gaussian distributions.

- The components $X_1, \cdots, X_m$ are standard $N(0,1)$ Gaussian random variables when $\mu = (0, \cdots, 0)$ and Σ is the $m \times m$ identity matrix I_m.. To generate a sample of size N from this

special m-variate distribution, one generates a sample of size mN from the distribution $N(0,1)$ and one regroup the entries m by m;
- $\mathbf{X} \sim N_m(\mu, \Sigma)$ if and only if $\mathbf{X} - \mu \sim N_m(\mathbf{0}, \Sigma)$
- $\mathbf{X} \sim N_m(\mathbf{0}, \Sigma)$ if $\mathbf{X} = \Sigma^{1/2}\mathbf{Z}$ and $\mathbf{Z} \sim N_m(\mathbf{0}, I_k)$

In other words, in to generate a sample of size N from the m-variate distribution $N_m(\mu, \Sigma)$, one generates a sample $\mathbf{Z}_1, \cdots, \mathbf{Z}_N$ from the distribution $N_m(\mathbf{0}, I_m)$ as explained in the first bullet point above, and then one sets

$$\mathbf{X}_1 = \mu + \Sigma^{1/2}\mathbf{Z}_1, \cdots, \mathbf{X}_N = \mu + \Sigma^{1/2}\mathbf{Z}_N.$$

Application: Pricing a Spread Option

A European call (resp. put) spread option on two assets $S_1(t)$ and $S_2(t)$ is a plain vanilla European call (resp. put) option on the difference $S_1(t) - S_2(t)$ of the two assets. Such options are of interest in the *presence of correlation* between the two assets. Hence, the price of such an option with maturity T and strike K is given by: the risk neutral expectation

$$e^{-rT}\mathbb{E}\{(S_1(t) - S_2(t) - K)^+\}, \qquad \text{(resp.} \quad e^{-rT}\mathbb{E}\{(K - (S_1(t) - S_2(t)))^+\}).$$

Notice that, even in the case of assets log-normally distributed, there no closed form formula for such prices, except in the case $K = 0$ of an option to exchange an asset for another. Indeed the difference of two correlated log-normally distributed is not log-normally distributed, and we do not have simple expressions for its distribution and the values of the above expectations. The special case $K = 0$ was treated by Margrabe who derived a formula of the Black-Scholes type for the price of exchange options.

For the purposes of the present discussion, we assume that the dynamics of the two assets are given by geometric Brownian motions

$$dS_i(t) = S_i(t)[\mu_i dt + \sigma_i dW_t^{(i)}], \qquad i = 1, 2$$

where the fact that the correlation coefficient of the two Wiener processes is ρ is symbolically denoted by

$$[dW^{(1)}, dW^{(2)}]_t = \rho dt,$$

for some $\rho \in [-1, +1]$. So when we say that the two assets are correlated, we mean that the infinitesimal returns are correlated and

$$[\frac{dS^{(1)}}{S^{(1)}}, \frac{dS^{(2)}}{S^{(2)}}]_t = [d\log S^{(1)}, d\log S^{(2)}]_t = \sigma_1\sigma_2[dW^{(1)}, dW^{(2)}]_t = \sigma_1\sigma_2\rho t.$$

This formal notation is in fact a rigorous statement in the framework of Ito's stochastic calculus. Since

$$S_i(T) = S_i(0)\exp[(\mu_i - \sigma_i^2/2)T + \sigma_i W_T^{(i)}], \qquad i = 1, 2$$

and since the two random variables $W_T^{(1)}$ and $W_T^{(2)}$ are mean-zero Gaussian random variables with variances T and correlation ρ, we produce a sample of size n from the difference $D(T) = S_1(T) - S_2(T)$ by generating a sample

$$(\epsilon_1^{(1)}, \epsilon_1^{(2)}), \cdots, (\epsilon_n^{(1)}, \epsilon_n^{(2)})$$

from the bivariate Gaussian distribution $N_2(\mu, \Sigma)$ with mean $\mu = (0, 0)$ and variance/covariance matrix

$$\Sigma = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}$$

and computing

$$D_i = S_1(0)\exp[(\mu_1 - \sigma_1^2/2)T + \sigma_1\sqrt{T}\epsilon_i^{(1)}] - S_2(0)\exp[(\mu_2 - \sigma_2^2/2)T + \sigma_2\sqrt{T}\epsilon_i^{(2)}],$$

for $i = 1, 2, \cdots, n$. Finally the Monte Carlo approximation of the price of the spread option is given by the average

$$\frac{e^{-rT}}{n}\sum_{j=1}^{n} D_j.$$

The Multivariate t-Distribution

The m-variate Student (or t-) distribution with ν degrees of freedom is the distribution of the m-dimensional random vector $\mathbf{X} = \xi/\sqrt{\chi/\nu}$ where $\xi = (\xi_1, \cdots, \xi_m)$ is a multivariate Gaussian vector with distribution $N_m(0, \Sigma)$, and where χ has the χ_ν chi-square distribution with ν degrees of freedom, ξ and χ being statistically independent. The parameters of the distributions are 1) the dimension m, 2) the number of degrees of freedom ν, and 3) the variance/covariance matrix Σ of ξ.

Notice that such a random vector $\mathbf{X}$ can be written in the form

$$\mathbf{X} = \frac{A\tilde{\xi}}{\sqrt{\chi/\nu}} = A\tilde{\mathbf{X}}$$

where $\tilde{\xi}$ is a mean-zero Gaussian random vector with independent components, and where A is any matrix satisfying $AA^t = \Sigma$. Notice that the random vector $\tilde{\mathbf{X}}$ has the multivariate t-distribution with ν degrees of freedom and its variance/covariance matrix is the m-dimensional identity matrix. Its marginal components have univariate t-distributions, they are uncorrelated but *they are not independent!*

As in the univariate case, the density of the t-distribution can be computed in closed form.

$$f_{m,\nu(x)} = \frac{\Gamma((\nu+m)/2)}{(\nu\pi)^{m/2}\Gamma(\nu/2)\text{det}(\Sigma)^{1/2}}\frac{1}{(1+\mathbf{x}^t\Sigma^{-1}\mathbf{x}/\nu)^{(\nu+m)/2}} \tag{2.26}$$

Our interest in the multivariate t-distribution is triggered by the discussion of copulas.

2.4.3 **VaR Computations for Risk Analysis**

The risk of a portfolio of holdings is often quantified by the so-called Value at Risk (VaR for short) of the profit and loss (P&L for short) distribution of the future values of the portfolio. Let us denote by P_t the portfolio value at time t, and for a given time horizon Δt, VaR is defined for each risk tolerance α, . For the sake of definiteness, we shall choose $\alpha = 1\%$ in the following discussion. The value at risk at the level α over the period $[t, t + \Delta t]$ is defined as the negative of the α-quantile of the distribution of the log-return of the portfolio over that period. In other words,

$$\mathbb{P}\{\log\frac{P_{t+\Delta t}}{P_t} \le -VaR_\alpha\} = \alpha$$

Despite widespread use, the value at risk of a portfolio is not a satisfactory measure of the risk of the portfolio. A naive explanation is that VaR gives the minimal loss occurring with a fixed probability (i.e. a loss of size at least VaR_α occurs with probability α) without giving the actual size of these losses. However, a more serious shortcoming of VaR as a measure of risk is that it is not *sub-additive* and hence, does not encourage diversification.

For this reason another quantity is often used to measure the risk of a portfolio. It is called the expected shortfall distribution and it is defined as the expected loss assuming (conditioning by) losses of size greater than VaR_α.In other words, the expected shortfall, ES_α is given by:

$$ES_\alpha = \mathbb{E}\{L|L > VaR_\alpha\} = \frac{1}{\alpha}\int_{x>VaR_\alpha} x\, dF_L(x). \tag{2.27}$$

if we denote by L the loss and by F_L its cdf. We proceed to show how to compute these measures of risk by Monte Carlo methods. In so doing we shall illustrate the dramatic influence of the loss distribution on the values of these measures.

Monte Carlo Estimations

Given a sample $L_1, \cdots, L_n$ of Monte Carlo realizations of the possible losses of a portfolio over a given period the value at risk at level α is estimated by the empirical 100α percentile of the sample. It is obtained by ordering the losses

$$L_{(1)} < L_{(2)} < \cdots < L_{(n)}$$

and by using $\widehat{VaR_\alpha} = L_{([n\alpha])}$ as estimate. We use the notation $[x]$ for the integer part of x. Once the value at risk is estimated, the expected shortfall is estimated by replacing the conditional expectation entering in its definition by the corresponding empirical analog. In other words

$$\widehat{ES_\alpha} = \frac{1}{[n\alpha]}\sum_{i=1}^{[n\alpha]} L_{(i)}$$

2.4.4 **Copulas and more VaR Computations**

By polarization of the expression (2.25) for the variance, it is plain to see that if $\mathbf{X} = (X_1, \cdots, X_k)$ is a Gaussian vector with variance/covariance matrix Σ, then

$$\Sigma_{i,j} = \text{cov}\{X_i, X_j\}$$

and that conversely, it is possible to generate samples of couples of random variables (X_1, X_2) jointly Gaussian, with marginal distributions $N(\mu_1, \sigma_1^2)$ and $N(\mu_2, \sigma_2^2)$ and correlation ρ. Indeed, these prescriptions completely determine the joint distribution of (X_1, X_2) which is necessarily $N_2((\mu, \Sigma)$ with

$$\mu = (\mu_1, \mu_2), \qquad \text{and} \qquad \Sigma = \begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix}$$

So correlating Gaussian random variables does not seem to be too difficult. How about coupling together random variables which are not necessarily Gaussian? The problem is much more difficult, and we shall address only elements of partial answers to these difficult problems.

NOTES & COMPLEMENTS

A detailed discussion of the various discretization schemes and a thorough analysis of their respective convergence properties can be found in the book by Kloeden and Platen [**?**].

Tree methods can be used to simulate and price instruments on underlying assets with stochastic volatility. Unfortunately, the size of the tree becomes prohibitive very fast and Monte Carlo methods are usually preferred.

For an exhaustive review of the properties of spread options and detailed discussions of the various approximation proposed for their prices and their Greeks, and in particular for a proof of Margrabe's formula, the interested reader is referred to the review article [**?**] of Carmona and Durrleman.

The statistical concept of copula has seen a recent renewal of interest due to the growing need for understanding and controlling the dependencies between non-Gaussian random variates. A clean mathematical account of copulas can be found in Nelsen's lecture notes [**?**]. The `S-Plus` implementation together with financial applications are presented in Carmona's textbook [**?**].

3

Monte Carlo Valuations based on Dynamic Programming Principle

In this chapter we review some of the most popular methods of American option pricing by Monte Carlo techniques. We believe that despite the enormous popularity of some of these algorithms, major confusion still persists about the very nature of these algorithms. The first algorithm we report uses Monte Carlo techniques to compute plain expectations. The other ones are written in terms of conditional expectations and the Monte Carlo computations are coupled with non-parametric regression ideas.

3.1 THE TSITSIKLIS-VAN ROY ALGORITHM

The version of the Tsitsiklis-van Roy algorithm which we present is based on the form (1.18) of the dynamic programming backward induction written in terms of the continuation value functions constructed from the discounted transition kernel $\tilde{P}_n$ defined earlier in (**??**).

The algorithm is based on the following idea: in (1.18), replace the discounted transition kernels $\tilde{P}_n$ by finite dimensional approximations $\Pi_n \tilde{P}_n$, and use the resulting approximate value of continuation functions $\tilde{q}_n$ instead of q_n. According to this prescription, $\tilde{q}_n$ is defined recursively by:

$$\begin{cases} \tilde{q}_{N-1} = \Pi \tilde{P}_{N-1} \varphi \\ \tilde{q}_n = \Pi \tilde{P}_n \max\{\varphi, \tilde{q}_{n+1}\}, \qquad n = 0, 1, \cdots, N-2 \end{cases} \tag{3.1}$$

As explained earlier, the rationale for using this form of the backward induction is that the discounted transition kernel $\tilde{P}_n$'s enters linearly in the equations, and following them with linear projections is very convenient: the projection operator Π acts directly on the discounted transition operator $\tilde{P}_n$, so Π and $\tilde{P}_n$ can be bundled together to perform the induction. This is not possible if the value function form of the dynamic programming principle is used as Π and $\tilde{P}_n$ would be on both sides of the maximum operator.

It is important to emphasize that the subsequent algorithm is not based on Monte Carlo estimation of regression functions. Indeed, as we are about to see, the dynamic programming principle is left written in terms of expectations only, as opposed to conditional expectations, and plain Monte Carlo computations of these expectations is possible very much in the spirit of the original Monte Carlo prescriptions.

3.1.1 **Numerical Algorithm**

In this form of the algorithm, a set of basis (feature) functions $\{f_\ell\}_\ell$ is chosen and all the value and continuation functions are approximated by linear combinations of the first L basis elements:

$$\sum_{\ell=1}^{L} \alpha_\ell f_\ell$$

where the coefficients $\{\alpha_\ell\}_{\ell=1,\cdots,L}$ are chosen to minimize some form of sum of squared errors. The choice of the functions f_ℓ is made for convenience of numerical computations, and possibly to tailor the numerics to the pay-off function φ, and the specifics of the dynamics of the underlying assets as given by the Markov chain $\{X_n\}_n$.

A natural procedure would be to choose a Hilbert space structure and for each function f, to use its orthogonal projection on the span of $\{f_1, \cdots, f_L\}$ as approximation. Two important features of (3.1) need to be considered. The natural Hilbert structure which should be used at time n is given by the inner product

$$< g, h >_n = \mathbb{E}\{g(X_n)h(X_n)\}.$$

Unfortunately, this Hilbertian structure changes with n and a fixed set of functions f_ℓ is not likely to form an orthonormal system simultaneously for all inner products. Consequently, Gramm-Schmidt ortho-normalization needs to be used to compute the orthogonal projections providing the approximations. This procedure requires computing the matrix $\Gamma^{(n)} = \left[\Gamma^{(n)}_{\ell\ell'}\right]_{\ell,\ell'}$ with

$$\Gamma^{(n)}_{\ell\ell'} = \mathbb{E}\{f_\ell(X_n)f_{\ell'}(X_n)\}.$$

As we know from classical least squares theory, the inverse $\Gamma^{(n)-1}$ and its square root $\Gamma^{(n)-1/2}$ play a crucial role. Indeed, the functions $f^{(n)}_\ell$ defined by:

$$f^{(n)}_\ell = \sum_{\ell'=1}^{L} [\Gamma^{(n)-1/2}]_{\ell\ell'} f_{\ell'} \tag{3.2}$$

form an orthonormal system, and the orthogonal projection $\Pi^{(n)} f$ of a function f onto the span of $f_1, \cdots, f_L$ can be written as

$$\Pi^{(n)} f = \sum_{\ell=1}^{L} < f, f^{(n)}_\ell > f^{(n)}_\ell = \sum_{\ell=1}^{L} \alpha^{(n)}_\ell f_\ell$$

where the $\alpha^{(n)}_\ell$ are computed from (3.2). It holds

$$\alpha^{(n)}_\ell = \sum_{\ell'=1}^{L} [\Gamma^{(n)-1}]_{\ell\ell'} < f, f_{\ell'} >_n$$

which has a nice expression in terms of expectations over the process when f is actually of the form $\tilde{P}_n f$. Indeed, in this case we have

$$\alpha^{(n)}_\ell = D(n, n+1) \sum_{\ell'=1}^{L} [\Gamma^{(n)-1}]_{\ell\ell'} \mathbb{E}\{f_{\ell'}(X_n)\varphi(X_{n+1})\}$$

We summarize the backward induction in terms of the coefficients $\alpha_\ell^{(n)}$ of the decomposition of the approximation of the continuation value function on the basis functions f_ℓ introduced above:

- Compute $\alpha_\ell^{(N-1)} = D(N-1,N)\sum_{\ell'=1}^{L}[\Gamma^{(N-1)-1}]_{\ell\ell'}\mathbb{E}\{f_{\ell'}(X_{N-1})\varphi(X_N)\}$ for $\ell = 1,\cdots,L$;
- for $n = N-2,\cdots,1,0$ compute

$$\alpha_\ell^{(n)} = D(n,n+1)\sum_{\ell'=1}^{L}[\Gamma^{(n)-1}]_{\ell\ell'}\mathbb{E}\left\{f_{\ell'}(X_n)\max\left\{\varphi(X_{n+1}),\sum_{\ell''=1}^{L}\alpha_{\ell''}^{(n+1)}f_{\ell''}(X_{n+1})\right\}\right\}$$

  for $\ell = 1,\cdots,L$;

The price of the option is given by

$$\max\{\varphi(x_0),\tilde{q}_0(x_0)\},$$

where the value of $\tilde{q}_0(x_0)$ is given by:

$$\sum_{\ell=1}^{L}\alpha_\ell^{(0)}f_\ell(x_0).$$

3.1.2 **Monte Carlo Implementation**

We now give the gory details of the implementation of the above algorithm when we choose a Monte Carlo approach to the computation of the expectations appearing in the various steps of the algorithm described in the above bullet points. As most of these expectations involve the underlying Markov chain at two successive instants, it is economical to use the same set of Monte Carlo samples of the entire paths for all these expectations. So we generate M samples $\omega^{(1)},\cdots,\omega^{(M)}$ from the distribution of the sample path $\{X_n\}_{n=0,1,\cdots,N}$, and replace all the expectations in the above bullet points by the corresponding Monte Carlo sample averages. Moreover, as we do not have the exact values for the entries of the matrices $\Gamma^{(n)}$, we use the empirical estimate $\hat{\Gamma}^{(n)}$ computed from the Monte Carlo observation proxies by

$$\hat{\Gamma}_{\ell\ell'}^{(n)} = \frac{1}{M}\sum_{m=1}^{M}f_\ell(X_n(\omega^{(m)}))f_{\ell'}(X_n(\omega^{(m)})) \tag{3.3}$$

- To start
  - Compute the matrix $\hat{\Gamma}^{(N-1)}$ using formula (3.3) with $n = N-1$;
  - Compute the numbers

$$\hat{\alpha}_\ell^{(N-1)} = \frac{D(N-1,N)}{M}\sum_{\ell'=1}^{L}[\hat{\Gamma}^{(N-1)-1}]_{\ell\ell'}\sum_{m=1}^{M}f_{\ell'}(X_{N-1}(\omega^{(m)}))\varphi(X_N(\omega^{(m)}))$$

    for $\ell = 1,\cdots,L$;
- for $n = N-2,\cdots,1,0$
  - Compute the matrix $\hat{\Gamma}^{(n)}$ using formula (3.3);

– Compute

$$\hat{\alpha}_\ell^{(n)} = \frac{D(n,n+1)}{M} \sum_{\ell'=1}^{L} [\hat{I}^{(n)-1}]_{\ell\ell'} \sum_{m=1}^{M} f_\ell(X_n(\omega^{(m)})) \max\left\{ \varphi(X_{n+1}(\omega^{(m)})), \sum_{\ell''=1}^{L} \hat{\alpha}_{\ell''}^{(n+1)} f_{\ell''}(X_{n+1}(\omega^{(m)})) \right\}$$

for $\ell = 1, \cdots, L$;

Finally, the price of the option is approximated by

$$\max\left\{ \varphi(x_0), \sum_{\ell=1}^{L} \hat{\alpha}_\ell^{(0)} f_\ell(x_0) \right\}.$$

As explained before, the smallest optimal exercise time can be approximated on each Monte Carlo simulated path $\omega^{(m)}$ by computing the first time that the pay-off equals the continuation value, i.e.

$$\hat{\tau}^*(\omega^{(m)}) = \min\left\{ n \geq 0;\ \varphi(X_n(\omega^{(m)})) = \sum_{\ell=1}^{L} \hat{\alpha}_\ell^{(n)} f_\ell(X_n \omega^{(m)}) \right\}.$$

3.2 General Monte Carlo Regression Strategy

We now present a general approach based on the classical form (1.13) of the dynamic programming principle. An approximation $\hat{v}_n(x)$ of the value $v_n(x)$ is computed for all $0 \leq n \leq N$ and $x \in E$. Then the true value of the option is approximated by $\hat{v}_0(x_0)$ (we denote by x_0 the value of the underlying at time $n = 0$). Then, the minimal optimal time of exercise is approximated along each sample path $X_\cdot$ by

$$\hat{\tau}^* = \inf\{n \geq 0;\ \hat{v}_n(X_n) = \varphi(X_n)\}.$$

The approximation $\hat{v}_n(x)$ is computed in the following way. Expressing the Markovian transition operator as a conditional expectation, the backward induction formula (1.13) giving the value function can be rewritten in the following way:

$$\begin{cases} v_N(x) = \varphi(x) \\ v_n(x) = \max\{\varphi(x), \mathbb{E}\{D(n,n+1) v_{n+1}(X_{n+1}) | X_n = x\}\}, \end{cases} \tag{3.4}$$

for $n = 0, 1, \cdots, N-1$. So, if we choose a method to compute approximations for the above conditional expectations, we got ourselves a way to compute an approximation for the value function, and hence of the price of the option as well as an approximate optimal exercise strategy. The above conditional expectation is the regression of the univariate random variable $D(n,n+1)v_{n+1}(X_{n+1})$ against the random variable X_n which is most of the time multivariate. Classical statistics provides all sorts of methods to compute approximations of regression functions from sample observations. They are classified as parametric, semi-parametric and non-parametric. Because we want to handle underlying processes $\{X_n\}_n$ in relatively high dimension, the Monte Carlo approach to American option pricing relies on non-parametric regression procedures. Moreover, samples created from Monte Carlo simulations of these couples of random variables are used in lieu of observation data. In these lecture notes, the Monte Carlo approach is based on the following strategy:

- Generate M samples $\omega^{(1)}, \cdots, \omega^{(M)}$ from the distribution of the sample path $\{X_n\}_{n=0,1,\cdots,N}$;
- Choose a non-parametric regression method, and for each $n = N-1, N-2, \cdots, 1, 0$, compute an approximation of the theoretical regression function

$$\mathbb{E}\{D(n,n+1)v_{n+1}(X_{n+1})|X_n = x\}$$

from the *observation data proxies*

$$\begin{aligned}(X_n(\omega^{(1)}), D(n,n+1)v_{n+1}(X_{n+1}(\omega^{(1)}))), \cdots, \\ \cdots (X_n(\omega^{(M)}), D(n,n+1)v_{n+1}(X_{n+1}(\omega^{(M)})))\end{aligned}$$

of M couples computed from the Monte Carlo samples generated in the first bullet point;
- Plug these approximations in formula (3.4) to compute the approximated value function $\hat{v}_n(x)$ for all x of the form $x = X_n(\omega^{(M)})$ and for $n = N$ down to $n = 0$.

Remarks .

1. Each sample $\omega^{(m)}$ is characterized by a sequence

$$X_0(\omega^{(m)}) = x_0, X_1(\omega^{(m)}), \cdots, X_N(\omega^{(m)})$$

of random elements of $\mathbb{R}^p$. Obviously, we should try to use exact simulation whenever the model for the underlying Markov chain permits. This is the case when the underlying dynamics are given by a binomial tree or a geometric Brownian motion (see for example Subsection 1.4.2), or a mean reverting Ornstein-Uhlenbeck or CIR process. However, this is not the case when these samples are obtained from the implementation of a numerical scheme for stochastic differential equations, e.g. the Euler's scheme or any other higher order scheme.
2. As we already mentioned, the regression method is most often chosen to be of the non-parametric type. This is due in part to the fact that we do not know a priori the functional form of the regression function and to the typically high dimension p of the underlying process. However, things will be different when we know or suspect that the regression functions are of a certain form. This is for example the case in the case of the regressions used for the convertible bond pricing where we used hockey-stick functions mimicking the European option pay-offs.
3. Glasserman random tree offers a smooth transition between the T vR method which requires only expectations and more general methods based on full fledge non-parametric regressions.

3.3 LONGSTAFF-SCHWARTZ ALGORITHM

Instead of implementing directly the backward induction of the value function version of the dynamic programming principle, this algorithm reformulates the dynamic programming principle in terms of the minimal optimal stopping times τ_n^*. In this more subtle form of the backward induction, the conditional expectations do not involve consecutive times any more. We shall see that the numerical performance of the approximation so-obtained are different, quite possibly better, than the naive implementation discussed in the previous section.

The thrust of the method is to compute inductively the optimal stopping times path by path. Concentrating on the optimal times of exercise does not take anything away from pricing needs. Indeed, as emphasized by formulae (1.15) and (1.16), the various value functions are expectations involving the optimal stopping times. Hence they can be approximated by Monte Carlo averages as:

$$\hat{v}_n(x) = \frac{1}{M}\sum_{m=1}^{M} D(n, \tau_n^*(\omega^{(m)}))\varphi(X_{\tau_k^*(\omega^{(m)})}(\omega^{(m)})) \tag{3.5}$$

In particular, the approximate price of the option will be computed as

$$\hat{v}_0(x) = \frac{1}{M}\sum_{m=1}^{M} D(0, \tau^*(\omega^{(m)}))\varphi(X_{\tau^*(\omega^{(m)})}(\omega^{(m)}) \tag{3.6}$$

3.3.1 **Dynamic Programming on Exercise Times**

Notice that the smallest optimal stopping times are characterized by

$$\begin{cases} \tau_N^* = N \\ \tau_n^* = n\mathbf{1}_{\{\varphi(X_n)\geq v(n,X_n)\}} + \tau_{n+1}^*\mathbf{1}_{\{\varphi(X_n)<v(n,X_n)\}} \qquad \text{for } n = 0, 1, \cdots, N-1 \end{cases}$$

This expression seems to depend on the value functions. However, the definitions of the two complementary sets appearing above can be rewritten in terms of the stopping times only. Indeed:

$$\{\varphi(X_n) < v_n(X_n)\} = \{\varphi(X_n) < [\tilde{P}_n v_{n+1}](X_n)\} \tag{3.7}$$

$$= \{\varphi(X_n) < \mathbb{E}\{D(n, \tau_{n+1}^*)\varphi(X_{\tau_{n+1}^*})|X_n\}\} \tag{3.8}$$

Equality (3.7) is due to the fact that the value $v_n(X_n)$ at time n is the maximum of the immediate pay-off $\varphi(X_n)$ and the continuation value, while and equality (3.8) is due to the fact that along the path of the underlying process, the Snell envelop (and hence the value function) is a martingale up until the smallest optimal time, τ_{n+1}^* in this case, time at which it is first equal to the discounted pay-off function, namely $D(n, \tau_{n+1}^*)\varphi(X_{\tau_{n+1}^*})$.

Once re-expressed in this way, the backward induction can be done on the stopping times instead of the value functions.

Remarks. 1. The form of the algorithm presented above is an abstraction, together with a mathematical justification of the procedure described on a simple numerical example in Section 1 of the original Longstaff-Schwartz paper [**?**]. Section 2 of this same paper is devoted to a formal description of the algorithm in the abstract framework used here. However, the very nature of the backward induction on exercise times seems to be lost in the process, and what is presented relates more to the general regression algorithm alluded to in the previous section. It is only by reading between the lines that one can recognize it.

2. The authors make a big fuss out of the use of a rejection method which at each stage of the induction, discards all the sample scenarios which are not in the money. We shall address this issue later in our discussion of the numerical results.

3.3.2 **Numerical Implementation**

A numerical implementation of the method requires the computation of the conditional expectations

$$\mathbb{E}\{D(n, \tau_{n+1}^*)\varphi(X_{\tau_{n+1}^*})|X_n\} \tag{3.9}$$

As in the previous section, we choose to use a Monte Carlo approach to the computation of these conditional expectations. As before, we first generate M samples $\omega^{(1)}, \cdots, \omega^{(M)}$ from the distribution of the sample path $\{X_n\}_{n=0,1,\cdots,N}$, and then we can use any of the the non-parametric regression methods discussed earlier and in the appendix.

NOTES & COMPLEMENTS

The confusion alluded to in the abstract of the chapter lies in the fact that many people call the general approach presented in Section 3.2 by the name of Tsitsiklis - van Roy method, and other call it Longstaff - Schwarz method. This confusion is partly due to the fact that these two papers have very many similarities and appeared almost simultaneously. We thank Bernard Lapeyre for enlightening discussions on the differences between these methods. In fact, the presentation in the layout of this chapter owes enormously to these private discussion. Both methods rely on the decomposition of regression functions on special bases of feature functions. The only convergence results apply when one increases the number of paths (scenarios) while keeping the number of basis functions constant. See for example the papers of Tsitsiklis and van Roy [**?**] and Clément, Lamberton and Protter [**?**] Recent studies by Egloff and Min-oo [**?**] and Glasserman and Yu [**?**] show that the number of paths should grow super-exponentially in the number of basis functions for simultaneous convergence to take place.

We stated clearly in the text that any method producing approximations of conditional expectations from Monte Carlo samples could be used to price American options. In a separate appendix , we review some of the most common nonparametric regression procedures which have been used in this context. A method based on Malliavin calculus is presumably the only method not familiar to classical statisticians.

A quantization algorithm based on classical ideas of signal analysis was proposed in [**?**] by Bally and Pages. The stochastic dynamics of the risky assets underlying the option are approximated by a Markov chain computed on a tesselation of the space. The major overhead in terms of computing time comes from the generation of the tesselation. So, the quantization algorithm is bound to be efficient if one has to price many American options written on the same underlying risky assets. Indeed in this case, the tesselation needs to be generated only once. For this reason, the authors made sample tesselations available on their web site. Still pricing is practially limited to underlying assets of dimension $p \leq 8$. Notice that for the quantization method, Monte Carlo techniques contribute in two different ways. They can be used to generate the tesselation and compute the transition probabilities. But once the tesselation is set and the transition probabilities are computed, as we demonstrated in this chapter, Monte Carlo methods can be used to price options on such a Markov chain. As we just explained, performance of the quantization method is seriously limited by the dimension of the underlying asset. This was known to many authors, in particular to Barraquand and Martineau who proposed five years earlier an approximate algorithm where quantization was used only in one dimension. The gist of their method is to compute an optimal exercise strategy which would only depend on the knowledge of the time evolution of the (one dimensional) reward process instead of the (multi-dimensional) underlying asset process. See [**?**] for details on this original use of the quantization idea.

The iterative construction mentioned in the last section of the chapter has been proposed by Kolodko and Schoenmakers in [**?**] and was extended to the multiple stopping problem by Bender and Schoenmakers in [**?**].

4

Dynkin Games and Monte Carlo Valuation of Convertible Bonds

Summary.

4.1 INTRODUCTION

Issuing bonds is a typical way for corporation to raise money. The level of corporate borrowing and the potentially disastrous effect of default have led to the birth and the growth of a vibrant credit derivatives market. Most corporate bond covenants are extremely complex (a typical prospectus is several hundred pages long), and conversion provisions are often found among the many intricate conditions for exercise. Convertible bonds are corporate bonds, and as such they are subject to default of the issuer. Embedded in a convertible bond is an option for the holder of the bond, to exchange the security for a given number of shares of the company stock or of an index chosen to underly the convertible component of the bond. This number of shares is determined by a conversion ratio identified in the indenture of the bond.

So convertible bonds are hybrid derivatives with a fixed income component (the interest coupon payments) and an equity component (underlier shares in case of conversion). They are very attractive to investors interested in the *upside* of a stock price appreciating significantly) with little or no *downside* due to the bond protection. Convertible bonds were extremely popular and their market volume increased very fast until the credit crunch of May 2005 due to the credit downgrade of GM and Ford.

The realization that convertible bonds are nothing more than Dynkin games of timing has spurred a recent wave of papers in the framework of the theory of mathematical models of continuous time finance. Starting with the introduction of game option by Kifer in [**?**], this culminated in the works of Kallsen and Kühn [**?**] and [**?**] who proved that a no-arbitrage price was given by the value function of a Dynkin game. These results set the stage for further theoretical developments. See for example the series of works by Bielecki et al. [**?**, **?**, **?**] or [**?**] for examples of applications of the theory of doubly reflected backward stochastic differential equations. We shall not need any of these theoretical developments as the purpose of these lecture is to discuss numerical valuation methods and their implementations. We work in the discrete time setting, and for that reason, we only need very simple facts from the theory of Dynkin games, for example as they are presented in the classical textbooks [**?**] and [**?**].

The structure of this chapter is very straightforward. We first describe the main features of convertible bonds as they are traded. The formalism which we introduce naturally brings to bare games of timing. Then, we review the basic facts of the theory of Dynkin games of timing in the discrete time setting, and we highlight the form of the dynamic programming principle which holds in this situation. This recursive form of the dynamic programming principle is recast as a backward induction not much different than in the case of the optimal stopping, and as such it is easily amenable to implementation. The resulting computation algorithms are riddled with computations of conditional expectations, and many of the considerations discussed earlier in the case of American options will come handy.

4.2 Convertible Bonds as Dynkin Games

4.2.1 Typical Corporate Bond Scenario

A corporate bond is a contract between two counterparties who agree to the following terms:
In general, the *seller* or *issuer* of the bond

- *collects the nominal* (namely the loan amount) at inception of the contract
- and *pays interest coupons* at regular time intervals
- and finally *returns* the nominal amount of the bond at maturity

In general the *buyer* or *holder* of the bond

- *pays* the nominal amount (loan amount) upfront
- and *receives* the interest coupon payments according to the tenor schedule (this is the fixed income part of her investment)
- and finally *retrieves* the nominal amount at maturity.

This description of what a corporate bond is highly incomplete as it lacks any reference to the possibility of default, the main feature differentiating a corporate bond from a Treasury which is usually assumed to be default free. Indeed, on the top of the items described in the bullet points above, the holder of the bond gets a *recovery* amount expressed as a proportion of the nominal of the bond in case of *default* before maturity.

4.2.2 Extra Features of Convertible Bonds

The issuer or seller of the bond can at a time of his *choosing* (namely at a stopping time τ_s that she can choose) put the bond namely

- she can return the nominal (loan amount) to the buyer of the bond and
- stop paying interest coupons

On the other hand, the buyer of the bond, i.e. the bond holder can at a time of his *choosing* (i.e. at a stopping time τ_b that she chooses)

- request the nominal of the bond and walk away (in which case the game is over) or
- convert the loan title into shares of the underlying specified in the contract in an amount also specified in the contract.

The contract **ends the first time** one of the two counterparties exercises her right. But as corporate bonds, convertible bonds are subject to default of the issuer. The precise conditions under which default is declared are specified in the indenture of the bond. Moreover, what should happen in case of default, i.e. what kind of recovery if any should the buyer of the bond expect, is also discuss in the prospectus of the bond. So a convertible bond can be viewed as a corporate bond with an option to exchange the title for a given number of shares of an agreed upon index (most often the stock of the issuing company) with the caveat of a few complex indentures including

- a Put / Redemption provision
- a Call provision
- a Put / Redemption protection
- a Call protection
- a Call notice
- specifics for the settlement at maturity
- definition of what constitutes default
- recovery in case of default
- . . . . . .

4.2.3 Mathematical Problem

In more mathematical terms, the settlement of a typical convertible bond contract can be described after

- the bond holder chooses an exercise strategy in the form of a stopping time τ_b
- the bond seller chooses an exercise strategy in the form of a stopping time τ_s

from which the time - zero value of the payments to the bond holder B from the bond seller S are defined by the quantity

$$R(\tau_b, \tau_s) = \begin{cases} L_{\tau_b}, & \text{whenever } \tau_b \leq \tau_s \text{ or } \tau_b = \tau_s < N \\ \xi, & \text{whenever } \tau_b = \tau_s = N \\ U_{\tau_s}, & \text{whenever } \tau_b > \tau_s \end{cases} \tag{4.1}$$

In other words, this present value of the cash flow is given by

- an adapted stochastic process $\mathbf{L} = \{L_n\}_n$ if B converts first,
- an adapted stochastic process $\mathbf{L} = \{U_n\}_n$ if S calls the bond first,
- and a random variable ξ if neither party exercises the option before maturity

In this setting,

- The bond holder B tries to **maximize** $\mathbb{E}\{R(\tau_b, \tau_s)\}$, so the holder of the bond can be view as a player trying to maximize the expected cash flow, while
- the issuer S tries to **minimize** the expected cash flow $\mathbb{E}\{R(\tau_b, \tau_s)\}$.

The question we address in this chapter is

What is the value of such a contract?

and how can we practically implement the computation of the value of the contract.

4.2.4 Current Practice in the Industry

We first remark why the methods used most frequently in pricing convertible bonds are not Satisfactory. For an exhaustive review of the implementation details of these methods, we refer the interested reader to Andersen and Buffum's review [**?**].

- On one hand, most current models in use are based on restrictive assumptions such that
  - Traditional split between a straight bond and the option to convert, but this form of the convertible bond is not exactly equivalent but merely an approximation
  - Structural approach based on precise information on the term structure of assets and liabilities of the firm and modeling of fundamental factors such as stock price, interest rate, ...
- On the other hand, most current implementations are based on tree models and PDE solvers which limit the bond models to low dimensional driving factors and rather short maturities. Moreover, the numerical results produced by these models
  - do not match market prices
  - do not match market deltas
  - .....

Obviously, there is room for experimentation with new ideas and new pricing algorithms. The Monte Carlo approach which we propose is in this spirit. But before we proceed, it is a good idea to check the wish-list of all parties involved, academics and practitioners, including traders and quants.

From an academic point of view, a versatile model with simple implementations is needed for the systematic analysis of many of the features of convertible bonds which are often mentioned and not studied rigorously by lack of the proper tools. See for example the discussion of possible extensions in Subsection **??** below for some of the claims that need to be substantiated. From a practitioner point of view, every day transactions require

- a program for quick pricing with prices in line with those from third party providers, but also
- a robust model implementation including more of the bond indentures and with implementations for risk managers, market makers and possibly proprietary traders. Such a program will be quite likely slower than the program alluded to in the first bullet point.

4.3 Proposed Monte Carlo Approach

We now describe the various components of our Monte Carlo approach to the pricing of the convertible bonds described in the previous section.

4.3.1 Monte Carlo Simulation

The first step is to generate sample paths (Monte Carlo scenarios) for the underlier. Indeed even though we will model default via an intensity, we still need the dynamics of the index into which the bond ownership can be converted. So even though we are not working within the structural approach à la Merton, we need a model for equity dynamics. We shall start from a geometric Brownian motion for the sake of definiteness, but the versatility of Monte Carlo methods allows for any kind of dynamics given by stochastic differential equations (including stochastic volatility, local volatility, ... models). As already emphasized, we use exact simulation whenever possible. Otherwise, we use Euler's scheme or any higher order scheme to generate Monte Carlo scenarios.

The next step is to choose a model for default intensity. We work in the framework of Cox processes in the spirit of the reduced form approach. If we denote by τ_d the time of default, assuming that it has intensity $\{\lambda_t\}_t$ with respect to the filtration $\{\mathcal{F}_t\}_t$ is essentially assuming that

$$\mathbb{P}\{\tau_d > t|\mathcal{F}_t\} = e^{-\int_0^t \lambda_s ds}, \qquad t > 0.$$

Since the intensity λ is non-negative, samples from such a stopping time are easily obtained as the first times the running integral of the intensity crosses over a level given by an independent exponential random variable with unit rate. For the sake of convenience, we restrict ourselves to intensities of the form

$$\lambda_t = \lambda(X_t)$$

for some deterministic function λ of the underlying factor X_t. Accordingly, including the possibility of default in our model will require that we

- draw an exponential random variable, say ϵ with rate 1.
- generate scenarios for the intensity λ_s for $0 \leq s \leq T$,
- compute the running integral $\int_0^t \lambda_s \, ds$
- in which case the scenarios for the time of default will be given by the first time the intensity running integral crosses above the exponential variate,.

We are now left with the implementation of the bond indenture on the Monte Carlo scenarios so obtained. This is done via the construction of **two random sequences**

- $\{L_j\}_j$ gives the present value at time $t = 0$ of the cumulative cash flows to the holder, before-and-including time $t = j\Delta t$, should she decide to exercise her right(s) (**conversion, redemption/put**, ...) at that time, while the issuer has not exercised her option yet.
- $\{U_j\}_j$ gives the present value of the cumulative cash flows still to the holder from the issuer before and including time $t = j\Delta t$ should the issuer decide to exercise her right to call the bond at time $j\Delta t$ while the holder has not exercised any of her options yet.

We denote by C_t the time zero value of all the coupon payments having occurred at time t or before. Typically

$$C_t = \sum i \geq 1, T_i \leq t D(0, T_i) c(T_i) \tag{4.2}$$

where as usual the nominal of the bond has been normalized to 1, and where $c(s)$ is the coupon rate applicable at time s as agreed upon in the indenture of the bond. In most cases, the function $s \hookrightarrow c(s)$ is piecewise constant if not constant throughout the life of the bond. Remember that we use the notation $T_0 = 0 < T_1 < \cdots < Tn < \cdots$ for the dates of the coupon payments. Their schedule is also specified in the bond prospectus, but for the sake of definiteness, we shall assume that the coupons are quarterly. Next we denote by A_t the interests accrued since the last coupon payment. It is given by the simple formula

$$A_t = c\frac{t - T_i}{T_{i+1} - T_i}, \qquad T_i \leq t < T_{i+1} \tag{4.3}$$

when the coupon rate is a constant independent of time.

Typical Monte Carlo Scenario for $\{L_j\}_j$

For the sake of illustration, we give several examples of what the present value of cumulative cash flows to the holder of the bond should be when she exercises her right (whatever right that is) before the seller.

- **Before default** (i.e. whenever $t < \tau$)
  - $L_t = A_t + C_t$ if $t = j\Delta t < T_{conv}$ since conversion is not allowed during the **conversion protection** period,, i.e. as long as t is smaller than the threshold T_{conv};
  - $L_t = A_t + C_t + \alpha S_t$ if $t > T_{conv}$ and no redemption is possible, i.e. if $t < T_{put}$;
  - $L_t = A_t + C_t + \max\{\alpha S_t, P\}$ if $t > T_{conv}$ and put is possible, i.e. $t > T_{put}$;
  - . . . . . .
- **After default** (i.e. when $t \geq \tau$)
  - $L_t = A_t + C_\tau + R_\tau$ where R_t denotes the present value of recovery if default occurs at time t

Typical Monte Carlo Scenario for $\{U_j\}_j$

Similarly, we describes the cash flows to the bond holder when the issuer exercises first by giving their present values.

- **Before default** (i.e. whenever $t < \tau$)
  - $U_t \equiv \infty$ up to the time T_{call} announcing the end of the **hard call protection**;
  - $U_t = A_t + C_t + \max\{\alpha S_t, P_{call}\}$ if the issuer is allowed to call the bond, i.e. if $t > T_{call}$;
  - Include **Make Whole** provisions
  - Implement **Soft Call Protection**
  - Include **Call Notice of Redemption Period** provisions
  - . . . . . .
- **After default** (i.e. if $t \geq \tau$)
  - $U_t = A_t + C_\tau + R_\tau$

Notice that in all cases we have

$$L_t \leq U_t \quad \text{and} \quad L_T = U_T = \xi.$$

The inequality between the processes L and U justifies the notation L for lower bound and U for upper bound, and it plays a crucial role in the derivation of the classical results of the theory of Dynkin games which we recall below and which we use in our valuation implementation.

Typical Convertible Bond Scenario

In words, *the present value of the cumulative payments to the bond holder B from the seller S* is given by

- L is B converts (strictly) before maturity and no later than S calls the bond
- U if S calls the bond first
- ξ if neither party exercise their respective options before maturity

In more mathematical terms,

- Given an exercise strategy τ_b (stopping time) chosen by the bond holder

- Given an exercise strategy τ_s (stopping time) chosen by the issuer of the bond

the present value $R(\tau_b, \tau_s)$ of all the cash flows from the issuer to the holder of the bond during the life of the bond is given by

$$R(\tau_b, \tau_s) = \begin{cases} L_{\tau_b}, & \text{whenever } \tau_b \leq \tau_s \text{ or } \tau_b = \tau_s < N \\ \xi, & \text{whenever } \tau_b = \tau_s = N \\ U_{\tau_s}, & \text{whenever } \tau_b > \tau_s \end{cases} \tag{4.4}$$

The holder of the bond tries of maximize the expected value of R while the issuer tries to minimize it. The analysis of this problem is part by the classical theory of Dynkin games of timing which we review in the next section.

4.4 Dynkin Games of Timing

Throughout this section we work within the mathematical framework of a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with a filtration $\{\mathcal{F}_n\}_{n\geq 0}$ and we denote by $\mathcal{S}$ the set of stopping times for the filtration. We only consider times n smaller than or equal to a final horizon N. We could allow for $N = \infty$ in order to include infinite horizon problems (and perpetual options and games) but we shall restrict ourselves to finite maturities N for the sake of definiteness.

4.4.1 Definition

The rules of such a game are completely determined by a *terminal* (i.e. $\mathcal{F}_N$-measurable) random variable ξ, and two adapted sequences $\{L_n\}_{0\leq n\leq N}$ and $\{U_n\}_{0\leq n\leq N}$ of integrable random variables satisfying

$$L_n \leq U_n \quad \mathbb{P} - a.s. \qquad n = 0, \cdots, N. \tag{4.5}$$

The game is played in the following way. Independently of each other, each player chooses a stopping time, say τ_b and τ_s to use the same notation as in the previous section, and for each scenario, the first player receives from the second player the *reward amount* $R(\tau_b, \tau_s)$ given by

$$R(\tau_b, \tau_s) = L_{\tau_b}\mathbf{1}_{\{\tau_b \leq \tau_s,\, \tau_b < N\}} + U_{\tau_s}\mathbf{1}_{\{\tau_s < \tau_b\}} + \xi\mathbf{1}_{\{\tau_b = \tau_s = N\}}. \tag{4.6}$$

This is an instance of a zero-sum game as whatever the first player receives is paid by the second player. Notice that in this abstract setting, this pay-off can be negative, implying that cash can flow in both directions. In words, the second player pays the amount L_{τ_b} if the first player ends the game first before time N, U_{τ_s} if she stops the game first, or ξ if none of the players stop the game before N. According to these rules,

- the first player tries to **maximize** $\mathbb{E}\{R(\tau_b, \tau_s)\}$ while
- the second player tries to **minimize** $\mathbb{E}\{R(\tau_b, \tau_s)\}$,

and for this reason, the first player is often referred to as the maximizing player while the second player is called the maximizing player. In any case, it is now clear that framed this way, the mathematical set-up of convertible bond valuation fits in the framework of Dynkin games of timing.

4.4.2 Game Value Functions

We now review the results of the theory of Dynkin games in discrete time which we will need in our convertible bond pricing exercise. We emphasize the striking parallel with the discussion of Chapter 1 which cast the Snell envelop as the solution of the optimal stopping problem.

We first introduce the game upper value as the number:

$$\overline{V} = \sup_{\tau_b} \inf_{\tau_s} \mathbb{E}\{R(\tau_b, \tau_s)\} \tag{4.7}$$

and the game lower value as

$$\underline{V} = \inf_{\tau_s} \sup_{\tau_b} \mathbb{E}\{R(\tau_b, \tau_s)\} \tag{4.8}$$

whose interpretations are as follow.

- For any given minimizing player strategy τ_s, the maximizing player will choose a strategy τ_b which maximizes her expected reward

$$\sup_{\tau_b} \mathbb{E}\{R(\tau_b, \tau_s)\} \tag{4.9}$$

  and if the minimizing player is prudent, she will choose τ_s in order to minimize (4.9), explaining how $\underline{V}$ appears.
- This *min-max* strategy guarantees that the expected payment to the maximizing player is at least $\underline{V}$, which appears as a lower bound to the game value.
- By exchanging the roles of the maximizing and minimizing players in the above argument, we highlight a *max-min* strategy for which, independently of the strategy used by the maximizing player, the expected payment cannot exceed $\overline{V}$, which hence appears as an upper bound to the value of the game.

4.4.3 Obvious Bounds

Fixing the exercise time τ_s of the seller in the definition of the upper value of the game will give an upper bound. If we use $\tau_s \equiv N$ (which amounts to saying that the seller actually does not exercise his or her right), then we get:

$$\begin{aligned}
\overline{V} &= \sup_{\tau_b} \inf_{\tau_s} \mathbb{E}\{R(\tau_b, \tau_s)\} && (4.10)\\
&\leq \sup_{\tau_b} \mathbb{E}\{R(\tau_b, N)\} && (4.11)\\
&= \sup_{\tau} \mathbb{E}\{L_\tau)\} && (4.12)
\end{aligned}$$

which shows that the value of the game is bounded from above by the value of an American option with pay-off $\{L_n\}_n$. Obviously the buyer cannot expect more than what he would get should the seller act as if she did not have the option to stop the game.

Similarly, fixing the exercise time τ_b of the buyer in the definition of the lower value of the game will give a lower bound. If we use $\tau_b \equiv N$ (which amounts to saying that the buyer actually does not exercise her exercise right), then we get:

$$\underline{V} = \inf_{\tau_s} \sup_{\tau_b} \mathbb{E}\{R(\tau_b, \tau_s)\} \tag{4.13}$$

$$\geq \inf_{\tau_s} \mathbb{E}\{R(N, \tau_s)\} \tag{4.14}$$

$$= \inf_{\tau} \mathbb{E}\{U_\tau)\} \tag{4.15}$$

which shows that the value of the game is bounded from below by the negative of the value of an American option with pay-off $\{-U_n\}_n$. Obviously the seller cannot expect to do better than what she would do should the buyer act as if she did not have the option to stop the game.

4.4.4 Main Result

As in the case of our discussion of the optimal stopping problem and the introduction of the Snell envelop as a dynamic version of the original problem, for each time $n \leq N$ we define the random variables $\underline{V_n}$ and $\overline{V_n}$ by

$$\underline{V_n} = \inf_{\tau_b \in \mathcal{S}_n} \sup_{\tau_s \in \mathcal{S}_n} \mathbb{E}_n\{R(\tau_b, \tau_s)\} \tag{4.16}$$

and

$$\overline{V_n} = \sup_{\tau_s \in \mathcal{S}_n} \inf_{\tau_b \mathcal{S}_n} \mathbb{E}_n\{R(\tau_b, \tau_s)\}. \tag{4.17}$$

The first result of the theory is captured by

$$\boxed{\underline{V_n} = \overline{V_n}, \qquad 0 \leq n \leq N.}$$

We denote by V_n the common value of $\underline{V_n}$ and $\overline{V_n}$. The second main result of the theory is a constructive algorithm which gives a practical way to actually compute the value process $\{V_n\}_{0 \leq n \leq N}$.

$$V_n = \begin{cases} L_n, & \text{if } \mathbb{E}_n\{V_{n+1}\} < L_n \\ \mathbb{E}_n\{V_{n+1}\}, & \text{if } L_n \leq \mathbb{E}_n\{V_{n+1}\} \leq U_n \\ U_n & \text{, if } U_n < \mathbb{E}_n\{V_{n+1}\} \end{cases} \tag{4.18}$$

starting from the terminal condition $V_N = \xi$. The backward induction given by formula (4.18) is the analog of the dynamic programming principle (1.3) derived for the Snell envelop constructed as the value process of the optimal stopping problem in Chapter 1. The third and final theoretical result which we borrow from the theory of Dynkin games of timing is the characterization of a minimal set of optimal stopping times

$$\tau_s^* = \inf\{n \geq 0; \;\; V_n \leq L_n\} \tag{4.19}$$

and

$$\tau_b^* = \inf\{n \geq 0; \;\; V_n \geq U_n\}. \tag{4.20}$$

These three theoretical results will be proven in Chapter **??** as part of our discussion of the duality/pointwise approach to the problem.

4.5 MONTE CARLO IMPLEMENTATION

The conclusion of the discussion of the previous section is that we have a pricing algorithm for convertible bonds without call notice. Like in the case of American options, we shall not argue why the value function computed above is in fact a reasonable *no-arbitrage* price according to the classical models of financial mathematics. The interested reader is referred to the Notes and Complements at the end of the chapter for references on this subject.

The strategy to implement such a pricing algorithm goes as follows:

- Choose our favorite regression method
- Compute the value functions V_n *backward-in-time* starting from time $n = N$ down to $n = 0$.
- Read off the convertible bond price as the value function at time $n = 0$
- Compute the optimal exercise times, scenario-by-scenario, in a *forward-in-time* pass through the scenario

4.5.1 Choice of a Model with Equity dependent Spreads

As explained earlier, we model default in the reduced form approach by choosing a *default intensity* λ, and conditioned on the value of this intensity, the time of default becomes an exponential random variable ϵ with rate λ. We model the intensity as a function of the underlying spot value of the company's shares. So we choose an intensity of the form $\lambda_t = \lambda(S_t)$ for some deterministic function $x \hookrightarrow \lambda(x)$. As explained in Duffie's book [**?**], the default intensity should be a decreasing function of the equity as we expect that large values of the stock come with a low probability of default and small values of the stock with an increase in the probability of default. Several parametric families have been used for the function λ.

- The power family $\lambda(x) = \beta_0 x^{-\beta_1}$ was used in [**?**, **?**, **?**, **?**, **?**, **?**, **?**], and possibly other works which we are not aware of, while
- the exponential family $\lambda(x) = \beta_0 e^{-\beta_1 x}$ was used in [**?**, **?**].

Except possibly for the fact that calibration to the CDS spread curves available to us seemed to have been more robust with the exponential family, we did not find significant differences between the valuations within each model.

In any case, because of this particular form of the default intensity, we work in the framework of a one-factor model.

4.5.2 Calibration

The numerical results presented below were obtained from simulations of geometric Brownian motions with constant volatility σ and drift given by the short interest rate the day we value the bond. The drift was adjusted for the rate of dividend payments when such information was available, but as we did not have any option data, we used a wild guess for the implied volatility.

A first improvement would be to use a local volatility function $(t, x) \hookrightarrow \sigma(t, x)$ function of the underlying spot x in lieu of the constant σ.

The intensity parameters β_0 and β_1 were chosen in order to match market the CDS spread curve whenever the information was available, or the 1yr and the 5yr probabilities of default given by

Moody's transition probability matrices when the credit rating of the company was known. We used a Levenberg-Marquardt form of least squares calibration, the CDS spreads being computed by Monte Carlo methods from the model. This calibration procedure seems to have worked reasonably well, its major shortcoming being its lack of robustness.

4.5.3 Setting up the Regressions

The σ-field $\mathcal{F}_n$ giving the information available at time n contains the information about S_k and all the events $\{\tau \leq k\}$ for $k \leq n$. So even if we were to assume that S is Markovian, we could not replace the conditioning by $\mathcal{F}_n$ by a conditional expectation with respect to S_n. However, in order to avoid multivariate regressions we would like to replace conditional expectations with respect to $\mathcal{F}_n$ by conditional expectations with respect to S_n. In order to achieve this we propose to add a jump to 0 at default. In other words, we set $S_t = 0$ when $t \geq \tau$ or to be more specific we work with $\tilde{S}_n$ defined by

$$\tilde{S}_n = \begin{cases} S_n, & \text{if } n < \tau \\ 0 & \text{if } n \geq \tau \end{cases}$$

In the models of continuous time finance, this would amount to adding a jump martingale term to the stochastic differential equation for S_t.

Back to our discrete time model, the effect of the change we propose to the backward induction amounts to

Regress V_{n+1} against $\tilde{S}_n$ instead of $\mathcal{F}_n$!

In order to understand the effect of such an assumption we look at empirical data in the (S_n, V_{n+1}) - plane. We see

- a blob of points above $S_n = 0$ (corresponding to instances of default prior to or at time n)
- a mild "hockey-stick" shape or linear cloud above $S_n > 0$

Given these empirical facts, we propose an easy way-out to the regression problem:

- we compute a plain average for of the values of V_{n+1} corresponding to $S_n = 0$
- we perform a plain least squares piecewise linear regression of V_{n+1} against S_n when $S_n > 0$!

4.5.4 Looking at Data: Sirius

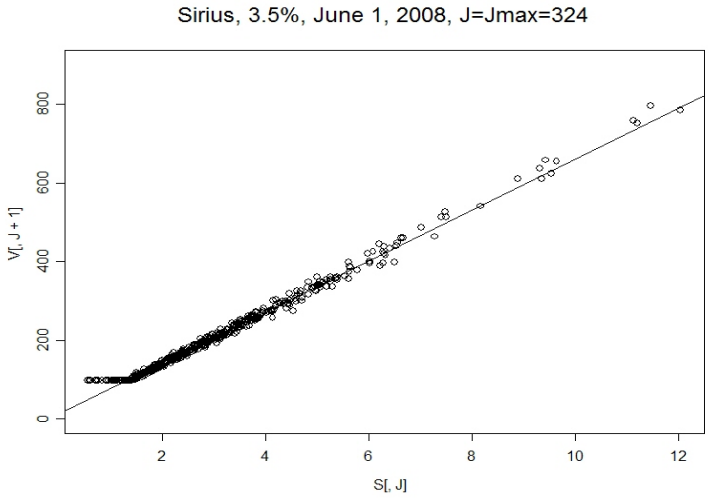

4.5.5 Looking at Data: Sirius

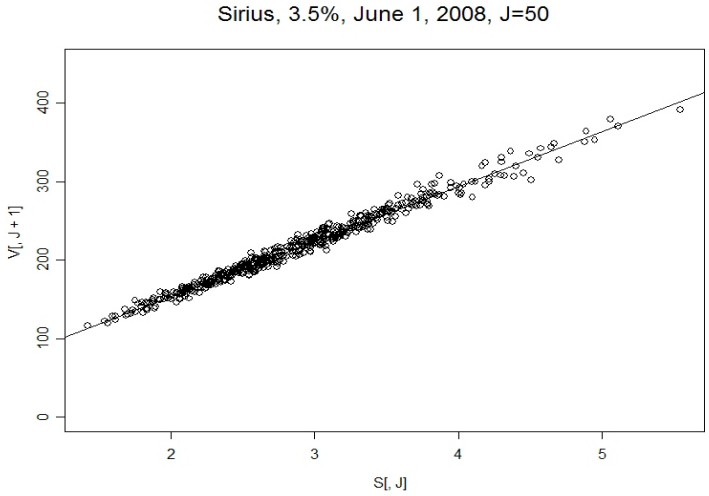

4.5.6 Looking at Data: Bearingpoint

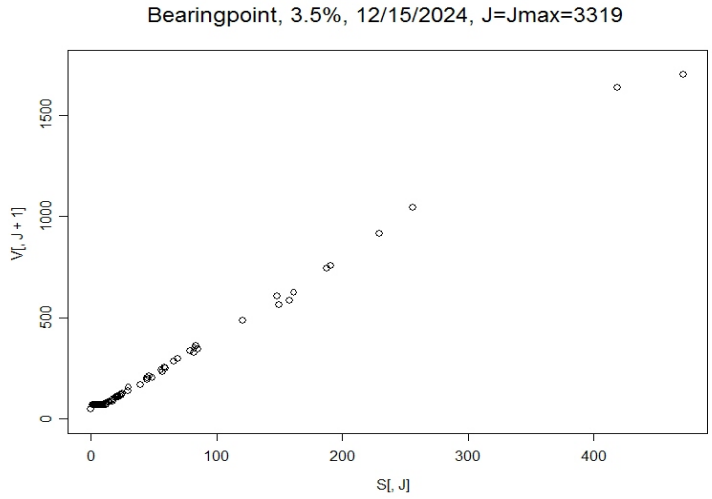

4.5.7 Looking at Data: Bearingpoint

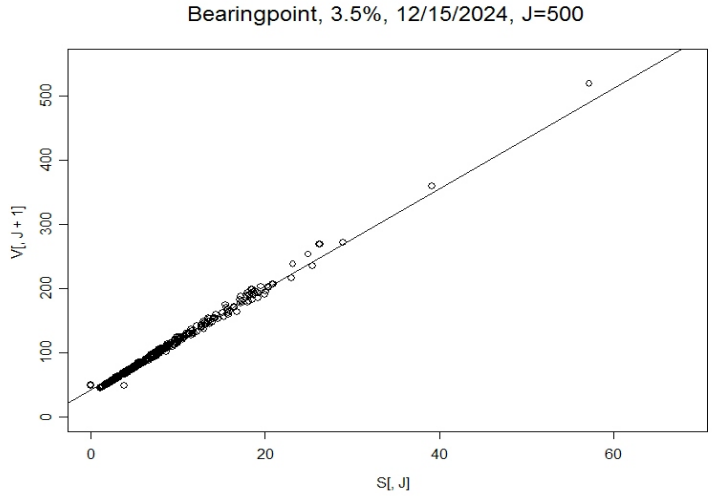

4.5.8 Looking at Data: Bearingpoint

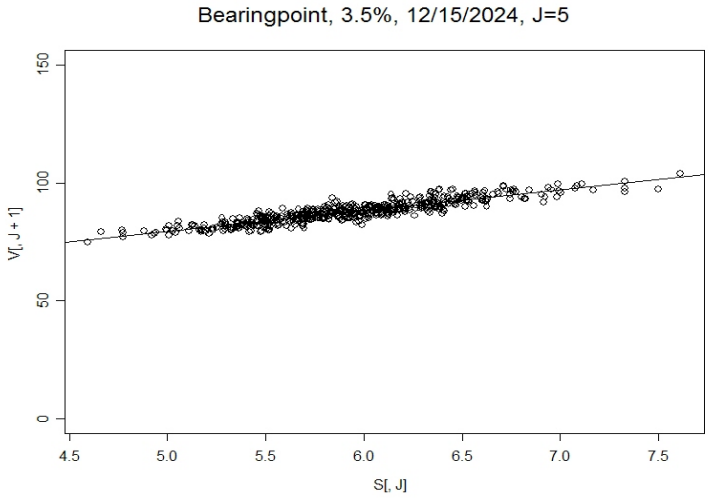

4.5.9 Looking at Data: Schlumberger

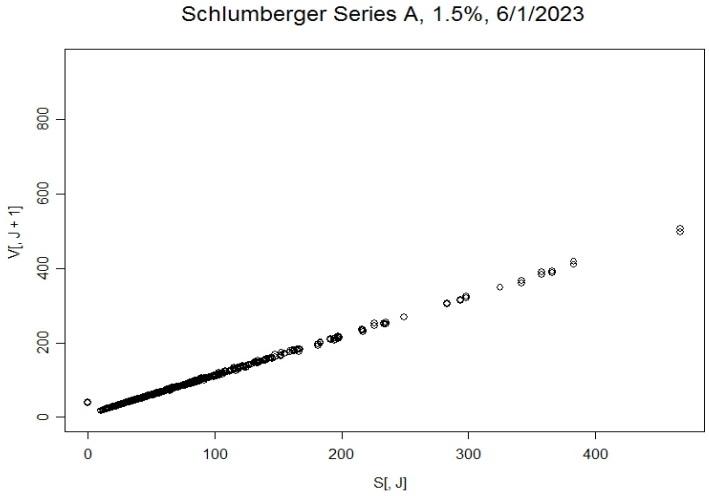

4.5.10 Prices and Model Deltas

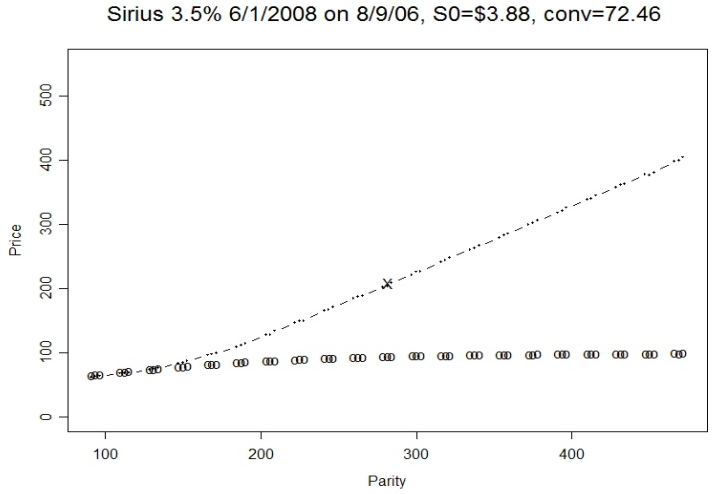

4.5.11 Prices and Model Rhos

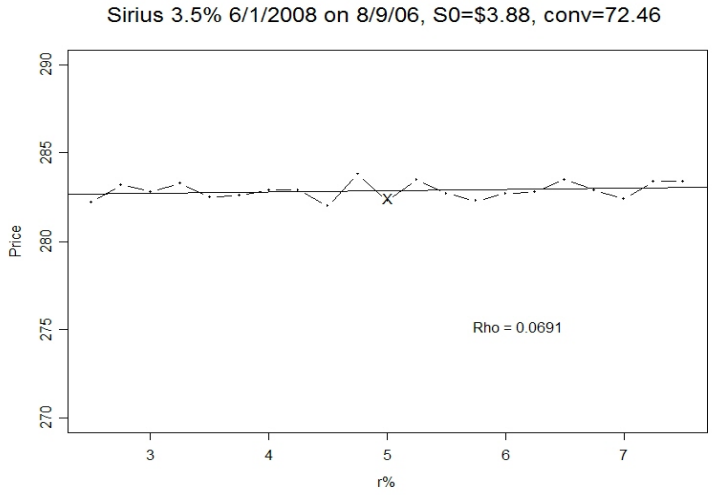

4.5.12 Prices and Model Vegas

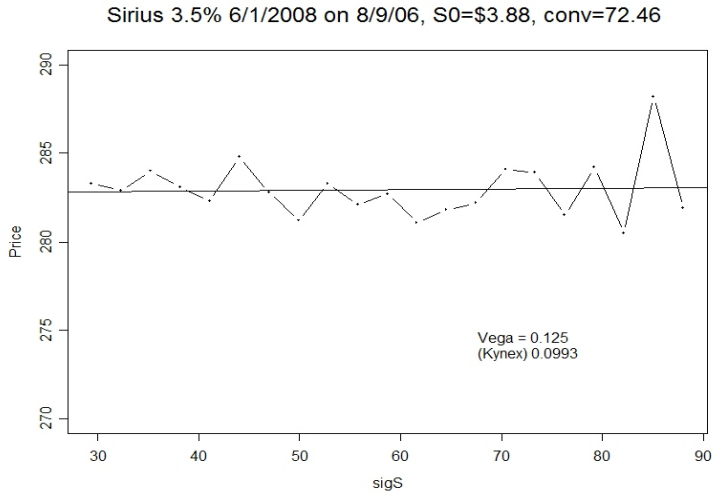

4.5.13 Implemention of Notices of Redemption Periods

This is usually done as an American option pricing problem.

4.5.14 Possible Extensions

Among the features which we did not test numerically and which could be of interest we list

- *Adding more stochastic factors* such as interest rate (for long dated bonds with so-called *borrow fees*), or stochastic volatility (e.g. Heston model) or even other underliers (stocks, indexes, baskets) for exchange conversions. Of these three possible extensions, the third one is imposed by the indenture of the contract. An alternative to the second one is discussed in the next bullet point below. Finally, concerning the first point, it has been argued repeatedly in the technical literature that the volatility of interest rates was significantly smaller than equity volatility and that, as a consequence, the correction for stochastic interest rates was of a much smaller order. See for example [**?**], citeGrimwoodHodges, or [**?**]. In any case, all these changes are easy to implement, the only problem being the increase in the dimension of the regression, and this is the only reason why we refrain from implementing any of these changes: we did not want to face the curse of dimensionality in this first numerical analysis.
- *Including local volatility* ("à la Dupire") in the underlier's model when the price of more than one liquidly traded option is available. Techniques to construct implied and local volatility surfaces have been developed for underliers with a large number of actively traded options. However, for most of the underliers of convertible bonds, only a small number of prices of traded options are available, and whether these constructions are based on parametric or non-parametric smoothing techniques, relying on their results can be very probematic, especially for long dated bonds. Prices of risky bonds should be included (when available) in the calibration of the volatility surface, as the use of calibration methods ignoring the possibility of default will create an artificial skew. See for example [**?**] for a discussion of this last point.
- Developing *importance sampling* methods for variance reduction

- Identification of the *optimal conversion time* as the first crossing time of an *Exercise Boundary*. This is easily achieved by computing a Monte Carlo of the support of $S(\tau)$.
- Identification of the *optimal call time* and statistical analysis of the *call lag*
- *Exogeneous modelling of tax effects and so-called fundamental changes*.

NOTES & COMPLEMENTS

The realization that convertible bonds are nothing more than Dynkin games of timing has spurred a recent wave of papers in the framework of the theory of mathematical models of continuous time finance. Starting with the introduction of game option by Kifer in [**?**], this culminated in the works of Kallsen and Kühn [**?**] and [**?**] who proved that a no-arbitrage price was given by the value function of a Dynkin game. These results set the stage for further theoretical developments. See for example the series of works by Bielecki, Crepey, Jeanblanc, and Rutkowski [**?**, **?**, **?**] or [**?**] for examples of applications of the theory of doubly reflected backward stochastic differential equations. We shall not need any of these theoretical results as the purpose of this paper is numerical. The facts from the theroy of Dynkin games in discrete time used in this chapter can be found in the classical books of Dynkin and Yuschevich [**?**], and Neveu [**?**].

The discussion of the Monte Carlo implementation presented in this chapter is borrowed from Carmona's technical report [**?**].

5

Valuation of Instruments with Multiple American Exercises

Summary. This chapter is devoted to the numerical analysis of a set of financial contracts involving multiple American exercises. We describe large classes of financial instruments with multiple American exercises before we set up the mathematical framework in which we tackle the problem of multiple American options valuation. Some of the Monte Carlo valuation algorithms described in the previous chapters are implemented in several multiple exercise cases and we discuss the numerical results so obtained. Motivated by the structure of energy swing contracts, we consider the generalization of our Monte Carlo approach to multiple American options to the case of multiple Dynkin games and we give numerical results in this case as well.

5.1 INTRODUCTION

Due to a renewal of interest in exotic interest rate derivatives and a growing involvement of academic researchers in problems of energy trading, options with multiple American exercises have caught the attention of many financial engineers. Extensive references are given in the Notes & Complements at the end of the chapter. We first present a set of instruments for which optionality appears in multiple instances. In preparation for the numerical analysis of these instruments, we outline the stylized facts of a mathematical set up in which simplified versions of these derivatives can be studied. Roughly speaking, we revert to the case of a single underlying risky asset, and we consider options which can be exercised several times. In such simplified models, these instruments could be priced by solving partial differential equations like in the classical Black-Scholes theory. However, non-linearities rule out solutions in the classical sense, and complex boundary conditions make the design and the control of numerical schemes quite a challenge. In the spirit of this set of lecture notes, we approach the pricing problem by Monte Carlo methods. We show how to expand the regression based procedures used in the previous chapters for single exercise instruments in order to apply them to the case of multiple exercises. A natural application of this generalization can be used to solve *multiple Dynkin games*, providing a tool for the pricing of two-sided swing contracts.

5.2 FIRST MOTIVATING EXAMPLES

This section is devoted to the description of financial instruments with multiple embedded American exercises. These instruments will serve as motivation for the mathematical problems we solve numerically in the following sections.

5.2.1 **Chooser Caps**

The first example we propose is borrowed from the fixed income markets. The reason for the popularity of these markets when it comes to implementing models of instruments with American and Bermudan exercises is the fact that the typical time interval separating two successive possible exercise times is typically a quarter (if not a semester) and as a consequence, the total number of possible exercise times is generally much smaller than in the case of the equity or energy markets. Following Meinshausen and Hambly (see [**?**]) we choose to discuss the *Chooser Cap* also known as *Chooser Flexible Cap* as our first example. We refer the interested reader to Section 9.7 pp.125-127 of Pelsser book [**?**] for details.

A cap contract with strike K on a 3-months interest rate is a portfolio of options on the quarterly interest payments that have to be made. Let us assume that the life of the contracts is N quarters, and the nominal is X. Each individual option is called a *caplet*. It covers one of the N quarters, and its pay-off is

$$X\Delta t(L_3 - K)^+ \tag{5.1}$$

where Δt is the daycount fraction (Act/360 is typically used for 3-months US dollar interest rate derivatives), and where L_3 denotes the 3-months LIBOR rate observed at the end of the quarter. An *Auto Cap* comprises a total of N quarters together with a limit $N_c \leq N$ on the number of caplets which are actually settled. The holder of such a contract receives a normal caplet payment for each of the first N_r caplets in the money (i.e. for which $L_3 > K$ at the end of the quarter). Obviously, an auto cap reduces to a regular cap when $N_c = N$. A *Chooser Cap* (also called a *Chooser Flexible Cap*) is a contract with the same premises as an auto cap, except for the fact that at each reset date, the holder of the contract has the right to decide whether to exercise that particular caplet and count it as part of the N_c allowable ones, or spare it for later use, each decision being final. Also note that the N_c caplet rights expire worthless if not used before the end of the N periods. As before, a chooser cap reduces to a regular cap when $N_c = N$, while it can be viewed as a standard American/Bermudan option when $N_c = 1$.

From now on we will fix $X = 1$ and drop the nominal of the contract from our notation.

The payoff of the chooser flexible cap under an exercise decision is maxfRt ..K; 0g, where Rt is the interest rate at the beginning of the t-th quarter. The constants are set to

We use exact simulation to simulate the values of r_{t_i} for $t_i = t_0 + i\Delta t$ for $i = 0, 1, \cdots, N$ with $\Delta t = .25$.

5.2.2 **Executive Option Programs**

The discussion in this subsection is modeled after the works of [**?**] and [**?**] and some of the references cited in these two papers.

5.2.3 **Optimal Dividend Payment Schedules**

This example is borrowed from [**?**]. We assume that the cash reservoir of a financial corporation follows a mean reverting process. The firm must decide the optimal dividend strategy, which consists of the optimal times and the optimal amounts to pay as dividends. We model this as an stochastic impulse control problem.

5.2.4 Energy Delivery Contracts

The following discussion is motivated by the structure of most *base-load* contracts used for the delivery of oil, gas and electricity. For the length of the contract, say one year, a fixed amount of commodity is to be delivered on a regular basis, say daily, in return for a periodic payment, say monthly, from the receiving party to the delivering counter-party.

Plain vanilla contracts of this type are essentially restricted to specific financially settled deals. Contracts physically settled by the actual delivery of the commodity, and most financially settled contracts are sprinkled with clauses increasing dramatically the complexity of the covenants. For the purpose of illustration, we give two stylized examples which we introduce as motivation for the developments of this chapter.

5.2.4.1 ***Recall Option***

Let N_r be an integer (we assume that N_r is not greater than the number N of days in the life of the contract), and let us assume that the seller of the commodity can withhold delivery on at most N_r days of his choice over the life of the contract provided he or she gives notice to the other counter-party the day before withholding delivery. Let us denote by L_n the present value of the net gain beyond the income expected from the base-load contract due to the recall of the commodity on day n. Assuming that there exists a liquid spot market for the commodity (we ignore delivery location and basis issues in this naive discussion), a simple (though naive) way to conceptualize L_n is to think of L_n as of the present value of the net gain or loss due to forfeiting the income from the base-load contract, and selling the daily amount of commodity contracted on the spot market. In other words,

$$L_n = e^{-rn}\alpha(S_n - K) \tag{5.2}$$

where α is the quantity to be delivered daily, S_n is the spot price of one unit of the commodity, and K is the fixed price agreed upon in the contract and where as usual, we denote by r the short interest rate which we use for discounting and which we assume to be constant over the life of the contract. According to such a contract, the holder of a recall option can choose N_r stopping times $\tau_1 < \tau_2 < \cdots < \tau_{N_r}$, and the present value of his or her expected profit and loss (P&L) is

$$\mathbb{E}\left\{\sum_{i=1,\,\tau_i\le T}^{N_r} L_{\tau_i}\right\}. \tag{5.3}$$

The holder of the recall option is thus interested in maximizing the value of the contract he holds, and for that, he will solve the optimization problem:

$$\sup_{\tau_1<\tau_2<\cdots<\tau_{N_r}\in\mathcal{S}} \mathbb{E}\left\{\sum_{i=1,\,\tau_i\le T}^{N_r} L_{\tau_i}\right\} \tag{5.4}$$

Notice that, because we do not assume that the reward L_n is non-negative, exercising all the N_r recall rights can be sub-optimal in some cases.

This type of instrument with multiple American exercises was the motivation behind the analyzes of [**?**] and [**?**].

5.2.4.2 *Swing Option*

We now assume that the optionality added to the base-load delivery contract is on the side of the buyer of the commodity. The terminology *swing* was motivated by the buyer's right to increase or decrease (i.e. swing) the amount to be delivered the day following the decision to exercise one of the swing rights. We assume that N_s times in the life of the contract, the buyer of the commodity can request with one day *notice* that the amount delivered the following day be *increased* or *decreased* by a fixed percentage amount. Let us denote by U_n the present value of the net cash flow beyond the income expected from the base load contract, due to the exercise of a swing right on day n. As before we assume the existence of a liquid spot market, and we denote by u and d the relative amounts (percentages) of commodity by which the base-load amount can be modified. In other words, on any given day, the buyer can choose to receive either $(1+u)\alpha$ or $(1-d)\alpha$ units of the commodity the following day still at the unit price of K, if he or she chooses to exercise one of his or her N_r swing rights. Obviously, the swing is up or down depending on whether the price is above or below K respectively. The holder of the swing option can sell the excess amount $u\alpha$, or buy the complement $d\alpha$ on the spot market. In other words, the present value U_n of the net gain to the holder of the swing option is

$$U_n = \begin{cases} e^{-rn} u\alpha(S_n - K) & \text{if } S_n > K \\ e^{-rn} d\alpha(K - S_n) & \text{if } S_n \le K \end{cases} \tag{5.5}$$

if he or she decides to exercise one of the swing rights on day n. Consequently, the holder of a swing contract can choose N_s stopping times $\tau_1 < \tau_2 < \cdots < \tau_{N_s}$ and the present value of his or her expected profit and loss (P&L) is

$$\mathbb{E}\left\{ \sum_{i=1,\, \tau_i \le T}^{N_s} U_{\tau_i} \right\}.$$

He or she should maximize the value of the contract, and for that, solve the optimization problem:

$$\sup_{\tau_1 < \tau_2 < \cdots < \tau_{N_s} \in \mathcal{S}} \mathbb{E}\left\{ \sum_{i=1,\, \tau_i \le T}^{N_s} U_{\tau_i} \right\}.$$

As before, the fact that the reward U_n can be negative implies that exercising all of the N_s swing rights can be sub-optimal.

It is important to notice that the definitions of the cash-flows L_n and U_n is not completely consistent with our claim that the exercise of a recall or swing right should be announced one day earlier. Indeed, either the action takes place the same day as the notice of exercise, or else the party exercising an option needs a crystal ball for the cash-flows to be defined in this way. Indeed, the conditions $S_n > K$ and $S_n \le K$ appearing in the definition of U_n should involve S_{n-1} instead of S_n in the absence of a crystal ball. We shall ignore the problem of the timing of the notice of exercise here, and consider the problem as defined above. It makes perfect sense mathematically, and it is a good first order approximation to the actual real life problem. This type of contract with multiple American exercises can be treated in exactly the same way as the base-load contract with finitely many recall rights.

5.2.4.3 ***Impulse Control for Plant Valuation***

5.3 MULTIPLE AMERICAN OPTIONS VALUATION

We now introduce a single mathematical framework in which we can solve at once a;; the valuation problems introduced in this chapter. As before we work on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with a filtration $\{\mathcal{F}_n\}_n$ and we denote by $\mathbb{E}_n\{\ \cdot\ \}$ the conditional expectation with respect to the σ-field $\mathcal{F}_n$ giving the information available at time n. The set of stopping times for this filtration is denoted by $\mathcal{S}$, and $\mathcal{S}_{\alpha,\beta}$ denotes the set of stopping times greater than or equal to α and smaller than or equal to β. We then fix an integer $p \geq 0$ not greater than the finite horizon N. The integer p represents the number of exercise rights while $\mathcal{S}^{(p)}_{\alpha,\beta}$ denotes the set of p-tuples of stopping times used to model the multiple exercises:

$$\mathcal{S}^{(p)}_{\alpha,\beta} = \{\boldsymbol{\tau} = (\tau_1, \cdots, \tau_p) \in \mathcal{S}^p; \alpha \leq \tau_1 < \tau_2 < \cdots < \tau_{p-1} < \tau_p \leq \beta\}$$

We now introduce the reward process $\{\xi_n\}_{0\leq n\leq N}$ as a random sequence on integrable random variables ξ_n. The random variable ξ_n can be interpreted as the present value (so ξ_n includes the discount factor e^{-rn}) of the pay-off to the holder of the option should he or she decide to exercise one of his or her rights at time n. Notice that we do not assume that ξ_n is non-negative, so ξ_n can represent a profit as well as a loss. With this notation out of the way, the *Optimal Multiple Stopping Problem* can be formulated as the optimization problem

$$\sup_{\boldsymbol{\tau}\in\mathcal{S}^{(p)}_{0,T}} \mathbb{E}\{\xi_{\boldsymbol{\tau}}\} \qquad \text{with} \qquad \xi_{\boldsymbol{\tau}} = \sum_{i=1}^{p} \xi_{\tau_i}$$

In analogy with the classical optimal stopping problem, we search for the value of the supremum, we try to find out if this supremum is actually attained, and if yes, we look for a computational algorithm giving a "minimal" set of optimal $\boldsymbol{\tau}$'s. We solve the problem of multiple exercises by an inductive procedure based on the solution of a sequence of single exercise American options with pay-off functions $\xi^{(i)}$ and value function $V^{(i)}$ defined inductively by:

$$\begin{array}{c|c} \xi^{(0)} \equiv 0 & V^{(0)} \equiv 0 \\ \xi^{(1)}_n = \xi_n & V^{(1)} = \hat{\xi}^{(1)} \\ \xi^{(2)}_n = \xi_n + \mathbb{E}_n\{V^{(1)}_{n+1}\} & V^{(2)} = \hat{\xi}^{(2)} \\ \cdots\cdots & \cdots\cdots \\ \xi^{(i)}_n = \xi_n + \mathbb{E}_n\{V^{(i-1)}_{n+1}\} & V^{(i)} = \hat{\xi}^{(i)} \\ \cdots\cdots & \cdots\cdots \\ \xi^{(p)}_n = \xi_n + \mathbb{E}_n\{V^{(p-1)}_{n+1}\} & V^{(p)} = \hat{\xi}^{(p)} \end{array}$$

where we use the hat-notation $\hat{}$ introduced in Chapter 1 for the Snell envelop of a random sequence. Let us explain some of these formulae in words. $\xi^{(i)}$ represents the reward process of exercising one right when the holder of the option still holds i exercise rights. So at time n, $\xi^{(i)}_n$ is equal to the reward ξ_n for immediate exercise of one right, plus the expected reward he or she should expect if from time $n+1$ on, the remaining $i-1$ rights are exercised optimally. In other words, one should have $\xi^{(i)}_n = \xi_n + \mathbb{E}_n\{V^{(i-1)}_{n+1}\}$. Using the results of Chapter 1, we easily get that the maximum expected reward is given by

$$V_0^{(p)} = \mathbb{E}\{\xi_{\boldsymbol{\tau}^*}\}$$

with $\boldsymbol{\tau}^* = (\tau_1^*, \cdots, \tau_p^*) \in \mathcal{S}_{0,T}^{(p)}$ defined by:

$$\begin{aligned}
\tau_1^* &= \inf\{n \geq 0;\ \xi_n^{(p)} = V_n^{(p)}\} \\
\cdots &= \cdots\cdots\cdots \\
\tau_i^* &= \inf\{n > \tau_{i-1}^*;\ \xi_n^{(p-i+1)} = V_n^{(p-i+1)}\}
\end{aligned}$$

Remark. In the Markovian case we can speak of value function instead of value process because, as in the case of the Snell envelop we have:

$$\begin{array}{c|c}
\xi_n^{(1)} = \xi_n = f(Z_n) & V^{(1)} = \hat{\xi}^{(1)} = f^*(Z_n) \\
\xi_n^{(2)} = (f + Pf^*)(Z_n) & V^{(2)} = \hat{\xi}^{(2)} = (f + Pf^*)^*(Z_n) \\
\cdots\cdots\cdots & \cdots\cdots\cdots\cdots \\
\xi_n^{(i)} = (f + P(f + Pf^*)^*)(Z_n) & V^{(i)} = (f + P(f + Pf^*)^*)^*(Z_n) \\
\cdots\cdots\cdots & \cdots\cdots\cdots\cdots
\end{array}$$

5.3.1 **Mathematical Models**

We shall consider several specific models in order to test our numerical implementation. They are both discrete time approximations of the continuous time models which we now describe.

American Put Option

We set ourselves in the Black-Scholes framework of an underlyer $\{S_t\}_t$ following a geometric Brownian motion driven by an $\mathbb{R}$-valued Wiener process $\{W_t\}_t$.

$$S_t = S_0 \exp\left[\left(r - \frac{\sigma^2}{2}\right) t + \sigma W_t\right], \quad t \geq 0 \tag{5.6}$$

The reward process $\{\xi_t\}_t$ is given by the pay-off of the put option $\xi_t = e^{-rt}\phi(S_t)$ where $\phi : \mathbb{R}_+ \longrightarrow \mathbb{R}_+$ is the function $\phi(x) = (K - x)^+$. Under these conditions, and after proper discretization of time, the desired value is given by the optimum

$$V_0^{(p)} = \sup_{(\tau_1,\ldots,\tau_p)\in\mathcal{S}_\delta^{(p)}} \mathbb{E}\left\{\sum_{i=1}^{p} e^{-r\tau_i}\phi(S_{\tau_i})\right\} .$$

Gaussian Mean Reverting Underlyer

Motivated by well known properties of commodity underlyers we consider the example of a call option written on an underlying process $\{S_t\}_t$ following a mean-zero Ornstein-Uhlenbeck process driven by an $\mathbb{R}$-valued Wiener process $\{W_t\}_t$. In this case

$$S_t = e^{-\lambda t}S_0 + \sigma \int_0^t e^{-\lambda(t-s)} dW_t , \quad t \geq 0 \tag{5.7}$$

where λ and σ positive constant used for the mean reversion rate and the volatility. The reward process $\{\xi_t\}_t$ is given by the pay-off of a call option with strike K, say $\xi_t = \phi(X_t)$, where ϕ :

$\mathbb{R}_+ \longrightarrow \mathbb{R}_+$ is now the hockey-stick function $\phi(x) = (x-K)^+$. Under these conditions, and after proper discretization of time, the desired value is given by the same optimum

$$V_0^{(p)} = \sup_{(\tau_1,\ldots,\tau_p)\in\mathcal{S}_\delta^{(p)}} \mathbb{E}\left\{\sum_{i=1}^{p} e^{-r\tau_i}\phi(X_{\tau_i})\right\} .$$

as before.

Mean Reverting CIR Underlyer

As explained in the description of the first example of this chapter, many interest rate models are based on the dynamics of short term interest rates. In order to analyze chooser cap prices, we denote the value of the 3-months LIBOR rate L_3 at time t by r_t and we model the dynamics of the stochastic process $\{r_t\}_{t\geq 0}$ as a mean reverting square root diffusion given by the stochastic differential equation

$$dr_t = -\lambda(r_t - \overline{r})dt + \sigma\sqrt{r_t}dW_t \ , \qquad t \geq 0 \tag{5.8}$$

where λ, $\overline{r}$, and σ are positive constants used for the mean reversion rate, the asymptotic rate, and the volatility, and where $\{W_t\}_t$ is an $\mathbb{R}$-valued Wiener process. We assume that $\sigma^2 \leq 2\overline{r}/\lambda$ which guarantees that r_t remains positive at all times.

Even though we do not have a closed form formula like (2.9) or (**??**) for the solution r_t of (2.12), it is still possible to implement exact simulation procedures to generate values at finitely many times T_i of Monte Carlo samples from (2.12).

The reward process $\{\xi_t\}_t$ is given by the pay-off of a call option with strike K, say $\xi_t = e^{-rt}\Delta t(r_t - K)^+$, given by the hockey-stick function $\phi(x) = (x-K)^+$. Under these conditions, and after proper discretization of time, the desired value is given by the optimum

$$V_0^{(p)} = \sup_{(\tau_1,\ldots,\tau_p)\in\mathcal{S}_\delta^{(p)}} \mathbb{E}\left\{\sum_{i=1}^{p} e^{-r\tau_i}(r_{\tau_i} - K)^+\right\} .$$

5.3.2 **Numerical Results**

This subsection gathers numerical results obtained by the implementation of natural extensions to the multiple exercises case of the Monte Carlo regression method described in the previous chapters.

Multiple American Put Options

We implemented the Monte Carlo computations described in the previous section with a simple piecewise linear least squares regression, and we recomputed the price and the exercise boundaries first analyzed by Carmona and Touzi in [**?**] using Malliavin calculus. See the Appendix at the end of these lecture notes for further details on this approach. As in [**?**] we used $p = 5$ exercise rights on a 1 year American put striked at the money with $S_0 = K = 100$, volatility $\sigma = 0.3$ and short interest rate $r = 5\%$. The exercise boundaries are given in Figure **??**.

Notice the differences with the boundaries given in Figure **??** of the Appendix. We believe that the Monte Carlo results given in Figure **??** are more reliable, though we cannot justify this claim at this stage.

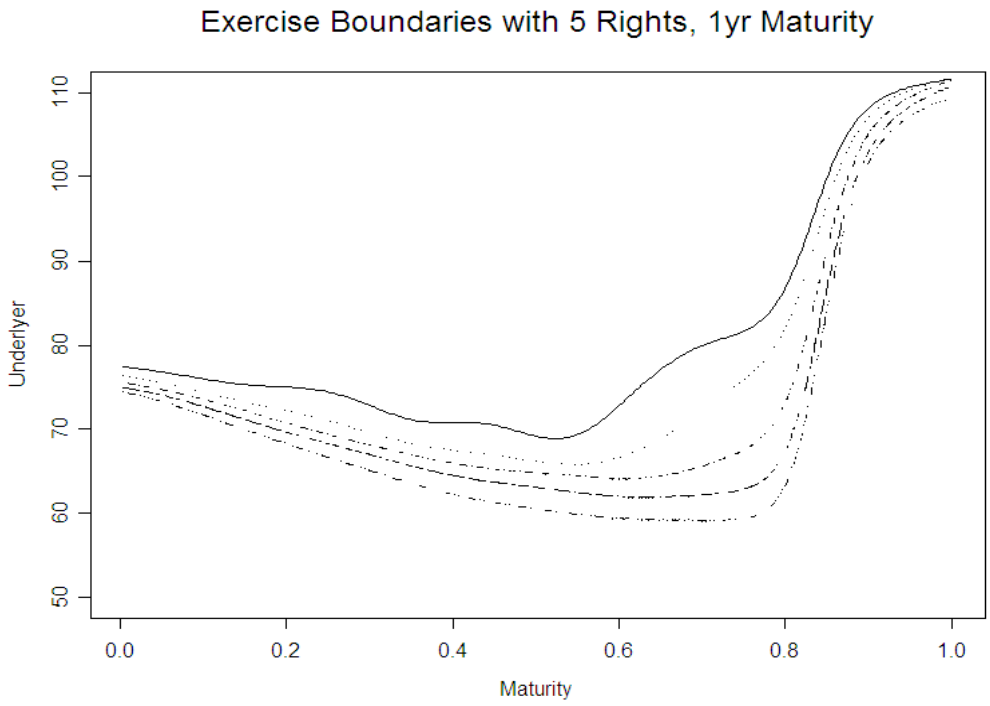

Fig. 5.1. Estimates of 5 exercise boundaries for the 1 year multiple American put option described in the text.

Gaussian Mean Reverting Underlyer

Chooser Cap

We concentrate on 10-years contracts, so the number of possible exercises is $N = 40$, and we specify the (risk neutral) dynamics of the 3-months LIBOR interest rate as discrete time observations of a continuous time stochastic process $\{r_t\}_{t\geq 0}$ chosen to be a mean reverting square root diffusion (known as a CIR model):

$$dr_t = -\lambda(r_t - \overline{r})dt + \sigma\sqrt{r_t}dW_t, \qquad t \geq 0. \tag{5.9}$$

Numerical results are presented in the next figures. Figure **??** for the values $\lambda = 1$, $\overline{r} = 0.05$ and $\sigma = 0.3$.

We varied the number of caplets that can be exercised from $N_c = 1$ to the maximal possible number $N_c = 40$. The values we obtained for the chooser cap are compared to the corresponding values of the auto cap (with the same number N_c of caplets) and the full-fledged cap for two values of the initial interest rate and strike couples, $r_0 = K = 0.01$ and $r_0 = K = 0.05$.

Remark. As we already mentioned, for $N_c = 40$ the price of the chooser cap reduces to the price of a regular cap. In the case of the Vasicek model, the latter can be computed explicitly by a change of numeraire technique as explained in the book [**?**] by Brigo and Mercurio. See Theorem 4.2.1 and the interest rate dynamics in the T-forward measure given in Lemma 4.2.2.

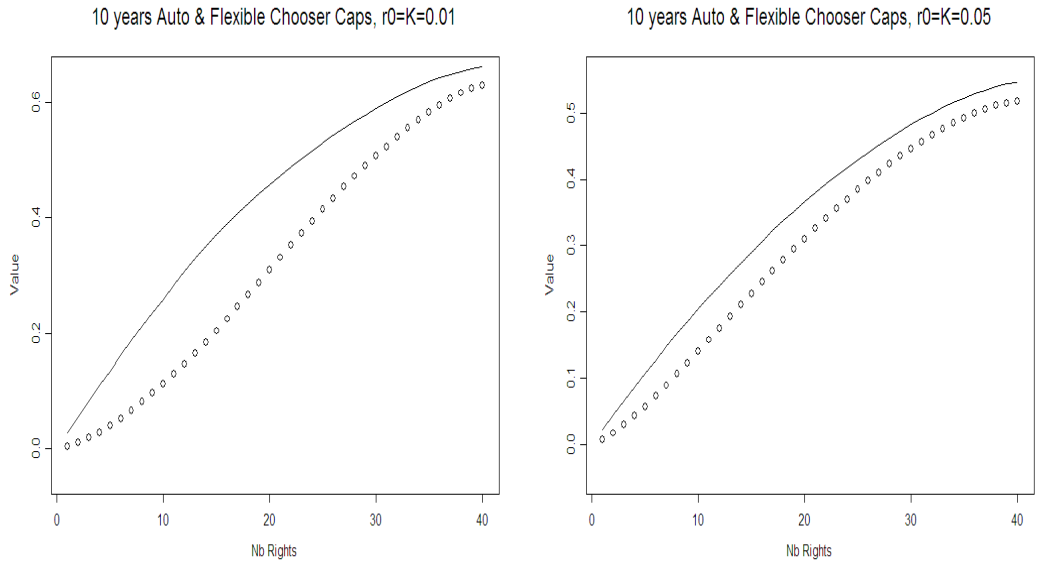

Fig. 5.2. Values of 10 years auto and chooser caps with $r_0 = K = 0.01$ (left) and $r_0 = K = 0.05$ (right).

It appears that as expected, the prices of chooser caps are much higher than the prices of auto caps. An intuitive explanation goes as follows. When a caplet is exercised in an auto cap, it is likely that it is barely in the money. However, in a chooser cap, when a caplet is not far enough in the money, it is possible to wait until a later date and exercise the caplet when it is much deeper in the money. The success of such a change in strategy obviously depends upon the current term structure of interest rates, and whether or not the yield curve is upward or downward sloping.

Figure **??** illustrate the sensitivities of the flexible chooser cap prices with respect to the mean reversion parameter λ (left) and the volatility factor σ (right).

The results reproduced in Figure **??** were produced with the following parameters. In both cases we fixed $r_0 = K = \overline{r} = 0.05$. For the dependence with respect to λ we chose $\sigma = 0.1$ and we varied the mean reversion parameter λ by choosing $\lambda = 5$, $\lambda = 2$, $\lambda = 0.8$, $\lambda = 0.4$ and $\lambda = 0.15$ successively. For the dependence with respect to σ we chose $\lambda = 0.5$ and we varied the volatility parameter σ by choosing $\sigma = 0.001$, $\sigma = 0.05$, $\sigma = 0.1$, $\sigma = 0.15$ and $\sigma = 0.2$ successively.

Notice that the various combinations of parameter values used in the numerical experiments reported above all satisfy the condition $\sigma^2 \leq 2\lambda\overline{r}$ which guarantees the positivity at all times of the square root diffusion, ie.e the solution of the stochastic differential equation (2.12).

5.4 MULTIPLE DYNKIN GAMES

As a motivation for the models considered in this section, we revisit the commodity delivery contract discussed earlier, embedding both the recall and swing options in the indenture of the contract.

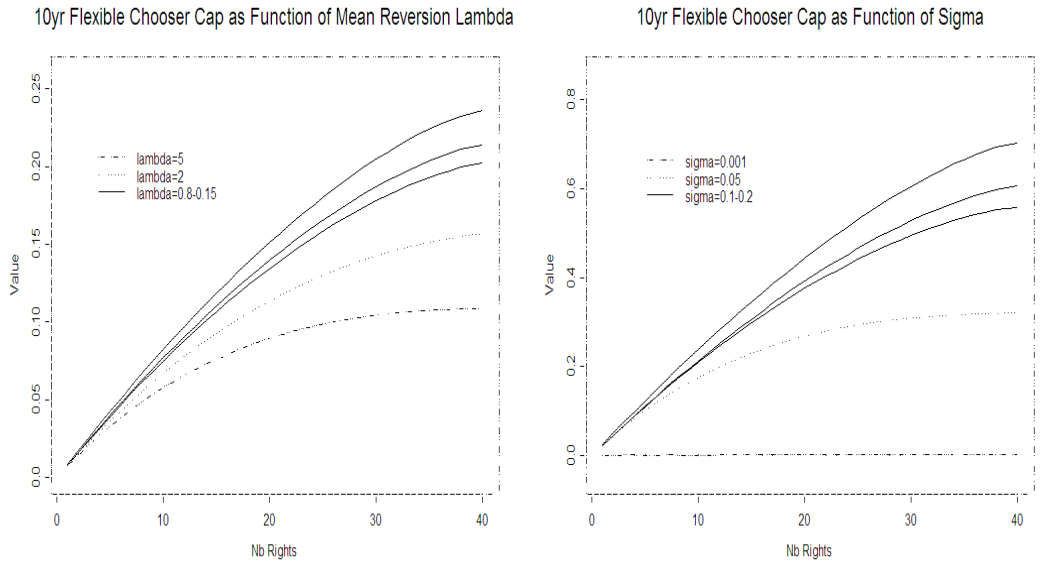

Fig. 5.3. Values of 10 years auto and chooser caps with $r_0 = K = 0.01$ (left) and $r_0 = K = 0.05$ (right).

5.4.1 **Two-sided Swing Contracts**

We now consider a base-load contract augmented with N_r recall rights for the party delivering the commodity, and N_s swing rights for the counter-party. Obviously, we are about to merge the discussions of the two preceding subsections. But before doing just that, we need to alter the cash-flow direction in one of the two examples as they are opposite one of the other. For the sake of consistency with our discussion of the applications of Dynkin games to the valuation of convertible bonds, in this subsection we consider only cash flows from the buyer of the commodity to the seller. For that reason, we keep the meaning of the notation L_n, but we replace U_n by its negative. Hence, the buyer of the commodity becomes the *minimizing* agent while seller of the commodity becomes the *maximizing* agent, and we denote by $\tau^{(r)} = (\tau_1^{(r)}, \tau_2^{(r)}, \cdots, \tau_{N_r}^{(r)})$ with $\tau_1^{(r)} < \tau_2^{(r)} < \cdots < \tau_{N_r}^{(r)}$ the stopping times chosen by the maximizing agent for the recall times, and by $\tau^{(s)} = (\tau_1^{(s)}, \tau_2^{(s)}, \cdots, \tau_{N_s}^{(s)})$ where $\tau_1^{(s)} < \tau_2^{(s)} < \cdots < \tau_{N_s}^{(s)}$ are the stopping times chosen by the minimizing agent for the swing times. We then define

$$J(\tau_1^{(r)}, \cdots, \tau_{N_r}^{(r)}, \tau_1^{(s)}, \cdots, \tau_{N_s}^{(s)})$$

as the expected present value of the pay-off to the buyer of the two sided swing contract. Finally we set

$$J^{(r)} = \sup_{\tau_1^{(r)}<\cdots<\tau_{N_r}^{(r)}} \inf_{\tau_1^{(s)}<\cdots<\tau_{N_s}^{(s)}} J(\tau_1^{(r)}, \cdots, \tau_{N_r}^{(r)}, \tau_1^{(s)}, \cdots, \tau_{N_s}^{(s)})$$

for the buyers value of the contract and

$$J^{(s)} = \inf_{\tau_1^{(s)}<\cdots<\tau_{N_s}^{(s)}} \sup_{\tau_1^{(r)}<\cdots<\tau_{N_r}^{(r)}} J(\tau_1^{(r)}, \cdots, \tau_{N_r}^{(r)}, \tau_1^{(s)}, \cdots, \tau_{N_s}^{(s)})$$

By construction we always have:

$$J^{(r)} \le J^{(s)}$$

Moreover we also have

Theorem 1.

$$J^{(r)} = J^{(s)}$$

We will denote by J or $J^{(N_r,N_s)}$ the common value of $J^{(r)}$ and $J^{(s)}$.

5.4.2 Recursive Hierarchy of Dynkin Games

We will rely on the notation of our first discussion of Dynkin games in Section 4.4. We only emphasize the changes due to the recusrive nature of the set-up of this subsection.

We start with two integrable stochastic processes $\xi^{(r)} = \{\xi_n^{(r)}\}_{0\le n\le N}$ and $\xi^{(s)} = \{\xi_n^{(s)}\}_{0\le n\le N}$ whose interpretations are as follows. The random variable $\xi_n^{(r)}$ represents the amount earned (or lost when $\xi_n^{(r)} < 0$) by the first player if she exercises one right at time n. Similarly the random variable $\xi_n^{(s)}$ represents the amount paid by the second player to the first player, if the second player exercises one right at time n. So $\xi_n^{(r)}$ and $\xi_n^{(s)}$ both represent cash flows from the second player to the first one.

Let us compare the direction of the cash-flows descried above with those in the case of the two-sided swing contracts: the regular payments go from the buyer of the commodity to the seller. So if we want to use the terminology of the Dynkin games of timing used in the case of convertible bonds models, the *maximizing player* should be the seller of the commodity, possibly a retailer, and the *minimizing player* should be the buyer of the commodity, possibly a consumer.

Next, we fix two integers $N_r \ge 0$ and $N_s \ge 0$ and we define inductively the random sequences $V^{(i,j)} = \{V_n^{(i,j)}\}_{0\le n\le N}$ starting with $V^{(0,0)} \equiv 0$. For each integer $1 \le i \le N_r$ and $1 \le j \le N_s$ we define $V^{(i,j)}$ as the value of the Dynkin game of timing with pay-offs

$$\begin{aligned} L_n^{(i,j)} &= \xi_n^{(r)} + \mathbb{E}_n\{V_{n+1}^{(i-1,j)}\} \\ U_n^{(i,j)} &= \xi_n^{(s)} + \mathbb{E}_n\{V_{n+1}^{(i,j-1)}\}. \end{aligned}$$

When $i = 0$ and $1 \le j \le N_s$, we define $V^{(0,j)}$ as the value function of a multiple American option with j exercise rights and single exercise pay-off $\xi^{(s)}$. Similarly, for $j = 0$ and $1 \le i \le M$, $V^{(i,0)}$ is defined as the value function of a multiple American option with i exercise rights and single exercise pay-off $\xi^{(r)}$. So according to our discussion of Section **??** on instruments with multiple American exercises, if $i \ge 1$, then $V^{(i,0)}$ is the value of an American option with pay-off

$$\xi_n^{(i,\cdot)} = \xi_n^{(r)} + \mathbb{E}_n\{V_{n+1}^{(i-1,0)}\},$$

while if $j \ge 1$, then $V^{(0,j)}$ is the value of an American option with pay-off

$$\xi_n^{(\cdot,j)} = \xi_n^{(s)} + \mathbb{E}_n\{V_{n+1}^{(0,j-1)}\}.$$

Theorem 2.

$$J^{(M,N)} \equiv V^{(M,N)}.$$

5.5 TWO-SIDED SWING CONTRACTS AS RECURSIVE DYNKIN GAMES

Let us assume that the stochastic process $\{S_t\}_{t\geq 0}$ gives the spot price of a commodity on which a two-sided swing contract is written. As before we denote by N_r and N_s the numbers of recall and swing exercise rights respectively, by $K > 0$ the strike of the base-load contract, and by d and u the quantities of down and up swings. We now introduce the reward processes $\xi^{(r)}$ and $\xi^{(s)}$ defining a recursive Dynkin game which is equivalent to the two-sided swing contract. As explained earlier, we consider cash flows in one direction only, and we choose flows from the seller to the buyer.

We first consider the commodity seller side, and we denote by $\xi_n^{(r)}$ the present value at time $n = 0$ of the cash flow from the seller to the buyer if the seller exercises one of his rights at time n. We have:

$$\xi_n^{(r)} = \phi^{(r)}(S_n) = e^{-rn}[\alpha(S_n - K) + f_r] \tag{5.10}$$

where f_r denotes the *recall fee* the commodity seller has to pay in order to be allowed to exercise one right. On the other hand, the cash flow from the commodity seller to the buyer when the latter decides to exercise one of his swing rights is given by

$$\xi_n^{(s)} = \phi^{(s)}(S_n) = e^{-rn}[u\alpha(S_n - K)^+ d\alpha(K - S_n)^+ - f_s] \tag{5.11}$$

where the constant f_s represents the swing fee the commodity buyer has to pay for each swing exercise.

5.5.1 Numerical Results

Base-load Delivery Contract

M-recalls Option

N-swings Option

A Two-sided Swing Contract

5.6 OPTIMAL SWITCHING AND REAL-OPTION ASSET-VALUATION

NOTES & COMPLEMENTS

The inductive procedure for multiple American exercises was implemented by Meinshausen and Hambly in [**?**] to value chooser interest rate swaps. These authors used an extension to the multiple exercise case of the duality upper bounds presented in Chapter **??**.

Detailed proofs of Theorem **??** and Theorem **??** can be found in the recent paper [**?**] by Carmona and Dayanik.

The optimal switching version of the impulse control problem discussed above was introduced by Carmona and Ludkovski in [**?**] and [**?**] to value power plants and gas storage facilities.

6

Duality and Pointwise Approach

In the first part of this last chapter, we review the duality approach to the optimal stopping problem and we explain how it can lead to Monte Carlo computations of approximate prices for American options. This method which was introduced approximately at the same time as the Tsitsiklis-van Roy and Longstaff-Schwartz methods. Despite the fact that it did not get the same *press*, it deserves consideration for two main reasons: the first one is that by reformulating the problem as an infimum over martingales (instead of a supremum over stopping times), it provides lower bounds for the value function of the problem, and these upper bounds can in principle be tight if the martingale is chosen appropriately. The second reason is the fact that these upper bounds can easily be computed by Monte Carlo method. Indeed, once a martingale is chosen, the upper bound is the expectation of a maximum which can be computed path by path at very little cost.

The last part of the chapter is devoted to the extension of this pointwise strategy to the case of Dynkin games.

6.1 INTRODUCTION

This chapter is motivated by recent developments in the numerics of American option valuation. Based on an idea of Davis and Karatzas [**?**] three groups of authors have proposed, essentially simultaneously, a duality formulation where the supremum over stopping times is replaced by an infimum over martingales. Even though the duality argument is the same in the three papers, the details of their implementations are slightly different. While Haug and Kogan use neural networks [**?**], Rogers uses potential theory to guess *a right martingale* [**?**], and Andersen and Broadie focus on unbiased estimates by embedding smaller trees within the Monte Carlo structure set up to compute the expectation once the martingale is chosen [**?**]. Combined with easily obtained lower bounds, the upper bounds provided by the duality approach proved to be very useful in squeezing the solution of the optimal stopping problem in tight intervals. These bounds were used in the multiple exercise case by Meinshausen and Hambly in [**?**].

What should the right analog of the duality theory for Dynkin games be is not clear at this stage. However the Davis and Karatzas' pathwise formulation of the optimal stopping problem was extended by Karatzas in [**?**] to cover the case of Dynkin games of timing. We give a detailed proof of this fact to lay the foundations for the development of efficient Monte Carlo procedure to compute bounds for the values of these games.

6.2 OPTIMAL STOPPING: DUALITY AND UPPER BOUNDS

We use freely the notation of Chapter 1. Let $\{M_n\}_n$ be any (bounded) martingale in the given filtration $\{\mathcal{F}_n\}_n$. We have (recall that $\{Y_n\}_n$ is the reward process and $D(0,n)$ is the discounting factor):

$$\begin{aligned}\hat{Y}_0 &= \sup_{\tau\in\mathcal{S}_{0,N}} \mathbb{E}\{D(0,\tau)Y_\tau\}\\ &= \sup_{\tau\in\mathcal{S}_{0,N}} \mathbb{E}\{M_\tau + D(0,\tau)Y_\tau - M_\tau\}\\ &= M_0 + \sup_{\tau\in\mathcal{S}_{0,N}} \mathbb{E}\{D(0,\tau)Y_\tau - M_\tau\}\\ &\leq M_0 + \mathbb{E}\{\max_{n=0,1,\cdots,N}[D(0,n)Y_n - M_n]\} \qquad (6.1)\end{aligned}$$

where we used the optional sampling theorem for martingales to derive the third equality. Hence, we proved that

$$\hat{Y}_0 \leq \inf_{\{M_n\}_n\in\mathcal{M}_0} \mathbb{E}\{\max_{n=0,1,\cdots,N}[D(0,n)Y_n - M_n]\}, \qquad (6.2)$$

where $\mathcal{M}_0$ denotes the set of $\mathcal{F}_n$-martingales $\{M_n\}_n$ vanishing at time 0 (i.e. for which $M_0 = 0$). This implies that, choosing a martingale $\{M_n\}_n$ starting from $M_0 = 0$, and computing the quantity

$$\mathbb{E}\{\max_{n=0,1,\cdots,N}[D(0,n)Y_n - M_n]\} \qquad (6.3)$$

provides an upper bound for the value of the American option. The computation of such an upper bound is very simple as it does not require the computation of conditional expectations. Moreover as it is expressed as an expectation, it is screaming for a Monte Carlo evaluation. Indeed, if one can generate sample scenarios $\omega^{(m)}$ for the discounted pay-off and the martingale, one can compute the difference $D(0,n)Y_n(\omega^{(m)}) - M_n(\omega^{(m)})$ for each time n, then compute the maximum over n, and finally average out the resulting maxima over m to get an approximation of the expectation.

The ease with which one should be able to compute these upper bounds is one of the main attractions of the method. But the main question remains: how good are these upper bound? We now show that, at least theoretically, they could be as good as we need. Indeed the duality gap collapses as we have:

$$\hat{Y}_0 = \inf_{\{M_n\}_n\in\mathcal{M}_0} \mathbb{E}\{\max_{n=0,1,\cdots,N}[D(0,n)Y_n - M_n]\}. \qquad (6.4)$$

To prove this result, we identify a martingale for which the upper bound (**??**) is tight, i.e. for which equality holds. The stochastic process $\{D(0,n)\hat{Y}_n\}_n$ giving the discounted price of the American option is a super-martingale, so it has a Doob-Meyer decomposition

$$D(0,n)\hat{Y}_n = \hat{Y}_0 + \tilde{M}_n - \tilde{A}_n \qquad (6.5)$$

where $\{\tilde{M}_n\}_n$ is a martingale in $\mathcal{M}_0$ and $\{\tilde{A}_n\}_n$ is a non-decreasing predictable process (i.e. $\tilde{A}_n$ is $\mathcal{F}_{n-1}$-measurable). The Doob-Meyer decomposition is a rather deep result for continuous time stochastic processes, but it is quite simple in the discrete case. Indeed one can write any adapted

process $\{Z_n\}_n$ as $Z_n = Z_0 + M_n + A_n$ where $\{M_n\}_n$ is a martingale such as $M_0 = 0$ and $\{A_n\}_n$ is a predictable process such as $A_0 = 0$. Indeed, it is enough to choose

$$M_n = \sum_{j=1}^{n} \left(Z_j - \mathbb{E}_{j-1}\{Z_j\}\right), \qquad \text{and} \qquad A_n = \sum_{j=1}^{n} (\mathbb{E}_{j-1}\{Z_j\} - Z_{j-1}). \tag{6.6}$$

This fact is intuitively clear: A_n is merely whatever is needed to compensate for the fact that Z_n may not be equal to the conditional expectation of its next value. Moreover, it is easy to check that A_n is non-negative and non-decreasing in n if and only if Z_n is a super-martingale. Now, using the martingale part of the Doob-Meyer decomposition of $D(0,n)\hat{Y}_n$ as a test martingale in (**??**) we get:

$$\hat{Y}_0 \leq \hat{Y}_0 + \mathbb{E}\left\{\max_{0\leq n\leq N}\left(D(0,n)Y_n - D(0,n)\hat{Y}_n - \tilde{A}_n\right)\right\} \leq \hat{Y}_0$$

where we used the fact that $Y_n \leq \hat{Y}_n$ and $\tilde{A}_n \geq 0$. So the martingale of the Doob-Meyer decomposition of the discounted value of the American option closes the duality gap!

Remark. It is important to emphasize the fact that the duality gap is closed because the value function V_n (i.e. the discounted Snell's envelop) is a super-martingale, fact which is true whether or not the pay-off process Y_n is non-negative.

6.2.1 Implementation Issues

- In order to take full advantage of the path - by - path form of the duality theory, one needs to identify *good martingales*. In the Markovian case, martingales can be obtained by computing harmonic functions for the infinitesimal generator of the process along the sample paths. If a basis $\{\varphi_1, \varphi_2\}$ can be computed explicitly (by solving an ordinary differential equation) it is reasonable to look for the coefficients λ_1 and λ_2 for which $M_n = \lambda_1\varphi_1(X_n) + \lambda_2\varphi_2(X_n)$ gives the best (i.e. the smallest) upper bound. This gives **FINISH**
- When looking for this elusive martingale capable of providing a good upper bound to the value of the optimum, a reasonable strategy could be to use the martingale part of the Doob-Meyer decomposition of an easy to compute approximation to the value function. Indeed, such approximations are easy to come by, for example by choosing a guestimate of the exercise boundary, and approximating the true value function by the expected discounted pay-off when the underlying process first hits the exercise boundary proxy. However, as one sees from formula (**??**), the computation of the martingale part of a Doob-Meyer decomposition involves the computation of many conditional expectations. Hence this route may not bring any relief to the computational burden. Moreover, computing conditional expectations by non-parametric methods always carries a bias, and in order to avoid this bias, Andersen and Broadie proposed to compute these conditional expectations by a subsampling procedure, embedding a fixed size smaller tree at each node of the Monte Carlo samples of the underlyer.

In our opinion, the direct search for a good martingale is still more of an art than a science.

6.2.2 Extension to Multiple Exercises

As explained in in [**?**] by Meinshausen and Hambly, the duality bounds can be used to approximate the values of instruments with multiple American exercises. **FINISH**

6.3 DUALITY SQUEEZES

Throughout these lectures, we derive approximations to the value of instruments with one or several American exercise opportunities. But not only didn't we worry about how good or bad these approximations were, but we did not attempt to identify possible biases. The possibility to squeeze the true value in between a lower and an upper bound is one of the main reasons for the success of the duality approach, especially since the *confidence* intervals straddling the unknown optimal value happen to be quite short in many practical applications without the need for an excessive computation budget.

6.3.1 Lower Bounds

As we already mentioned, lower bounds are relatively easy to come by. From any reasonable candidate for an exercise strategy, say τ, the expectation

$$\mathbb{E}\{D(0,\tau)\varphi(X_\tau)\}$$

is obviously a lower bound to the optimum. Given a Monte Carlo sample $\omega^{(1)}, \cdots, \omega^{(m)}$, an unbiased estimate of this expectation is given by:

$$\ell(\tau) = \frac{1}{M}\sum_{m=1}^{M} D(0,\tau(\omega^{(m)}))f(X_{\tau(\omega^{(m)})}(\omega^{(m)})).$$

This random number may or may not be below the true optimum, but its expectation is. So it is plain to use Monte Carlo scenarios to produce a downward biased estimate if we can identify such a reasonable candidate for the exercise strategy. This can be done in several different ways using the tools developed in Chapter 3. For example, one can use the last stopping time produced by the Longstaff-Schwartz backward induction. But we can also use an approximate value function, say the result of the dynamic programming backward induction performed with any kind of regression method, in which case the exercise strategy can by identified as the first time the Monte Carlo sample paths hit the set where the approximate Snell's envelop (i.e. the value function computed along the paths) coincides with the reward.

Remark. At this stage, we should not call such an estimate a lower bound as it is only in expectation that we can be sure that it lies below the unknown true value. However if we construct a confidence interval around this negatively biased estimate, say with confidence level $\alpha/2$, the lower limit of this confidence interval provides a lower bound for the true value with probability $1 - \alpha/2$.

6.3.2 Upper Bounds

Even though the search for an optimal martingale is still elusive in general, the duality theory presented in this chapter offers a straightforward tool to turn the lower bound computed above in Subsection **??** into an upper bound. Indeed, one can use the upper bound provided by the martingale of the Doob-Meyer decomposition (**??**) of the approximate Snell envelop. But before doing so, we need to tackle the following problem: as explained earlier, this computation requires the evaluation of a large number of conditional expectations, and if these conditional expectations are estimated with a bias (as we would do using plain non-parametric procedures) the resulting process may not be a real martingale and the resulting computation of what should be an upper bound in expectation may

not be so. For this reason, it is recommended to estimate these conditional expectations by regular averages (i.e. plain estimates of regular expectations) over new independent samples generated according to the *conditional distribution* at the node where the estimate of the conditional expectation is needed. This extra sampling adds to the complexity of the algorithm, but its effect on the complexity is only a constant factor independent of the number M of original Monte Carlo samples and of the number of time steps as the number of branches of the extra trees is constant over the nodes.

A direct application of the strategy described above is the analog of the Tsitsiklis - van Roy approach to the implementation of the dynamic programming backward induction. But as in Chapter 3, there is an alternative backward induction in terms of approximately optimal stopping times. This provided a Longstaff-Schwartz version of the argument. This idea was formulated and tested by Andersen and Broadie in [**?**]

Remark. As in the case of the upper bound, the above estimate is only an upper bound in expectation. However if as before we construct a confidence interval around this positively biased estimate, say with confidence level $\alpha/2$, then the random interval limited by the lower bound of the confidence interval computed earlier in Subsection **??** and the upper limit of the present confidence interval, contains the true value of the optimum with probability $1 - \alpha$. This idea was promoted by Broadie and Glasserman. See for example [**?**] Chapter 8.

6.3.3 Numerical Examples

6.4 PATHWISE APPROACH TO DYNKIN GAMES

In this section, we use the notation and the results of Section 4.4 on Dynkin games of timing. In particular, we use the notation L_n and U_n for the pay-offs to the maximizer player when respectively, the maximizer and the minimizer player stop the game, and V_n for the value function. One of the main results of Section 4.4 is a form of the dynamic programming principle which can be summarized as

$$V_n = \text{med}\{L_n, \mathbb{E}_n\{V_{n+1}\}, U_n\} \tag{6.7}$$

where we use the notation $\text{med}\{x_1, x_2, \cdots, x_n\}$ for the median of the sample $x_1, x_2, \cdots, x_n$.

Recalling our discussion of the Doob-Meyer deconposition for discrete time processes, and especially formulae (**??**) and (**??**), the process $\{M_n\}_n$ defined by

$$V_n = V_0 + M_n + \sum_{j=1}^{n} (\mathbb{E}_{j-1}\{V_j\} - V_{j-1})$$

is a martingale in $\mathcal{M}_0$. Again the compensation is only here to correct for what prevents V_n to be equal to $\mathbb{E}_n\{V_{n+1}\}$. Looking back at the form of the dynamic programming principle given by (**??**), the summation above can be split into two separate sums: a first sum over the indexes j for which $\mathbb{E}_{j-1}\{V_j\} > V_{j-1}$ (in which case $V_{j-1} = U_{j-1}$) and a second sum over the indexes j for which $\mathbb{E}_{j-1}\{V_j\} < V_{j-1}$ (in which case $V_{j-1} = L_{j-1}$). Consequently we have the decomposition

$$V_n = V_0 + M_n + A_n - B_n \tag{6.8}$$

for some martingale $\{M_n\}_n$ in $\mathcal{M}_0$, and for predictable processes $\{A_n\}_n$ and $\{B_n\}_n$ defined by:

$$A_n = \sum_{j=1}^{n} (\mathbb{E}_{j-1}\{V_j\} - U_{j-1})^+, \qquad n = 1, 2, \cdots, N \tag{6.9}$$

and

$$B_n = \sum_{j=1}^{n} (L_{j-1} - \mathbb{E}_{j-1}\{V_j\})^+, \qquad n = 1, 2, \cdots, N. \tag{6.10}$$

Notice also that both predictable processes $\{A_n\}_{1\leq n\leq N}$ and $\{B_n\}_{1\leq n\leq N}$ are non-negative and non-decreasing. Also, by the very definition of these processes we have

$$A_n = A_{\tau_n^S}, \quad \text{and} \quad B_n = B_{\tau_n^B}, \qquad n = 1, 2, \cdots, N, \tag{6.11}$$

where the stopping times τ_n^S and τ_n^B are the minimal optimal exercise times of the buyer and the seller as defined in (**??**) and (**??**).

We state the result of the decomposition which we just proved as a lemma for future reference.

The fact that the upper bound (**??**) is tight follows from Karatzas analysis [**?**]. We reproduce the proof for the sake of completeness.

For any fixed $0 \leq n \leq N$ we have

$$R(n, \tau_s^*) - M^*_{n\wedge\tau_s^*} = (L_n - M_n)\mathbf{1}_{n\leq\tau_s^*, n<N} + (U_{\tau_s^*} - M_{\tau_s^*})\mathbf{1}_{n>\tau_s^*} + (\xi - M_N)\mathbf{1}_{n=\tau_s^*=N}$$

6.4.1 **More Bounds**

We now prove the easy upper and lower bounds for the value of the game.

$$\sup_{M^-\in\mathcal{M}_0} \mathbb{E}\{\min_{0\leq n\leq N}(U_n - M_n^-)\} \leq V_0 \leq \inf_{M^+\in\mathcal{M}_0} \mathbb{E}\{\max_{0\leq n\leq N}(L_n - M_n^+)\} \tag{6.12}$$

In order to prove the upper bound we set $\tau_s = N$ (in which case $R(\tau_b, \tau_s) = L_{\tau_B}$) and we apply the duality theory of Section **??** to the reward process $\{U_n\}_n$. We get:

$$\begin{aligned} V_0 &= \sup_{\tau_b}\inf_{\tau_s} \mathbb{E}\{R(\tau_b, \tau_s)\} \\ &\leq \sup_{\tau_b} \mathbb{E}\{R(\tau_b, N)\} \\ &\leq \sup_{\tau_b} \mathbb{E}\{L_{\tau_b}\} \\ &= \inf_{M^+\in\mathcal{M}_0} \mathbb{E}\{\max_{0\leq n\leq N}(L_n - M_n^+)\} \end{aligned}$$

Similarly, in order to prove the lower bound we set $\tau_b = N$ (in which case $R(\tau_b, \tau_s) = U_{\tau_s}$) and we use duality for the optimal stopping problem with reward process $\{-L_n\}_n$. We get:

$$
\begin{aligned}
V_0 &= \sup_{\tau_b} \inf_{\tau_s} \mathbb{E}\{R(\tau_b, \tau_s)\} \\
&\geq \inf_{\tau_s} \mathbb{E}\{R(N, \tau_s)\} \\
&\geq -\sup_{\tau_s} \mathbb{E}\{-U_{\tau_s}\} \\
&= -\inf_{M^- \in \mathcal{M}_0} \mathbb{E}\{\max_{0 \leq n \leq N} (-U_n + M_n^-)\} \\
&= \sup_{M^- \in \mathcal{M}_0} \mathbb{E}\{\min_{0 \leq n \leq N} (M_n^- - U_n)\}
\end{aligned}
$$

NOTES & COMPLEMENTS

The duality approach to American option pricing was proposed simultaneously by Haug and Kogan in [**?**] and Rogers [**?**]. The fundamental idea can be traced back to an earlier paper by Davis and Karatzas [**?**]. This method was used in the multiple exercise case by Meinshausen and Hambly in [**?**]. The implementation of the duality approach was more recently refined by Andersen and Broadie in [**?**]. In this chapter, we used the additive form of the Doob-Meyer decomposition. There exists a multiplicative form as well, and a duality theory based on this form was developed by Jamshidian in [**?**]. It is unclear which method is more efficient numerically, so we chose to present and work with the additive version as it is slightly simple to present and implement, and most importantly, because it has been studied most extensively.

7

Appendix: Monte Carlo Regression

As evidenced by our discussion of the dynamic programming principle, the practical implementations of this backward procedure as well as some of the implementations of the duality approach require the computation of a large number of conditional expectations. This appendix is devoted to a review of some of the regression methods which have been used in this context. We first review the basic facts of parametric regression by restricting the function f to families of functions which can be characterized by a small number of parameters. But since in some cases we do not want to assume anything a priori on the regression function, except possibly for requiring that is (at least piecewise) smooth, we then concentrate on non-parametric regression techniques. Also, because of the applications we are interested in, we mostly work in the multidimensional case, and especially in cases where the dimension p of the explanatory variables is relatively large in the sense that we would like to have methods working reasonably well for values of p up to 100 and possibly higher.

7.1 REVIEW OF THE REGRESSION SET UP

We review the notation and the general set up of regression. This will give us a chance to stress the main differences between the parametric point of view and the nonparametric approach.

Given a random variable $\mathbf{X}$ in a measurable space $(E, \mathcal{E})$ and a real valued random variable Y, the regression of Y against $\mathbf{X}$ is the function

$$E \ni \mathbf{x} \hookrightarrow f(\mathbf{x}) = \mathbb{E}\{Y|\mathbf{X} = \mathbf{x}\} \tag{7.1}$$

which is defined almost surely on E for the distribution of $\mathbf{X}$. Obviously this definition can be generalized to the case where Y takes values in a more general vector space as long as the notion of expected value for E-valued random variables is clearly defined. For the purpose of this appendix, $\mathbf{X}$ will be a real valued random variable or a multivariate random vector taking values in a Euclidean space $\mathbb{R}^p$. Statisticians have developed practical methods of estimating regression functions from sample data. From such a statistical point of view, the starting point is a sample of n observations

$$(\mathbf{x}_1, y_1), \cdots\cdots, (\mathbf{x}_n, y_n)$$

where for each $i = 1, 2, \cdots, n$, $\mathbf{x}_j$ is a vector of p numerical components $\mathbf{x}_i = [x_{i,1}, x_{i,2}, \cdots, x_{i,p}]$, and y_i is a real number. The components of the $\mathbf{x}_i$'s are observations of the p explanatory (scalar) variables, while the y_i are observations of the response variable. The theory justifying the estimation procedures is developed under the assumption that the data are sample observations of random

couples $(\mathbf{X}_1, Y_1), \cdots, (\mathbf{X}_n, Y_n)$ having the same distribution, say the distribution of the couple $(\mathbf{X}, Y)$ having the regression function (**??**). For the sake of notation, we use the notation $(\mathbf{X}, Y)$ for a generic couple of random vector, random variable with the same joint distribution. The statistical dependence of the Y-component upon the $\mathbf{X}$-components is determined by the knowledge of the (common) joint distribution of the couples $(\mathbf{X}_i, Y_i)$'s. This distribution *lives* in $(p+1)$ dimensions. It is determined by the marginal distribution of the p-variate random vector $\mathbf{X}$ (which lives in p dimensions), and the conditional distribution of Y which gives, for each possible value $\mathbf{x}$ of $\mathbf{X}$, the conditional distribution of Y given that $\mathbf{X} = \mathbf{x}$. The regression function of Y against $\mathbf{X}$ gives the expectation of this conditional distribution, so it does not contain all the information about the conditional distribution, just its mean. The graph of the regression function is a one-dimensional curve when $p = 1$, it is a 2-dimensional surface when $p = 2$, and it becomes a hyper-surface more difficult to visualize for larger values of the number p of explanatory variables.

7.1.1 Practical Objective

In most practical applications the search for the regression estimate $\hat{f}$ is done by

- computing the real numbers $\hat{f}(\mathbf{x})$ for all the values $\mathbf{x} = \mathbf{x}_i$ for $i = 1, \cdots, n$, or by
- computing the real numbers $\hat{f}(\mathbf{x})$ for the values $\mathbf{x} = \overline{\mathbf{x}}_i$ for a grid $\overline{\mathbf{x}}_1, \cdots, \overline{\mathbf{x}}_{\overline{n}}$.

The option described in the second bullet point is used when the sample size n is too large for the function $\hat{f}$ to be computed at all the observations $\mathbf{x}_i$, In this case a grid of values of $\mathbf{x}$ with cardinality $\overline{n}$ much smaller than n is chosen to be as fine enough for the computations of $\hat{f}(\mathbf{x})$ to be computed by simple linear interpolation from the values of $\hat{f}$ at the points of the grid without loosing too much information.

7.2 PARAMETRIC REGRESSION

The idea of parametric regression is to search for an estimate $\hat{f}(\mathbf{x})$ in a specific class of functions $\{f_\theta\}_{\theta \in \Theta}$ parameterized by a parameter θ varying in a parameter set Θ, the chosen estimate $\hat{\theta}$ minimizing a criterion quantifying the fit of the function f_θ to the data. For example, the least squares estimate would be a value $\hat{\theta}$ minimizing the least squares error criterion, i.e.

$$\hat{\theta} = \arg \inf_{\theta \in \Theta} \sum_{i=1}^{n} |y_i - f_\theta(\mathbf{x}_i)|^2. \tag{7.2}$$

Similarly, the least absolute deviations estimate would be a value $\hat{\theta}$ minimizing the least absolute deviations criterion, i.e.

$$\hat{\theta} = \arg \inf_{\theta \in \Theta} \sum_{i=1}^{n} |y_i - f_\theta(\mathbf{x}_i)|. \tag{7.3}$$

A more systematic procedure is to assume that the joint distribution of $(\mathbf{X}, Y)$ belongs to a family $f_\theta(\mathbf{x}, y)$, and then to use the maximum likelihood procedure to estimate the parameter θ from the sample observations. This procedure recovers the least squares estimate when the conditional distributions are Gaussian, and the least absolute deviations estimate in the case of double exponential distributions.

7.2.1 **Standard Examples**

Except possibly for the case of simple and multiple linear regressions, we shall not work with least absolute deviations regression, and we shall concentrate on least squares regression. The most commonly used least squares regression procedures are

- simple or multiple regression
- polynomial regression
- piecewise linear and splines regression

7.3 NON-PARAMETRIC POINT OF VIEW

Let us assume for example that the set of the sample values $\mathbf{x}_1, \cdots, \mathbf{x}_n$ is equal to a smaller set $\tilde{\mathbf{x}}_1$, $\cdots$, $\tilde{\mathbf{x}}_{\tilde{n}}$ of different values, each of them being taken by many equal $\mathbf{x}_i$. To be more specific, let us assume for example that $n = k\tilde{n}$ and that

$$\mathbf{x}_1 = \tilde{\mathbf{x}}_1, \cdots, \mathbf{x}_k = \tilde{\mathbf{x}}_1, \quad \cdots\cdots \quad \mathbf{x}_{k(\tilde{n}-1)+1} = \tilde{\mathbf{x}}_{\tilde{n}}, \cdots, \mathbf{x}_n = \tilde{\mathbf{x}}_{\tilde{n}}$$

In such a case, it is natural to compute the regression function $f(\mathbf{x}) = \mathbb{E}\{Y|\mathbf{X} = \mathbf{x}\}$ at the points $\tilde{\mathbf{x}}_j$ for $j = 1, \cdots, k$. For each such $\tilde{\mathbf{x}}_j$, the observations $y_{k(j-1)+1}, \cdots, y_{kj}$ of the response variable form a sample of size k of the conditional distribution of Y given $\mathbf{X} = \tilde{\mathbf{x}}_j$. So the regression, which is by definition the expectation of this conditional distribution, can be estimated by the sample average of these observations, i.e.

$$\hat{f}(\tilde{\mathbf{x}}_j) = \frac{1}{k} \sum_{i=1}^{k} y_{k(j-1)+i},$$

and the larger the value of k, the better the estimate.

Unfortunately, such a situation is rarely encountered, in other words, in most cases, $k = 1$. The philosophy of most of the non-parametric regression procedures is to overcome this fact.

7.4 BASIS EXPANSION REGRESSION

Let us assume that the regression function

$$f(\mathbf{x}) = \mathbb{E}\{Y|\mathbf{X} = \mathbf{x}\}$$

belongs to a function space for which we know a basis, say $\{f_\ell\}_\ell$. For the sake of simplicity, we assume that the function space is a Hilbert space and that the basis is orthonormal. In this way

$$f(\mathbf{x}) = \sum_{\ell \geq 1} < f_\ell, f > f_\ell \tag{7.4}$$

Non-parametric regression typically involves infinite dimensional function spaces, and hence, infinite sums. Indeed, if the expansion involves only a small number of terms, estimating f amounts to estimating these parameters and we are back in the realm of parametric statistics. Notice that the above expansion in (**??**) contains finitely many terms when the regressor X can only take finitely

many values. A natural first step in estimating the function f is to approximate expansion (**??**) by a finite sum. Indeed, the above decomposition means that

$$f(\mathbf{x}) = \lim_{L\to\infty} \sum_{\ell=1}^{L} < f_\ell, f > f_\ell.$$

We use the notation $f^{(L)} = \Pi_L f$ for the finite sum appearing in the above right hand side. If we use this finite sum $f^{(L)}$ as an approximation for f, the squared error is given

$$\|f - f^{(L)}\|^2 = \sum_{\ell \geq L+1} < f_\ell, f >^2$$

and classical numerical analysis results can be used to estimate the rate of convergence to 0 of this error. This rate will of course depend upon the specifics of the Hilbert space and the smoothness properties of the (unknown) regression function f, but precise convergence rates can be obtained for very general classes of functions f.

We assume from now on that the integer L has been chosen. Notice that the difference between non-parametric and parametric statistics is more at the level of L large versus L small, rather than at the level of the duality L finite versus L infinite.

Notice also that because of the Hilbert space structure, the coefficients $\{< f_\ell, f >\}_{1\leq\ell\leq L}$ of the basis decomposition are the arguments $\{\alpha_\ell\}_{1\leq\ell\leq L}$ of the minimization problem

$$\inf_{\alpha_1,\cdots,\alpha_L} \|f - \sum_{\ell=1}^{\ell} \alpha_\ell f_\ell\|^2 \tag{7.5}$$

So estimating the function f by the finite sum $f^{(L)}$ can be done by looking for real numbers $\alpha_1, \cdots, \alpha_L$ solving (**??**), and using the approximation

$$F^{(L)} = \sum_{\ell=1}^{L} \alpha_\ell f_\ell. \tag{7.6}$$

7.4.1 **Sample estimation**

We now assume that we have observations

$$(\mathbf{x}_1, y_1), \cdots\cdots, (\mathbf{x}_M, y_M)$$

where for each $m = 1, 2, \cdots, M$, $\mathbf{x}_m$ is a vector of p numerical components and y_m is a real number. We use these data to estimate the approxiamtion $f^{(L)}$ of regression function f. Each y_m can be viewed as a noisy observation of $f(\mathbf{x}_m)$ since by definition

$$f(\mathbf{x}_m) = \mathbb{E}\{Y|\mathbf{X} = \mathbf{x}_m\}.$$

Informally:

$$\begin{aligned} y_m &= f(\mathbf{x}_m) + \text{noise} \\ &\approx f^{(L)}(\mathbf{x}_m) + \text{noise} \\ &= \sum_{\ell=1}^{\ell} \alpha_\ell f_\ell(\mathbf{x}_m) + \text{noise} \end{aligned}$$

hence, the least squares paradigm suggests to look for the α_m's minimizing the least squares objective

$$\inf_{\alpha_1,\cdots,\alpha_L} \sum_{m=1}^{M} |y_m - \sum_{\ell=1}^{\ell} \alpha_\ell f_\ell(\mathbf{x}_m)|^2. \tag{7.7}$$

This is now a problem which can be solve by *plain* least squares methods, possibly of high dimension, but definitely with standard tools!

7.5 KERNEL REGRESSION

We first discuss the case $p = 1$ of univariate regression.

7.5.1 The Kernel Smoother

The idea of the kernel smoother is to rely on the observed responses to neighboring values of x to predict the response $f(x)$. The only difference is that, instead of relying on a limited number of observations y_i's of the response, the local character of the averaging is realized by a weighted average of all the observed values y_i's, the weights being decreasing with the distance between x and the corresponding value x_i of the explanatory variable. To be more specific, the weights are computed by means of a *kernel* function $x \hookrightarrow K(x)$, and our good old enemy, the smoothing parameter. The latter is called *bandwidth* in the case of the kernel method, and it will be denoted by $b > 0$. By now, we should be familiar with the terminology and the notation associated with the kernel method. Indeed, we already introduced them in our discussion of the kernel density estimation method. We give a lucid account of the relationship between the applications of the kernel method to density estimation and regression in the Appendix at the end of this chapter. The actual formula giving the kernel scatterplot smoother $f(x)$ is:

$$f(x) = f_{b,K}(x) = \frac{\sum_{i=1}^{n} y_i K\left(\frac{x-x_i}{b}\right)}{\sum_{j=1}^{n} K\left(\frac{x-x_j}{b}\right)}. \tag{7.8}$$

Notice that the formula giving $f(x)$ can be rewritten in the form:

$$f(x) = \sum_{i=1}^{n} w_i(x) y_i \tag{7.9}$$

provided we define the weights $w_i(x)$ by the formula:

$$w_i(x) = \frac{K\left(\frac{x-x_i}{b}\right)}{\sum_{j=1}^{n} K\left(\frac{x-x_j}{b}\right)}. \tag{7.10}$$

Understanding the properties of these weights is crucial to understanding the very nature of kernel regression. These properties will be clear once we define what we mean by a kernel function. A nonnegative function $x \hookrightarrow K(x)$ is called a kernel function if it is integrable and if its integral is equal to 1, i.e. if it satisfies:

$$\int_{-\infty}^{+\infty} K(x)\, dx = 1.$$

Kernel function	Formula
`box`	$K_{box}(x) = \begin{cases} 1, & \text{if } \quad |x| \leq .5 \\ 0 & \text{otherwise} \end{cases}$
`triangle`	$K_{triangle}(x) = \begin{cases} 1 - \lvert x\rvert, & \text{if } \quad |x| \leq 1 \\ 0 & \text{otherwise} \end{cases}$
`parzen`	$K_{parzen}(x) = \begin{cases} \frac{9}{8} - \frac{3}{2}\lvert x\rvert + \frac{1}{2}x^2, & \text{if } \quad \frac{1}{2} \leq \lvert x\rvert \leq \frac{3}{2} \\ \frac{3}{4} - x^2, & \text{if } \quad \lvert x\rvert \leq \frac{1}{2} \\ 0 & \text{otherwise} \end{cases}$
`normal`	$K_{normal}(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$

Table 7.1. Table of the four kernel functions used by the `Splus` function `ksmooth`.

In other words, K is a probability density. The fact that the integral of $K(x)$ is equal to one is merely a normalization condition useful in the applications to density estimation. It will not be of any consequence in the case of the applications to regression since K always appear simultaneously in the numerator and the denominator: indeed, as one can easily see from formulae (**??**) and (**??**), multiplying K by a constant does not change the value of the regression function f as defined in (**??**). But in order to be useful, the kernel $K(x)$ has to take relatively large values for small values of x, and relatively small values for large values of x. In fact, it is also often assumed that K is symmetric in the sense that:

$$K(-x) = K(x)$$

and that $K(x)$ decreases as x goes from 0 to $+\infty$. The above symmetry condition implies that:

$$\int_{-\infty}^{+\infty} xK(x)\,dx = 0 \tag{7.11}$$

which will be used in our discussion of the connection with kernel density estimation in appendix. They are also some of the most commonly used kernel functions when it comes to regression. Notice that the first three of them vanish outside a finite interval, while the fourth one (theoretically) never vanishes. Nevertheless, since its computation involves evaluating exponentials, it will not come as a surprise that such a kernel can be (numerically) zero because of the evaluation of exponentials of large negative numbers: indeed for all practical purposes, there is no significant difference between e^{-60} and 0, and exponents that negative can appear very often!

Except for the choice of the kernel function `box` which leads to the crudest results, the other three kernels give essentially the same results in most applications. The situation is different when it comes to the choice of the bandwidth parameter b. Indeed, the choice of the bandwidth is the *Achilles heel* of kernel regression. This choice can have an enormous impact, and the results can vary dramatically: small values of b give rough graphs which fit the data too closely, while too large a value of b produces a flatter graph as b increases. By experimenting with the choice of the bandwidth, one can easily see that as b tends to ∞, the graph of the kernel smoother converges toward the horizontal straight line with intercept the mean $\overline{y}$ of the observed responses y_j's. As we explained earlier, this means that regression is meaningless since the explanatory variable does not have any influence on the value of the prediction of the response variable.

7.5.2 The Multivariate Case

Multivariate kernel regression is a typical example of multivariate nonparametric nonlinear regression, but it can also be viewed as a high dimensional generalization of the procedure described in the subsection on the kernel scatterplot smoother, and especially the discussion of the function `ksmooth`. Indeed, most of what we said then, can be generalized to the case where the dimension p of the explanatory variables is not necessarily equal to 1. Indeed, formula (**??**) can be used in the form:

$$f(\mathbf{x}) = f_{b,K}(\mathbf{x}) = \frac{\sum_{j=1}^{n} y_j K\left(\frac{\mathbf{x}-\mathbf{x}_j}{b}\right)}{\sum_{j=1}^{n} K\left(\frac{\mathbf{x}-\mathbf{x}_j}{b}\right)} \tag{7.12}$$

provided the function $\mathbf{x} \hookrightarrow K(\mathbf{x})$ is a kernel function in p dimensions, in the sense that it is a nonnegative function of p variables which integrates to one.

The simplest example of p-dimensional kernel function is given by a function of the form:

$$K(\mathbf{x}) = k(\text{dist}(\mathbf{x}, 0)) \tag{7.13}$$

for some nonnegative and non-increasing function $k(\,\cdot\,)$ of one variable and some choice of a notion of distance from the origin in p dimensions. Possible choices for this notion of distance include the usual Euclidean norms in $\mathbb{R}^p$:

$$\text{dist}(\mathbf{x}, 0)) = \left(\sum_{j=1}^{p} x_j^2\right)^{1/2} \qquad \text{or} \qquad \text{dist}(\mathbf{x}, 0)) = \left(\sum_{j=1}^{p} w_j x_j^2\right)^{1/2}$$

or non Euclidean norms such as:

$$\text{dist}(\mathbf{x}, 0)) = \sum_{j=1}^{p} |x_j| \qquad \text{or} \qquad \text{dist}(\mathbf{x}, 0)) = \sup_{j=1\cdots p} |x_j|.$$

These choices are popular because of their convenient scaling properties. If one excepts the Euclidean distance computed with different weights w_j for the different components x_j of the explanatory vector $\mathbf{x}$, all these kernel functions share the same shortcoming: all the components of the explanatory vector are treated equally, and this may be very inappropriate if the numerical values are on different scales. Indeed, in such a case, the value of the distance is influenced mostly (if not exclusively) by the variables having the largest values. We illustrate this point with a short discussion of an example which we will study in detail in the later part of the chapter. Let us imagine for example that the first explanatory variable is an annualized interest rate. Its values are typically of the order of a few percentage points. Let us also imagine that the second explanatory variable is a time to maturity. If for some strange reason this second variable is expressed in days instead of years, its values will be in the hundreds on a regular basis, and a distance of the type given above will ignore the small changes in interest rate, and report only on the differences in maturity. A change in unit in one of the variables can dramatically change the qualitative properties measured by these notions of distance, and consequently affect strongly the results of the kernel regression. This effect is highly undesirable. We discuss below alternative choices of kernel functions which can overcome this difficulty, as well as a standardization procedure which re-scales all the explanatory variables in an attempt to balance their relative contributions to the regression results.

Another very popular class of kernel functions is given by direct products (sometimes called tensor products) of one dimensional kernel functions. Indeed, if K_1, K_2, $\cdots$, K_p are one-dimensional kernel functions (possibly equal to each other), then the function:

$$K(\mathbf{x}) = K(x_1, x_2, \cdots, x_p) = K_1(x_1)K_2(x_2)\cdots K_p(x_p) \tag{7.14}$$

is obviously a p-dimensional kernel function. For these kernel functions, the weight multiplying the i-th response y_i is proportional to:

$$K\left(\frac{\mathbf{x}-\mathbf{x}_i}{b}\right) = K_1\left(\frac{x_1-x_{i,1}}{b}\right)K_2\left(\frac{x_2-x_{i,2}}{b}\right)\cdots K_p\left(\frac{x_p-x_{i,p}}{b}\right)$$

and from this expression one sees that there is no harm in choosing different values for the p occurrences of the bandwidth b in the right hand side. In other words, it is possible to choose p different bandwidths b_1, b_2, $\cdots$, b_p, one for each component of the explanatory variable. This feature of the direct product kernels makes them very attractive. In some sense, normalizing the scalar explanatory variables and using one single bandwidth amounts to the same as using different bandwidths for the components of the explanatory vector. See Subsection **??** for an example of standardization before running a kernel regression.

We now recast some of the most important properties of kernel regression as elementary remarks which apply as well to the one dimensional case of the kernel scatterplot smoother `ksmooth` discussed earlier.

• The kernel regression estimate $f(\mathbf{x})$ is a linear function of the observations. Indeed, the definition formula (**??**) can be rewritten in the form:

$$f(\mathbf{x}) = \sum_{i=1}^{n} w_i(\mathbf{x})y_i$$

where the weights $w_i(\mathbf{x})$ are defined by:

$$w_i(\mathbf{x}) = \frac{K\left(\frac{\mathbf{x}-\mathbf{x}_i}{b}\right)}{\sum_{j=1}^{n} K\left(\frac{\mathbf{x}-\mathbf{x}_j}{b}\right)}.$$

Notice that these weights are nonnegative and they do sum up to one. Because the kernel function is typically very small when its argument is large and relatively large when its argument is small, the weight $w_i(\mathbf{x})$ is (relatively) large when $\mathbf{x}$ is close (i.e. similar) to the observation $\mathbf{x}_i$ and small otherwise. This shows that the kernel regression function f given by (**??**) is a weighted average of the observed values y_i's of the response (and hence it is linear in the y_i's) with weights which give more importance to the responses from values of $\mathbf{x}_i$ close to the value $\mathbf{x}$ of the explanatory variables under consideration.

• The choice of the bandwidth is a very touchy business. Many proposals have been made for an automatic (i.e. data driven) choice of this smoothing parameter. Whether one uses the results of difficult asymptotic analyses to implement bootstrap or cross validation procedures or simple rules of thumb, our advise is to be wise and to rely on experience to detect distortions due to a poor choice of the bandwidth.

• As we explained earlier, it is tempting to use a separate bandwidth for each explanatory variable. This is especially the case when the kernel function is of the product type as given in (**??**) and when

the dynamic ranges of these variables are very different. For example, if a variable is expressed in a physical unit, changing the unit system may change dramatically the range of the actual values of the measurements, and small numbers can suddenly become very large and as a consequence of the change of units. Accordingly, the influence of this variable on the computation of the kernel regression can increase dramatically. This undesirable effect is often overcome by normalizing the variables. See details in the discussions of the practical examples presented in Subsection **??** below.

• The sample observations of the explanatory vector form a cloud of points in the p-dimensional Euclidean space $\mathbb{R}^p$. The larger the dimension p, the further apart these points appear. Filling up space with points is more difficult in higher dimension, and in any given neighborhood of a point $\mathbf{x} \in \mathbb{R}^p$, we are less likely to find points from the cloud of sample observations when p is large. This fact is known as Bellman's *curse of dimensionality.* When the number n of observations is not excessively large, the kernel regression has proven to be very powerful when the number of explanatory variables (i.e. the number p) is reasonably small, typically 2 or 3. How small should this number be obviously depends upon the sample size n, and the more observations we have, the larger the number of explanatory variables we may include. This form of Bellman's *curse of dimensionality* can easily be illustrated by heuristic arguments but it can also be quantified by rigorous asymptotic results which show that n should grow exponentially with p. This is a serious hindrance.

7.6 PROJECTION PURSUIT REGRESSION

Projection pursuit searches for an approximation of the true regression function $f(\mathbf{x})$ in the form:

$$f^{(L)}(\mathbf{x}) = \alpha + \sum_{\ell=1}^{m} \phi_\ell(\mathbf{a}_\ell \cdot \mathbf{x}). \tag{7.15}$$

where α is a real number, the $\mathbf{a}_\ell$ are unit vectors in $\mathbb{R}^p$ and the ϕ_ℓ's are functions of a real variable. Remember that we are working in the usual regression setting:

$$(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \cdots\cdots, (\mathbf{x}_M, y_M)$$

where the *explanatory variables* $\mathbf{x}_1, \mathbf{x}_2, \cdots, \mathbf{x}_M$ are p-dimensional. This algorithm was designed in order to cope with the *curse of dimensionality* inherent with large values of p, by replacing the p-dimensional explanatory vectors $\mathbf{x}_i$ by suitably chosen one-dimensional projections $\mathbf{a}_j \cdot \mathbf{x}_i$, hence the term *projection* in the name of the method. We now explain how the quantities α, $\phi_1(\mathbf{a}_1 \cdot \mathbf{x}), \cdots$, and $\phi_L(\mathbf{a}_L \cdot \mathbf{x})$ appearing in formula (**??**) are estimated in practice. The projection pursuit algorithm is based on an inductive procedure in which residuals are recomputed and fitted at each iteration. To start, we assume that the observed values y_i's as the starting residuals, i.e. the residuals of order zero. Next, each time one of the terms in the sum appearing in the right hand side of (**??**) is estimated, the actual estimates are subtracted from the current values of the residuals, providing in this way a new set of residuals from which we proceed to estimate the next term in (**??**). This recursive fitting of the residuals justifies the term *pursuit* in the name of the method.

The constant α is naturally estimated by the mean of the observations of the response:

$$\hat{\alpha} = \overline{y} = \frac{1}{n} \sum_{i=1}^{n} y_i.$$

This sample mean is subtracted from the observations (i.e. the residuals of order zero) to get the residuals of order one. Next we proceed to the estimation of the first direction $\mathbf{a}_1$ and the first function ϕ_1. Because of computational considerations, it is important to choose a specific procedure capable of selecting the best function ϕ for each choice of the direction given by the unit vector $\mathbf{a}$.

The creators of the projection pursuit algorithm proposed to fit recursively the residuals (starting from the values of the response variable) with terms of the form $\phi_{\mathbf{a}_j}(\mathbf{a}_j \cdot \mathbf{x})$ and to associate to each such optimal term a figure of merit, for example the proportion of the total variance of the response actually explained by such a term. In this way, the whole procedure would depend only upon one parameter. One could choose this *tolerance* parameter first (typically a small number) and one would fit the response and the successive residuals until the figure of merit (i.e. the proportion of the variance explained by $\phi_{\mathbf{a}_j}(\mathbf{a}_j \cdot \mathbf{x})$) would drop below the tolerance parameter and this would automatically take care of the choice of the order m of the model (**??**).

7.7 MONTE CARLO MALLIAVIN REGRESSION

This final section is devoted to a recently developed method to compute conditional expectations when random variables/vectors are functions of a (possibly multivariate) Wiener process. Se we propose to compute

$$\mathbb{E}\{Y|\mathbf{X}=\mathbf{x}\} \;=\; \frac{\mathbb{E}\{Y\delta_{\mathbf{x}}(\mathbf{X})\}}{\mathbb{E}\{\delta_{\mathbf{x}}(\mathbf{X})\}} \tag{7.16}$$

when both Y and $\mathbf{X}$ are functions of an underlying Wiener process $\{W(t)\}_t$. The above formula is not always rigorous as it involves Dirac delta functions, but its intuitive power will be extremely useful. The idea of the method is to use an integration by parts to replace the delta function by its anti-derivative which is a bona fide function. The idea is very simple. However, its implementation is very intricate as the integration by parts in question needs to be done in function space (the path space of the Wiener process). This is how and why the so-called Malliavin calculus (calculus on Wiener space) is brought to bear.

Malliavin Calculus Based Simulation Method

Using Malliavin integration by parts formula to get rid of the Dirac point masses in (**??**), one gets a ratio of the form:

$$\mathbb{E}\{Y|\mathbf{X}=\mathbf{x}\} \;=\; \frac{\mathbb{E}\{Y H_{\mathbf{x}}(\mathbf{X})S\}}{\mathbb{E}\{H_{\mathbf{x}}(\mathbf{X})S\}} \tag{7.17}$$

where $H_b(\mathbf{x}') = \prod_{i=1}^{p} \mathbf{1}_{[x_i,\infty)}(x_i')$, and S is some non-negative random variable. An important consequence of this formula is the fact that the associated Monte Carlo estimator:

$$\hat{\mathbb{E}}\{Y|\mathbf{X}=\mathbf{x}\} = \frac{\frac{1}{N}\sum_{n=1}^{N} Y^{(n)} H_{\mathbf{x}}(\mathbf{X}^{(n)})S^{(n)}}{\frac{1}{N}\sum_{n=1}^{N} H_{\mathbf{x}}(\mathbf{X}^{(n)})S^{(n)}} \;,$$

constructed from an independent sample $\{(Y^{(n)}, \mathbf{X}^{(n)}, S^{(n)})\}_{n=1,\cdots,N}$ of size N, converges at the $\sqrt{N}-$rate by the classical central limit theorem. The following subsection is devoted to a self-contained derivation of these facts. We use a pedestrian approach based on the log-normality of our Gaussian framework, without ever appealing to results of the Malliavin calculus.

7.7.1 Integration-by-Parts based Regression Estimation

We first concentrate on a regression function of the form:

$$r_h(x) \;:=\; \mathbb{E}\{g(W_{t+h})|W_t = x\} \;=\; \frac{\mathbb{E}\{g(W_{t+h})\delta_x(W_t)\}}{\mathbb{E}\{\delta_x(W_t)\}} \,.$$

Integration by Parts

Let us denote by φ the density of the standard one dimensional normal distribution, and let us assume that g is a smooth function with a bounded derivative. By the independence of the increments of the Brownian motion, we have:

$$\mathbb{E}\{g(W_{t+h})\delta_x(W_t)\} \;=\; \iint g(w_1+w_2)\delta_x(w_1)\varphi\left(\frac{w_1}{\sqrt{t}}\right)\varphi\left(\frac{w_2}{\sqrt{h}}\right)dw_1dw_2$$

where we use the notation φ for the density of the standard normal distribution, i.e.

$$\varphi(x) = \frac{1}{\sqrt{2\pi}}e^{-x^2/2}, \qquad x \in \mathbb{R}.$$

Integrating by parts with respect to the w_1 variable, we get :

$$\begin{aligned}\mathbb{E}\{g(W_{t+h})\delta_x(W_t)\} = &\iint g(w_1+w_2)\mathbf{1}_{[x,\infty)}(w_1)\frac{w_1}{t}\varphi\left(\frac{w_1}{\sqrt{t}}\right)\varphi\left(\frac{w_2}{\sqrt{h}}\right)dw_1dw_2\\ &-\iint g'(w_1+w_2)\mathbf{1}_{[x,\infty)}(w_1)\varphi\left(\frac{w_1}{\sqrt{t}}\right)\varphi\left(\frac{w_2}{\sqrt{h}}\right)dw_1dw_2\,.\end{aligned}$$

Next, we compute the second integral by integrating by parts with respect to the w_2 variable. We get:

$$\begin{aligned}\mathbb{E}\{g(W_{t+h})\delta_x(W_t)\} &= \iint g(w_1+w_2)\mathbf{1}_{[x,\infty)}(w_1)\frac{w_1}{t}\varphi\left(\frac{w_1}{\sqrt{t}}\right)\varphi\left(\frac{w_2}{\sqrt{h}}\right)dw_1dw_2\\ &\quad-\iint g(w_1+w_2)\mathbf{1}_{[x,\infty)}(w_1)\frac{w_2}{h}\varphi\left(\frac{w_1}{\sqrt{t}}\right)\varphi\left(\frac{w_2}{\sqrt{h}}\right)dw_1dw_2\\ &= \mathbb{E}\{g(W_{t+h})\mathbf{1}_{[x,\infty)}(W_t)S_h\}\,,\end{aligned}$$

where the random variable

$$S_h = \frac{W_t}{t} - \frac{W_{t+h}-W_t}{h} \tag{7.18}$$

is independent of the function g. Notice that formula (**??**) is established for a function $g \in C_b^1$. However, since it does not involve the regularity of g, we can conclude by a classical density argument that it is valid whenever $g(W_{t+h}) \in L^2$.

Actual Simulation

Let $(W^{(s)})_{1\le s\le S}$ be n independent samples of the Wiener process W. Then, the Monte Carlo estimator suggested by the above formula is defined by

$$\hat{r}_N(x) \;:=\; \frac{\hat{q}_N[g](x)}{\hat{q}_N[1](x)} \text{ where } \hat{q}_N[g](x) \;:=\; \frac{1}{N}\sum_{s=1}^{N} g(W_{t+h}^{(s)})\mathbf{1}_{[x,\infty)}(W_t^{(s)})S_h^{(s)}\,,$$

where $S_h^{(s)}$ is computed from the sample $W^{(s)}$ using formula (**??**). Its asymptotic properties are directly deduced from the law of large numbers and the central limit theorem for independent identically distributed random variables. In particular, the rate of convergence is of the order $\sqrt{N}$.

The price to pay in order to recover the $\sqrt{N}$ rate of convergence is that the variance of the estimator $\hat{q}_N[g](x)$ explodes as h shrinks to zero since

$$\lim_{h \searrow 0} S_h = 0$$

in L^2. Since our objective is to send the time step h to zero, it is necessary to find a remedy to this variance explosion problem.

Localization

In order to do so, we introduce a localization function. Let χ be an arbitrary smooth function with $\chi(0) = 1$. Following the computations leading to formula (**??**) we get:

$$\begin{aligned}
\mathbb{E}\{g(W_{t+h})\delta_x(W_t)\} &= \mathbb{E}\{g(W_{t+h})\delta_x(W_t)\chi(W_t - x)\} \\
&= \iint g(w_1 + w_2)\delta_x(w_1)\chi(w_1 - x)\varphi\left(\frac{w_1}{\sqrt{t}}\right)\varphi\left(\frac{w_2}{\sqrt{h}}\right)dw_1 dw_2 \\
&= \iint g(w_1 + w_2)\mathbf{1}_{[x,\infty)}(w_1)\chi(w_1 - x)\frac{w_1}{t}\varphi\left(\frac{w_1}{\sqrt{t}}\right)\varphi\left(\frac{w_2}{\sqrt{h}}\right)dw_1 dw_2 \\
&\quad - \iint g'(w_1 + w_2)\mathbf{1}_{[x,\infty)}(w_1)\chi(w_1 - x)\varphi\left(\frac{w_1}{\sqrt{t}}\right)\varphi\left(\frac{w_2}{\sqrt{h}}\right)dw_1 dw_2 \\
&\quad - \iint g(w_1 + w_2)\mathbf{1}_{[x,\infty)}(w_1)\chi'(w_1 - x)\varphi\left(\frac{w_1}{\sqrt{t}}\right)\varphi\left(\frac{w_2}{\sqrt{h}}\right)dw_1 dw_2 \\
&= \iint g(w_1 + w_2)\mathbf{1}_{[x,\infty)}(w_1)\chi(w_1 - x)\frac{w_1}{t}\varphi\left(\frac{w_1}{\sqrt{t}}\right)\varphi\left(\frac{w_2}{\sqrt{h}}\right)dw_1 dw_2 \\
&\quad - \iint g(w_1 + w_2)\mathbf{1}_{[x,\infty)}(w_1)\chi(w_1 - x)\frac{W_2}{h}\varphi\left(\frac{w_1}{\sqrt{t}}\right)\varphi\left(\frac{w_2}{\sqrt{h}}\right)dw_1 dw_2 \\
&\quad - \iint g(w_1 + w_2)\mathbf{1}_{[x,\infty)}(w_1)\chi'(w_1 - x)\varphi\left(\frac{w_1}{\sqrt{t}}\right)\varphi\left(\frac{w_2}{\sqrt{h}}\right)dw_1 dw_2 \\
&= \mathbb{E}\{g\left(W_{t+h}\right)\mathbf{1}_{[x,\infty)}(W_t)S_{h,\chi}\}\,,
\end{aligned}$$

where the random variable $S_{h,\chi}$ is defined by:

$$\begin{aligned}
S_{h,\chi} &= \chi(W_t - x)\left(\frac{W_t}{t} - \frac{W_{t+h} - W_t}{h}\right) - \chi'(W_t - x) \\
&= \chi(W_t - x)S_h - \chi'(W_t - x)
\end{aligned} \tag{7.19}$$

is again independent of the function g. For each localization function χ, one can now define a new Monte Carlo estimator as before. All these estimators share the nice convergence property at the $\sqrt{N}$-rate. Therefore, the natural question is whether one can reduce the variance of the Monte Carlo estimator by some convenient choice of localization function χ.

Variance Reduction by Localization. Set $G := g(W_{t+h})^2$, and let us consider the integrated mean square error

$$J(\chi) \ := \ \int_{\mathbb{R}} \mathbb{E}\left\{G\mathbf{1}_{W_t>x} S_{h,\chi}^2\right\} dx \ .$$

We are interested in the integrated mean square error minimization :

$$V := \min\left\{J(\chi) \ : \ \chi \text{ smooth, bounded and } \chi(0)=1\right\} \ .$$

Using Fubini's theorem and a path by path substitution, we get:

$$\begin{aligned} J(\chi) &= \mathbb{E}\left\{G\int_{-\infty}^{W_t} \left|\chi(W_t-x)S_h - \chi'(W_t-x)\right|^2 dx\right\} \\ &= \mathbb{E}\left\{G\int_0^{+\infty} \left|\chi(y)S_h - \chi'(y)\right|^2 dy\right\} \ . \end{aligned}$$

Observing that $\mathbb{E}\{GS_h\}=0$, this provides

$$J(\chi) \ = \ \int_0^{+\infty} \left[\mathbb{E}\{GS_h^2\}|\chi(y)|^2 + \mathbb{E}\{G\}|\chi'(y)|^2\right] dy \ .$$

Hence the integrated mean square error minimization is reduced to a classical problem of calculus of variations, which can be solved explicitly. The optimal localization function is then given by

$$\chi_h(x) \ := \ e^{-\eta_h x} \text{ where } \eta_h \ := \ \left(\frac{\mathbb{E}\{GS_h^2\}}{\mathbb{E}\{G\}}\right)^{1/2} \ . \tag{7.20}$$

In particular, this shows that

$$\eta_h \ = \ \mathrm{O}\left(h^{-1/2}\right) \ .$$

7.7.2 Monte Carlo Estimation for the Finite Maturity Problem

We now return to the problem of the optimal multiple stopping problem, and more precisely to the pricing of swing options in the framework of the discrete time approximation set up in Subsection **??**. Let N_n be some integer depending on the time step parameter n, and let $\{W^{(s)}, 1 \le s \le N_n\}$ be N_n independent samples of the Wiener process. For each integer s, we denote by $X^{(s)}$ the process X associated to the Brownian motion $W^{(s)}$ via formula (**??**). Also, we set:

$$\begin{aligned} R_h^{(s)}(t_j,x) &:= S_{h,\chi_h}^{(s)}(t_j)\mathbf{1}_{[x,\infty)}(W_{t_j}^{(s)}) \\ &= \chi_h(W_t^{(s)}-x)\left[\eta_h + h^{-1}\left(2W_{t_j}^{(s)} - W_{t_j-h}^{(s)} - W_{t_j+h}^{(s)}\right)\right]\mathbf{1}_{[x,\infty)}(W_{t_j}^{(s)}) \ , \end{aligned}$$

where S_h and χ_h are defined respectively in (**??**) and (**??**). Following the discussion of the previous paragraph, we define the estimators :

$$\begin{aligned} \tilde{\rho}_n^{(i)}\left(t_j, X_{t_j}\right) &:= \frac{\frac{1}{N_n}\sum_{s=1}^S \hat{v}_n^{(i)}\left(t_{j+1}, X_{t_{j+1}}^{(s)}\right) R_{1/n}^{(s)}\left(t_j, X_{t_j}\right)}{\frac{1}{N_n}\sum_{s=1}^S R_{1/n}^{(s)}\left(t_j, W_{t_j}\right)} \ , \\ \tilde{\phi}_n^{(i)}\left(t_j, X_{t_j}\right) &:= \phi\left(X_{t_j}\right) \\ &\quad + e^{-r\delta}\ \frac{\frac{1}{N_n}\sum_{s=1}^S \hat{v}_n^{(i-1)}\left(t_j+\delta, X_{t_j+\delta}^{(s)}\right) R_\delta^{(s)}\left(t_j, X_{t_j}\right)}{\frac{1}{N_n}\sum_{s=1}^S R_\delta^{(s)}\left(t_j, W_{t_j}\right)} \mathbf{1}_{t_j \le 1-\delta} \end{aligned}$$

of $\rho_n^{(i)}\left(t_j, X_{t_j}\right)$ and $\phi_n^{(i)}\left(t_j, X_{t_j}\right)$ respectively. These estimators are defined inductively, given the estimators $\hat{v}^{(i-1)}(.,.)$ and the previous estimator $\hat{v}_n^{(i)}\left(t_{j+1}, X_{t_{j+1}}^{(s)}\right)$ in the backward procedure. Finally, we observe that $\rho_n^{(i)} \leq iK$ and $\phi_n^{(i)} \leq iK$. Hence, in order to ovoid an explosion of the algorithm, we define the truncated estimators (see Bouchard and Touzi (2002)):

$$\hat{\rho}_n^{(i)}\left(t_j, X_{t_j}\right) := (iK) \wedge \tilde{\rho}_n^{(i)}\left(t_j, X_{t_j}\right)^+ ,$$
$$\hat{\phi}_n^{(i)}\left(t_j, X_{t_j}\right) := (iK) \wedge \tilde{\phi}_n^{(i)}\left(t_j, X_{t_j}\right)^+ ,$$

and

$$\hat{v}_n^{(i)}\left(t_j, X_{t_j}\right) := \max\left\{\hat{\rho}_n^{(i)}\left(t_j, X_{t_j}\right), \hat{\phi}_n^{(i)}\left(t_j, X_{t_j}\right)\right\} .$$

According to the error estimate of Bouchard and Touzi (2002), in order for the approximation error to be of the order of $n^{-1/2}$, one has to choose a number N_n of simulated trajectories such that

$$N_n \;=\; \mathrm{O}\left(n^{7/2}\right) .$$

The Value Functions

The above algorithm was implemented and tested in the case of an American put option with four exercise rights and the following characteristics: maturity $T = 1$ year, $r = .05$, $\sigma = .30$, maximal number of exercise rights $\ell = 5$, $n = 50$. We also used a refraction period $\delta = 0.1$ to separate successive exercises. This feature is not so much of an extra difficulty when dealing with discrete time models, however, it can be a serious nuisance with continuous time models. We added the refraction time to our computations as this feature is typical in the energy markets.

The left pane of Figure **??** gives the plots of the graphs of the functions $x \hookrightarrow v^{(1)}(t,x)$ for $t = .59, .58, \cdots, .02, .01$. Two remarks are in order. First, these graphs are not computed over the same range of values of x. Essentially, we computed the values of $v^{(k)}(t,x)$ for the values of x which can be reached by the sample paths of the diffusion process X_t, and we determined this range of values of x from the results of our simulations. The second remark concerns the noise in the numerical results. Obviously, we should expect zero in the right hand side of the plots, and we see quite significant departures from this expectation. The right pane of the figure gives the plots of the graphs of the functions $x \hookrightarrow v^{(3)}(t,x)$ for $t = .59, .58, \cdots, .02, .01$. for $t = .49, .47, \cdots, .02, .01$.

Figure **??** gives the same plot as the left pane of Figure **??**, but instead of super-imposing the one-dimensional graphs on the same plot, we use both the t and the x variables to produce surface plots, or to be more specific the scaled time to maturity $\tau = 100(T - t)$ and x. The fact that the range of x varies with t is obvious from this surface plot, and as expected, it is limited by some form of parabola. Plotting the graphs of the other value functions $v^{(k)}$ would produce very similar results and we refrain from producing them.

Number of Monte Carlo Scenarios

We present some partial numerical results to illustrate the effect of the number of trajectories N_n. According to the result of Bouchard and Touzi (2002) which we re-derived above, the number N_n should be of the order of $n^{7/2}$. The results collected in the following Table 1 show that a very high precision can be achieved even with a significantly smaller number of simulated trajectories.

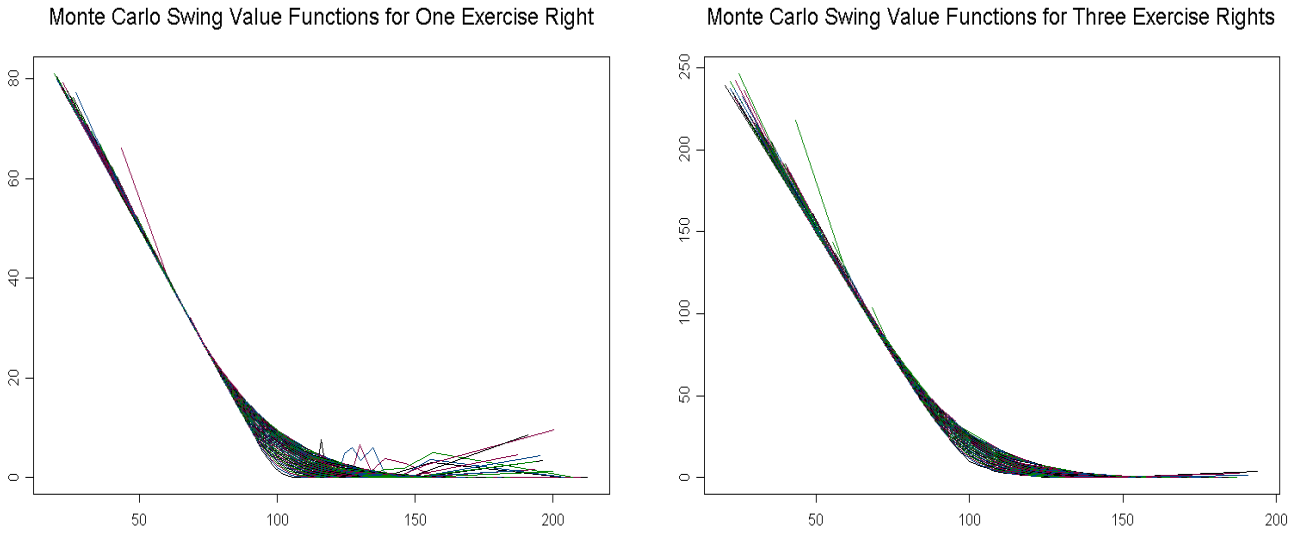

Fig. 7.1. Graphs of the functions $v^{(1)}(t,\cdot)$ for $t = .59, .58, \cdots, .02, .01$ (left) and of the functions $v^{(3)}(t,\cdot)$ for $t = .49, .47, \cdots, .02, .01$ (right).

Table 1. *Swing put option values for various numbers of simulations*
$T = 1$ year, $\delta = 0.1$ year, $S_0 = K = 100$, $r = .05$, $\sigma = .30$, $n = 50$

	$N = 8,192$	$N = 16,384$
$v^{(1)}$ [stand. dev.]	9.84 [.22%]	9.85 [.12%]
$v^{(2)}$ [stand. dev.]	19.21 [.56%]	19.26 [.30%]
$v^{(3)}$ [stand. dev.]	28.69 [.68%]	28.80 [.30%]
$v^{(4)}$ [stand. dev.]	38.34 [.57%]	38.48 [.27%]
$v^{(5)}$ [stand. dev.]	48.17 [.50%]	48.32 [.30%]

Exercise Regions

Next we identify an estimate of the exercise region for each of the value functions $v^{(i)}$ considered as a single stopping problem associated to the reward function $\phi^{(i)}(t,x)$. The corresponding exercise boundaries are given by the graphs of the functions $t \hookrightarrow \hat{x}_i^*(t)$. Estimates of these boundaries computed with the Monte Carlo method described in this section are plotted in Figure **??**. The computations were performed with the following parameters: maturity $T = 10$ months, refraction period $\delta = 2$ months, $r = .05$, $\sigma = .30$, maximal number of exercise rights $\ell = 5$, $n = 50$, $N = 8192$. As expected these exercise boundaries are increasing functions of the time-to-maturity variable. We also verify that $\hat{x}_i^*(t) \geq \hat{x}_{i-1}^*(t)$. This property is consistent with the intuition. We proved rigorously

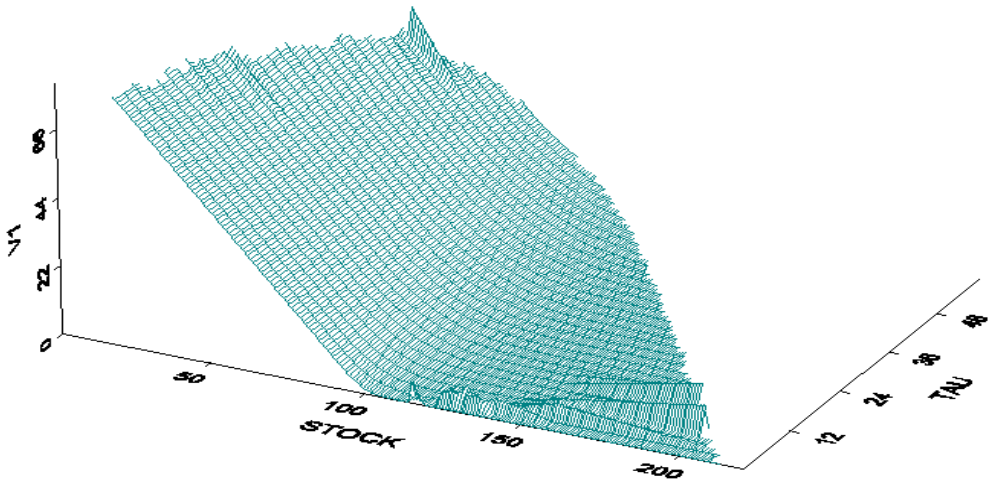

Fig. 7.2. Surface plot of the graph of the function $v^{(1)}$ when regarded as function of both t and x. The variable TAU represents $100 * (T - t)$.

this result in the case of the perpetual put options in Lemma **??**, but a proof of this fact in the finite maturity case is still lacking: this monotonicity remains an interesting open problem.

NOTES & COMPLEMENTS

The use of kernel regression in the context of American option pricing was first suggested by Carrière in [**?**]. If the bandwidth is chosen to go to zero as the sample size increases without bound, the kernel estimate can be made to converge toward the true value of the conditional expectation. However, despite some freedom in the choice of the rate of convergence of the bandwidth toward 0, the bias introduced by the approximation of the point mass at $\mathbf{x}$ by the kernel function is responsible for the fact that the classical $\sqrt{N}$ rate of convergence of the central limit theorem is lost. We refer to Bosq's monograph [**?**] for a detailed analysis of the rate of convergence of the kernel estimator. There the interested reader will find extensions to dependent samples and proofs that this rate decreases dramatically when the dimension of the random variable B increases.

The use of basis function expansions in the context of American option pricing can be traced back to the work of Tsitsiklis - van Roy [**?**] and Longstaff and Schwartz [**?**] , As in the case of the kernel method, the $\sqrt{N}-$rate of convergence is lost because of the bias introduced by the finite dimensional approximation. The choice of the orthonormal basis can drastically influence the rate of convergence. For example, it was shown by Egloff and Min-oo in [**?**] (see also the paper by Glasserman and Yu quoted in the previous chapter) that the rate of convergence of this algorithm could be exponentially slow. See for example their Theorem 6.15.

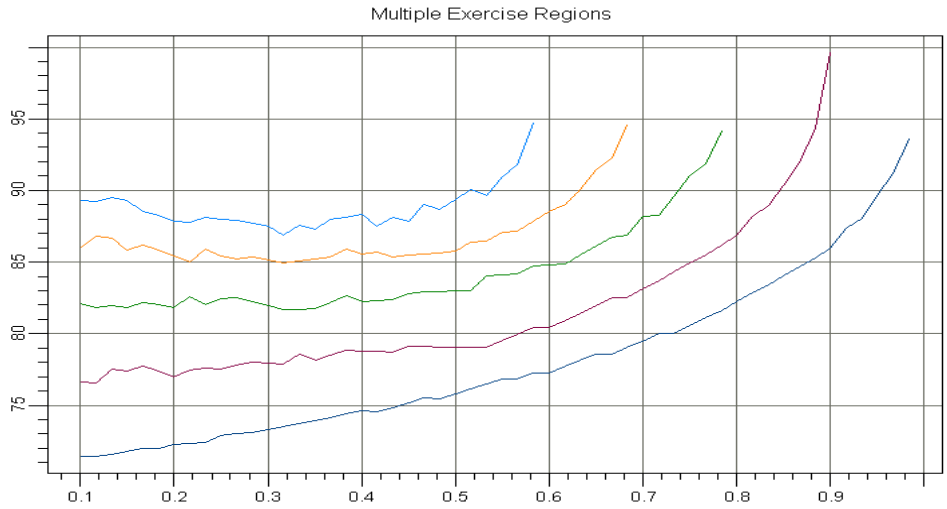

Fig. 7.3. Estimates of the boundaries of the exercise regions of swing options with $\ell = 5$ exercise rights, as given by the graphs of the functions $t \hookrightarrow x_i^*(t)$ computed via the Monte Carlo procedure described in the text.

The use of Malliavin calculus in financial derivative pricing has been proposed by Fournié, Lasry, Lebuchoux and Lions in [**?**], and further developed by Bouchard, Ekeland and Touzi in [**?**]. The asymptotic properties of the resulting numerical algorithm for the computation of the price of American put options (and more generally, for the expected value of functions of the solutions of reflected backward stochastic differential equations) have been analyzed in Bouchard and Touzi [**?**]. The presentation given in the text is borrowed from a work of Carmona and Touzi [**?**] where a self contained approached was given in the Black-scholes framework of geometric Brownian motion without any requirement other than basis probability calculus.

Bibliography

References

1. J. Cvitanic A. Cadenillas and F. Zapatero, *Executive stock options with effort disutility and choice of volatility*.
2. L. Andersen and M. Broadie, *A primal-dual simulation algorithm for pricing multidimensional american options*, Management Science **50** (2004), 1222–1234.
3. I. Ekeland B. Bouchard and N. Touzi, *On the malliavin approach to monte carlo approximation of conditional expectations*, 2002.
4. V. Bally and C. Pages, *A quantization algorithm for solving multidimensional discrete time optimal stopping problems*, Bernoulli **9** (2003), 1003 – 1049.
5. J. Barraquand and D. Martineau, *Numerical valuation of high dimensional multivariate american securities*, Journal of Financial and Quantitative Analysis **30** (1995), 383405.
6. C. Bender and J. Schoenmakers, *An iterative algorithm for multiple stopping: Convergence and stability*, 2004.
7. D. Bosq, *Non-parametric statistics for stochastic processes*, Springer Verlag, 1998.
8. B. Bouchard and N. Touzi, *Discrete time approximation and monte-carlo simulation of backward stochastic differential equations*, 2002.
9. D. Brigo and F. Mercurio, *Interest rate models: Theory and practice*, Springer Finance, Springer Verlag, 2001.
10. A. Cadenillas, *Optimal dividend policy with mean-reverting cash reservoir*.
11. R. Carmona, *Statistical analysis of financial data in* `S-plus`, Springer Verlag, 2004.
12. R. Carmona and S. Dayanik, *Optimal multiple stopping of linear diffusions and swing options*, Tech. report, Princeton University, 2003.
13. _____, *Multiple dynkin games of timing and swing options*, Tech. report, Princeton University, 2006.
14. R. Carmona and V. Durrleman, *Pricing and hedging spread options*, SIAM Review **45** (2003), 627 – 685.
15. R. Carmona and M. Ludkovski, *Gas storage and supply guarantees: an optimal switching approach*, Tech. report, Princeton University, 2004.
16. _____, *Optimal switching with applications to energy tolling agreements*, Tech. report, Princeton University, 2004.
17. R. Carmona and N. Touzi, *Optimal multiple stopping and valuation of swing options*, 2003, p. (to appear).
18. J. F. Carrière, *Valuation of the early-exercise price for options using simulations and nonparametric regression*, Insurance, Mathematics and Economics **19** (1996), 19–30.
19. L. Clewlow and C. Strickland, *Implementing derivatives modesl*, John wiley & Sons, 1998.
20. M.H.A. Davis and I. Karatzas, *A deterministic approach to optimal stopping*, Probability, Statistics and Optimization: A Tribute to Peter Wittle, Wiley, 1994, pp. 455 – 466.
21. D. Lamberton E. Clément and P. Protter, *An analysis of a least squares regression method for american option pricing*, **6** (2002), 449–472.
22. D. Egloff and M. Min-Oo, *Convergence of monte carlo algorithms for pricing american options*, 2002.

23. P. Glasserman, *Monte carlo methods in financial engineering*, Springer-Verlag, 2004.
24. P. Glasserman and B. Yu, *Number of paths versus number of basis functions in american option pricing*, Annals of Applied Probability **14** (2004), 2090 – 2119.
25. S. Hamadene, *Mixed zero-sum stochastic differential game and american game options*, SIAM J. Control Optim. **45** (2006), no. 2, 496–518.
26. M. Haug and L. Kogan, *Pricing american options: a duality approach*, Operations Research **52** (2004), 258 – 270.
27. F. Jamshidian, *Numeraire-invariant option pricing and american, bermudan and trigger stream rollover*, Tech. report, University of Twente, 2004.
28. S.A. Ross J.C. Cox and M. Rubinstein, *Option pricing: a simplified approach*, Journal of Financial Economics', voluem = (1979).
29. J. Kallsen and C. Kuhn, *Pricing derivatives of american and game type in incomplete markets*, Finance and Stochastics **8** (2004), 261–284.
30. ———, *Convertible bonds: Financial derivatives of game type*, Exotic Option Pricing and Advanced Lévy Models, Wiley, 2005, pp. 277 – 288.
31. I. Karatzas, *A pathwise approach to dynkin games*, Statistics, Probability and Game Theory, IMS Lecture Notes - Monographs Series, vol. 30, IMS, 1996.
32. Y. Kifer, *Game options*, Finance and Stochastics **4** (2000), 443–463.
33. A. Kolodko and J. Shoenmakers, *Iterative construction of the optimal bermudan stopping time*, Finance Stochastics **10** (2005), 27 – 49.
34. D. Lamberton and B. Lapeyre, *Introduction to stochastic calculus applied to finance*, Chapman & Hall, 1996.
35. F. A. Longstaff and R. S. Schwartz, *Valuing american options by simulation: A simple least-square approach*, Review of Financial Studies **14** (2001), 113–147.
36. N. Meinshausen and B.M. Hambly, *Monte carlo methods for the valuation of multiple exercise options*, Mathematical Finance (2004).
37. J. Neveu, *Discrete parameter martingales*, North-Holland, 1975.
38. A. Pelsser, *Efficient methods for valuing interest rate derivatives*, Springer Finance, Springer Verlag, 2000.
39. L.C.G. Rogers, *Monte carlo valuation of american options*, Mathematical Finance **12** (2002), 271 – 286.
40. B. Van Roy and J. N. Tsitsiklis, *Regression methods for pricing complex american-style options*, IEEE Trans. on Neural Networks **12** (2000), 694–703.
41. R. Sircar and W. Xiong, *Evaluating incentive options*.
42. L. De Vroye, *Non-uniform variates*, Springer Verlag, 1984.
43. J.B. Walsh, *Convergence of the binomial tree model to black-scholes*, 2003, pp. 337 – 361.

Indexes

AUTHOR INDEX

NOTATION INDEX

SUBJECT INDEX

Z