

INTEREST RATES AND FX MODELS

7. Backward Induction and Monte Carlo Simulations

Andrew Lesniewski
Courant Institute of Mathematical Sciences
New York University
New York

March 13, 2013

Contents

1	Introduction	2
2	Bermudan swaptions	3
3	Valuation of American options	5
3.1	Principle of optimality and backward induction	5
3.2	Snell envelope	7
3.3	Conditional expected values	8
3.4	Monte Carlo simulations and estimation	9
3.5	Optimal exercise	11
A	Hermite polynomials and their properties	13
A.1	Classical Hermite polynomials	13
A.2	Normalized Hermite polynomials	14
A.3	The addition theorem	15

1 Introduction

Fixed income securities come in a great variety of shapes and sizes. They have often built in various termination options which allow a counterparty to exit the transaction during a contractually specified period of time. From the point of view of interest rates these exit options may be (i) endogenous, driven by rates only, or (ii) exogenous, when, in addition to interest rates, they are driven by other risk factors such as credit or prepayment events. The presence of termination options in an instrument has an impact on its valuation and risk profile. In this lecture, we focus on pure interest rate options whose modeling does not require methodologies going beyond interest rate models.

As an example, an early termination clause may allow a bond issuer to early repay the principal and cancel all future coupon payments (such bonds are called *callable*), or it may allow the bond holder to request early repayment of the principal (such bonds are called *puttable*) and forfeit the future coupon stream. Typically, these options are American style, i.e. they allow for multiple exercise dates. Such pure interest rate options are complicated enough to merit detailed understanding, and this is the subject of this chapter.

LIBOR market model is a powerful interest rate modeling methodology, and its practical value lies in the suitability to reliably model complex interest rate

options. We will cover here a limited number of topics only, and focus the discussion on the valuation of Bermudan swaptions by means of LMM style Monte Carlo simulations.

2 Bermudan swaptions

Synthetic instruments corresponding to these embedded termination options trade in financial markets in a variety of forms.

A *callable swap* is a fixed for LIBOR swap which gives the fixed rate payer a periodic right to cancel the remaining cash flows on the swap (“call the swap”). Typically, the cancelation dates follow an initial *lockout period*, when no calls are allowed. A typical example is a 10 year swap, which can be called every 6 months, two business days preceding a fixed rate roll date, following a 2 year lockout period. In the market lingo, this swap is called a 10 no call 2 swap.

As in the case of vanilla swaps discussed in Lecture 1, callable swaps are often created on the back of bond issuance. Consider the situation depicted in Figure 1, where an issuer issues a callable bond which is purchased by an investor. In addition to paying a periodic coupon and the principal at the bond’s maturity, the issuer has the right to call the bond following an initial lockout period. Because of the ability to finance themselves at a favorable spread to LIBOR and unwillingness to manage option risk, the issuer enters into a callable swap with an interest rate derivatives dealer. The dealer assumes the right to call the swap at the expense of paying an above the market fixed rate on the swap. Should the dealer call the swap, the issuer notifies the investor about it, and calls the bond.




Figure 1: Swapping a callable bond

A *puttable swap* is a similar structure on which the fixed rate receiver has a periodic right to cancel the swap, following an initial lockout period. In exchange for that right, the fixed rate on the swap is below the market swap rate for the term

of the swap. For example, on a 10 no put 1 swap, the fixed rate receiver can cancel the swap every quarter, two business days prior the fixed leg roll dates.

Multiple exercise swap options are often referred to as *Bermudan swaptions*. Unlike American options in the equity markets, a Bermudan swaption is a *multi-underlying* derivative. A Bermudan swaption can be exercised periodically, say semiannually or quarterly, or at any time excluding an initial lockout period. For example, a semiannual 5.50% 1 year Bermudan receiver swaption into a 5 year swap (a “1 into 5” Bermudan receiver) gives the holder the right to receive 5.50% on a 5 year swap starting in 1 year, or on a 4.5 year swap starting 1.5 year from now, or ..., or on a 6 month swap starting 5.5 year from now. Note that the total maturity of the transaction is $1 + 5 = 6$ years. More precisely, the option holder has the right to exercise the option on the 1 year anniversary of today (with the usual business day convention adjustments) and every six months thereafter, in which case they enter into a receiver swap starting two business days later. The maturity of the swap plus the number of years past the first exercise is equal to 5 years. Similarly, a quarterly 5.50% “5 into 10” Bermudan payer swaption gives the holder the right to pay 5.50% on a 10 year swap starting in 5 year or on a $9\frac{3}{4}$ year swap starting $5\frac{1}{4}$ years from now, or Bermudan swaptions trade over the counter.

It is convenient to think about a callable swap as a basket consisting of a vanilla swap and a Bermudan swaption. Consider, for example a 10 no call 1 callable swap paying 3%. From the point of view of the fixed rate receiver, this is a long position in a 10 year swap paying 3%, and a short position in a 1 into 9 Bermudan receiver swaption struck at 3%. Calling the swap by the fixed rate payer is equivalent to exercising that swaption. For example, upon exercise two years into the swap, the fixed rate payer enters into a swap on which they receives a fixed rate of 3% for the the remaining 8 years thus effectively canceling the original swap.

Oftentimes, exotic swaps are structured as callable swaps. As an example, consider a LIBOR range accrual. If, in addition, the coupon payer has a periodic right (following, as usual, a lockout period) to call the swap, this structure is called a callable range accrual. In analogy to a callable swap, a callable range accrual is a basket of a range accrual and a Bermudan option to enter into the trade offsetting the remainder of the underlying range accrual. This option is a complicated multi-underlying option where each of the underlyings is itself a basket of digital options on LIBOR.

3 Valuation of American options

Monte Carlo based valuation of options with multiple exercises poses more of a challenge. The reason is the apparent incompatibility between the forward moving process of path generation and the backward moving optimal exercise decision process. In recent years, a number of efficient approximate Monte Carlo algorithms for pricing American options have been proposed, notably by Longstaff and Schwartz [3]. The method we discuss here is a version of the Longstaff - Schwartz methodology.

3.1 Principle of optimality and backward induction

We consider an American option which can be exercised on a discrete set of dates. Specifically, let $T_0 = 0$ denote the valuation date, and let $T_1 < T_2 < \dots < T_n$, where $T_1 > 0$, denote the possible exercise dates. Upon exercise at time T_j the option seller delivers the underlying instrument. For example, in the case of a 1 into 10 Bermudan receiver swaption struck at K , the underlying instrument is the swap receiving coupon K which starts at T_j and matures in 11 years. Let $g_j(T_j)$ denote the payoff at the exercise time T_j .

The act of exercising the American option is modeled by a stopping time¹ τ , a random variable which represents an early exercise policy for the option. For convenience, we shall work under the terminal forward measure which will be denoted by \mathbb{Q} . Today's value $V_0 = V(T_0)$ of the option is the maximum over all possible stopping times:

$$V_0 = \max_{\tau \in \mathcal{T}(T_1, \dots, T_n)} \mathbb{E}^{\mathbb{Q}}[P(0, \tau) g(\tau)], \quad (1)$$

where $\mathcal{T}(T_1, \dots, T_n)$ denotes the set of stopping times taking values in the set $\{T_1, \dots, T_n\}$, and where $P(t, T)$ is the discount factor. We seek to determine the *optimal stopping time* τ_* (or, the *optimal exercise policy*) such that

$$V_0 = \mathbb{E}^{\mathbb{Q}}[P(0, \tau_*) g(\tau_*)]. \quad (2)$$

In order to find the optimal exercise policy, and thus the value of an American option, we shall invoke Bellman's *principle of optimality* which reads (in his own

¹Recall that a stopping time is a random variable τ such that $\{\tau \leq t\} \in \mathcal{F}_t$ for all t . In our context, this means that the exercise decision at time t is made solely based on the information prior to t .

words): *An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision.* Bellman's principle leads us thus to require that, given a state of the world, the value $V_j = V(T_j)$ of the option at any exercise date T_j , $j = 1, \dots, n - 1$, satisfies

$$V_j = \max_{\tau \in \mathcal{T}(T_{j+1}, \dots, T_n)} E_j[P(T_j, \tau) g(\tau)]. \quad (3)$$

where $E_j[\cdot]$ is a shorthand for $E^Q[\cdot | \mathcal{F}_{T_j}]$. Clearly, we have

$$V_n = E_n[g(T_n)]. \quad (4)$$

This can be expressed algorithmically in terms of *backward induction*. At each future scenario, and at each future time the option holder compares the value of the immediate exercise versus the remaining value of the option. Given a future scenario, he thus starts at the latest possible exercise date, where the value of the option is simply equal to its intrinsic value (as there are no further exercise opportunities left). He then moves backward in time, computes the current value of the remaining option, and compares it to its intrinsic value. He continues this processes until reaching the first possible exercise date. The earliest exercise date, along the scenario, at which the intrinsic value exceeds the continuation value of the option is the optimal exercise date. He repeats this analysis for all future scenarios.

This algorithm is easy to implement for models which allow for recombining tree implementation, such as the Hull - White model. LMM does not allow for a natural implementation based on recombining trees, and thus all valuations have to be performed via Monte Carlo simulations. Because backward induction is an essential part of the valuation logic, calculating prices of American options by means of naive Monte Carlo simulations is inefficient. The key reason why backward induction is difficult to combine with Monte Carlo simulations lies in the fact that it is hard to estimate conditional expected values given a future scenario. Suppose that initially we have generated N Monte Carlo scenarios. At each decision point along a scenario, one would have to generate a new set of N Monte Carlo paths in order to calculate the continuation value of the option and the value of the immediate exercise. This process would result in a number of paths that grows fast in the number of exercises and prove computationally extremely expensive.

3.2 Snell envelope

For simplicity of exposition we assume that our model is a one factor model driven by the Wiener process $W(s)$, such as the 1 factor LMM model. Having d factors is merely a notational complication which would obfuscate the main ideas of the method. We introduce the notation:

$$X_j = \frac{W(T_j)}{\sqrt{T_j}}, \quad j = 1, \dots, n,$$

so that each X_j is a standard normal random variable,

$$X_j \sim N(0, 1), \quad (5)$$

and we let $X_{1:j}$ denote the vector:

$$X_{1:j} = (X_1, \dots, X_j). \quad (6)$$

Finally, we denote by $g_j = g(T_j)$ the payoff function at time T_j . Note that g_j is a function of the path $X_{1:j}$ and, when appropriate, we will make this fact explicit by writing, somewhat inconsistently, $g_j = g_j(X_{1:j})$.

We now proceed in a number of steps. The first step is to construct an auxiliary process $S_j = S_j(X_{1:j})$, called the *Snell envelope* of g_j . We begin with $j = n$, and set

$$S_n(X_{1:n}) = g_n(X_{1:n}). \quad (7)$$

As stated before, the time T_n value of the option is equal to the value of the immediate payoff, as no further opportunities to exercise exist, and thus $V_n = S_n$. On the other hand, for $j = n - 1, \dots, 1$, the value of the option is the greater of the immediate payoff and the *continuation value* $C_j = C_j(X_{1:j})$ of the option. We set

$$S_j(X_{1:j}) = \max(g_j(X_{1:j}), C_j(X_{1:j})). \quad (8)$$

The option should be exercised if the maximum in (8) equals g_j ; otherwise the option does not get exercised. Clearly, it is optimal to exercise at the earliest opportunity when $\max(g_j, C_j) = g_j$. Later we will utilize this observation in order to relate S_j to V_j .

The continuation value of the option at time T_j is the current (i.e. time T_j) value of the remaining options conditioned on the current state of the world. In mathematical terms,

$$C_j(X_{1:j}) = \begin{cases} P_{j,j+1} \mathbf{E}_j[S_{j+1}], & \text{for } j = 0, \dots, n-1, \\ 0, & \text{for } j = n, \end{cases} \quad (9)$$

where $P_{j,k} = P(T_j, T_k)$ is the discount factor from T_j to T_k , $j < k$. The value S_0 of the Snell envelope on the value date is

$$\begin{aligned} S_0 &= P_{0,1} \mathbb{E}[S_1] \\ &= C_0. \end{aligned} \tag{10}$$

The Snell envelope has the property that $S_j \geq g_j$, for all j . In fact, one can show that S_j is the smallest *supermartingale* with this property². Note also that S_j is not, in general, equal to V_j , as it does not encode the information when it is optimal to exercise the option.

The key to implementation of the above procedure is the ability to compute, at each exercise date T_j and for each scenario $X_{1:j}$, (i) the value of the immediate payoff of the option $g_j(X_{1:j})$, and (ii) the continuation value of the option $C_j(X_{1:j})$. In the following two sections we present a methodology for such computations.

3.3 Conditional expected values

Consider a function $f(X_{1:j})$. Our goal is to compute algorithmically the conditional expected value $\mathbb{E}_i[f(X_{1:j})]$, where $i < j$.

Let us first take $i = j - 1$. In order to compute the conditional expectation, we proceed as follows.

Step 1. We write for $j = 1, \dots, n$,

$$W(T_j) = W(T_{j-1}) + (W(T_j) - W(T_{j-1})),$$

and note that this identity implies

$$\sqrt{T_j} X_j = \sqrt{T_{j-1}} X_{j-1} + \sqrt{T_j - T_{j-1}} \xi_j,$$

with $\xi_j \sim N(0, 1)$ independent of X_{j-1} . Dividing by $\sqrt{T_j}$, we obtain the decomposition

$$X_j = \sqrt{\alpha_j} X_{j-1} + \sqrt{1 - \alpha_j} \xi_j, \quad \text{with } \xi_j \sim N(0, 1), \tag{11}$$

where the innovation ξ_j is independent of X_{j-1} , and α_j is defined as the ratio of the consecutive exercise times,

$$\alpha_j = \frac{T_{j-1}}{T_j}. \tag{12}$$

²Recall that a supermartingale is a process X such that $X(s) \geq \mathbb{E}[X(t) | \mathcal{F}_s]$, for $s \leq t$.

This decomposition is key in the following calculations.

Step 2. Viewing $f(X_{1:j})$ as a function of X_j , we can, according to (42) - (43), represent it as a series of Hermite polynomials,

$$f(X_{1:j}) = \sum_{0 \leq k < \infty} c_k(X_{1:j-1}) h_k(X_j). \quad (13)$$

The Fourier coefficients c_k are calculated by

$$c_k(X_{1:j-1}) = \int_{-\infty}^{\infty} f(X_{1:j}) h_k(X_j) d\mu(X_j). \quad (14)$$

Step 3. Using the decomposition (11) and the conditioning rule (47) for Hermite polynomials we find that

$$\begin{aligned} \mathbb{E}_{j-1}[h_k(X_j)] &= \mathbb{E}_{j-1}[h_k(\sqrt{\alpha_j} X_{j-1} + \sqrt{1 - \alpha_j} \xi_j)] \\ &= \alpha_j^{k/2} h_k(X_{j-1}). \end{aligned}$$

Applying this identity to the expansion (13), we finally obtain the following result:

$$\mathbb{E}_{j-1}[f(X_{1:j})] = \sum_{0 \leq k < \infty} \alpha_j^{k/2} c_k(X_{1:j-1}) h_k(X_{j-1}). \quad (15)$$

This formula allows us, in principle, to compute all relevant conditional expectations in an algorithmic manner.

Let us now turn to the general case of $i < j$. Since taking conditional expected values can be nested,

$$\mathbb{E}_i[f(X_{1:j})] = \mathbb{E}_i[\dots \mathbb{E}_{j-2}[\mathbb{E}_{j-1}[f(X_{1:j})]] \dots], \quad (16)$$

$\mathbb{E}_i[f(X_{1:j})]$ can be computed in $j - i$ iterations of the steps described above.

3.4 Monte Carlo simulations and estimation

The problem with the above calculation is, of course, that it requires taking complicated integrals and summing infinite series. A good numerical methodology is needed in order to make the calculation practical and accurate. Assume that we have generated N Monte Carlo paths $X_{1:n}^i$, $i = 1, \dots, n$, representing samples from the underlying market dynamics.

Numerical implementation of the methodology above is done in the three following steps.

1. *Truncation.* The expansions with respect to the Hermite polynomials are truncated at some finite order κ , i.e.:

$$f(X_{1:j}) \approx \sum_{0 \leq k \leq \kappa} c_k(X_{1:j-1}) h_k(X_j). \quad (17)$$

The conditional expected value has thus the truncated form:

$$\mathbb{E}_{j-1}[f(X_{1:j})] \approx \sum_{0 \leq k \leq \kappa} \alpha_j^{k/2} c_k(X_{1:j-1}) h_k(X_{j-1}). \quad (18)$$

Remarkably, relatively low values of κ , $\kappa \sim 5$, lead to good numerical results.

2. *Estimation.* The Fourier coefficients c_k are replaced by their estimates \hat{c}_k . The naive Monte Carlo estimate, i.e.:

$$\hat{c}_k \approx \frac{1}{N} \sum_{1 \leq i \leq N} f(X_{1:j}^i) h_k(X_j^i), \quad (19)$$

does not perform well, and should not be used. Instead, we shall choose the vector \hat{c} of coefficients \hat{c}_k so as to minimize the square error:

$$\hat{c} = \arg \min_{c_0, \dots, c_\kappa} \sum_{1 \leq i \leq N} \left(f(X_{1:j}^i) - \sum_{0 \leq k \leq \kappa} c_k h_k(X_j^i) \right)^2.$$

This amounts to regressing the values the function f takes on the Monte Carlo paths $X_{1:j}^i$ on the first few Hermite polynomials evaluated on X_j^i . The Fourier coefficients are thus estimated as regression coefficients, and are explicitly given by

$$\hat{c} = G^{-1}v. \quad (20)$$

Here G is a $\kappa \times \kappa$ matrix whose components are:

$$G_{kl} = \frac{1}{N} \sum_{1 \leq i \leq N} h_k(X_j^i) h_l(X_j^i), \text{ for } k, l = 1, \dots, \kappa, \quad (21)$$

and v is a vector of dimension κ with the components:

$$v_k = \frac{1}{N} \sum_{1 \leq i \leq N} f(X_{1:j}^i) h_k(X_j^i), \text{ for } k = 1, \dots, \kappa. \quad (22)$$

In the limit as $N \rightarrow \infty$, G tends to the identity matrix, and so (20) and (19) coincide in that limit. We thus arrive at the following *estimated conditional expected value* $\widehat{E}_{j-1}[f(X_{1:j})]$:

$$\widehat{E}_{j-1}[f(X_{1:j})] = \sum_{0 \leq k \leq \kappa} \alpha_j^{k/2} \widehat{c}_k h_k(X_{j-1}). \quad (23)$$

It is obtained from (19) by replacing the c'_k s with their estimated values \widehat{c}_k s.

3. Acceleration. From the performance point of view, is worthwhile to try to accelerate the convergence of the series (17). Payoff functions of options are, typically, not smooth. This causes slowdowns of the rate of convergence at the points where the payoff has kinks (as in the usual calls or puts) or discontinuities (as in digital options). This behavior resembles somewhat of the Gibbs phenomenon encountered in the practice (and theory) of Fourier series, and can be remedied by using an accelerating filter. This step is really optional and we leave out the technical details.

3.5 Optimal exercise

Let us now go back to the issue of numerically determining the optimal exercise policy for an American option. Suppose that we are given N Monte Carlo paths $X_{1:n}^i$, $i = 1, \dots, N$, simulating the underlying dynamics. We shall proceed in two steps.

Step 1. Assume for simplicity that the immediate payoff of the option is easy to compute, i.e. it does not require calculating expected values. For example, for a Bermudan swaption, the underlying instrument is a swap which can be priced based on the multicurve built off the data available at T_j . If all the $g_j(X_{1:j}^i)$ are known, the first step consists of computing the Snell envelope only.

We proceed inductively. Start with $j = n$, and set

$$\begin{aligned} C_n(X_{1:n}^i) &= 0, \\ S_n(X_{1:n}^i) &= g_n(X_{1:n}^i). \end{aligned} \quad (24)$$

Assume that we have already computed $C_j(X_{1:j}^i)$ and $S_j(X_{1:j}^i)$, for some $1 < j < n$, and all $i = 1, \dots, N$. Then, we set

$$\begin{aligned} C_{j-1}(X_{1:j-1}^i) &= P_{j-1,j} \widehat{E}_{j-1}[S_j(X_{1:j})], \\ S_{j-1}(X_{1:j-1}^i) &= \max(g_j(X_{1:j-1}^i), C_j(X_{1:j-1}^i)), \end{aligned} \quad (25)$$

where $P_{j-1,j}$ is the relative discount factor. We continue this process until we reach $j = 1$.

If you do not want to get lost in the technical details that follow, you may want, on the first reading, to skip directly to Step 2.

On the other hand, if the underlying instrument itself is a contingent claim, its price is a conditional expected value. In this case, the first step requires computing the payoffs $g_j(X_{1:j}^i)$ along with the continuation values. In order to prepare for it, we first pre-compute all $p_j(X_{1:n}^i)$, for $j = 1, \dots, n, i = 1, \dots, N$, where $p_j(X_{1:n}^i)$ denotes the present value at T_j of the cash flows of the underlying instrument along the Monte Carlo path $X_{1:n}^i$. The payoff $g_j(X_{1:j}^i)$ is then equal to the conditional expected value of $p_j(X_{1:n}^i)$ given $X_{1:j}^i$, namely $g_j(X_{1:j}^i) = E_j[p_j(X_{1:n}^i)]$. Using the nesting property (16) of conditional expectation, we write this as

$$g_j(X_{1:j}^i) = E_j[E_{j+1}[\dots E_n[p_j(X_{1:n}^i)] \dots]], \quad (26)$$

and calculate $g_j(X_{1:j}^i)$ in a sequence of steps starting at the last exercise date.

We thus arrive at the following inductive procedure. Start with $j = n$, and set

$$\begin{aligned} g_l(X_{1:n}^i) &= p_l(X_{1:n}^i), \text{ for } l = 1, \dots, n, \\ C_n(X_{1:n}^i) &= 0, \\ S_n(X_{1:n}^i) &= g_n(X_{1:n}^i). \end{aligned} \quad (27)$$

Assume that we have computed $p_1(X_{1:j}^i), \dots, p_j(X_{1:j}^i), C_j(X_{1:j}^i)$, and $S_j(X_{1:j}^i)$, for some $1 < j < n$, and all $i = 1, \dots, N$. Then, we set

$$\begin{aligned} p_l(X_{1:j-1}^i) &= \widehat{E}_{j-1}[p_l(X_{1:j}^i)], \text{ for } l = 1, \dots, j-1, \\ g_{j-1}(X_{1:j-1}^i) &= p_{j-1}(X_{1:j-1}^i), \\ C_{j-1}(X_{1:j-1}^i) &= P_{j-1,j} \widehat{E}_{j-1}[S_j(X_{1:j}^i)], \\ S_{j-1}(X_{1:j-1}^i) &= \max(g_j(X_{1:j-1}^i), C_j(X_{1:j-1}^i)), \end{aligned} \quad (28)$$

where, as before, $P_{j-1,j}$ is the relative discount factor. We continue this process until we reach $j = 1$.

Having estimated the conditional expected values (and thus calculated the consecutive continuation values), we construct the optimal stopping time as follows. **Step 2.** Let now $X^i = (X_0^i, X_1^i, \dots, X_n^i)$, $i = 1, \dots, N$, be a Monte Carlo path. For each i , we initialize the value of the stopping time to the last possible option expiration, $\tau_*(X^i) = T_n$. As we move backward along X^i , we update the value

of $\tau_*(X^i)$ according to the following rule

$$\begin{aligned} & \mathbf{for } j = n \mathbf{ to } j = 1 \\ & \quad \mathbf{if } g_j(X_{1:j}^i) \geq C_j(X_{1:j}^i) \\ & \quad \quad \mathbf{then } \tau_*(X_{1:n}^i) = T_j \end{aligned} \quad (29)$$

In other words,

$$\tau_*(X_{1:n}^i) = \min \{j : S_j(X_{1:j}^i) = g_j(X_{1:j}^i)\}. \quad (30)$$

The value of the option is the arithmetic mean over all Monte Carlo paths:

$$V_0 \approx \frac{1}{N} \sum_{1 \leq i \leq N} P_{0,j_*} g_{j_*}(X_{1:j_*}^i), \quad (31)$$

where $j_* = j_*(X_{1:n}^i)$ is the index j for which $\tau_*(X_{1:n}^i) = T_j$. This formula represents a Monte Carlo approximation to the valuation formula (2).

A Hermite polynomials and their properties

A.1 Classical Hermite polynomials

The *classical Hermite polynomials* $H_n(x)$, $n = 0, 1, \dots$, are defined by means of the expansion:

$$\Psi(t, x) = \sum_{n \geq 0} \frac{t^n}{n!} H_n(x), \quad (32)$$

where $\Psi(t, x)$, the *generating function*, is given by

$$\Psi(t, x) = e^{-t^2+2tx}. \quad (33)$$

This expansion can be viewed as the Taylor expansion of the function $t \rightarrow \Psi(t, x)$ around $t = 0$, and it thus implies that

$$\begin{aligned} H_n(x) &= \left. \frac{d^n}{dt^n} e^{-t^2+2tx} \right|_{t=0} \\ &= (-1)^n e^{x^2} \frac{d^n}{dx^n} e^{-x^2}. \end{aligned} \quad (34)$$

Explicitly, the first few Hermite polynomials are:

$$\begin{aligned} H_0(x) &= 1, \\ H_1(x) &= 2x, \\ H_2(x) &= 4x^2 - 2, \\ H_3(x) &= 8x^3 - 12x, \\ &\vdots \end{aligned} \tag{35}$$

Simple manipulations using formula (34) show that the Hermite polynomials satisfy the following recurrence relations:

$$\begin{aligned} H_{n+1}(x) &= 2xH_n(x) - 2nH_{n-1}(x), \\ H'_n(x) &= 2nH_{n-1}(x). \end{aligned} \tag{36}$$

These in turn lead to the following differential equation obeyed by $H_n(x)$:

$$H''_n(x) - 2xH'_n(x) + 2nH_n(x) = 0. \tag{37}$$

A.2 Normalized Hermite polynomials

It will be convenient for our purposes to use the *normalized Hermite polynomials*, denoted here by $h_n(x)$, $n = 0, 1, \dots$, which are defined as follows:

$$h_n(x) = \frac{1}{2^{n/2}\sqrt{n!}} H_n\left(\frac{x}{\sqrt{2}}\right). \tag{38}$$

The reader will verify that the recurrence relations (36) restated in terms of the normalized Hermite polynomials read:

$$\begin{aligned} \sqrt{n+1} h_{n+1}(x) &= x h_n(x) - \sqrt{n} h_{n-1}(x) \\ h'_n(x) &= \sqrt{n} h_{n-1}(x). \end{aligned} \tag{39}$$

From these relations, one verifies readily that the normalized Hermite polynomials form an orthonormal system,

$$\int_{-\infty}^{\infty} h_m(x) h_n(x) d\mu(x) = \delta_{mn}, \tag{40}$$

with respect to the Gaussian measure

$$d\mu(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx. \tag{41}$$

In fact, they form an orthonormal basis for the Hilbert space $L^2(\mathbb{R}, d\mu)$ of functions on the real line which are square integrable with respect to the Gaussian measure (41). Namely, any function $f \in L^2(\mathbb{R}, d\mu)$ can be expressed as a convergent series,

$$f(x) = \sum_{n \geq 0} a_n h_n(x), \quad (42)$$

where the *Fourier coefficient* a_n is given by:

$$a_n = \int_{-\infty}^{\infty} h_n(x) f(x) d\mu(x). \quad (43)$$

A.3 The addition theorem

A useful property of the Hermite polynomials is the addition theorem. For any $0 \leq \alpha \leq 1$, the following identity holds

$$\begin{aligned} H_n(\alpha^{1/2}x + (1-\alpha)^{1/2}y) \\ = \sum_{0 \leq k \leq n} \binom{n}{k} \alpha^{k/2} (1-\alpha)^{(n-k)/2} H_k(x) H_{n-k}(y). \end{aligned} \quad (44)$$

The proof is an immediate consequence of the following identity for the generating function (33):

$$\Psi(t, \alpha^{1/2}x + (1-\alpha)^{1/2}y) = \Psi(\alpha^{1/2}t, x) \Psi((1-\alpha)^{1/2}t, y). \quad (45)$$

This identity is easy to establish by manipulating the properties of the exponential function. The normalized version of (44) reads:

$$\begin{aligned} h_n(\alpha^{1/2}x + (1-\alpha)^{1/2}y) \\ = \sum_{0 \leq k \leq n} \sqrt{\binom{n}{k}} \alpha^{k/2} (1-\alpha)^{(n-k)/2} h_k(x) h_{n-k}(y). \end{aligned} \quad (46)$$

A consequence of the identity above is the following *conditioning rule*:

$$\int_{-\infty}^{\infty} h_n(\alpha^{1/2}x + (1-\alpha)^{1/2}y) d\mu(y) = \alpha^{n/2} h_n(x). \quad (47)$$

References

- [1] Andersen, L., and Piterbarg, V.: *Interest Rate Modeling*, Vol. 3, Atlantic Financial Press (2010).
- [2] Glasserman, P.: *Monte Carlo Methods in Financial Engineering*, Springer (2004).
- [3] Longstaff, F., and Schwartz, E.: Valuing American options by simulation: a simple least squares approach, *Rev. Financial Stud.*, **14**, 113 - 147 (2001).