MONTE CARLO METHODS IN FINANCE: AN INTRODUCTORY TUTORIAL

Sandeep Juneja

School of Technology and Computer Science
Tata Institute of Fundamental Research
Mumbai, 400 005, INDIA

ABSTRACT

In this introductory tutorial we discuss the problem of pricing financial derivatives, the key application of Monte Carlo in finance. We review the mathematics that uses no-arbitrage principle to price derivatives and expresses derivatives price as an expectation under the equivalent martingale measure. In the presentation at the conference we will also elaborate on the use of Monte Carlo methods for pricing American options and in portfolio risk measurement.

1 INTRODUCTION

A key application of Monte Carlo methods in modern finance is in pricing and hedging complex financial derivatives especially when these derivatives have early exercise 'American' features built into them. Monte Carlo methods are also used in risk management for credit, market as well as operational risks. These methods increasingly find utility in portfolio optimization as well as in model calibration. In this introductory tutorial we spell out the underlying mathematics that makes derivatives pricing amenable to Monte Carlo. We briefly review the framework for pricing American/Bermudan options and some related literature. At the presentation in the conference we will also review the techniques developed to price American/Bermudan options as well as the use of Monte Carlo methods in portfolio risk measurement.

Recall that options are popular financial derivatives that give the buyer an option but not an obligation to engage in specific transactions related to the underlying assets. For instance, a call option allows the buyer an option but not an obligation to purchase an underlying asset at a specified price at a particular time in future, referred to as time to maturity. Similarly, a put option gives the buyer the option but not an obligation to sell an asset at a specified price at the time to maturity. Broadly speaking, payoffs from options are a function of the trajectory followed by the underlying asset prices up to a specified time. European options can be exercised only at a fixed time, while American options can be exercised at any time up to a fixed time. This early exercise opportunity embedded in American options makes the problem of pricing them far more challenging. Bermudan options are intermediate between European and American options in that they can be exercised at any of the finite number of specified set of times. If the number of exercise opportunities are large enough and well spaced, the associated Bermudan option price closely approximates the price of an American option. Typically, American options are numerically priced by pricing a Bermudan option that closely approximates it.

The no-arbitrage principle for pricing European derivatives, introduced by Black and Scholes (1973), relies on replicating the pay-off from the derivative in every possible scenario by a series of intricate trades involving the underlying more basic securities, carried out every infinitesimal interval till option

maturity. Then, since the derivative and this trading strategy have identical payoffs, the no-arbitrage principle posits that they must have the same price, otherwise an arbitrage is created (positive profits are generated while taking zero risk). This reasoning allows European derivative's price to be expressed as an expectation of its discounted payoff under a probability measure different from the physical measure, referred to as risk neutral or equivalent martingale measure. This representation makes the pricing problem amenable to Monte Carlo. We briefly review the associated mathematics in Section 2 (see, e.g., Shreve (2004), Duffie (1996), Steele (2001) for a comprehensive yet accessible account).

American or Bermudan options have the early exercise feature that add further non-trivial challenges to pricing them. Essentially, one may associate with such an option an exercise policy, i.e., a prescription that specifies in every scenario when to exercise the option. Using the no-arbitrage principle, the value of the option under each such policy can again be expressed as an expectation of a random variable. The rational price for the American option equals that of the policy having the maximum value (otherwise, an arbitrage can be created, see Duffie (1996)). Finding this policy and hence the value of the option is difficult. We outline the mathematical framework to analyze American options and briefly review some of the literature that addresses this problem using Monte Carlo, in Section 3.

2 Mathematical Framework

We now briefly outline the argument to construct a replicating portfolio for a derivative that shows that the no-arbitrage derivative security price is the expectation of its discounted payoff under an equivalent martingale measure. We do this in a simple setting of two securities. One risky security or stock price. The other a risk free security or the money market account. First step in this effort is to find an equivalent martingale measure under which both the securities when discounted are martingales. This ensures that any portfolio that dynamically trades in the two securities is also a martingale. In particular, if the derivative security can be replicated by such a portfolio, its price process is a martingale. This provides a conditional expectation representation of the option price process.

This argument requires two broad steps. First involves identifying the equivalent martingale measure. Crucial to this is the use of Girsanov Theorem. Second requires identifying the replicating portfolio for a given derivative. Here the Martingale Representation Theorem plays a key role. Note that a market where all derivative securities can be replicated are referred to as complete (else, incomplete).

2.1 The Model

Suppose that the stock price $(S_t : 0 \leq t \leq T)$ follows the stochastic differential equation (SDE)

$$dS_t = \mu_t S_t dt + \sigma_t S_t dW_t$$

under the probability measure $\mathscr{P}$. This is a differential and a convenient shorthand notation for the rigorous form

$$S_t = S_0 + \int_0^t \mu_s S_s ds + \int_0^t \sigma_s S_s dW_s. \tag{1}$$

In the above equation, $(\mu_t : 0 \leq t \leq T)$ and $(\sigma_t : 0 \leq t \leq T)$ are allowed to be a random processes where for each t, μ_t and σ_t are known by time t (that is, they are *adapted* processes). $(W_t : 0 \leq t \leq T)$ is Brownian motion: It has continuous paths, its increments are stationary and independent, $W(0) = 0$ and $W(t)$ has a Gaussian distribution with mean zero and variance one. $\int_0^t \mu_s S_s ds$ is the usual Lebesgue integral, however $\int_0^t \sigma_s S_s dW_s$ is the stochastic Ito integral, which under mild conditions can be seen to be zero mean martingales.

Equation (1) may be best viewed through its discretized Euler's approximation at times t and $t + \delta t$:

$$\frac{S_{t+\delta t} - S_t}{S_t} = \mu_t \delta t + \sigma_t (W_{t+\delta t} - W_t).$$

This suggests that μ_t captures the drift in the instantaneous return from the security at time t. Similarly, σ_t captures the sensitivity to the independent noise $W_{t+\delta t} - W_t$ present in the instantaneous return at time t.

The money market $(R_t : 0 \leq t \leq T)$ is governed by a *short rate* adapted stochastic process $(r_t : 0 \leq t \leq T)$ and satisfies the differential equation

$$dR_t = r_t R_t dt$$

with $R_0 = 1$. Here short rate r_t corresponds to instantaneous return on investment in the money market at time t and is known at time t. In particular, \$1 invested at time zero in the money market equals $R_t = \exp(\int_0^t r_s ds)$ at time t.

2.2 Equivalent martingale measure

We first look for a probability measure under which the discounted security process $R_t^{-1} S_t$ is martingale. For this purpose, first recall the Ito's lemma

$$df(t, S_t) = f_t(t, S_t)dt + f_x(t, S_t)dS_t + \frac{1}{2} f_{xx}(t, S_t)\sigma_t^2 dt \qquad (2)$$

where f_t denotes the partial derivative w.r.t. the first argument of $f(\cdot, \cdot)$ and f_x and f_{xx} denote the first and the second order partial derivatives with respect to its second argument.

Using the Ito's lemma, the discounted stock price process satisfies the relation

$$d(R_t^{-1} S_t) = -r_t R_t^{-1} S_t + R_t^{-1} dS_t = R_t^{-1} S_t \left((\mu_t - r_t)dt + \sigma_t dW_t \right).$$

(This relation follows from (2) if we assume that the short rate process is deterministic. However, this result is true more generally as can be seen by applying the multi-dimensional Ito's lemma).

Now we look for a probability measure that makes this a martingale, that is, under which the drift coefficient associated with the dt term is zero.

Girsanov theorem is useful here. For that consider the positive process

$$Y_t = \exp \left(\int_0^t v_s dW_s - \frac{1}{2} \int_0^t v_s^2 ds \right)$$

for a well behaved process $(v_t : 0 \leq t \leq T)$. Using Ito's lemma it can be seen that

$$dY_t = Y_t v_t dW_t$$

or in its meaningful form

$$Y_t = 1 + \int_0^t Y_s v_s dW_s.$$

Under technical conditions the stochastic integral is a mean zero martingale so that $(Y_t : 0 \leq t \leq T)$ is a positive mean 1 martingale. Let $\mathscr{P}^v(A) = E_{\mathscr{P}}[Y_T I(A)]$. This is clearly a probability measure that is equivalent to the original probability measure $\mathscr{P}$ (Two probability measures are equivalent if they assign positive probability to same sets). Girsanov Theorem states that under the probability measure

$\mathscr{P}^v$, under technical conditions on $(v_t : 0 \leq t \leq T)$, the process $(W_t^v : 0 \leq t \leq T)$ where

$$W_t^v = W_t - \int_0^t v_s ds$$

(or $dW_t^v = dW_t - v_t dt$ in the differential notation) is a standard Brownian motion. Equivalently, $(W_t : 0 \leq t \leq T)$ is standard Brownian motion plus the drift process $(\int_0^t v_s ds : 0 \leq t \leq T)$.

It is now easy to see from Girsanov Theorem that if $\sigma_t > 0$ almost surely, then under the probability measure $\mathscr{P}^v$ with $v_t = \frac{r_t - \mu_t}{\sigma_t}$ the discounted stock price process is a martingale, so that $\mathscr{P}^v$ is the equivalent martingale measure. This is true as

$$d(R_t^{-1} S_t) = R_t^{-1} S_t \sigma_t (dW_t - v_t dt) = R_t^{-1} S_t \sigma_t dW_t^v$$

is a stochastic integral and hence a martingale (modulo technical conditions). It is easy to see by applying the Ito's formula (Note that $S_t = R_t X_t$ where $X_t = R_t^{-1} S_t$ satisfies the SDE above) that

$$dS_t = r_t S_t dt + S_t \sigma_t dW_t^v. \tag{3}$$

Hence, under the equivalent martingale measure $\mathscr{P}^v$, the drift of the stock price process changes from $\{\mu_t\}$ to $\{r_t\}$. Therefore, $\mathscr{P}^v$ is also referred to as the risk neutral measure.

2.3 Replicating Portfolio Process

Now consider the problem of creating the replicating process for an option with pay-off H. For instance, for a call option that matures at time T with strike price K, we have $H = \max(S_T - K, 0)$.

Define V_t for $0 \leq t \leq T$ so that

$$R_t^{-1} V_t = E_{\mathscr{P}^v}[R_T^{-1} H | \mathscr{F}_t].$$

Here $\mathscr{F}_t$ denotes the information available by time t. Heuristically, this corresponds to knowing $(W_s : 0 \leq s \leq t)$. Also note that $V_T = H$.

Our plan is to construct a replicating portfolio process $(P_t : 0 \leq t \leq T)$ such that $P_t = V_t$ for all t. Then, since $P_T = H$ we have replicated the option with this portfolio process and P_t then denotes the price of the option at time t, i.e.,

$$P_t = E_{\mathscr{P}^v} \left[\exp \left(- \int_t^T r_s ds \right) H | \mathscr{F}_t \right]. \tag{4}$$

Then, the option price is simply the conditional expectation of the discounted option payoff under the risk neutral or the equivalent martingale measure.

To this end, it is easily seen from the law of iterated conditional expectations that for $s < t$,

$$R_s^{-1} V_s = E_{\mathscr{P}^v} \left(E_{\mathscr{P}^v}[R_T^{-1} V_T | \mathscr{F}_t] | \mathscr{F}_s \right) = E_{\mathscr{P}^v}[R_t^{-1} V_t | \mathscr{F}_s],$$

that is, $(R_t^{-1} V_t : 0 \leq t \leq T)$ is a martingale. Martingale Representation Theorem essentially states that if all the noise in the system is due to the Brownian motion then every martingale is a stochastic integral. Hence, there exists an adapted process $(w_t : 0 \leq t \leq T)$ such that

$$d(R_t^{-1} V_t) = w_t dW_t^v. \tag{5}$$

Now consider a portfolio process $(P_t : 0 \leq t \leq T)$ with the associated recipe process $(b_t : 0 \leq t \leq T)$. This means that we start with wealth P_0. At any time t, the number of stocks kept in the portfolio equals b_t. The portfolio value is denoted by P_t. The amount $P_t - b_t S_t$ is invested in the money market.

Then, the portfolio process evolves as in

$$dP_t = (P_t - b_t S_t) r_t dt + b_t dS_t. \tag{6}$$

The discounted portfolio process evolves as

$$d(R_t^{-1} P_t) = R_t^{-1} b_t S_t ((\mu_t - r_t) dt + \sigma_t dW_t) = R_t^{-1} b_t S_t \sigma_t dW_t^\nu. \tag{7}$$

Therefore, under technical conditions, the discounted portfolio process being a stochastic integral, is a martingale under $\mathscr{P}^\nu$.

From (5) and (7) its clear that if we set $P_0 = V_0$ and $b_t = \frac{w_t R_t}{S_t \sigma_t}$ then $P_t = V_t$ for all t. In particular we have constructed a replicating portfolio. In particular, under technical conditions, primarily that $\sigma_t > 0$ almost surely, for almost every t every option can be replicated so that this market is complete.

2.4 Multiple Underlying Securities

The analysis can be similarly extended to d risky assets or stocks driven by m independent sources of noise modeled by independent Brownian motions $(W_t^1, \ldots, W_t^m : 0 \le t \le T)$. Specifically, we assume that d assets $(S_t^1, \ldots, S_t^d : 0 \le t \le T)$ satisfy the SDE

$$dS_t^i = \mu_t^i S_t^i dt + \sum_{j=1}^m \sigma_t^{ij} S_t^i dW_t^j$$

for $i = 1, \ldots, d$. Here, we assume that each μ^i and σ^{ij} is an adapted process and satisfy restrictions so that the integrals associated with the SDEs are well defined. In addition we let $R_t = \exp(\int_0^t r(s)ds)$ as before.

Then it can be similarly seen that (see, e.g., Duffie (1996), Karatzas and Shreve (1998)) under the no-arbitrage and some additional regularity conditions, the market is complete and there exists a unique equivalent martingale measure under which the asset prices evolve according to the SDE's

$$dS_t^i = r_t S_t^i dt + \sum_{j=1}^m \sigma_t^{ij} S_t^i dB_t^j \tag{8}$$

so that each security has an instantaneous drift equal to r_t. Here, $(B^j : j \le m)$ are again a collection of independent standard Brownian motions under the equivalent martingale measure. Let $\tilde{E}$ denote the expectation operator under this measure. Then, the price of a derivative security at time t that pays amount H at time T equals

$$\tilde{E} \left(\exp(-\int_t^T r_t) H | \mathscr{F}_t \right).$$

This can be estimated via Monte-Carlo simulation. To keep the discussion simple, we henceforth assume that $r_t = r$ a constant.

Suppose that the payoff of the derivative security is a function of the asset prices at times $0 = t_0 < t_1 < t_2 < \ldots < t_n = T^1$. For example, consider an Asian call option on an asset i whose discounted payoff equals

$$\exp(-rT)(\frac{1}{n} \sum_{k=1}^n S_{t_k}^i - K)^+$$

---

[1] Some of the discussion here also appeared in Bolia and Juneja (2005). We refer the reader to Bolia and Juneja (2005) where Monte Carlo methods for pricing European and Bermudan options is covered in greater depth.

where K denotes the strike price of the option. In such cases, if each σ_t^{ij} is not a constant, it may be difficult to generate samples of $(S_{t_k} : k \leq n)$ via simulation such that their joint distribution matches the corresponding joint distribution of the solution of the SDE's.

However, when $\sigma_t^{ij} = \sigma^{ij}$, a constant, then this is feasible. In that setting, we may set (by applying Ito's lemma to $\log S_t^i$)

$$S_{t_{k+1}}^i = S_{t_k}^i \exp(r - \frac{1}{2} \sum_{j \leq m} (\sigma^{ij})^2)(t_{k+1} - t_k) + \sqrt{t_{k+1} - t_k} \sum_{j \leq m} \sigma^{ij} N_{k+1,j}) \tag{9}$$

where $N_{k,j}$ for each k and j are independent Gaussian random variables with mean zero and variance one. In this case, the $(S_{t_k} : k \leq n)$ so generated have the joint distribution that matches the corresponding joint distribution of the solution of the SDE's (8).

In a general case, where $\{r_t\}$ and $\{\sigma_t^{ij}\}$ are adapted processes, generating exact distributions at the selected times may not be practical and one resorts to generating the approximate time-discretized version of the SDE using Euler's scheme. This amounts to setting

$$S_{t_{k+1}}^i = S_{t_k}^i + S_{t_k}^i r_{t_k}(t_{k+1} - t_k) + S_{t_k}^i \sqrt{t_{k+1} - t_k} \sum_{j=1}^m \sigma_t^{ij} N_{k+1,j}. \tag{10}$$

This method will have a discretization error as the generated joint distributions will differ from those corresponding to the SDE. To reduce the error larger number of well spaced times may be chosen so that the accuracy is achieved at the cost of greater computational effort. We refer the reader to Asmussen and Glynn (2007) for a discussion on improved discretization methods. Recently, there has been some evolving literature on generating exact samples from one dimensional SDEs (see, e.g., Beskos and Roberts (2005), Chen (2010)).

The Monte-Carlo simulation in this setting may be implemented as follows: Generate samples $N_{k,j}$ for $k = 1, \ldots, n$ and $j = 1, \ldots, m$ (see any standard simulation text for algorithms for generating $N(0,1)$ random variable (rv), e.g., Glasserman (2004)) and use them to generate samples of $(S_{t_k}^i : i = 1, \ldots, d; k = 1, \ldots, n)$ using (9) or (10), whichever is appropriate. These are then used to arrive at the sample value of the discounted payoff from the option, call it $\hat{H}_1$. This process is repeated independently, say l times so that independent samples $(\hat{H}_i : i \leq l)$ are obtained. The point estimate of the (un-discounted) option price is given by

$$\bar{H} = \frac{1}{l} \sum_{i \leq l} \hat{H}_i.$$

Let σ_X^2 denote the variance of rv X and σ_X its standard deviation. Then from the central limit theorem, an approximate $100(1 - \alpha)\%$ confidence interval for H equals

$$\bar{H} \pm z_{1-\alpha/2} \frac{\sigma_{\hat{H}}}{\sqrt{l}},$$

where $z_{1-\alpha}$ is the solution to the equation

$$P(N(0,1) > z) = \alpha.$$

In practice $\sigma_{\hat{H}}^2$ is not known and is estimated from the generated samples by

$$\bar{\sigma}_{\hat{H}}^2 = \frac{1}{l-1} \sum_{i \leq l} (\hat{H}_i - \bar{H})^2$$

and $\bar{\sigma}_{\hat{H}}$ replaces $\sigma_{\hat{H}}$ in the approximate confidence interval. This is asymptotically valid as $\frac{\bar{\sigma}_{\hat{H}}}{\sigma_{\hat{H}}} \to 1$ a.s. as $l \to \infty$.

3 Pricing Bermudan Options

3.1 Mathematical Framework

We now construct a mathematical framework for pricing Bermudan options. Recall that American options are approximately priced by discretizing the times of exercise and estimating the price of the resulting Bermudan option.

Again, without essential loss of generality suppose that the option can be exercised only at $T+1$ times $0, 1, 2, \ldots, T$. Denote the underlying security prices by $(S_t \in \mathfrak{R}^d : t = 0, 1, \ldots, T)$. In addition, the description of the state at time t may include variables such as the value of stochastic interest rates and volatilities, and supplementary path dependent information so that the resulting process is Markov. Thus, each S_t may take values in a more general space denoted by $\mathscr{S}$. The value of the option at time t if exercised at that time, is denoted by $g_t : \mathscr{S} \to \mathfrak{R}^+$ (i.e., its exercise value or intrinsic value). Let $\mathscr{T}_k$ denote the set of stopping times taking value in $\{k, k+1, \ldots, T\}$ (recall that τ is a stopping time w.r.t. $\{S_k\}$ if $\{\tau = k\}$ is determined by observing $\{S_1, \ldots, S_k\}$). Note that each stopping time represents an exercise strategy by the owner of the option at time period k. Let

$$J_k(s) = \sup_{\tau \in \mathscr{T}_k} E[g_\tau(S_\tau)|S_k = s], \ s \in \mathscr{S}, \tag{11}$$

where, the expectation is taken under the equivalent martingale measure. Then $J_k(s)$ is the value of the option at time k given that the option is not exercised before time k. The initial state $S_0 = s_0$ is fixed and known. So, our pricing problem is to evaluate $J_0(s_0)$. This formulation is sufficiently general to include discounted payoffs through appropriate definition of the $\{S_k\}$ and $\{g_k\}$ (see Glasserman 2004, p.425), and hence these are not explicitly stated. It can be shown that (see, e.g., Duffie 1996) there exists an optimal exercise policy specified by a collection of increasing stopping times $(\tau_k^* : k \le T)$ where

$$\tau_k^* = \inf\{m \ge k : g_m(S_m) \ge J_m(S_m)\}.$$

Then,

$$J_k(s) = E[g_{\tau_k^*}(S_{\tau_k^*})|S_k = s]$$

a.s. Note that knowing the optimal policy corresponds to knowing at each state at each time whether it is optimal to exercise the option or to hold on to it.

Suppose that the pdf of S_{k+1} conditioned on $S_k = s$ evaluated at y is given by $f_k(s, y)$ under the risk-neutral measure (for example, the expression for $f_k(s, y)$ can be easily derived using (9) or (10) if either is appropriate). Let $Q_k(s)$ denote the conditional expectation

$$E[J_{k+1}(S_{k+1})|S_k = s] = \int_{\mathscr{S}} J_{k+1}(y) f_k(s, y) dy. \tag{12}$$

We refer to $Q = (Q_k(s) : s \in \mathscr{S}, k \le T - 1)$ as continuation value functions as $Q_k(s)$ denotes the value of the option at time k at state s if it is not exercised at time k.

It is not feasible to evaluate $J_k(s)$ by evaluating expectation $E[g_\tau(S_\tau)|S_k = s]$ for each stopping time τ in (11). Fortunately, it can be shown that the value functions $J = (J_k(s) : s \in \mathscr{S}, k \le T)$ satisfy the following intuitively plausible backward recursions (see, e.g., Shirayev (1978)):

$$J_T(s) = g_T(s)$$
$$J_k(s) = \max\{(g_k(s), Q_k(s)\} \tag{13}$$

Define $P_k H(\cdot)$ as

$$(P_k H)(s) \triangleq E[H(S_{k+1})|S_k = s] = \int_{\mathscr{S}} H(u) f_k(s,u) du.$$

Then, an alternative set of recursions satisfied by the continuation value function $Q = (Q_k(s) : s \in \mathscr{S}, k \leq T-1)$ is given by:

$$Q_{T-1}(s) = (P_{T-1} g_T)(s)$$
$$Q_k(s) = (P_k \max(g_{k+1}, Q_{k+1}))(s) \tag{14}$$

for $k = 0, 1, 2, ..., T-2$.

Evaluating the value functions using these recursions would require discretizing state space and then solving these recursions approximately. However, even this becomes computationally unviable due to state-space blow-up when the dimension of the underlying process is large.

In the presentation we shall review some simulation methods in the literature to approximately solve these backward recursions to estimate the value function and the associated optimal exercise policy (τ_k^* : $k \leq T$). These include the popular regression based methods introduced by Tsitsiklis and Roy (2001) and Longstaff and Schwartz (2001). Haugh and Kogan (2004) and Rogers (2002) propose additive duality for American options. Jamishidian (2007) develops multiplicative duality. Additive duality is exploted by Andersen and Broadie (2004) to develop accurate and efficient pricing algorithms (also see Broadie and Cao (2008)). Chen and Glasserman (2007) further exploit the additive and multiplicative duality to develop iterative schemes that converge to the true option value. Kolodko and Schoenmakers (2006) Kolodko and Schoemakers (2006) develop an iterative procedure for constructing stopping times that converge to the optimal stopping time. Juneja and Kalra (2009) and Ehrlichman and Henderson (2007) relate additive and multiplicative duality approach to control variate and importance sampling variance reduction techniques.

REFERENCES

Andersen, L., and M. Broadie. 2004. A primal-dual simulation algorithm for pricing multi-dimensional American options. *Management Science* 50:1222–1234.

Asmussen, S., and P. Glynn. 2007. *Stochastic simulation: Algorithms and analysis*. New York: Springer.

Beskos, A., and G. Roberts. 2005. Exact simulaton of diffusions. *Ann. Appl. Probab.* 15 (4): 2422–2444.

Black, F., and M. Scholes. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81 (3): 637–654.

Bolia, N., and S. Juneja. 2005. Monte carlo methods in pricing financial options. *Sadhana*:347–385.

Broadie, M., and M. Cao. 2008. Improved lower and upper bound algorithms for pricing American options by simulation. *Quantitative Finance* 8 (8): 845–861.

Chen, N. 2010. Exact simulation of brownian motion driven stochastic differential equations. *Preprint*.

Chen, N., and P. Glasserman. 2007. Additive and multiplicative duals for American option pricing. *Finance and Stochastics* 11:153–179.

Duffie, D. 1996. *Dynamic asset pricing theory*. Princeton NJ: Princeton university press.

Ehrlichman, S. M. T., and S. Henderson. 2007. Adaptive control variates for pricing multi-dimensional American options. *Journal of Computational Finance* 11 (1): 65–91.

Glasserman, P. 2004. *Monte carlo methods in financial engineering*. New York: Springer.

Haugh, M. B., and L. Kogan. 2004. Pricing American options: A dual approach. *Operations Research* 52 (2): 258–270.

Jamishidian, F. 2007. The duality of optimal exercise and domineering claims: a DoobMeyer decomposition approach to the snell envelope. *Stochastics: An International Journal of Probability and Stochastics Processes* 79 (1-2): 27–60.

Juneja, S., and H. Kalra. 2009. Variance reduction techniques for pricing American options. *Journal of Computational Finance* 12 (3): 79–102.

Karatzas, I., and S. Shreve. 1998. *Methods of mathematical finance*. New York: Springer.

Kolodko, A., and J. Schoenmakers. 2006. Iterative construction of the optimal Bermudan stopping time. *Finance and Stochastics* 10:27–49.

Longstaff, A. F., and E. S. Schwartz. 2001. Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies* 14:113–147.

Rogers, L. C. G. 2002. Monte Carlo valuation of American options. *Math. Finance* 12:271–286.

Shirayev, A. 1978. *Optimal stopping rules*. Springer-Verlag.

Shreve, S. E. 2004. *Stochastic calculus for finance II: Continuous-time models*. New York: Springer.

Steele, J. M. 2001. *Stochastic calculus and financial applications*. New York: Springer.

Tsitsiklis, J., and B. V. Roy. 2001. Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks* 12:694–703.

AUTHOR BIOGRAPHY

SANDEEP JUNEJA is an Associate Professor School of Technology and Computer Science at the Tata Institute of Fundamental Research. His Ph.D. was from Department of Operations Research at Stanford University (1994). His research interests include applied probability, computational finance, simulation and rare-event analysis. He serves on the editorial board on Mathematics of Operations Research. His email address is <juneja@tifr.res.in> and his web page is www.tcs.tifr.res.in/~sandeepj>.