金融数学丛书

影印版

Monte Carlo Methods in Financial Engineering

金融工程中的蒙特卡罗方法

Paul Glasserman 著

高等教育出版社
HIGHER EDUCATION PRESS

金融数学丛书

影印版

Monte Carlo Methods in Financial Engineering

金融工程中的蒙特卡罗方法

Paul Glasserman 著

高等教育出版社
HIGHER EDUCATION PRESS

图书在版编目(CIP)数据

金融工程中的蒙特卡罗方法 = Monte Carlo Methods in Financial Engineering：英文/(美)格拉瑟曼(Glasserman,P.)著. —北京：高等教育出版社，2008.6
(金融数学丛书)
ISBN 978 -7 -04 -024752 -7

Ⅰ.金… Ⅱ.格… Ⅲ.蒙特卡罗法 - 应用 - 金融学 - 英文 Ⅳ.F830

中国版本图书馆 CIP 数据核字(2008)第 066615 号

出版发行	高等教育出版社	**购书热线**	010 - 58581118
社　　址	北京市西城区德外大街 4 号	**免费咨询**	800 - 810 - 0598
邮政编码	100120	**网　　址**	http://www.hep.edu.cn
总　　机	010 - 58581000		http://www.hep.com.cn
		网上订购	http://www.landraco.com
经　　销	蓝色畅想图书发行有限公司		http://www.landraco.com.cn
印　　刷	北京铭成印刷有限公司	**畅想教育**	http://www.widedu.com
开　　本	787 × 1092　1/16	**版　　次**	2008 年 6 月第 1 版
印　　张	38.25	**印　　次**	2008 年 6 月第 1 次印刷
字　　数	800 000	**定　　价**	79.00 元

本书如有缺页、倒页、脱页等质量问题，请到所购图书销售部门联系调换.

物 料 号　24752 - 00

Preface

This is a book about Monte Carlo methods from the perspective of financial engineering. Monte Carlo simulation has become an essential tool in the pricing of derivative securities and in risk management; these applications have, in turn, stimulated research into new Monte Carlo techniques and renewed interest in some old techniques. This is also a book about financial engineering from the perspective of Monte Carlo methods. One of the best ways to develop an understanding of a model of, say, the term structure of interest rates is to implement a simulation of the model; and finding ways to improve the efficiency of a simulation motivates a deeper investigation into properties of a model.

My intended audience is a mix of graduate students in financial engineering, researchers interested in the application of Monte Carlo methods in finance, and practitioners implementing models in industry. This book has grown out of lecture notes I have used over several years at Columbia, for a semester at Princeton, and for a short course at Aarhus University. These classes have been attended by masters and doctoral students in engineering, the mathematical and physical sciences, and finance. The selection of topics has also been influenced by my experiences in developing and delivering professional training courses with Mark Broadie, often in collaboration with Leif Andersen and Phelim Boyle. The opportunity to discuss the use of Monte Carlo methods in the derivatives industry with practitioners and colleagues has helped shaped my thinking about the methods and their application.

Students and practitioners come to the area of financial engineering from diverse academic fields and with widely ranging levels of training in mathematics, statistics, finance, and computing. This presents a challenge in setting the appropriate level for discourse. The most important prerequisite for reading this book is familiarity with the mathematical tools routinely used to specify and analyze continuous-time models in finance. Prior exposure to the basic principles of option pricing is useful but less essential. The tools of mathematical finance include Itô calculus, stochastic differential equations, and martingales. Perhaps the most advanced idea used in many places in this book is the concept of a change of measure. This idea is so central both to derivatives pricing and to Monte Carlo methods that there is simply no avoiding it. The prerequisites to understanding the statement of the Girsanov theorem should suffice for reading this book.

Whereas the language of mathematical finance is essential to our topic, its technical subtleties are less so for purposes of computational work. My use of mathematical tools is often informal: I may assume that a local martingale is a martingale or that a stochastic differential equation has a solution, for example, without calling attention to these assumptions. Where convenient, I take derivatives without first assuming differentiability and I take expectations without verifying integrability. My intent is to focus on the issues most important to Monte Carlo methods and to avoid diverting the discussion to spell out technical conditions. Where these conditions are not evident and where they are essential to understanding the scope of a technique, I discuss them explicitly. In addition, an appendix gives precise statements of the most important tools from stochastic calculus.

This book divides roughly into three parts. The first part, Chapters 1–3, develops fundamentals of Monte Carlo methods. Chapter 1 summarizes the theoretical foundations of derivatives pricing and Monte Carlo. It explains the principles by which a pricing problem can be formulated as an integration problem to which Monte Carlo is then applicable. Chapter 2 discusses random number generation and methods for sampling from nonuniform distributions, tools fundamental to every application of Monte Carlo. Chapter 3 provides an overview of some of the most important models used in financial engineering and discusses their implementation by simulation. I have included more discussion of the models in Chapter 3 and the financial underpinnings in Chapter 1 than is strictly necessary to run a simulation. Students often come to a course in Monte Carlo with limited exposure to this material, and the implementation of a simulation becomes more meaningful if accompanied by an understanding of a model and its context. Moreover, it is precisely in model details that many of the most interesting simulation issues arise.

If the first three chapters deal with running a simulation, the next three deal with ways of running it better. Chapter 4 presents methods for increasing precision by reducing the variance of Monte Carlo estimates. Chapter 5 discusses the application of deterministic *quasi*-Monte Carlo methods for numerical integration. Chapter 6 addresses the problem of discretization error that results from simulating discrete-time approximations to continuous-time models.

The last three chapters address topics specific to the application of Monte Carlo methods in finance. Chapter 7 covers methods for estimating price sensitivities or "Greeks." Chapter 8 deals with the pricing of American options, which entails solving an optimal stopping problem within a simulation. Chapter 9 is an introduction to the use of Monte Carlo methods in risk management. It discusses the measurement of market risk and credit risk in financial portfolios. The models and methods of this final chapter are rather different from those in the other chapters, which deal primarily with the pricing of derivative securities.

Several people have influenced this book in various ways and it is my pleasure to express my thanks to them here. I owe a particular debt to my frequent collaborators and co-authors Mark Broadie, Phil Heidelberger, and Perwez Shahabuddin. Working with them has influenced my thinking as well as the book's contents. With Mark Broadie I have had several occasions to collaborate on teaching as well as research, and I have benefited from our many discussions on most of the topics in this book. Mark, Phil Heidelberger, Steve Kou, Pierre L'Ecuyer, Barry Nelson, Art Owen, Philip Protter, and Jeremy Staum each commented on one or more draft chapters; I thank them for their comments and apologize for the many good suggestions I was unable to incorporate fully. I have also benefited from working with current and former Columbia students Jingyi Li, Nicolas Merener, Jeremy Staum, Hui Wang, Bin Yu, and Xiaoliang Zhao on some of the topics in this book. Several classes of students helped uncover errors in the lecture notes from which this book evolved.

Paul Glasserman
New York, 2003

Contents

1 Foundations

This chapter's two parts develop key ideas from two fields, the intersection of which is the topic of this book. Section 1.1 develops principles underlying the use and analysis of Monte Carlo methods. It begins with a general description and simple examples of Monte Carlo, and then develops a framework for measuring the efficiency of Monte Carlo estimators. Section 1.2 reviews concepts from the theory of derivatives pricing, including pricing by replication, the absence of arbitrage, risk-neutral probabilities, and market completeness. The most important idea for our purposes is the representation of derivative prices as expectations, because this representation underlies the application of Monte Carlo.

1.1 Principles of Monte Carlo

1.1.1 Introduction

Monte Carlo methods are based on the analogy between probability and volume. The mathematics of measure formalizes the intuitive notion of probability, associating an event with a set of outcomes and defining the probability of the event to be its volume or measure relative to that of a universe of possible outcomes. Monte Carlo uses this identity in reverse, calculating the volume of a set by interpreting the volume as a probability. In the simplest case, this means sampling randomly from a universe of possible outcomes and taking the fraction of random draws that fall in a given set as an estimate of the set's volume. The law of large numbers ensures that this estimate converges to the correct value as the number of draws increases. The central limit theorem provides information about the likely magnitude of the error in the estimate after a finite number of draws.

A small step takes us from volumes to integrals. Consider, for example, the problem of estimating the integral of a function f over the unit interval. We may represent the integral

$$\alpha = \int_0^1 f(x)\, dx$$

as an expectation $\mathsf{E}[f(U)]$, with U uniformly distributed between 0 and 1. Suppose we have a mechanism for drawing points $U_1, U_2, \ldots$ independently and uniformly from $[0,1]$. Evaluating the function f at n of these random points and averaging the results produces the Monte Carlo estimate

$$\hat{\alpha}_n = \frac{1}{n} \sum_{i=1}^{n} f(U_i).$$

If f is indeed integrable over $[0,1]$ then, by the strong law of large numbers,

$$\hat{\alpha}_n \to \alpha \quad \text{with probability 1 as } n \to \infty.$$

If f is in fact square integrable and we set

$$\sigma_f^2 = \int_0^1 (f(x) - \alpha)^2\, dx,$$

then the error $\hat{\alpha}_n - \alpha$ in the Monte Carlo estimate is approximately normally distributed with mean 0 and standard deviation $\sigma_f/\sqrt{n}$, the quality of this approximation improving with increasing n. The parameter σ_f would typically be unknown in a setting in which α is unknown, but it can be estimated using the sample standard deviation

$$s_f = \sqrt{\frac{1}{n-1} \sum_{i=1}^{n} (f(U_i) - \hat{\alpha}_n)^2}.$$

Thus, from the function values $f(U_1), \ldots, f(U_n)$ we obtain not only an estimate of the integral α but also a measure of the error in this estimate.

The form of the standard error $\sigma_f/\sqrt{n}$ is a central feature of the Monte Carlo method. Cutting this error in half requires increasing the number of points by a factor of four; adding one decimal place of precision requires 100 times as many points. These are tangible expressions of the square-root convergence rate implied by the $\sqrt{n}$ in the denominator of the standard error. In contrast, the error in the simple trapezoidal rule

$$\alpha \approx \frac{f(0) + f(1)}{2n} + \frac{1}{n} \sum_{i=1}^{n-1} f(i/n)$$

is $O(n^{-2})$, at least for twice continuously differentiable f. Monte Carlo is generally not a competitive method for calculating one-dimensional integrals.

The value of Monte Carlo as a computational tool lies in the fact that its $O(n^{-1/2})$ convergence rate is not restricted to integrals over the unit interval.

Indeed, the steps outlined above extend to estimating an integral over $[0,1]^d$ (and even $\Re^d$) for all dimensions d. Of course, when we change dimensions we change f and when we change f we change σ_f, but the standard error will still have the form $\sigma_f/\sqrt{n}$ for a Monte Carlo estimate based on n draws from the domain $[0,1]^d$. In particular, the $O(n^{-1/2})$ convergence rate holds for all d. In contrast, the error in a product trapezoidal rule in d dimensions is $O(n^{-2/d})$ for twice continuously differentiable integrands; this degradation in convergence rate with increasing dimension is characteristic of all deterministic integration methods. Thus, Monte Carlo methods are attractive in evaluating integrals in high dimensions.

What does this have to do with financial engineering? A fundamental implication of asset pricing theory is that under certain circumstances (reviewed in Section 1.2.1), the price of a derivative security can be usefully represented as an expected value. Valuing derivatives thus reduces to computing expectations. In many cases, if we were to write the relevant expectation as an integral, we would find that its dimension is large or even infinite. This is precisely the sort of setting in which Monte Carlo methods become attractive.

Valuing a derivative security by Monte Carlo typically involves simulating paths of stochastic processes used to describe the evolution of underlying asset prices, interest rates, model parameters, and other factors relevant to the security in question. Rather than simply drawing points randomly from $[0,1]$ or $[0,1]^d$, we seek to sample from a space of paths. Depending on how the problem and model are formulated, the dimension of the relevant space may be large or even infinite. The dimension will ordinarily be at least as large as the number of time steps in the simulation, and this could easily be large enough to make the square-root convergence rate for Monte Carlo competitive with alternative methods.

For the most part, there is nothing we can do to overcome the rather slow rate of convergence characteristic of Monte Carlo. (The quasi-Monte Carlo methods discussed in Chapter 5 are an exception — under appropriate conditions they provide a faster convergence rate.) We can, however, look for superior sampling methods that reduce the implicit constant in the convergence rate. Much of this book is devoted to examples and general principles for doing this.

The rest of this section further develops some essential ideas underlying Monte Carlo methods and their application to financial engineering. Section 1.1.2 illustrates the use of Monte Carlo with two simple types of option contracts. Section 1.1.3 develops a framework for evaluating the efficiency of simulation estimators.

1.1.2 First Examples

In discussing general principles of Monte Carlo, it is useful to have some simple specific examples to which to refer. As a first illustration of a Monte Carlo method, we consider the calculation of the expected present value of the payoff of a call option on a stock. We do not yet refer to this as the option *price*; the connection between a price and an expected discounted payoff is developed in Section 1.2.1.

Let $S(t)$ denote the price of the stock at time t. Consider a call option granting the holder the right to buy the stock at a fixed price K at a fixed time T in the future; the current time is $t = 0$. If at time T the stock price $S(T)$ exceeds the strike price K, the holder exercises the option for a profit of $S(T) - K$; if, on the other hand, $S(T) \leq K$, the option expires worthless. (This is a *European* option, meaning that it can be exercised only at the fixed date T; an *American* option allows the holder to choose the time of exercise.) The payoff to the option holder at time T is thus

$$(S(T) - K)^+ = \max\{0, S(T) - K\}.$$

To get the present value of this payoff we multiply by a discount factor e^{-rT}, with r a continuously compounded interest rate. We denote the expected present value by $\mathsf{E}[e^{-rT}(S(T) - K)^+]$.

For this expectation to be meaningful, we need to specify the distribution of the random variable $S(T)$, the terminal stock price. In fact, rather than simply specifying the distribution at a fixed time, we introduce a model for the dynamics of the stock price. The Black-Scholes model describes the evolution of the stock price through the stochastic differential equation (SDE)

$$\frac{dS(t)}{S(t)} = r\,dt + \sigma\,dW(t), \tag{1.1}$$

with W a standard Brownian motion. (For a brief review of stochastic calculus, see Appendix B.) This equation may be interpreted as modeling the percentage changes dS/S in the stock price as the increments of a Brownian motion. The parameter σ is the volatility of the stock price and the coefficient on dt in (1.1) is the mean rate of return. In taking the rate of return to be the same as the interest rate r, we are implicitly describing the *risk-neutral* dynamics of the stock price, an idea reviewed in Section 1.2.1.

The solution of the stochastic differential equation (1.1) is

$$S(T) = S(0)\exp\left([r - \tfrac{1}{2}\sigma^2]T + \sigma W(T)\right). \tag{1.2}$$

As $S(0)$ is the current price of the stock, we may assume it is known. The random variable $W(T)$ is normally distributed with mean 0 and variance T; this is also the distribution of $\sqrt{T}Z$ if Z is a standard normal random variable (mean 0, variance 1). We may therefore represent the terminal stock price as

$$S(T) = S(0)\exp\left([r - \tfrac{1}{2}\sigma^2]T + \sigma\sqrt{T}Z\right). \tag{1.3}$$

The logarithm of the stock price is thus normally distributed, and the stock price itself has a lognormal distribution.

The expectation $\mathsf{E}[e^{-rT}(S(T)-K)^+]$ is an integral with respect to the lognormal density of $S(T)$. This integral can be evaluated in terms of the standard normal cumulative distribution function Φ as $\mathrm{BS}(S(0),\sigma,T,r,K)$ with

$$\mathrm{BS}(S,\sigma,T,r,K) = S\Phi\left(\frac{\log(S/K)+(r+\frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}\right) - e^{-rT}K\Phi\left(\frac{\log(S/K)+(r-\frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}\right). \quad (1.4)$$

This is the Black-Scholes [50] formula for a call option.

In light of the availability of this formula, there is no need to use Monte Carlo to compute $\mathsf{E}[e^{-rT}(S(T)-K)^+]$. Moreover, we noted earlier that Monte Carlo is not a competitive method for computing one-dimensional integrals. Nevertheless, we now use this example to illustrate the key steps in Monte Carlo. From (1.3) we see that to draw samples of the terminal stock price $S(T)$ it suffices to have a mechanism for drawing samples from the standard normal distribution. Methods for doing this are discussed in Section 2.3; for now we simply assume the ability to produce a sequence $Z_1, Z_2, \ldots$ of independent standard normal random variables. Given a mechanism for generating the Z_i, we can estimate $\mathsf{E}[e^{-rT}(S(T)-K)^+]$ using the following algorithm:

for $i = 1, \ldots, n$
  generate Z_i
  set $S_i(T) = S(0)\exp\left([r-\frac{1}{2}\sigma^2]T + \sigma\sqrt{T}Z_i\right)$
  set $C_i = e^{-rT}(S(T)-K)^+$
set $\hat{C}_n = (C_1 + \cdots + C_n)/n$

For any $n \geq 1$, the estimator $\hat{C}_n$ is *unbiased,* in the sense that its expectation is the target quantity:

$$\mathsf{E}[\hat{C}_n] = C \equiv \mathsf{E}[e^{-rT}(S(T)-K)^+].$$

The estimator is *strongly consistent*, meaning that as $n \to \infty$,

$$\hat{C}_n \to C \quad \text{with probability 1.}$$

For finite but at least moderately large n, we can supplement the point estimate $\hat{C}_n$ with a confidence interval. Let

$$s_C = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}(C_i - \hat{C}_n)^2} \quad (1.5)$$

denote the sample standard deviation of $C_1, \ldots, C_n$ and let z_δ denote the $1-\delta$ quantile of the standard normal distribution (i.e., $\Phi(z_\delta) = 1-\delta$). Then

$$\hat{C}_n \pm z_{\delta/2}\frac{s_C}{\sqrt{n}} \quad (1.6)$$

is an asymptotically (as $n \to \infty$) valid $1-\delta$ confidence interval for C. (For a 95% confidence interval, $\delta = .05$ and $z_{\delta/2} \approx 1.96$.) Alternatively, because the standard deviation is estimated rather than known, we may prefer to replace $z_{\delta/2}$ with the corresponding quantile from the t distribution with $n-1$ degrees of freedom, which results in a slightly wider interval. In either case, the probability that the interval covers C approaches $1-\delta$ as $n \to \infty$. (These ideas are reviewed in Appendix A.)

The problem of estimating $\mathsf{E}[e^{-rT}(S(T)-K)^+]$ by Monte Carlo is simple enough to be illustrated in a spreadsheet. Commercial spreadsheet software typically includes a method for sampling from the normal distribution and the mathematical functions needed to transform normal samples to terminal stock prices and then to discounted option payoffs. Figure 1.1 gives a schematic illustration. The Z_i are samples from the normal distribution; the comments in the spreadsheet illustrate the formulas used to transform these to arrive at the estimate $\hat{C}_n$. The spreadsheet layout in Figure 1.1 makes the method transparent but has the drawback that it requires storing all n replication in n rows of cells. It is usually possible to use additional spreadsheet commands to recalculate cell values n times without storing intermediate values.

Replication	Normals	Stock Price	Option Payoff
1	Z_1	S_1	C_1
2	Z_2	S_2	C_2
3	Z_3	S_3	C_3
4	Z_4	S_4	
5	Z_5	S_5	
6	Z_6	S_6	C_6
7	Z_7	S_7	C_7
8	Z_8	S_8	C_8
9	Z_9	S_9	C_9
10	Z_10	S_10	C_10
11	Z_11	S_11	C_11
⋮	⋮	⋮	⋮
n	Z_n	S_n	C_n

$\hat{C}_n$ = AVERAGE(C_1,....,C_n)

s_C = STDEV(C_1,...,C_n)

S_1=S(0)*exp((r-0.5*σ^2)*T+σ*sqrt(T)*Z_1)

C_8=exp(-rT)*max(0,S_8-K)

Fig. 1.1. A spreadsheet for estimating the expected present value of the payoff of a call option.

This simple example illustrates a general feature of Monte Carlo methods for valuing derivatives, which is that the simulation is built up in layers: each of the transformations

$$Z_i \longrightarrow S_i(T) \longrightarrow C_i$$

exemplifies a typical layer. The first transformation constructs a path of underlying assets from random variables with simpler distributions and the second calculates a discounted payoff from each path. In fact, we often have additional