

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 $E[f(U)]$, with U uniformly distributed between 0 and 1. Suppose we have a mechanism for drawing points U_1, U_2, \dots 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,

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

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 σ_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 $f(U_1), \dots, f(U_n)$ 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 σ_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 \mathfrak{R}^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 σ_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