Contents

Preface

The aim of this book is to provide a wide class of readers with a low- to intermediate-level treatment of Monte Carlo methods for applications in finance and economics. The target audience consists of students and junior practitioners with a quantitative background, and it includes not only students in economics and finance, but also in mathematics, statistics, and engineering. In fact, this is the kind of audience I typically deal with in my courses. Not all of these readers have a strong background in either statistics, financial economics, or econometrics, which is why I have also included some basic material on stochastic modeling in the early chapters, which is typically skipped in higher level books. Clearly, this is not meant as a substitute for a proper treatment, which can be found in the references listed at the end of each chapter. Some level of mathematical maturity is assumed, but the prerequisites are rather low and boil down to the essentials of probability and statistics, as well as some basic programming skills.[1] Advanced readers may skip the introductory chapters on modeling and estimation, which are also included as a reminder that no Monte Carlo method, however sophisticated, will yield useful results if the input model is flawed. Indeed, the power and flexibility of such methods may lure us into a false sense of security, making us forget some of their inherent limitations.

Option pricing is certainly a relevant application domain for the techniques we discuss in the book, but this is not meant to be a book on financial engineering. I have also included a significant amount of material on optimization in its many guises, as well as a chapter related to computational Bayesian statistics. I have favored a wide scope over a deeply technical treatment, for which there are already some really excellent and more demanding books. Many of them, however, do not quite help the reader to really "feel" what she is learning, as no ready-to-use code is offered. In order to allow anyone to run the code, play with it, and hopefully come up with some variations on the theme, I have chosen to develop code in R. Readers familiar with my previous book written in MATLAB might wonder whether I have changed my mind. I did not: I never use R in research or consulting, but I use it a lot for teaching. When I started writing the book, I was less than impressed by the lack of an adequate development environment, and some design choices of the language itself left me a bit puzzled. As an example, the `*` operator in MATLAB multiplies matrices row by column; whenever you want to work elementwise, you use the `.` operator, which has a clear and uniform meaning when applied to other operators. On the

[1] In case of need, the mathematical prerequisites are covered in my other book: *Quantitative Methods: An Introduction for Business Management*. Wiley, 2011.

contrary, the operator `*` works elementwise in R, and row-by-column matrix product is accomplished by the somewhat baroque operator `%*%`. Furthermore, having to desperately google every time you have to understand a command, because documentation is a bit poor and you have to make your way in a mess of packages, may be quite frustrating at times. I have also found that some optimization functions are less accurate and less clever in dealing with limit cases than the corresponding MATLAB functions. Having said that, while working on the book, I have started to appreciate R much more. Also my teaching experience with R has certainly been fun and rewarding. A free tool with such a potential as R is certainly most welcome, and R developers must be praised for offering all of this. Hopefully, the reader will find R code useful as a starting point for further experimentation. I did *not* assemble R code into a package, as this would be extremely misleading: I had no plan to develop an integrated and reliable set of functions. I just use R code to illustrate ideas in concrete terms and to encourage active learning. When appropriate, I have pointed out some programming practices that may help in reducing the computational burden, but as a general rule I have tried to emphasize clarity over efficiency. I have also avoided writing an introduction to R programming, as there are many freely available tutorials (and a few good books[2]). A reader with some programming experience in any language should be able to make her way through the code, which has been commented on when necessary. My assumption is that a reader, when stumbling upon an unknown function, will take advantage of the online help and the example I provide in order to understand its use and potentiality. Typically, R library functions are equipped with optional parameters that can be put to good use, but for the sake of conciseness I have refrained from a full description of function inputs.

Book structure

The book is organized in five parts.

1. Part I, Overview and Motivation, consists of two chapters. Chapter 1 provides an introduction to Monte Carlo methods and applications. The different classes of dynamic models that are encountered in simulation are outlined, and due emphasis is placed on pitfalls and limitations of Monte Carlo methods. Chapter 2 deals with numerical integration methods. Numerical integration is quite relevant, as it provided most of the historical motivation to develop Monte Carlo methods; furthermore, there are cases in which one is much better off using good quadrature formulas than throwing random numbers around. Finally, framing Monte Carlo methods within numerical integration provides the necessary discipline to understand and properly use low-discrepancy sequences, sometimes referred to as quasi–Monte Carlo methods.

[2]Unfortunately, I have also run into very bad books using R; hopefully, this one will not contribute to the list.

2. Part II, Input Analysis: Modeling and Estimation, is specifically aimed at students and newcomers, as it includes two introductory chapters dealing with stochastic model building (Chapter 3) and model fitting (Chapter 4). Essentially, in this part of the book we are concerned with the modeling of inputs of a Monte Carlo simulation. Many advanced books on Monte Carlo methods for finance skip and take for granted these concepts. I have preferred to offer a limited treatment for the sake of unfamiliar readers, such as students in engineering or practitioners without an econometrics background. Needless to say, space does not allow for a deep treatment, but I believe that it is important to build at least a framework for further study. In order to make this part useful to intermediate readers, too, I have taken each topic as an excuse for a further illustration of R functionalities. Furthermore, some more advanced sections may be useful to students in economics and finance as well, such as those on stochastic calculus, copulas, and Bayesian statistics.
3. Part III, Sampling and Path Generation, is more technical and consists of two chapters. In Chapter 5 we deal with pseudorandom number and variate generation. While it is certainly true that in common practice one takes advantage of reliable generators provided by software tools like R, and there is no need for an overly deep treatment, some basic knowledge is needed in order to select generators and to manage simulation runs properly. We also outline scenario generation using copulas. In Chapter 6 we deal with sample path generation for continuous-time models based on stochastic differential equations. This is an essential tool for any financial engineer and is at the heart of many derivative pricing methods. It is important to point out that this is also relevant for risk managers, insurers, and some economists as well.
4. Part IV, Output Analysis and Efficiency Improvement, looks at the final step of the simulation process. Monte Carlo methods are extremely powerful and flexible; yet, their output may not be quite reliable, and an unreasonable computational effort may be called for, unless suitable countermeasures are taken. Chapter 7 covers very simple, and possibly overlooked, concepts related to confidence intervals. Counterexamples are used to point out the danger of forgetting some underlying assumptions. Chapter 8 deals with variance reduction strategies that are essential in many financial engineering and risk management applications; indeed, the techniques illustrated here are applied in later chapters, too. Chapter 9 deals with low-discrepancy sequences, which are sometimes gathered under the quasi–Monte Carlo nickname. Actually, there is nothing stochastic in low-discrepancy sequences, which should be regarded as deterministic numerical integration strategies. For a certain range of problem dimensionality, they are a good alternative to pseudorandom sampling.
5. Part IV, Miscellaneous Applications, includes five more or less interrelated chapters dealing with:

- The interplay between optimization and Monte Carlo methods, including stochastic methods for deterministic global optimization, scenario generation for stochastic programming with recourse, and stochastic dynamic programming (Chapter 10)
- Option pricing, with an emphasis on variance reduction methods (Chapter 11)
- Obtaining sensitivity of performance measures with Monte Carlo simulation (Chapter 12)
- Risk measurement and management, with an emphasis on value-at-risk and related risk measures for financial risk management (Chapter 13)
- Markov chain Monte Carlo (MCMC) methods, which are relevant for different applications, most notably computational Bayesian statistics (Chapter 14)

There are some logical precedences among these final chapters, but they need not be read in a strict sequential order. Chapter 14 is independent of the others, and the only link is represented by the possibility of regarding simulated annealing, a stochastic approach for both global and combinatorial optimization, as an MCMC method. Stochastic dynamic programming, dealt with in Chapter 10, is needed to understand American-style option pricing in Chapter 11. Measuring option price sensitivity is used as a motivating example in Chapter 12, but the methods outlined there have a much more general applicability, as they can also be used within optimization procedures. Finally, there is certainly a natural link between option pricing and the risk management examples discussed in Chapter 13.

Supplements and R code

The code has been organized in two files per chapter. The first one contains all of the function definitions and the reference to required packages, if any; this file should be sourced before running the scripts, which are included as chunks of code in a second file. An archive including all of the R code will be posted on a webpage. My current URL is

- `http://staff.polito.it/paolo.brandimarte/`

A hopefully short list of errata will be posted there as well. One of the many corollaries of Murphy's law states that my URL is going to change shortly after publication of the book. An up-to-date link will be maintained on the Wiley webpage:

- `http://www.wiley.com/`

For comments, suggestions, and criticisms, all of which are quite welcome, my e-mail address is

- paolo.brandimarte@polito.it

PAOLO BRANDIMARTE

Turin, February 2014