Numerical Methods in Finance and Economics

A MATLAB-Based Introduction

Second Edition

Paolo Brandimarte

A Wiley-Interscience Publication
JOHN WILEY & SONS, INC.
New York / Chichester / Weinheim / Brisbane / Singapore / Toronto

This book is dedicated to Commander Straker, Lieutenant Ellis, and all SHADO operatives. Thirty-five years ago they introduced me to the art of using both computers and gut feelings to make decisions.

Contents

Part II Numerical Methods

Preface to the Second Edition

After the publication of the first edition of the book, about five years ago, I have received a fair number of messages from readers, both students and practitioners, around the world. The recurring keyword, and the most important thing to me, was *useful*. The book had, and has, no ambition of being a very advanced research book. The basic motivation behind this second edition is the same behind the first one: providing the newcomer with an easy, but solid, entry point to computational finance, without too much sophisticated mathematics and avoiding the burden of difficult C++ code, also covering relatively non-standard optimization topics such as stochastic and integer programming. See also the excerpt from the preface to the first edition. However, there are a few new things here:

- a slightly revised title;
- completely revised organization of chapters;
- significantly increased number of pages.

The title mentions both Finance *and* Economics, rather than just Finance. To avoid any misunderstanding, it should be made quite clear that this is essentially a book for students and practitioners working in Finance. Nevertheless, it *can* be useful to Ph.D. students in Economics as well, as a complement to more specific and advanced textbooks. In the last four years, I have been giving a course on numerical methods within a Ph.D. program in Economics, and I typically use other available excellent textbooks covering advanced algorithms[1] or offering well-thought MATLAB toolboxes[2] which can be used to solve a wide array of problems in Economics. From the point of view of my students in such a course, the present book has many deficiencies: For instance, it does not cover ordinary differential equations and it does not deal with computing equilibria or rational expectations models; furthermore, practically all of the examples deal with option pricing or portfolio management. Nevertheless, given my experience, I believe that they can benefit from a more detailed and elementary treatment of the basics, supported by simple examples. Moreover, I believe that students in Economics should also get

[1] K.L. Judd, *Numerical Methods in Economics*, MIT Press, 1998.
[2] M.J. Miranda and P.L. Fackler, *Applied Computational Economics and Finance*, MIT Press, 2002.

at least acquainted with topics from Operations Research, such as stochastic programming and integer programming. Hence, the "*and Economics*" part of the title suggests potential use of the book as a complement, and by no means as a substitute.

The book has been reorganized in order to ease its use within standard courses on numerical methods for financial engineering. In the first edition, optimization applications were dealt with extensively, in chapters preceding those related to option pricing. This was a result of my personal background, which is mainly Computer Science and Operations Research, but it did not fit very well with the common use of a book on computational finance. In the present edition, advanced optimization applications are left to the last chapters, so they do not get into the way of most financial engineering students. The book consists of twelve chapters and three appendices.

- Chapter 1 provides the reader with motivations for the use of numerical methods, and for the use of MATLAB as well.
- Chapter 2 is an overview of financial theory. It is aimed at students in Engineering, Mathematics, or Operations Research, who may be interested in the book, but have little or no financial background.
- Chapter 3 is devoted to the basics of classical numerical methods. In some sense, this is complementary to chapter 2 and it is aimed at people with a background in Economics, who typically are not exposed to numerical analysis. To keep the book to a reasonable size, a few classical topics were omitted because of their limited role in the following chapters. In particular, I do not cover computation of eigenvalues and eigenvectors and ordinary differential equations.
- Chapter 4 is devoted to numerical integration, both by quadrature formulas and Monte Carlo methods. In the first edition, quadrature formulas were dealt with in the chapter on numerical analysis, and Monte Carlo was the subject of a separate chapter. I preferred giving a unified treatment of these two approaches, as this helps understanding their respective strengths and weaknesses, both for option pricing and scenario generation in stochastic optimization. Regarding Monte Carlo as a tool for integration rather than simulation is also helpful to properly frame the application of low-discrepancy sequences (which is also known under the more appealing name of quasi-Monte Carlo simulation). There is some new material on Gaussian quadrature, an extensive treatment of variance reduction methods, and some application to vanilla options to illustrate simple but concrete applications immediately, leaving more complex cases to chapter 8.
- Chapter 5 deals with basic finite difference schemes for partial differential equations. The main theme is solving the heat equation, which

is the prototype example of the class of parabolic equations, to which Black–Scholes equation belongs. In this simplified framework we may understand the difference between explicit and implicit methods, as well as the issues related to convergence and numerical stability. With respect to the first edition, I have added an outline of the Alternating Direction Implicit method to solve the two-dimensional heat equation, which is useful background for pricing multidimensional options.

- Chapter 6 deals with finite-dimensional (static) optimization. This chapter can be safely skipped by students interested in the option pricing applications described in chapters 7, 8, and 9. However, it may be useful to students in Economics. It is also necessary background for the relatively advanced optimization models and methods which are covered in chapters 10, 11, and 12.

- Chapter 7 is a new chapter which is devoted to binomial and trinomial lattices, which were not treated extensively in the first edition. The main issues here are proper implementation and memory management.

- Chapter 8 is naturally linked to chapter 4 and deals with more advanced applications of Monte Carlo and low-discrepancy sequences to exotic options, such as barrier and Asian options. We also deal briefly with the estimation of option sensitivities (the Greeks) by Monte Carlo methods. Emphasis is on European-style options; pricing American options by Monte Carlo methods is a more advanced topic which must be analyzed within an appropriate framework, which is done in chapter 10.

- Chapter 9 applies the background of chapter 5 to option pricing by finite difference methods.

- Chapter 10 deals with numerical dynamic programming. The main reason for including this chapter is pricing American options by Monte Carlo simulation, which was not covered in the first edition but is gaining more and more importance. I have decided to deal with this topic within an appropriate framework, which is dynamic stochastic optimization. In this chapter we just cover the essentials, which means discrete-time and finite-horizon dynamic programs. Nevertheless, we try to offer a reasonably firm understanding of these topics, both because of their importance in Economics and because understanding dynamic programming is helpful in understanding stochastic programming with recourse, which is the subject of the next chapter.

- Chapter 11 deals with linear stochastic programming models with recourse. This is becoming a standard topic for people in Operations Research, whereas people in Economics are much more familiar with dynamic programming. There are good reasons for this state of the matter, but from a methodological point of view I believe that it is very important to compare this approach with dynamic programming; from a practical point of view, stochastic programming has an interesting potential both for dynamic portfolio management and for option hedging in incomplete markets.

- Chapter 12 also deals with the relatively exotic topic of non-convex optimization. The main aim here is introducing mixed-integer programming, which can be used for portfolio management when practically relevant constraints call for the introduction of logical decision variables. We also deal, very shortly, with global optimization, i.e., continuous non-convex optimization, which is important when we leave the comfortable domain of easy optimization problems (i.e., minimizing convex cost functions or maximizing concave utility functions). We also outline heuristic principles such as local search and genetic algorithms. They are useful to integrate simulation and optimization and are often used in computational economics.

- Finally, we offer three appendices on MATLAB, probability and statistics, and AMPL. The appendix on MATLAB should be used by the unfamiliar reader to get herself going, but the best way to learn MATLAB is by trying and using the online help when needed. The appendix on probability and statistics is just a refresher which is offered for the sake of convenience. The third appendix on AMPL is new, and it reflects the increased role of algebraic languages to describe complex optimization models. AMPL is a modeling system offering access to a wide array of optimization solvers. The choice of AMPL is just based on personal taste (and the fact that a demo version is available on the web). In fact, GAMS is probably much more common for economic applications, but the concepts are actually the same. This appendix is only required for chapters 11 and 12.

Finally, there are many more pages in this second edition: more than 600 pages, whereas the first edition had about 400. Actually, I had a choice: either including many more topics, such as interest-rate derivatives, or offering a more extended and improved coverage of what was already included in the first edition. While there is indeed some new material, I preferred the second option. Actually, the original plan of the book included two more chapters on interest-rate derivatives, as many readers complained about this lack in the first edition. While writing this increasingly long second edition, I switched to plan B, and interest-rate derivatives are just outlined in the second chapter to point out their peculiarities with respect to stock options. In fact, when planning this new edition, many reviewers warned that there was little hope to cover interest-rate derivatives thoroughly in a limited amount of pages. They require a deeper understanding of risk-neutral pricing, interest rate modeling, and market practice. I do believe that the many readers interested in this topic can use this book to build a solid basis in numerical methods, which is helpful to tackle the more advanced texts on interest-rate derivatives.

Interest-rate derivatives are not the only significant omission. I could also mention implied lattices and financial econometrics. But since there are excellent books covering those topics and I see this one just as an entry point or a complement, I felt that it was more important to give a concrete understanding of the basics, including some less familiar topics. This is also why I prefer using MATLAB, rather than C++ or Visual Basic. While there is no doubt that C++ has many merits for developing professional code, both in terms of efficiency and object orientation, it is way too complex for newcomers. Furthermore, the heavy burden it places on the reader tends to overshadow the underlying concepts, which are the real subject of the book. Visual Basic would be a very convenient choice: It is widespread, and it does not require yet another license, since it is included in software tools that almost everyone has available. Such a choice would probably increase my royalties as well. Nevertheless, MATLAB code can exploit a wide and reliable library of numerical functions and it is much more compact. To the very least, it can be considered a good language for fast prototyping. These considerations, as well as the introduction of new MATLAB toolboxes aimed at financial applications, are the reasons why I am sticking to my original choice. The increasing number of books using MATLAB seems to confirm that it was a good one.

Acknowledgments. I have received much appreciated feedback and encouragement from readers of the first edition of the book. Some pointed out typos, errors, and inaccuracies. Offering apologies for possible omissions, I would like to thank I-Jung Hsiao, Sandra Hui, Byunggyoo Kim, Scott Lyden, Alexander Reisz, Ayumu Satoh, and Aldo Tagliani.

Supplements. As with the first edition, I plan to keep a web page containing the (hopefully short) list of errata and the (hopefully long) list of supplements, as well as the MATLAB code described in the book. My current URL is:

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

For comments, suggestions, and criticisms, my e-mail address is

- `paolo.brandimarte@polito.it`

One of the many corollaries of Murphy's law says that my URL is going to change shortly after publication of the book. An up-to-date link will be maintained both on Wiley Web page:

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

and on The MathWorks' web page:

- `http://www.mathworks.com/support/books/`

PAOLO BRANDIMARTE
Turin, March 2006

From the Preface to the First Edition

Crossroads are hardly, if ever, points of arrival; but neither are they points of departure. In some sense, crossroads may be disappointing, indeed. You are tired of driving, you are not at home yet, and by Murphy's law there is a far-from-negligible probability of taking the wrong turn. In this book, different paths cross, involving finance, numerical analysis, optimization theory, probability theory, Monte Carlo simulation, and partial differential equations. It is not a point of departure, because although the prerequisites are fairly low, some level of mathematical maturity on the part of the reader is assumed. It is not a point of arrival, as many relevant issues have been omitted, such as hedging exotic options and interest-rate derivatives.

The book stems from lectures I give in a Master's course on numerical methods for finance, aimed at graduate students in Economics, and in an optimization course aimed at students in Industrial Engineering. Hence, this is not a research monograph; it is a textbook for students. On the one hand, students in Economics usually have little background in numerical methods and lack the ability to translate algorithmic concepts into a working program; on the other hand, students in Engineering do not see the potential application of quantitative methods to finance clearly.

Although there is an increasing literature on high-level mathematics applied to financial engineering, and a few books illustrating how cookbook recipes may be applied to a wide variety of problems through use of a spreadsheet, I believe there is some need for an intermediate-level book, both interesting to practitioners and suitable for self-study. I believe that students should:

- Acquire *reasonably* strong foundations in order to appreciate the issues behind the application of numerical methods
- Be able to translate and check ideas quickly in a computational environment
- Gain confidence in their ability to apply methods, even by carrying out the apparently pointless task of using relatively sophisticated tools to pricing a vanilla European option
- Be encouraged to pursue further study by tackling more advanced subjects, from both practical and theoretical perspectives

The material covered in the book has been selected with these aims in mind. Of course, personal tastes are admittedly reflected, and this has something to do with my Operations Research background. I am afraid the book will not please statisticians, as no econometric model is developed; however, there is a wide and excellent literature on those topics, and I tried to come up with a complementary textbook.

The text is interspersed with MATLAB snapshots and pieces of code, to make the material as lively as possible and of immediate use. MATLAB is a flexible high-level computing environment which allows us to implement non-trivial algorithms with a few lines of code. It has also been chosen because of its increasing potential for specific financial applications.

It may be argued that the book is more successful at raising questions than at giving answers. This is a necessary evil, given the space available to cover such a wide array of topics. But if, after reading this book, students will want to read others, my job will have been accomplished. This was meant to be a crossroads, after all.

PS1. Despite all of my effort, the book is likely to contain some errors and typos. I will maintain a list of errata, which will be updated, based on reader feedback. Any comment or suggestion on the book will also be appreciated. My e-mail address is: `paolo.brandimarte@polito.it`.

PS2. The list of errata will be posted on a Web page which will also include additional material and MATLAB programs. The current URL is

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

An up-to-date link will be maintained on Wiley Web page:

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

PS3. And if (what a shame ...) you are wondering who Commander Straker is, take a look at the following Web sites:

- `http://www.ufoseries.com`
- `http://www.isoshado.org`

PAOLO BRANDIMARTE
Turin, June 2001