Introduction to Quantitative Finance

A Math Tool Kit

Robert R. Reitano

The MIT Press
Cambridge, Massachusetts
London, England

MIT Press books may be purchased at special quantity discounts for business or sales promotional use. For information, please email special_sales@mitpress.mit.edu or write to Special Sales Department, The MIT Press, 55 Hayward Street, Cambridge, MA 02142.

This book was set in Times New Roman on 3B2 by Asco Typesetters, Hong Kong and was printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Reitano, Robert R., 1950–
Introduction to quantitative finance : a math tool kit / Robert R. Reitano.
p. cm.
Includes index.
ISBN 978-0-262-01369-7 (hardcover : alk. paper) 1. Finance—Mathematical models. I. Title.
HG106.R45 2010
332.01′5195—dc22 2009022214

10 9 8 7 6 5 4 3 2 1

Introduction

This book provides an accessible yet rigorous introduction to the fields of mathematics that are needed for success in investment and quantitative finance. The book's goal is to develop mathematics topics used in portfolio management and investment banking, including basic derivatives pricing and risk management applications, that are essential to quantitative investment finance, or more simply, investment finance. A future book, *Advanced Quantitative Finance: A Math Tool Kit*, will cover more advanced mathematical topics in these areas as used for investment modeling, derivatives pricing, and risk management. Collectively, these latter areas are called quantitative finance or mathematical finance.

The mathematics presented in this book would typically be learned by an undergraduate mathematics major. Each chapter of the book corresponds roughly to the mathematical materials that are acquired in a one semester course. Naturally each chapter presents only a subset of the materials from these traditional math courses, since the goal is to emphasize the most important and relevant materials for the finance applications presented. However, more advanced topics are introduced earlier than is customary so that the reader can become familiar with these materials in an accessible setting.

My motivation for writing this text was to fill two current gaps in the financial and mathematical literature as they apply to students, and practitioners, interested in sharpening their mathematical skills and deepening their understanding of investment and quantitative finance applications. The gap in the mathematics literature is that most texts are focused on a single field of mathematics such as calculus. Anyone interested in meeting the field requirements in finance is left with the choice to either pursue one or more degrees in mathematics or expend a significant self-study effort on associated mathematics textbooks. Neither approach is efficient for business school and finance graduate students nor for professionals working in investment and quantitative finance and aiming to advance their mathematical skills. As the diligent reader quickly discovers, each such book presents more math than is needed for finance, and it is nearly impossible to identify what math is essential for finance applications. An additional complication is that math books rarely if ever provide applications in finance, which further complicates the identification of the relevant theory.

The second gap is in the finance literature. Finance texts have effectively become bifurcated in terms of mathematical sophistication. One group of texts takes the recipe-book approach to math finance often presenting mathematical formulas with only simplified or heuristic derivations. These books typically neglect discussion of the mathematical framework that derivations require, as well as effects of assumptions by which the conclusions are drawn. While such treatment may allow more discussion of the financial applications, it does not adequately prepare the student who will inevitably be investigating quantitative problems for which the answers are unknown.

The other group of finance textbooks are mathematically rigorous but inaccessible to students who are not in a mathematics degree program. Also, while rigorous, such books depend on sophisticated results developed elsewhere, and hence the discussions are incomplete and inadequate even for a motivated student without additional classroom instruction. Here, again, the unprepared student must take on faith referenced results without adequate understanding, which is essentially another form of recipe book.

With this book I attempt to fill some of these gaps by way of a reasonably economic, yet rigorous and accessible, review of many of the areas of mathematics needed in quantitative investment finance. My objective is to help the reader acquire a deep understanding of relevant mathematical theory and the tools that can be effectively put in practice. In each chapter I provide a concluding section on finance applications of the presented materials to help the reader connect the chapter's mathematical theory to finance applications and work in the finance industry.

What Does It Take to Be a "Quant"?

In some sense, the emphasis of this book is on the development of the math tools one needs to succeed in mathematical modeling applications in finance. The imagery implied by "math tool kit" is deliberate, and it reflects my belief that the study of mathematics is an intellectually rewarding endeavor, and it provides an enormously flexible collection of tools that allow users to answer a wide variety of important and practical questions.

By tools, however, I do not mean a collection of formulas that should be memorized for later application. Of course, some memorization is mandatory in mathematics, as in any language, to understand what the words mean and to facilitate accurate communication. But most formulas are outside this mandatorily memorized collection. Indeed, although mathematics texts are full of formulas, the memorization of formulas should be relatively low on the list of priorities of any student or user of these books. The student should instead endeavor to learn the mathematical frameworks and the application of these frameworks to real world problems.

In other words, the student should focus on the thought process and mathematics used to develop each result. These are the "tools," that is, the mathematical methods of each discipline of explicitly identifying assumptions, formally developing the needed insights and formulas, and understanding the relationships between formulas and the underlying assumptions. The tools so defined and studied in this book will equip the student with fairly robust frameworks for their applications in investment and quantitative finance.

Despite its large size, this book has the relatively modest ambition of teaching a very specific application of mathematics, that being to finance, and so the selection of materials in every subdiscipline has been made parsimoniously. This selection of materials was the most difficult aspect of developing this book. In general, the selection criterion I used was that a topic had to be either directly applicable to finance, or needed for the understanding of a later topic that was directly applicable to finance. Because my objective was to make this book more than a collection of mathematical formulas, or just another finance recipe book, I devote considerable space to discussion on how the results are derived, and how they relate to their mathematical assumptions. Ideally the students of this book should never again accept a formulaic result as an immutable truth separate from any assumptions made by its originator.

The motivation for this approach is that in investment and quantitative finance, there are few good careers that depend on the application of standard formulas in standard situations. All such applications tend to be automated and run in companies' computer systems with little or no human intervention. Think "program trading" as an example of this statement. While there is an interesting and deep theory related to identifying so-called arbitrage opportunities, these can be formulaically listed and programmed, and their implementation automated with little further analyst intervention.

Equally, if not more important, with new financial products developed regularly, there are increased demands on quants and all finance practitioners to apply the previous methodologies and adapt them appropriately to financial analyses, pricing, risk modeling, and risk management. Today, in practice, standard results may or may not apply, and the most critical job of the finance quant is to determine if the traditional approach applies, and if not, to develop an appropriate modification or even an entirely new approach. In other words, for today's finance quants, it has become critical to be able to think in mathematics, and not simply to do mathematics by rote.

The many finance applications developed in the chapters present enough detail to be understood by someone new to the given application but in less detail than would be appropriate for mastering the application. Ideally the reader will be familiar with some applications and will be introduced to other applications that can, as needed, be enhanced by further study. On my selection of mathematical topics and finance applications, I hope to benefit from the valuable comments of finance readers, whether student or practitioner. All such feedback will be welcomed and acknowledged in future editions.

Plan of the Book

The ten chapters of this book are arranged so that each topic is developed based on materials previously discussed. In a few places, however, a formula or result is introduced that could not be fully developed until a much later chapter. In fewer places, I decided to not prove a deep result that would have brought the book too far afield from its intended purpose. Overall, the book is intended to be self-contained, complete with respect to the materials discussed, and mathematically rigorous. The only mathematical background required of the reader is competent skill in algebraic manipulations and some knowledge of pre-calculus topics of graphing, exponentials and logarithms. Thus the topics developed in this book are interrelated and applied with the understanding that the student will be motivated to work through, with pen or pencil and paper or by computer simulation, any derivation or example that may be unclear and that the student has the algebraic skills and self-discipline to do so.

Of course, even when a proof or example appears clear, the student will benefit in using pencil and paper and computer simulation to clarify any missing details in derivations. Such informal exercises provide essential practice in the application of the tools discussed, and analytical skills can be progressively sharpened by way of the book's formal exercises and ultimately in real world situations. While not every derivation in the book offers the same amount of enlightenment on the mathematical tools studied, or should be studied in detail before proceeding, developing the habit of filling in details can deepen mathematical knowledge and the understanding of how this knowledge can be applied.

I have identified the more advanced sections by an asterisk (*). The beginning student may find it useful to scan these sections on first reading. These sections can then be returned to if needed for a later application of interest. The more advanced student may find these sections to provide some insights on the materials they are already familiar with. For beginning practitioners and professors of students new to the materials, it may be useful to only scan the reasoning in the longer proofs on a first review before turning to the applications.

There are a number of productive approaches to the chapter sequencing of this book for both self-study and formal classroom presentation. Professors and practitioners with good prior exposure might pick and choose chapters out of order to efficiently address pressing educational needs. For finance applications, again the best approach is the one that suits the needs of the student or practitioner. Those familiar with finance applications and aware of the math skills that need to be developed will focus on the appropriate math sections, then proceed to the finance applications to better understand the connections between the math and the finance. Those less familiar with finance may be motivated to first review the applications section of each chapter for motivation before turning to the math.

Some Course Design Options

This book is well suited for a first-semester introductory graduate course in quantitative finance, perhaps taken at the same time as other typical first-year graduate courses for finance students, such as investment markets and products, portfolio theory, financial reporting, corporate finance, and business strategy. For such students the instructor can balance the class time between sharpening mathematical knowledge and deepening a level of understanding of finance applications taken in the first term. Students will then be well prepared for more quantitatively focused investment finance courses on fixed income and equity markets, portfolio management, and options and derivatives, for example, in the second term.

For business school finance students new to the subject of finance, it might be better to defer this book to a second semester course, following an introductory course in financial markets and instruments so as to provide a context for the finance applications discussed in the chapters of this book.

This book is also appropriate for graduate students interested in firming up their technical knowledge and skills in investment and quantitative finance, so it can be used for self-study by students soon to be working in investment or quantitative finance, and by practitioners needing to improve their math skill set in order to advance their finance careers in the "quant" direction. Mathematics and engineering departments, which will have many very knowledgeable graduate and undergraduate students in the areas of math covered in this book, may also be interested in offering an introductory course in finance with a strong mathematical framework. The rigorous math approach to real world applications will be familiar to such students, so a balance of math and finance could be offered early in the students' academic program.

For students for whom the early chapters would provide a relatively easy review, it is feasible to take a sequential approach to all the materials, moving faster through the familiar math topics and dwelling more on the finance applications. For nonmathematical students who risk getting bogged down by the first four chapters in their struggle with abstract notions, and are motivated to learn the math only after recognizing the need in a later practical setting, it may be preferable to teach only a subset of the math from chapters 1 through 4 and focus on the intuition behind these chapters' applications. For example, an instructor might provide a quick overview of logic and proof from chapter 1, choose selectively from chapter 2 on number systems, then skip ahead to chapter 4 for set operations. After this topical tour the instructor could finally settle in with all the math and applications in chapter 5 on sequences and then move forward sequentially through chapters 6 to 10. The other mathematics topics of chapters 1 through 4 could then be assigned or taught as required to supplement the materials of these later chapters. This approach and pace could keep the students motivated by getting to the more meaningful applications sooner, and thus help prevent math burnout before reaching these important applications.

Chapter Exercises

Chapter exercises are split into practice exercises and assignment exercises. Both types of exercises provide practice in mathematics and finance applications. The more challenging exercises are accompanied by a "hint," but students should not be constrained by the hints. The best learning in mathematics and in applications often occurs in pursuit of alternative approaches, even those that ultimately fail. Valuable lessons can come from such failures that help the student identify a misunderstanding of concepts or a misapplication of logic or mathematical techniques. Therefore, if other approaches to a problem appear feasible, the student is encouraged to follow at least some to a conclusion. This additional effort can provide reinforcement of a result that follows from different approaches but also help identify errors and misunderstandings when two approaches lead to different conclusions.

Solutions and Instructor's Manuals

For the book's practice exercises, a *Solutions Manual* with detailed explanations of solutions is available for purchase by students. For the assignment exercises, solutions are available to instructors as part of an *Instructor's Manual*. This Manual also contains chapter-by-chapter suggestions on teaching the materials. All instructor materials are also available online.

Organization of Chapters

Few mathematics books today have an introductory chapter on mathematical logic, and certainly none that address applications. The field of logic is a subject available to mathematics or philosophy students as a separate course. To skip the material on logic is to miss an opportunity to acquire useful tools of thinking, in drawing appropriate conclusions, and developing clear and correct quantitative reasoning.

Simple conclusions and quantitative derivations require no formality of logic, but the tools of truth tables and statement analysis, as well as the logical construction of a valid proof, are indispensable in evaluating the integrity of more complicated results. In addition to the tools of logic, chapter 1 presents various approaches to proofs that follow from these tools, and that will be encountered in subsequent chapters. The chapter also provides a collection of paradoxes that are often amusing and demonstrate that even with careful reasoning, an argument can go awry or a conclusion reached can make no sense. Yet paradoxes are important; they motivate clearer thinking and more explicit identification of underlying assumptions.

Finally, for completeness, this chapter includes a discussion of the axiomatic formality of mathematical theory and explains why this formality can help one avoid paradoxes. It notes that there can be some latitude in the selection of the axioms, and that axioms can have a strong effect on the mathematical theory. While the reader should not get bogged down in these formalities, since they are not critical to the understanding of the materials that follow, the reader should find comfort that they exist beneath the more familiar frameworks to be studied later.

The primary application of mathematical logic to finance and to any field is as a guide to cautionary practice in identifying assumptions and in applying or deriving a needed result to avoid the risk of a potentially disastrous consequence. Intuition is useful as a guide to a result, but never as a substitute for careful analysis.

Chapter 2, on number systems and functions, may appear to be on relatively trivial topics. Haven't we all learned numbers in grade school? The main objective in reviewing the different number systems is that they *are* familiar and provide the foundational examples for more advanced mathematical models. Because the aim of this book is to introduce important concepts early, the natural numbers provide a relatively simple example of an axiomatic structure from chapter 1 used to develop a mathematical theory.

From the natural numbers other numbers are added sequentially to allow more arithmetic operations, leading in turn to integers, rational, irrational, real, and complex numbers. Along the way these collections are seen to share certain arithmetic structures, and the notions of group and field are introduced. These collections also provide an elementary context for introducing the notions of countable and uncountable infinite sets, as well as the notion of a "dense" subset of a given set. Once defined, these number systems and their various subsets are the natural domains on which functions are defined.

While it might be expected that only the rational numbers are needed in finance, and indeed the rational numbers with perhaps only 6 to 10 decimal point representations, it is easy to exemplify finance problems with irrational and even complex number solutions. In the former cases, rational approximations are used, and sometimes with reconciliation difficulties to real world transactions, while complex numbers are avoided by properly framing the interest rate basis. Functions appear everywhere in finance—from interest rate nominal basis conversions, to the pricing functions for bonds, mortgages and other loans, preferred and common stock, and forward contracts, and to the modeling of portfolio returns as a function of the asset allocation.

The development of number system structures is continued in chapter 3 on Euclidean and other spaces. Two-dimensional Euclidean space, as was introduced in chapter 2, provided a visual framework for the complex numbers. Once defined, the vector space structure of Euclidean space is discussed, as well as the notions of the standard norm and inner product on these spaces. This discussion leads naturally to the important Cauchy–Schwarz inequality relating these concepts, an inequality that arises time and again in various contexts in this book. Euclidean space is also the simplest context in which to introduce the notion of alternative norms, and the l_p-norms, in particular, are defined and relationships developed. The central result is the generalization of Cauchy–Schwarz to the Hölder inequality, and of the triangle inequality to the Minkowski inequality.

Metrics are then discussed, as is the relationship between a metric and a norm, and cases where one can be induced from the other on a given space using examples from the l_p-norm collection. A common theme in mathematics and one seen here is that a general metric is defined to have exactly the essential properties of the standard and familiar metric defined on $\mathbb{R}^2$ or generalized to $\mathbb{R}^n$. Two notions of equivalence of two metrics is introduced, and it is shown that all the metrics induced by the l_p-norms are equivalent in Euclidean space. Strong evidence is uncovered that this result is fundamentally related to the finite dimensionality of these spaces, suggesting that equivalence will not be sustained in more general forthcoming contexts. It is also illustrated that despite this general l_p-equivalence result, not all metrics are equivalent.

For finance applications, Euclidean space is seen to be the natural habitat for expressing vectors of asset allocations within a portfolio, various bond yield term structures, and projected cash flows. In addition, all the l_p-norms appear in the calculation of various moments of sample statistical data, while some of the l_p-norms, specifically $p = 1, 2$, and ∞, appear in various guises in constrained optimization problems common in finance. Sometimes these special norms appear as constraints and sometimes as the objective function one needs to optimize.

Chapter 4 on set theory and topology introduces another example of an axiomatic framework, and this example is motivated by one of the paradoxes discussed in chapter 1. But the focus here is on set operations and their relationships. These are important tools that are as essential to mathematical derivations as are algebraic manipulations. In addition, basic concepts of open and closed are first introduced in the familiar setting of intervals on the real line, but then generalized and illustrated making good use of the set manipulation results. After showing that open sets in $\mathbb{R}$ are relatively simple, the construction of the Cantor set is presented as an exotic example of a closed set. It is unusual because it is uncountable and yet, at the same time, shown to have "measure 0." This result is demonstrated by showing that the Cantor set is what is left from the interval $[0, 1]$ after a collection of intervals are removed that have total length equal to 1!

The notions of open and closed are then extended in a natural way to Euclidean space and metric spaces, and the idea of a topological space is introduced for completeness. The basic aim is once again to illustrate that a general idea, here topology, is defined to satisfy exactly the same properties as do the open sets in more familiar contexts. The chapter ends with a few other important notions such as accumulation point and compactness, which lead to discussions in the next chapter.

For finance applications, constrained optimization problems are seen to be naturally interpreted in terms of sets in Euclidean space defined by functions and/or norms. The solution of such problems generally requires that these sets have certain topological properties like compactness and that the defining functions have certain regularity properties. Function regularity here means that the solution of an equation can be approximated with an iterative process that converges as the number of steps increases, a notion that naturally leads to chapter 5. Interval bisection is introduced as an example of an iterative process, with an application to finding the yield of a security, and convergence questions are made explicit and seen to motivate the notion of continuity.

Sequences and their convergence are addressed in chapter 5, making good use of the concepts, tools, and examples of earlier chapters. The central idea, of course, is that of convergence to a limit, which is informally illustrated before it is formally defined. Because of the importance of this idea, the formal definition is discussed at some length, providing both more detail on what the words mean and justification as to why this definition requires the formality presented. Convergence is demonstrated to be preserved under various arithmetic operations. Also an important result related to compactness is demonstrated: that is, while a bounded sequence need not converge, it must have an accumulation point and contain a subsequence that convergences to that accumulation point. Because such sequences may have many—indeed infinitely many—such accumulation points, the notions of limit superior and limit inferior are introduced and shown to provide the largest and smallest such accumulation points, respectively.

Convergence of sequences is then discussed in the more general context of Euclidean space, for which all the earlier results generalize without modification, and metric spaces, in which some care is needed. The notion of a Cauchy sequence is next introduced and seen to naturally lead to the question of whether such sequences converge to a point of the space, as examples of both convergence and nonconvergence are presented. This discussion leads to the introduction of the idea of completeness of a metric space, and of its completion, and an important result on completion is presented without proof but seen to be consistent with examples studied.

Interval bisection provides an important example of a Cauchy sequence in finance. Here the sequence is of solution iterates, but again the question of convergence of the associated price values remains open to a future chapter. With more details on this process, the important notion of continuous function is given more formality.

Although the convergence of an infinite sequence is broadly applicable in its own right, this theory provides the perfect segue to the convergence of infinite sums addressed in chapter 6 on series and their convergence. Notions of absolute and conditional convergence are developed, along with the implications of these properties for arithmetic manipulations of series, and for re-orderings or rearrangements of the series terms. Rearrangements are discussed for both single-sum and multiple-sum applications.

A few of the most useful tests for convergence are developed in this chapter. The chapter 3 introduction to the l_p-norms is expanded to include l_p-spaces of sequences and associated norms, demonstrating that these spaces are complete normed spaces, or Banach spaces, and are overlapping yet distinct spaces for each p. The case of $p = 2$ gets special notice as a complete inner product space, or Hilbert space, and implications of this are explored. Power series are introduced, and the notions of radius of convergence and interval of convergence are developed from one of the previous tests for convergence. Finally, results for products and quotients of power series are developed.

Applications to finance include convergence of price formulas for various perpetual preferred and common stock models with cash flows modeled in different functional ways, and various investor yield demands. Linearly increasing cash flows provide an example of double summation methods, and the result is generalized to polynomial payments. Approximating complicated pricing functions with power series is considered next, and the application of the l_p-spaces is characterized as providing an accessible introduction to the generalized function space counterparts to be studied in more advanced texts.

An important application of the tools of chapter 6 is to discrete probability theory, which is the topic developed in chapter 7 starting with sample spaces and probability measures. By discrete, it is meant that the theory applies to sample spaces with a finite or countably infinite number of sample points. Also studied are notions of conditional probability, stochastic independence, and an n-trial sample space construction that provides a formal basis for the concept of an independent sample from a sample space. Combinatorics are then presented as an important tool for organizing and counting collections of events from discrete sample spaces.

Random variables are shown to provide key insights to a sample space and its probability measure through the associated probability density and distribution functions, making good use of the combinatorial tools. Moments of probability density functions and their properties are developed, as well as moments of sample data drawn from an n-trial sample space. Several of the most common discrete probability density functions are introduced, as well as a methodology for generating random samples from any such density function.

Applications of these materials in finance are many, and begin with loss models related to bond or loan portfolios, as well as those associated with various forms of insurance. In this latter context, various net premium calculations are derived. Asset allocation provides a natural application of probability methods, as does the modeling of equity prices in discrete time considered within either a binomial lattice or binomial scenario model. The binomial lattice model is then used for option pricing in discrete time based on the notion of option replication. Last, scenario-based option pricing is introduced through the notion of a sample-based option price defined in terms of a sampling of equity price scenarios.

With chapter 7 providing the groundwork, chapter 8 develops a collection of the fundamental probability theorems, beginning with a modest proof of the uniqueness of the moment-generating and characteristic functions in the case of finite discrete probability density functions. Chebyshev's inequality, or rather, Chebyshev's inequalities, are developed, as is the *weak law of large numbers* as the first of several results related to the distribution of the sample mean of a random variable in the limit as the sample size grows. Although the weak law requires only that the random variable have a finite mean, in the more common case where the variance is also finite, this law is derived with a sleek one-step proof based on Chebyshev.

The *strong law of large numbers* requires both a finite mean and variance but provides a much more powerful statement about the distribution of sample means in the limit. The strong law is based on a generalization of the Chebyshev inequality known as Kolmogorov's inequality. The De Moivre–Laplace theorem is investigated next, followed by discussions on the normal distribution and the *central limit theorem* (CLT). The CLT is proved in the special case of probability densities with moment-generating functions, and some generalizations are discussed.

For finance applications, Chebyshev is applied to the problem of modeling and evaluating asset adequacy, or capital adequacy, in a risky balance sheet. Then the binomial lattice model for stock prices under the real world probabilities introduced in chapter 7 is studied in the limit as the time interval converges to zero, and the probability density function of future stock prices is determined. This analysis uses the methods underlying the De Moivre–Laplace theorem and provides the basis of the next investigation into the derivation of the Black–Scholes–Merton formulas for the price of a European put or call option. Several of the details of this derivation that require the tools of chapters 9 and 10 are deferred to those chapters. The final application is to the probabilistic properties of the scenario-based option price introduced in chapter 7.

The calculus of functions of a single variable is the topic developed in the last two chapters. Calculus is generally understood as the study of functions that display various types of "smoothness." In line with tradition, this subject is split into a differentiation theory and an integration theory. The former provides a rigorous framework for approximating smooth functions, and the latter introduces in an accessible framework an important tool needed for a continuous probability theory.

Chapter 9 on the calculus of differentiation begins with the formal introduction of the notion of continuity and its variations, as well the development of important properties of continuous functions. These basic notions of smoothness provide the beginnings of an approximation approach that is generalized and formalized with the development of the derivative of a function. Various results on differentiation follow, as does the formal application of derivatives to the question of function approximation via Taylor series. With these tools important results are developed related to the derivative, such as classifying the critical points of a given function, characterizing the notions of convexity and concavity, and the derivation of Jensen's inequality. Not only can derivatives be used to approximate function values, but the values of derivatives can be approximated using nearby function values and the associated errors quantified. Results on the preservation of continuity and differentiability under convergence of a sequence of functions are addressed, as is the relationship between analytic functions and power series.

Applications found in finance include the continuity of price functions and their application to the method of interval bisection. Also discussed is the continuity of objective functions and constraint functions and implications for solvability of constrained optimization problems. Deriving the minimal risk portfolio allocation is one application of a critical point analysis. Duration and convexity of fixed income investments is studied next and used in an application of Taylor series to price function approximations and asset-liability management problems in various settings.

Outside of fixed income, the more common sensitivity measures are known as the "Greeks," and these are introduced and shown to easily lend themselves to Taylor series methods. Utility theory and its implications for risk preferences are studied as an application of convex and concave functions and Jensen's inequality, and then applied in the context of optimal portfolio allocation. Finally, details are provided for the limiting distributions of stock prices under the risk-neutral probabilities and special risk-averter probabilities needed for the derivation of the Black–Scholes–Merton option pricing formulas, extending and formalizing the derivation begun in chapter 8. The risk-averter model is introduced in chapter 8 as a mathematical artifact to facilitate the final derivation, but it is clear the final result only depends on the risk-neutral model.

The notion of Riemann integral is studied in chapter 10 on the calculus of integration, beginning with its definition for a continuous function on a closed and bounded interval where it is seen to represent a "signed" area between the graph of the function and the x-axis. A series of generalizations are pursued, from the weakening of the continuity assumption to that of bounded and continuous "except on a set of points of measure 0," to the generalization of the interval to be unbounded, and finally to certain generalizations when the function is unbounded. Properties of such integrals are developed, and the connection between integration and differentiation is studied with two forms of the fundamental theorem of calculus.

The evaluation of a given integral is pursued with standard methods for exact valuation as well as with numerical methods. The notion of integral is seen to provide a useful alternative representation of the remainder in a Taylor series, and to provide a powerful tool for evaluating convergence of, and estimating the sum of or rate of divergence of, an infinite series. Convergence of a sequence of integrals is included. The Riemann notion of an integral is powerful but has limitations, some of which are explored.

Continuous probability theory is developed with the tools of this chapter, encompassing more general probability spaces and sigma algebras of events. Continuously distributed random variables are introduced, as well as their moments, and an accessible result is presented on discretizing such a random variable that links the discrete and continuous moment results. Several continuous distributions are presented and their properties studied.

Applications to finance in chapter 10 include the present and accumulated value of continuous cash flow streams with continuous interest rates, continuous interest rate term structures for bond yields, spot and forward rates, and continuous equity dividends and their reinvestment into equities. An alternative approach to applying the duration and convexity values of fixed income investments to approximating price functions is introduced. Numerical integration methods are exemplified by application to the normal distribution.

Finally, a generalized Black–Scholes–Merton pricing formula for a European option is developed from the general binomial pricing result of chapter 8, using a "continuitization" of the binomial distribution and a derivation that this continuitization converges to the appropriate normal distribution encountered in chapter 9. As another application, the Riemann–Stieltjes integral is introduced in the chapter exercises. It is seen to provide a mathematical link between the calculations within the discrete and continuous probability theories, and to generalize these to so-called mixed probability densities.

Acknowledgments

I have had the pleasure and privilege to train under and work with many experts in both mathematics and finance. My thesis advisor and mentor, Alberto P. Calderón (1920–1998), was the most influential in my mathematical development, and to this day I gauge the elegance and lucidity of any mathematical argument by the standard he set in his work and communications. In addition I owe a debt of gratitude to all the mathematicians whose books and papers I have studied, and whose best proofs have greatly influenced many of the proofs presented throughout this book.

I also acknowledge the advice and support of many friends and professional associates on the development of this book. Notably this includes (alphabetically) fellow academics Zvi Bodie, Laurence D. Booth, F. Trenery Dolbear, Jr., Frank J. Fabozzi, George J. Hall, John C. Hull, Blake LeBaron, Andrew Lyasoff, Bruce R. Magid, Catherine L. Mann, and Rachel McCulloch, as well as fellow finance practitioners Foster L. Aborn, Charles L. Gilbert, C. Dec Mullarkey, K. Ravi Ravindran and Andrew D. Smith, publishing professionals Jane MacDonald and Tina Samaha, and my editor at the MIT Press, Dana Andrus.

I thank the students at the Brandeis University International Business School for their feedback on an earlier draft of this book and careful proofreading, notably Amidou Guindo, Zhenbin Luo, Manjola Tase, Ly Tran, and Erick Barongo Vedasto. Despite their best efforts I remain responsible for any remaining errors.

Last, I am indebted to my parents, Dorothy and Domenic, for a lifetime of advice and support. I happily acknowledge the support and encouragement of my wife Lisa, who also provided editorial support, and sons Michael, David, and Jeffrey, during the somewhat long and continuing process of preparing my work for publication.

I welcome comments on this book from readers. My email address is *rreitano@brandeis.edu.*

Robert R. Reitano
International Business School
Brandeis University