HANDBOOKS IN
OPERATIONS RESEARCH AND
MANAGEMENT SCIENCE

Volume 15

FINANCIAL ENGINEERING

John R. Birge Vadim Linetsky
Editors

NORTH-HOLLAND

HANDBOOKS IN OPERATIONS RESEARCH
AND MANAGEMENT SCIENCE
VOLUME 15

Handbooks in Operations Research and Management Science

Editors

J.K. Lenstra
Centrum voor Wiskunde en
Informatica, Amsterdam

G.L. Nemhauser
Georgia Institute of Technology

J.G. Dai
Georgia Institute of Technology

Volume 15

ELSEVIER

Amsterdam – Boston – Heidelberg – London – New York – Oxford
Paris – San Diego – San Francisco – Singapore – Sydney – Tokyo
North-Holland is an imprint of Elsevier

N·H

Financial Engineering

Edited by

John R. Birge
University of Chicago, IL, USA

Vadim Linetsky
Northwestern University, IL, USA

ELSEVIER

Amsterdam – Boston – Heidelberg – London – New York – Oxford
Paris – San Diego – San Francisco – Singapore – Sydney – Tokyo
North-Holland is an imprint of Elsevier

N·H

North-Holland is an imprint of Elsevier
Radarweg 29, PO Box 211, 1000 AE Amsterdam, The Netherlands
Linacre House, Jordan Hill, Oxford OX2 8DP, UK

First edition 2008

Library of Congress Cataloging-in-Publication Data
A catalog record for this book is available from the Library of Congress

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

ISBN: 978-0-444-51781-4
ISSN: 0927-0507

For information on all North-Holland publications
visit our web site at books.elsevier.com

Printed and bound in The Netherlands

08 09 10 11 12 10 9 8 7 6 5 4 3 2 1

Contents

PART I

Introduction

This page intentionally left blank

Introduction to the Handbook of Financial Engineering

John R. Birge
Graduate School of Business, University of Chicago, USA

Vadim Linetsky
Department of Industrial Engineering and Management Sciences, Northwestern University, USA

Financial engineering (FE) is an interdisciplinary field focusing on applications of mathematical and statistical modeling and computational technology to problems in the financial services industry. According to the report by the National Academy of Engineering (2003),[1] "Financial services are the foundation of a modern economy. They provide mechanisms for assigning value, exchanging payment, and determining and distributing risk, and they provide the essential underpinnings of global economic activity. The industry provides the wherewithal for the capital investment that drives innovation and productivity growth throughout the economy." Important areas of FE include mathematical modeling of market and credit risk, pricing and hedging of derivative securities used to manage risk, asset allocation and portfolio management.

Market risk is a risk of adverse changes in prices or rates, such as interest rates, foreign exchange rates, stock prices, and commodity and energy prices. Credit risk is a risk of default on a bond, loan, lease, pension or any other type of financial obligation. Modern derivatives markets can be viewed as a global marketplace for financial risks. The function of derivative markets is to facilitate financial risk transfer from risk reducers (hedgers) to risk takers (investors). Organizations wishing to reduce their risk exposure to a particular type of financial risk, such as the risk of increasing commodity and energy prices that will make future production more expensive or the risk of increasing interest rates that will make future financing more expensive, can offset those risks by entering into financial contracts that act as insurance, protecting the company against adverse market events. While the hedger comes to the derivatives market to reduce its risk, the counterparty who takes the other side

[1] National Academy of Engineering, The Impact of Academic Research on Industrial Performance, National Academies Press, Washington, DC, 2003, http://www.nap.edu/books/0309089735/html.

of the contract comes to the market to invest in that risk and expects to be adequately compensated for taking the risk. We can thus talk about buying and selling financial risks.

Global derivatives markets have experienced remarkable growth over the past several decades. According to a recent survey by the Bank for International Settlement in Basel (2006),[2] the aggregate size of the global derivatives markets went from about $50 trillion in notional amounts in 1995 to $343 trillion in notional amounts by the end of 2005 ($283 trillion in notional amounts in over-the-counter derivatives contracts and $58 trillion in futures and options traded on derivatives exchanges worldwide). Major segments of the global derivatives markets include interest rate derivatives, currency derivatives, equity derivatives, commodity and energy derivatives, and credit derivatives.

A *derivative* is a financial contract between two parties that specifies conditions, in particular, dates and the resulting values of underlying variables, under which payments or *payoffs* are to be made between parties (payments can be either in the form of cash or delivery of some specified asset). Call and put options are basic examples of derivatives used to manage market risk. A *call option* is a contract that gives its holder the right to buy some specified quantity of an underlying asset (for example, a fixed number of shares of stock of a particular company or a fixed amount of a commodity) at a predetermined price (called the *strike price*) on or before a specified date in the future (option *expiration*). A *put option* is a contract that gives its holder the right to sell some specified quantity of an underlying asset at a predetermined price on or before expiration. The holder of the option contract locks in the price for future purchase (in the case of call options) or future sale (in the case of put options), thus eliminating any price uncertainty or risk, at the cost of paying the *premium* to purchase the option. The situation is analogous to insurance contracts that pay pre-agreed amounts in the event of fire, flood, car accident, etc. In financial options, the payments are based on financial market moves (and credit events in the case of credit derivatives). Just as in the insurance industry the key problem is to determine the insurance premium to charge for a policy based on actuarial assessments of event probabilities, the option-pricing problem is to determine the premium or option price based on a stochastic model of the underlying financial variables.

Portfolio optimization problems constitute another major class of important problems in financial engineering. Portfolio optimization problems occur throughout the financial services industry as pension funds, mutual funds, insurance companies, university and foundation endowments, and individual investors all face the fundamental problem of allocating their capital across different securities in order to generate investment returns sufficient to achieve a particular goal, such as meeting future pension liabilities. These problems are

[2] Bank for International Settlement Quarterly Review, June 2006, pp. A103–A108, http://www.bis.org/publ/qtrpdf/r_qa0606.pdf.

often very complex owing to their dynamic and stochastic nature, their high dimensionality, and the complexity of real-world constraints.

The remarkable growth of financial markets over the past decades has been accompanied by an equally remarkable explosion in financial engineering research. The goals of financial engineering research are to develop empirically realistic stochastic models describing dynamics of financial risk variables, such as asset prices, foreign exchange rates, and interest rates, and to develop analytical, computational and statistical methods and tools to implement the models and employ them to evaluate financial products used to manage risk and to optimally allocate investment funds to meet financial goals. As financial models are stochastic, probability theory and stochastic processes play a central role in financial engineering. Furthermore, in order to implement financial models, a wide variety of analytical and computational tools are used, including Monte Carlo simulation, numerical PDE methods, stochastic dynamic programming, Fourier methods, spectral methods, etc.

The Handbook is organized in six parts: Introduction, Derivative Securities: Models and Methods, Interest Rate and Credit Risk Models and Derivatives, Incomplete Markets, Risk Management, and Portfolio Optimization. This division is somewhat artificial, as many chapters are equally relevant for several or even all of these areas. Nevertheless, this structure provides an overview of the main areas of the field of financial engineering.

A working knowledge of probability theory and stochastic processes is a prerequisite to reading many of the chapters in the Handbook. Karatzas and Shreve (1991) and Revuz and Yor (1999) are standard references on Brownian motion and continuous martingales. Jacod and Shiryaev (2002) and Protter (2005) are standard references on semimartingale processes with jumps. Shreve (2004) and Klebaner (2005) provide excellent introductions to stochastic calculus for finance at a less demanding technical level. For the financial background at the practical level, excellent overviews of derivatives markets and financial risk management can be found in Hull (2005) and McDonald (2005). Key texts on asset pricing theory include Bjork (2004), Duffie (2001), Jeanblanc et al. (2007), and Karatzas and Shreve (2001). These monographs also contain extensive bibliographies.

In Chapter 1 “A Partial Introduction to Financial Asset Pricing Theory,” Robert Jarrow and Philip Protter present a concise introduction to Mathematical Finance theory. The reader is first introduced to derivative securities and the fundamental financial concept of arbitrage in the binomial framework. The core asset pricing theory is then developed in the general semimartingale framework, assuming prices of risky assets follow semimartingale processes. The general fundamental theorems of asset pricing are formulated and illustrated on important examples. In particular, the special case when the risky asset price process is a Markov process is treated in detail, the celebrated Black–Scholes–Merton model is derived, and a variety of results on pricing European- and American-style options and more complex derivative securities are presented. This chapter summarizes the core of Mathematical Finance theory and is an essential reading.

Part II "Derivative Securities: Models and Methods," contains chapters on a range of topics in derivatives modeling and pricing. The first three chapters survey several important classes of stochastic models used in derivatives modeling. In Chapter 2 "Jump-Diffusion Models," Steven Kou surveys recent developments in option pricing in jump-diffusion models. The chapter discusses empirical evidence of jumps in financial variables and surveys analytical and numerical methods for the pricing of European, American, barrier, and lookback options in jump-diffusion models, with particular attention given to the jump-diffusion model with a double-exponential jump size distribution due to its analytical tractability.

In Chapter 3 "Modeling Financial Security Returns Using Levy Processes," Liuren Wu surveys a class of models based on time-changed Levy processes. Applying stochastic time changes to Levy processes randomizes the clock on which the process runs, thus generating stochastic volatility. If the characteristic exponent of the underlying Levy process and the Laplace transform of the time change process are known in closed form, then the pricing of options can be accomplished by inverting the Fourier transform, which can be done efficiently using the fast Fourier transform (FFT) algorithm. The combination of this analytical and computational tractability and the richness of possible process behaviors (continuous dynamics, as well as jumps of finite activity or infinite activity) make this class of models attractive for a wide range of financial engineering applications. This chapter surveys both the theory and empirical results.

In Chapter 4 "Pricing with Wishart Risk Factors," Christian Gourieroux and Razvan Sufana survey asset pricing based on risk factors that follow a Wishart process. The class of Wishart models can be thought of as multi-factor extensions of affine stochastic volatility models, which model a stochastic variance-covariance matrix as a matrix-valued stochastic process. As for the standard affine processes, the conditional Laplace transforms can be derived in closed form for Wishart processes. This chapter surveys Wishart processes and their applications to building a wide range of multi-variate models of asset prices with stochastic volatilities and correlations, multi-factor interest rate models, and credit risk models, both in discrete and in continuous time.

In Chapter 5 "Volatility," Federico Bandi and Jeff Russell survey the state of the literature on estimating asset price volatility. They provide a unified framework to understand recent advances in volatility estimation by virtue of microstructure noise contaminated asset price data and transaction cost evaluation. The emphasis is on recently proposed identification procedures that rely on asset price data sampled at high frequency. Volatility is the key factor that determines option prices, and, as such, better understanding of volatility is of key interest in options pricing.

In Chapter 6 "Spectral Methods in Derivatives Pricing," Vadim Linetsky surveys a problem of valuing a (possibly defaultable) derivative asset contingent on the underlying economic state modeled as a Markov process. To gain analytical and computational tractability both in order to estimate the model from empirical data and to compute the prices of derivative assets, financial models in applications are often Markovian. In applications, it is important to have a tool kit of analytically tractable Markov processes with known transition semigroups that lead to closed-form expressions for prices of derivative assets. The spectral expansion method is a powerful approach to generate analytical solutions for Markovian problems. This chapter surveys the spectral method in general, as well as those classes of Markov processes for which the spectral representation can be obtained in closed form, thus generating closed form solutions to Markovian derivative pricing problems.

When underlying financial variables follow a Markov jump-diffusion process, the value function of a derivative security satisfies a partial integro-differential equation (PIDE) for European-style exercise or a partial integro-differential variational inequality (PIDVI) for American-style exercise. Unless the Markov process has a special structure (as discussed in Chapter 6), analytical solutions are generally not available, and it is necessary to solve the PIDE or the PIDVI numerically. In Chapter 7 "Variational Methods in Derivatives Pricing," Liming Feng, Pavlo Kovalov, Vadim Linetsky and Michael Marcozzi survey a computational method for the valuation of options in jump-diffusion models based on converting the PIDE or PIDVI to a variational (weak) form, discretizing the weak formulation spatially by the Galerkin finite element method to obtain a system of ODEs, and integrating the resulting system of ODEs in time.

In Chapter 8 "Discrete Path-Dependent Options," Steven Kou surveys recent advances in the development of methods to price discrete path-dependent options, such as discrete barrier and lookback options that sample the minimum or maximum of the asset price process at discrete time intervals, including discrete barrier and lookback options. A wide array of option pricing methods are surveyed, including convolution methods, asymptotic expansion methods, and methods based on Laplace, Hilbert and fast Gauss transforms.

Part III surveys interest rate and credit risk models and derivatives. In Chapter 9 "Topics in Interest Rate Theory" Tomas Bjork surveys modern interest rate theory. The chapter surveys both the classical material on the Heath–Jarrow–Morton forward rate modeling framework and on the LIBOR market models popular in market practice, as well as a range of recent advances in the interest rate modeling literature, including the geometric interest rate theory (issues of consistency and existence of finite-dimensional state space realizations), and potentials and positive interest models.

Chapters 10 and 11 survey the state-of-the-art in modeling portfolio credit risk and multi-name credit derivatives. In Chapter 10 "Computational Aspects of Credit Risk," Paul Glasserman surveys modeling and computational issues associated with portfolio credit risk. A particular focus is on the problem of calculating the loss distribution of a portfolio of credit risky assets, such as corporate bonds or loans. The chapter surveys models of dependence, including structural credit risk models, copula models, the mixed Poisson model, and associated computational techniques, including recursive convolution, transform inversion, saddlepoint approximation, and importance sampling for Monte Carlo simulation.

In Chapter 11 "Valuation of Basket Credit Derivatives in the Credit Migrations Environment," Tomasz Bielecki, Stephane Crepey, Monique Jeanblanc and Marek Rutkowski present methods to value and hedge basket credit derivatives (such as collateralized debt obligations (CDO) tranches and *n*th to default swaps) and portfolios of credit risky debt. The chapter presents methods for modeling dependent credit migrations of obligors among credit classes and, in particular, dependent defaults. The focus is on specific classes of Markovian models for which computations can be carried out.

Part IV surveys incomplete markets theory and applications. In incomplete markets, dynamic hedging and perfect replication of derivative securities break down and derivatives are no longer redundant assets that can be manufactured via dynamic trading in the underlying primary securities. In Chapter 12 "Incomplete Markets," Jeremy Staum surveys, compares and contrasts many proposed approaches to pricing and hedging derivative securities in incomplete markets, from the perspective of an over-the-counter derivatives market maker operating in an incomplete market. The chapter discusses a wide range of methods, including indifference pricing, good deal bounds, marginal pricing, and minimum-distance pricing measures.

In Chapter 13 "Option Pricing: Real and Risk-Neutral Distributions," George Constantinides, Jens Jackwerth, and Stylianos Perrakis examine the pricing of options in incomplete and imperfect markets in which dynamic trading breaks down either because the market is incomplete or because it is imperfect due to trading costs, or both. Market incompleteness renders the risk-neutral probability measure nonunique and allows one to determine option prices only within some lower and upper bounds. Moreover, in the presence of trading costs, the dynamic replicating strategy does not exist. The authors examine modifications of the theory required to accommodate incompleteness and trading costs, survey testable implications of the theory for option prices, and survey empirical evidence in equity options markets.

In Chapter 14 "Total Risk Minimization Using Monte Carlo Simulation," Thomas Coleman, Yuying Li, and Maria-Cristina Patron study options hedging strategies in incomplete markets. While in an incomplete market it is generally impossible to replicate an option exactly, total risk minimization chooses an optimal self-financing strategy that best approximates the option payoff by its terminal value. Total risk minimization is a computationally challenging dynamic stochastic programming problem. This chapter presents computational approaches to tackle this problem.

In Chapter 15 "Queueing Theoretic Approaches to Financial Price Fluctuations," Erhan Bayraktar, Ulrich Horst, and Ronnie Sircar survey recent research on agent-based market microstructure models. These models of financial prices are based on queueing-type models of order flows and are capable of explaining many stylized features of empirical data, such as herding behavior, volatility clustering, and fat tailed return distributions. In particular, the chapter examines models of investor inertia, providing a link with behavioral finance.

Part V "Risk Management" contains chapters concerned with risk measurement and its application to capital allocation, liquidity risk, and actuarial risk. In Chapter 16 "Economic Credit Capital Allocation and Risk Contributions," Helmut Mausser and Dan Rosen provide a practical overview of risk measurement and management process, and in particular the measurement of economic capital (EC) contributions and their application to capital allocation. EC acts as a buffer for financial institutions to absorb large unexpected losses, thereby protecting depositors and other claim holders. Once the amount of EC has been determined, it must be allocated among the various components of the portfolio (e.g., business units, obligors, individual transactions). This chapter provides an overview of the process of risk measurement, its statistical and computational challenges, and its application to the process of risk management and capital allocation for financial institutions.

In Chapter 17 "Liquidity Risk and Option Pricing Theory," Robert Jarrow and Phillip Protter survey recent research advances in modeling liquidity risk and including it into asset pricing theory. Classical asset pricing theory assumes that investors' trades have no impact on the prices paid or received. In reality, there is a quantity impact on prices. The authors show how to extend the classical arbitrage pricing theory and, in particular, the fundamental theorems of asset pricing, to include liquidity risk. This is accomplished by studying an economy where the security's price depends on the trade size. An analysis of the theory and applications to market data are presented.

In Chapter 18 "Financial Engineering: Applications in Insurance," Phelim Boyle and Mary Hardy provide an introduction to the insurance area, the oldest branch of risk management, and survey financial engineering applications in insurance. The authors compare the actuarial and financial engineering approaches to risk assessment and focus on the life insurance applications in particular. Life insurance products often include an embedded investment component, and thus require the merging of actuarial and financial risk management tools of analysis.

Part VI is devoted to portfolio optimization. In Chapter 19 "Dynamic Portfolio Choice and Risk Aversion," Costis Skiadas surveys optimal consumption and portfolio choice theory, with the emphasis on the modeling of risk aversion given a stochastic investment opportunity set. Dynamic portfolio choice theory was pioneered in Merton's seminal work, who assumed that the investor maximizes time-additive expected utility and approached the problem using the Hamilton–Jacobi–Bellman equation of optimal control theory. This chapter presents a modern exposition of dynamic portfolio choice theory from a more advanced perspective of recursive utility. The mathematical tools include backward stochastic differential equations (BSDE) and forward–backward stochastic differential equations (FBSDE).

In Chapter 20 “Optimization Methods in Dynamic Portfolio Management,” John Birge describes optimization algorithms and approximations that apply to dynamic discrete-time portfolio models including consumption-investment problems, asset-liability management, and dynamic hedging policy design. The chapter develops an overall structure to the many methods that have been proposed by interpreting them in terms of the form of approximation used to obtain tractable models and solutions. The chapter includes the relevant algorithms associated with the approximations and the role that portfolio problem structure plays in enabling efficient implementation.

In Chapter 21 “Simulation Methods for Optimal Portfolios,” Jerome Detemple, Rene Garcia and Marcel Rindisbacher survey and compare Monte Carlo simulation methods that have recently been proposed for the computation of optimal portfolio policies. Monte Carlo simulation is the approach of choice for high-dimensional problems with large number of underlying variables. Simulation methods have recently emerged as natural candidates for the numerical implementation of optimal portfolio rules in high-dimensional portfolio choice models. The approaches surveyed include the Monte Carlo Malliavin derivative method, the Monte Carlo covariation method, the Monte Carlo regression method, and the Monte Carlo finite difference method. The mathematical tools include Malliavin’s stochastic calculus of variations, a brief survey of which is included in the chapter.

In Chapter 22 “Duality Theory and Approximate Dynamic Programming for Pricing American Options and Portfolio Optimization,” Martin Haugh and Leonid Kogan describe how duality and approximate dynamic programming can be applied to construct approximate solutions to American option pricing and portfolio optimization problems when the underlying state space is high-dimensional. While it has long been recognized that simulation is an indispensable tool in financial engineering, it is only recently that simulation has begun to play an important role in control problems in financial engineering. This chapter surveys recent advances in applying simulation to solve optimal stopping and portfolio optimization problems.

In Chapter 23 “Asset Allocation with Multivariate Non-Gaussian Returns,” Dilip Madan and Ju-Yi Yen consider a problem of optimal investment in assets with non-Gaussian returns. They present and back test an asset allocation procedure that accounts for higher moments in investment returns. The procedure is made computationally efficient by employing a signal processing technique known as independent component analysis (ICA) to identify long-tailed independent components in the vector of asset returns. The multivariate portfolio allocation problem is then reduced to univariate problems of component investment. They further assume that the ICs follow the variance gamma (VG) Levy processes and build a multivariate VG portfolio and analyze empirical results of the optimal investment strategy in this setting and compare it with the classical mean–variance Gaussian setting.

In Chapter 24 “Large Deviation Techniques and Financial Applications,” Phelim Boyle, Shui Feng and Weidong Tian survey recent applications of large deviation techniques in portfolio management (establishing portfolio selection criteria, construction of performance indexes), risk management (estimation of large credit portfolio losses that occur in the tail of the distribution), Monte Carlo simulation to better simulate rare events for risk management and asset pricing, and incomplete markets models (estimation of the distance of an incomplete model to a benchmark complete model). A brief survey of the mathematics of large deviations is included in the chapter.

A number of important topics that have recently been extensively surveyed elsewhere were not included in the Handbook. Statistical estimation of stochastic models in finance is an important area that has received limited attention in this volume, with the exception of the focused chapter on volatility. Recent advances in this area are surveyed in the forthcoming Handbook of Financial Econometrics edited by Ait-Sahalia and Hansen (2007). In the coverage of credit risk the Handbook is limited to surveying recent advances in multi-name credit portfolios and derivatives in Chapters 10 and 11, leaving out single-name credit models. The latter have recently been extensively surveyed in monographs Bielecki and Rutkowski (2002), Duffie and Singleton (2003), and Lando (2004). The coverage of Monte Carlo simulation methods is limited to applications to multi-name credit portfolios in Chapter 10, to hedging in incomplete markets in Chapter 14, and to portfolio optimization in Chapters 21 and 22. Monte Carlo simulation applications in derivatives valuation have recently been surveyed by Glasserman (2004). Our coverage of risk measurement and risk management is limited to Chapters 16, 17 and 18 on economic capital allocation, liquidity risk, and insurance risk, respectively. We refer the reader to the recently published monograph McNeil et al. (2005) for extensive treatments of Value-at-Risk and related topics. Modeling energy and commodity markets and derivatives is an important area of financial engineering not covered in the Handbook. We refer the reader to the recent monographs by Eydeland and Wolyniec (2002) and Geman (2005) for extensive surveys of energy and commodity markets.

References

Ait-Sahalia, Y., Hansen, L.P. (Eds.) (2007). *Handbook of Financial Econometrics*. Elsevier, Amsterdam, in press.

Bielecki, T., Rutkowski, M. (2002). *Credit Risk: Modeling, Valuation and Hedging*. Springer.

Bjork, T. (2004). *Arbitrage Theory in Continuous Time*, second ed. Oxford University Press, Oxford, UK.

Duffie, D. (2001). *Dynamic Asset Pricing Theory*, third ed. Princeton University Press, Princeton, NJ.

Duffie, D., Singleton, K. (2003). *Credit Risk*. Princeton University Press, Princeton, NJ.

Eydeland, A., Wolyniec, K. (2002). *Energy and Power Risk Management: New Developments in Modeling, Pricing and Hedging*. John Wiley & Sons, New Jersey.

Jacod, J., Shiryaev, A.N. (2002). *Limit Theorems for Stochastic Processes*, second ed. Springer.

Jeanblanc, M., Yor, M., Chesney, M. (2007). *Mathematical Methods for Financial Markets*. Springer, in press.

Geman, H. (2005). *Commodities and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals, and Energy*. John Wiley & Sons, New Jersey.

Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. Springer.
Hull, J. (2005). *Options, Futures, and Other Derivatives*, sixth ed. Prentice Hall.
Karatzas, I., Shreve, S.E. (1991). *Brownian Motion and Stochastic Calculus*, second ed. Springer.
Karatzas, I., Shreve, S.E. (2001). *Methods of Mathematical Finance*. Springer.
Klebaner, F.C. (2005). *Introduction to Stochastic Calculus with Applications*, second ed. Imperial College Press.
Lando, D. (2004). *Credit Risk Modeling*. Princeton University Press, Princeton, NJ.
McDonald, R.L. (2005). *Derivatives Markets*, second ed. Addison–Wesley.
McNeil, A.J., Frey, R., Embrechts, P. (2005). *Quantitative Risk Management: Concepts, Techniques, and Tools*. Princeton University Press, Princeton, NJ.
Protter, P.E. (2005). *Stochastic Integration and Differential Equations*, second ed. Springer.
Revuz, D., Yor, M. (1999). *Continuous Martingales and Brownian Motion*, third ed. Springer.
Shreve, S.E. (2004). *Stochastic Calculus for Finance II: Continuous-time Models*. Springer.

J.R. Birge and V. Linetsky (Eds.), *Handbooks in OR & MS, Vol. 15*

DOI: 10.1016/S0927-0507(07)15001-5

Chapter 1

An Introduction to Financial Asset Pricing

Robert A. Jarrow

Johnson Graduate School of Management, Cornell University, Ithaca, NY 14853, USA
E-mail: raj15@cornell.edu

Philip Protter[†]

ORIE – 219 Rhodes Hall, Cornell University, Ithaca, NY 14853-3801, USA
E-mail: pep4@cornell.edu

Abstract

We present an introduction to Mathematical Finance Theory, covering the basic issues as well as some selected special topics.

1 Introduction

Stock markets date back to at least 1531, when one was started in Antwerp, Belgium.[1] Today there are over 150 stock exchanges (see Wall Street Journal, 2000). The mathematical modeling of such markets however, came hundreds of years after Antwerp, and it was embroiled in controversy at its beginnings. The first attempt known to the authors to model the stock market using probability is due to L. Bachelier in Paris about 1900. Bachelier's model was his thesis, and it met with disfavor in the Paris mathematics community, mostly because the topic was not thought worthy of study. Nevertheless we now realize that Bachelier essentially modeled Brownian motion five years before the 1905 paper of Einstein [albeit twenty years after T.N. Thiele of Copenhagen (Hald, 1981)] and of course decades before Kolmogorov gave mathematical legitimacy to the subject of probability theory. Poincaré was hostile to Bachelier's thesis, remarking that his thesis topic was "somewhat remote from those our candidates are in the habit of treating" and Bachelier ended up spending his career in Besançon, far from the French capital. His work was then ignored and forgotten for some time.

[†] Supported in part by NSF grant DMS-0202958 and NSA grant MDA-904-03-1-0092

[1] For a more serious history than this thumbnail sketch, we refer the reader to the recent article (Jarrow and Protter, 2004).

Following work by Cowles, Kendall and Osborne, it was the renowned statistician Savage who re-discovered Bachelier's work in the 1950's, and he alerted Paul Samuelson (see Bernstein, 1992, pp. 22–23). Samuelson further developed Bachelier's model to include stock prices that evolved according to a geometric Brownian motion, and thus (for example) always remained positive. This built on the earlier observations of Cowles and others that it was the increments of the logarithms of the prices that behaved independently.

The development of financial asset pricing theory over the 35 years since Samuelson's 1965 article (Samuelson, 1965) has been intertwined with the development of the theory of stochastic integration. A key breakthrough occurred in the early 1970's when Black, Scholes, and Merton (Black and Scholes, 1973; Merton, 1973) proposed a method to price European options via an explicit formula. In doing this they made use of the Itô stochastic calculus and the Markov property of diffusions in key ways. The work of Black, Merton, and Scholes brought order to a rather chaotic situation, where the previous pricing of options had been done by intuition about ill defined market forces. Shortly after the work of Black, Merton, and Scholes, the theory of stochastic integration for semimartingales (and not just Itô processes) was developed in the 1970's and 1980's, mostly in France, due in large part to P.A. Meyer of Strasbourg and his collaborators. These advances in the theory of stochastic integration were combined with the work of Black, Scholes, and Merton to further advance the theory, by Harrison and Kreps (1979) and Harrison and Pliska (1981) in seminal articles published in 1979 and 1980. In particular they established a connection between complete markets and martingale representation. Much has happened in the intervening two decades, and the subject has attracted the interest and curiosity of a large number of researchers and of course practitioners. The interweaving of finance and stochastic integration continues today. This article has the hope of introducing researchers to the subject at more or less its current state, for the special topics addressed here. We take an abstract approach, attempting to introduce simplifying hypotheses as needed, and we signal when we do so. In this way it is hoped that the reader can see the underlying structure of the theory.

The subject is much larger than the topics of this article, and there are several books that treat the subject in some detail (e.g., Duffie, 2001; Karatzas and Shreve, 1998; Musiela and Rutkowski, 1997; Shiryaev, 1999), including the new lovely book by Shreve (2004). Indeed, the reader is sometimes referred to books such as (Duffie, 2001) to find more details for certain topics. Otherwise references are provided for the relevant papers.

2 Introduction to derivatives and arbitrage

Let $S = (S_t)_{0\leqslant t\leqslant T}$ represent the (nonnegative) price process of a risky asset (e.g., the price of a stock, a commodity such as "pork bellies," a currency exchange rate, etc.). The present is often thought of as time $t = 0$. One is interested in the unknown price at some future time T, and thus S_T constitutes a "risk." For example, if an American company contracts at time $t = 0$ to deliver machine parts to Germany at time T, then the unknown price of Euros at time T (in dollars) constitutes a risk for that company. In order to reduce this risk, one may use "derivatives": one can purchase – at time $t = 0$ – the right to buy Euros at time T at a price that is fixed at time 0, and which is called the "strike price." If the price of Euros is higher at time T, then one exercises this right to buy the Euros, and the risk is removed. This is one example of a derivative, called a *call option*.

A *derivative* is any financial security whose value is derived from the price of another asset, financial security, or commodity. For example, the call option just described is a derivative because its value is derived from the value of the underlying Euro. In fact, almost all traded financial securities can be viewed as derivatives.[2] Returning to the *call option* with strike price K, its payoff at time T can be represented mathematically as

$$C = (S_T - K)^+$$

where $x^+ = \max(x, 0)$. Analogously, the payoff to a *put option* with strike price K at time T is

$$P = (K - S_T)^+$$

and this corresponds to the right to *sell* the security at price K at time T. These are two simple examples of derivatives, called a *European call option* and *European put option*, respectively. They are clearly related, and we have

$$S_T - K = (S_T - K)^+ - (K - S_T)^+.$$

This simple equality leads to a relationship between the price of a call option and the price of a put option known as *put–call parity*. We return to this in Section 3.7.

We can also use these two simple options as building blocks for more complicated derivatives. For example, if

$$V = \max(K, S_T)$$

then

$$V = S_T + (K - S_T)^+ = K + (S_T - K)^+.$$

[2] A fun exercise is to try to think of a financial security whose value does not depend on the price of some other asset or commodity. An example is a precious metal itself, like gold, trading as a commodity. But, gold stocks are a derivative as well as gold futures!

More generally, if $f : \mathbb{R}_+ \to \mathbb{R}_+$ is convex then we can use the well-known representation

$$f(x) = f(0) + f'_+(0)x + \int_0^\infty (x-y)^+ \mu(\mathrm{d}y), \tag{1}$$

where $f'_+(x)$ is the right continuous version of the (mathematical) derivative of f, and μ is a positive measure on $\mathbb{R}$ with $\mu = f''$, where the mathematical derivative is in the generalized function sense. In this case if

$$V = f(S_T)$$

is our financial derivative, then V is effectively a portfolio consisting of a continuum of European call options, using (1) (see Brown and Ross, 1991):

$$V = f(0) + f'_+(0)S_T + \int_0^\infty (S_T - K)^+ \mu\,(\mathrm{d}K).$$

For the derivatives discussed so far, the derivative's time T value is a random variable of the form $V = f(S_T)$, that is, a function of the value of S at one fixed and prescribed time T. One can also consider derivatives of the form

$$V = F(S)_T = F(S_t; 0 \leqslant t \leqslant T)$$

which are functionals of the paths of S. For example if S has càdlàg paths (càdlàg is a French acronym for "right continuous with left limits") then $F : D \to \mathbb{R}_+$, where D is the space of functions $f : [0, T] \to \mathbb{R}_+$ which are right continuous with left limits.

If the derivative's value depends on a decision of its holder at only the expiration time T, then they are considered to be of the *European type*, although their analysis for pricing and hedging is more difficult than for simple European call and put options. The decision in the case of a call or put option is whether to exercise the right to buy or sell, respectively.[3] Hence, such decisions are often referred to as *exercise* decisions.

An *American type derivative* is one in which the holder has a decision to make with respect to the security at any time before or at the expiration time. For example, an *American call option* allows the holder to buy the security at a striking price K not only at time T (as is the case for a European call option), but at any time between times $t = 0$ and time T. (It is this type of option that is listed, for example, in the "Listed Options Quotations" in the Wall Street Journal.) Deciding when to exercise such an option is complicated. A strategy for exercising an American call option can be represented mathematically by

[3] This decision is explicitly represented by the maximum operator in the payoff of the call and put options.

a *stopping rule* τ. (That is, if $\mathbb{F} = (\mathcal{F}_t)_{t\geqslant 0}$ is the underlying filtration of S then $\{\tau \leqslant t\} \in \mathcal{F}_t$ for each t, $0 \leqslant t \leqslant T$.) For a given τ, the American call's payoff at time $\tau(\omega)$ is

$$C(\omega) = \big(S_{\tau(\omega)}(\omega) - K\big)^+.$$

We now turn to the *pricing of derivatives*. Let C be a random variable in $\mathcal{F}_T$ representing the time T payoff to a derivative. Let V_t be its *value* (or price) at time t. What then is V_0? From a traditional point of view based on an analysis of fair (gambling) games, classical probability tells us that[4]

$$V_0 = E\{C\}. \tag{2}$$

One should pay the expected payoff of participating in the gamble. But, one should also discount for the time value of money (the interest forgone or earned) and assuming a fixed spot interest rate r, one would have

$$V_0 = E\left\{\frac{C}{(1+r)^T}\right\} \tag{3}$$

instead of (2). Surprisingly, this value is not correct, because it ignores the impact of risk aversion on the part of the purchaser. For simplicity, we will take $r = 0$ and then show why the obvious price given in (2) does not work (!).

Let us consider a simple binary example. At time $t = 0$, 1 Euro = \$1.15. We assume at time $t = T$ that the Euro will be worth either \$0.75 or \$1.45. Let the probability that it goes up to \$1.45 be p and the probability that it goes down be $1 - p$.

Consider a European call option with exercise price $K = \$1.15$. That is, $C = (S_T - \$1.15)^+$, where $S = (S_t)_{0\leqslant t\leqslant T}$ is the price of one Euro in US dollars. The classical rules for calculating probabilities dating back to Huygens and Bernoulli give a fair price of C as

$$E\{C\} = (1.45 - 1.15)p = (0.30)p.$$

For example if $p = 1/2$ we get $V_0 = 0.15$.

The *Black–Scholes* method[5] for calculating the option's price, however, is quite different. We first replace p with a new probability p^* that (in the absence of interest rates) makes the security price $S = (S_t)_{t=0,T}$ a martingale. Since this is a two-step process, *we need only to choose p^* so that S has a constant expectation under P^*, the probability measure implied by the choice of p^**. Since

[4] This assumes, implicitly, that there are no intermediate cash flows from holding the derivative security.

[5] The "Black–Scholes method" dates back to the fundamental and seminal articles (Black and Scholes, 1973) and (Merton, 1973) of 1973, where partial differential equations were used; the ideas implicit in that (and subsequent) articles are now referred to as the Black–Scholes methods. More correctly, it should be called the Black–Merton–Scholes method. M.S. Scholes and R. Merton received the Nobel prize in economics for (Black and Scholes, 1973; Merton, 1973), and related work. (F. Black died and was not able to share in the prize.)

$S_0 = 1.15$, we need

$$E^*\{S_T\} = 1.45p^* + (1 - p^*)0.75 = 1.15, \tag{4}$$

where E^* denotes mathematical expectation with respect to the probability measure P^* given by $P^*(\text{Euro} = \$1.45 \text{ at time } T) = p^*$, and $P^*(\text{Euro} = \$0.75 \text{ at time } T) = 1 - p^*$. Solving for p^* gives

$$p^* = 4/7.$$

We get now

$$V_0 = E^*\{C\} = (0.30)p^* = \frac{6}{35} \simeq 0.17.$$

The change from p to p^* seems arbitrary. But, there is an *economics* argument to justify it. This is where the economics concept of *no arbitrage opportunities* changes the usual intuition dating back to the 16th and 17th centuries.

Suppose, for example, at time $t = 0$ you sell the call option, giving the buyer of the option the right to purchase 1 Euro at time T for \$1.15. He then gives you the price $v(C)$ of the option. Again we assume $r = 0$, so there is no cost to borrow money. You can then follow a safety strategy to prepare for the option you sold, as follows (calculations are to two decimal places):

Action at time $t = 0$	*Result*
Sell the option at price $v(C)$	$+v(C)$
Borrow $\$\frac{9}{28}$	$+0.32$
Buy $\frac{3}{7}$ euros at \$1.15	-0.49

The balance at time $t = 0$ is $v(C) - 0.17$.

At time T there are two possibilities:

What happens to the euro	*Result*
The euro has risen:	
Option is exercised	-0.30
Sell $\frac{3}{7}$ euros at \$1.45	$+0.62$
Pay back loan	-0.32
End balance:	0
The euro has fallen:	
Option is worthless	0
Sell $\frac{3}{7}$ euros at \$0.75	$+0.32$
Pay back loan	-0.32
End balance:	0

Since the balance at time T is zero in both cases, the balance at time 0 should also be 0; therefore we must have $v(C) = 0.17$. Indeed any price other than $v(C) = 0.17$ would allow either the option seller or buyer to make a sure profit

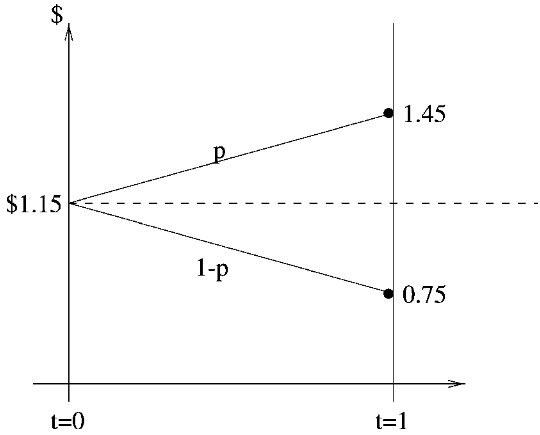

Fig. 1. Binary schematic.

without any risk. Such a sure profit with no risk is called an *arbitrage opportunity* in economics, and it is a standard assumption that such opportunities do not exist. (Of course if they were to exist, market forces would, in theory, quickly eliminate them.)

Thus we see that – at least in the case of this simple example – that the "no arbitrage price" of the derivative C is not $E\{C\}$, but rather it must be $E^*\{C\}$. We emphasize that this is contrary to our standard intuition based on fair games, since P is the probability measure governing the true laws of chance of the security, while P^* is an artificial construct.

Remark 1 *(Heuristic Explanation)*. We offer two comments here. The first is that the change of probability measures from P to P^* is done with the goal of keeping the expectation constant. (See Equation (4).) It is this property of constant expectation of the price process which excludes the possibility of arbitrage opportunities, when the price of the derivative is chosen to be the expectation under P^*. Since one can have many different types of processes with constant expectation, one can ask: what is the connection to martingales? The answer is that a necessary and sufficient condition for a process $M = (M_t)_{t \geqslant 0}$ to be a uniformly martingale is that $E(M_\tau) = E(M_0)$ for every stopping time τ. The key here is that it is required for every stopping time, and not just for fixed times. In words, the price process must have constant expectation at all random times (stopping times) under a measure P^* in order for the expectation of the contingent claim under P^* to be an arbitrage free price of the claim.

The second comment refers to Figure 1 (binary schematic). Intuition tells us that as $p \nearrow 1$, that the price of a call or put option must change, since as it becomes almost certain that the price will go up, the call might be worth less (or more) to the purchaser. And yet our no arbitrage argument tells us that it cannot, and that p^* is fixed for all p, $0 < p < 1$. How can this be? An economics explanation is that if one lets p increase to 1, one is implicitly perverting the economy. In essence, this perversion of the economy *a fortiori* reflects changes in participants' levels of *risk aversion*. If the price can change to only two prices, and it is near certain to go up, how can we keep the current price fixed at \$1.15? Certainly this change in perceived probabilities should affect the current price too. In order to increase p towards 1 and simultaneously keep the current price fixed at \$1.15, we are forced to assume that people's behavior has changed, and either they are very averse to even a small potential loss (the price going down to \$0.75), or they now value much less the near certain potential price increase to \$1.45.

This simple binary example can do more than illustrate the idea of using the lack of arbitrage to determine a price. We can also use it to approximate some continuous time models for the evolution of an asset's price. We let the time interval become small (Δt), and we let the binomial model already described become a recombinant tree, which moves up or down to a neighboring node at each time "tick" Δt. For an actual time "tick" of interest of length say δ, we can have the price go to 2^n possible values for a given n, by choosing Δt small enough in relation to n and δ. Thus for example if the continuous time process follows geometric Brownian motion:

$$\mathrm{d}S_t = \sigma S_t \,\mathrm{d}B_t + \mu S_t \,\mathrm{d}t$$

(as is often assumed in practice); and if the security price process S has value $S_t = s$, then it will move up or down at the next tick Δt to

$$s\exp(\mu\Delta t + \sigma\sqrt{\Delta t}\,) \quad \text{if up;} \qquad s\exp(\mu\Delta t - \sigma\sqrt{\Delta t}\,) \quad \text{if down;}$$

with p being the probability of going up or down (here take $p = \frac{1}{2}$). Thus for a time t, if $n = \frac{t}{\Delta t}$, we get

$$S_t^n = S_0 \exp\left(\mu t + \sigma\sqrt{t}\left(\frac{2X_n - n}{\sqrt{n}}\right)\right),$$

where X_n counts the number of jumps up. By the central limit theorem S_t^n converges, as n tends to infinity, to a log normal process $S = (S_t)_{t\geqslant 0}$; that is, $\log S_t$ has a normal distribution with mean $\log(S_0 + \mu t)$ and variance $\sigma^2 t$.

Next we use the absence of arbitrage to change p from $\frac{1}{2}$ to p^*. We find p^* by requiring that $E^*\{S_t\} = E^*\{S_0\}$, and we get p^* approximately equal to

$$p^* = \frac{1}{2}\left(1 - \sqrt{\Delta t}\left(\frac{\mu + \frac{1}{2}\sigma^2}{\sigma}\right)\right).$$

Thus under P^*, X_n is still binomial, but now it has mean np^* and variance $np^*(1 - p^*)$. Therefore $\left(\frac{2X_n - n}{\sqrt{n}}\right)$ has mean $-\sqrt{t}(\mu + \frac{1}{2}\sigma^2)/\sigma$ and a variance which converges to 1 asymptotically. The central limit theorem now implies that S_t converges as n tends to infinity to a log normal distribution: $\log S_t$ has mean $\log S_0 - \frac{1}{2}\sigma^2 t$ and variance $\sigma^2 t$. Thus

$$S_t = S_0 \exp\left(\sigma\sqrt{t}Z - \frac{1}{2}\sigma^2 t\right),$$

where Z is $N(0, 1)$ under P^*. This is known as the "binomial approximation." The binomial approximation can be further used to derive the Black–Scholes equations, by taking limits, leading to simple formulas in the continuous case. (We present these formulas in Section 3.10.) A simple derivation can be found in Cox et al. (1979) or in Duffie (2001, Chapter 12B, pp. 294–299).

3 The core of the theory

3.1 Basic definitions

Throughout this section we will assume that we are given an underlying probability space $(\Omega, \mathcal{F}, \mathbb{F}, P)$, where $\mathbb{F} = (\mathcal{F}_t)_{t\geqslant 0}$. We further assume $\mathcal{F}_s \subset \mathcal{F}_t$ if $s < t$; $\mathcal{F}_0$ contains all the P-null sets of $\mathcal{F}$; and also that $\bigcap_{s>t}\mathcal{F}_s \equiv \mathcal{F}_{t+} = \mathcal{F}_t$ by hypothesis. This last property is called the *right continuity of the filtration*. These hypotheses, taken together, are known as the *usual hypotheses*. (When the usual hypotheses hold, one knows that every martingale has a version which is càdlàg, one of the most important consequences of these hypotheses.)

3.2 The price process

We let $S = (S_t)_{t\geqslant 0}$ be a semimartingale[6] which will be the *price process* of a risky security. For simplicity, after the initial purchase or sale, we assume that the security has no cash flows associated with it (for example, if the security is a common stock, we assume that there are no dividends paid). This assumption is easily relaxed, but its relaxation unnecessarily complicates the notation and explanation, so we leave it to outside references.

3.3 Spot interest rates

Let r be a fixed spot rate of interest. If one invests 1 dollar at rate r for one year, at the end of the year one has $1 + r$ dollars. If interest is paid at n evenly spaced times during the year and compounded, then at the end of the year one has $(1 + \frac{r}{n})^n$. This leads to the notion of an *interest rate r compounded*

[6] One *definition of a semimartingale* is a process S that has a decomposition $S = M + A$, with M a local martingale and A an adapted process with càdlàg paths of finite variation on compacts. (See Protter, 2005.)

continuously:

$$\lim_{n\to\infty}\left(1+\frac{r}{n}\right)^n = e^r$$

or, for a fraction t of the year, one has \$$e^{rt}$ after t units of time for a spot interest rate r compounded continuously.

We define

$$R(t) = e^{rt};$$

then R satisfies the ODE (ODE abbreviates ordinary differential equation)

$$\mathrm{d}R(t) = rR(t)\,\mathrm{d}t; \quad R(0) = 1. \tag{5}$$

Using the ODE (5) as a basis for interest rates, one can treat a variable interest rate $r(t)$ as follows: ($r(t)$ can be random: that is $r(t) = r(t, \omega)$)[7]:

$$\mathrm{d}R(t) = r(t)R(t)\,\mathrm{d}t; \quad R(0) = 1 \tag{6}$$

and solving yields $R(t) = \exp(\int_0^t r(s)\,\mathrm{d}s)$. We think of the interest rate process $R(t)$ as the *time t value of a money market account*.

3.4 Trading strategies and portfolios

We will assume as given a risky asset with price process S and a money market account with price process R. Let $(a_t)_{t\geqslant 0}$ and $(b_t)_{t\geqslant 0}$ be our *time t holdings* in the security and the bond, respectively.

We call our holdings of S and R our *portfolio*. Note that for the model to make sense, we must have both the risky asset and the money market account present. When we receive money through the sale of risky assets, we place the cash in the money market account; and when we purchase risky assets, we use the cash from the money market account to pay for the expenditure. The money market account is allowed to have a negative balance.

Definition 1. The *value at time t*[8] *of a portfolio* (a, b) is

$$V_t(a, b) = a_t S_t + b_t R_t.$$

[7] An example is to take $r(t)$ to be a diffusion; one can then make appropriate hypotheses on the diffusion to model the behavior of the spot interest rate.

[8] This concept of value is a commonly used approximation. If one were to liquidate one's risky assets at time t all at once to realize this "value," one would find less money in the savings account, due to liquidity and transaction costs. For simplicity, we are assuming there are no liquidity and transaction costs. Such an assumption is not necessary, however, and we recommend the interested reader to Jarrow and Protter (2007) in this volume.

Now we have our first problem. Later we will want to change probabilities so that $V = (V_t(a, b))_{t \geqslant 0}$ is a martingale. One usually takes the right continuous versions of a martingale, so we want the right side of (4) to be at least càdlàg. Typically this is not a real problem. Even if the process a has no regularity, one can always choose b in such a way that $V_t(a, b)$ is càdlàg.

Let us next define two sigma algebras on the product space $\mathbb{R}_+ \times \Omega$. We recall that we are given an underlying probability space $(\Omega, \mathcal{F}, \mathbb{F}, P)$ with $\mathbb{F} = (\mathcal{F}_t)_{t \geqslant 0}$, satisfying the "usual hypotheses."

Definition 2. Let $\mathbb{L}$ denote the space of left continuous processes whose paths have right limits (càglàd), and which are adapted: that is, $H_t \in \mathcal{F}_t$, for $t \geqslant 0$. The *predictable σ-algebra* $\mathcal{P}$ on $\mathbb{R}_+ \times \Omega$ is

$$\mathcal{P} = \sigma\{H\colon\ H \in \mathbb{L}\}.$$

That is $\mathcal{P}$ is the smallest σ-algebra that makes all of $\mathbb{L}$ measurable.

Definition 3. The *optional σ-algebra* $\mathcal{O}$ on $\mathbb{R}_+ \times \Omega$ is

$$\mathcal{O} = \sigma\{H\colon\ H \text{ is càdlàg and adapted}\}.$$

In general we have $\mathcal{P} \subset \mathcal{O}$. In the case where $B = (B_t)_{t \geqslant 0}$ is a standard Wiener process (or "Brownian motion"), and $\mathcal{F}_t^0 = \sigma(B_s; s \leqslant t)$ and $\mathcal{F}_t = \mathcal{F}_t^0 \vee \mathcal{N}$ where $\mathcal{N}$ are the P-null sets of $\mathcal{F}$, then we have $\mathcal{O} = \mathcal{P}$. In general $\mathcal{O}$ and $\mathcal{P}$ are not equal. Indeed if they are equal, then every stopping time is predictable: that is, there are no totally inaccessible stopping times.[9] Since the jump times of (reasonable) Markov processes are totally inaccessible, any model which contains a Markov process with jumps (such as a Poisson Process) will have $\mathcal{P} \subset \mathcal{O}$, where the inclusion is strict.

Remark on filtration issues. The predictable σ-algebra $\mathcal{P}$ is important because it is the natural σ-field for which stochastic integrals are defined. In the special case of Brownian motion one can use the optional σ-algebra (since they are the same). There is a third σ-algebra which is often used, known as

---

[9] A *totally inaccessible stopping time* is a stopping time that comes with no advance warning: it is a complete surprise. A stopping time T is *totally inaccessible* if whenever there exists a sequence of nondecreasing stopping times $(S_n)_{n \geqslant 1}$ with $\Lambda = \bigcap_{n=1}^{\infty}\{S_n < T\}$, then

$$P\Big(\Big\{w\colon \lim_{n\to\infty} S_n = T\Big\} \cap \Lambda\Big) = 0.$$

A stopping time T is *predictable* if there exists a nondecreasing sequence of stopping times $(S_n)_{n \geqslant 1}$ as above with

$$P\Big(\Big\{w\colon \lim_{n\to\infty} S_n = T\Big\} \cap \Lambda\Big) = 1.$$

Note that the probabilities above need not be only 0 or 1; thus there are in general stopping times which are neither predictable nor totally inaccessible.

the progressively measurable sets, and denoted $\boldsymbol{\pi}$. One has, in general, that $\mathcal{P} \subset \mathcal{O} \subset \boldsymbol{\pi}$; however in practice one gains very little by assuming a process is $\boldsymbol{\pi}$-measurable instead of optional, if – as is the case here – one assumes that the filtration $(\mathcal{F}_t)_{t\geqslant 0}$ is right-continuous (that is $\mathcal{F}_{t+} = \mathcal{F}_t$, all $t \geqslant 0$). The reason is that the primary use of $\boldsymbol{\pi}$ is to show that adapted, right-continuous processes are $\boldsymbol{\pi}$-measurable and in particular that $S_T \in \mathcal{F}_T$ for T a stopping time and S progressive; but such processes are already optional if $(\mathcal{F}_t)_{t\geqslant 0}$ is right continuous. Thus there are essentially no "naturally occurring" examples of progressively measurable processes that are not already optional. An example of such a process, however, is the indicator function $1_G(t)$, where G is described as follows: let $\mathbb{Z} = \{(t, \omega)\colon\ B_t(\omega) = 0\}$. ($B$ is standard Brownian motion.) Then $\mathbb{Z}$ is a perfect (and closed) set on $\mathbb{R}_+$ for almost all ω. For fixed ω, the complement is an open set and hence a countable union of open intervals. $G(\omega)$ denotes the left end-points of these open intervals. One can then show (using the Markov property of B and P.A. Meyer's section theorems) that G is progressively measurable but not optional. In this case note that $1_G(t)$ is zero except for countably many t for each ω, hence $\int 1_G(s)\,\mathrm{d}B_s \equiv 0$. Finally we note that if $a = (a_s)_{s\geqslant 0}$ is progressively measurable, then $\int_0^t a_s\,\mathrm{d}B_s = \int_0^t \dot{a}_s\,\mathrm{d}B_s$, where $\dot{a}$ is the predictable projection of a.[10]

Let us now recall a few details of stochastic integration. First, let S and X be any two càdlàg semimartingales. The integration by parts formula can be used to define the quadratic co-variation of X and S:

$$[X, S]_t = X_t Y_t - \int_0^t X_{s-}\,\mathrm{d}S_s - \int_0^t S_{s-}\,\mathrm{d}X_s.$$

However if a càdlàg, adapted process H is not a semimartingale, one can still give the quadratic co-variation a meaning, by using a limit in probability as the definition. This limit always exists if both H and S are semimartingales:

$$[H, S]_t = \lim_{n\to\infty} \sum_{t_i \in \pi^n[0,t]} (H_{t_{i+1}} - H_{t_i})(S_{t_{i+1}} - S_{t_i}),$$

where $\pi^n[0, t]$ be a sequence of finite partitions of $[0, t]$ with $\lim_{n\to\infty} \mathrm{mesh}(\pi^n) = 0$.

[10] Let H be a bounded, measurable process. (H need not be adapted.) The *predictable projection* of H is the unique predictable process $\dot{H}$ such that

$$\dot{H}_T = E\{H \mid \mathcal{F}_{T-}\} \quad \text{a.s. on } \{T < \infty\}$$

for all predictable stopping times T. Here $\mathcal{F}_{T-} = \sigma\{A \cap \{t < T\};\ A \in \mathcal{F}_t\} \vee \mathcal{F}_0$. For a proof of the existence and uniqueness of $\dot{H}$ see Protter (2005, p. 119).

Henceforth let S be a (càdlàg) semimartingale, and let H be càdlàg and adapted, or alternatively $H \in \mathbb{L}$. Let $H_- = (H_{s-})_{s\geqslant 0}$ denote the left-continuous version of H. (If $H \in \mathbb{L}$, then of course $H = H_-$.) We have:

Theorem 1. *H càdlàg, adapted or $H \in \mathbb{L}$. Then*

$$\lim_{n\to\infty} \sum_{t_i \in \pi^n[0,t]} H_{t_i}(S_{t_{i+1}} - S_{t_i}) = \int_0^t H_{s-}\, \mathrm{d}S_s,$$

with convergence uniform in s on $[0, t]$ in probability.

We remark that it is crucial that we sample H at the left endpoint of the interval $[t_i, t_{i+1}]$. Were we to sample at, say, the right endpoint or the midpoint, then the sums would not converge in general (they converge for example if the quadratic covariation process $[H, S]$ exists); in cases where they do converge, the limit is in general different. Thus while the above theorem gives a pleasing "limit as Riemann sums" interpretation to a stochastic integral, it is not at all a perfect analogy.

The basic idea of the preceding theorem can be extended to bounded predictable processes in a method analogous to the definition of the Lebesgue integral for real-valued functions. Note that

$$\sum_{t_i \in \pi^n[0,t]} H_{t_i}(S_{t_{i+1}} - S_{t_i}) = \int_{0+}^t H_s^n\, \mathrm{d}S_s,$$

where $H_t^n = \sum H_{t_i} 1_{(t_i.t_{i+1}]}$ which is in $\mathbb{L}$; thus these "simple" processes are the building blocks, and since $\sigma(\mathbb{L}) = \mathcal{P}$, it is unreasonable to expect to go beyond $\mathcal{P}$ when defining the stochastic integral.

There is, of course, a maximal space of integrable processes where the stochastic integral is well defined and still gives rise to a semimartingale as the integrated process; without describing it [see any book on stochastic integration such as (Protter, 2005)], we define:

Definition 4. For a semimartingale S we let $L(S)$ denote the space of predictable processes a, where a is integrable with respect to S.

We would like to fix the underlying semimartingale (or vector of semimartingales) S. The process S represents the price process of our risky asset. A way to do that is to introduce the notion of a *model*. We present two versions. The first is the more complete, as it specifies the probability space and the underlying filtration. However it is also cumbersome, and thus we will abbreviate it with the second:

Definition 5. A sextuple $(\Omega, \mathcal{F}, \mathbb{F}, S, L(S), P)$, where $\mathbb{F} = (\mathcal{F}_t)_{t\geqslant 0}$, is called an *asset pricing model*; or more simply, the triple $(S, L(S), P)$ is called a *model*, where the probability space and σ-algebras are implicit: that is, $(\Omega, \mathcal{F}, \mathbb{F})$ is implicit.

We are now ready for a key definition.

A *trading strategy* in the risky asset is a predictable process $a = (a_t)_{t\geqslant 0}$ with $a \in L(S)$; its economic interpretation is that at time t one holds an amount a_t of the asset. We also remark that it is reasonable that a be predictable: a is the trader's holdings at time t, and this is based on information obtained at times strictly before t, but not t itself. Often one has in concrete situations that a is continuous or at least càdlàg or càglàd (left continuous with right limits). (Indeed, it is difficult to imagine a practical trading strategy with pathological path irregularities.) In the case a is adapted and càglàd, then

$$\int_0^t a_s \,\mathrm{d}S_s = \lim_{n\to\infty} \sum_{t_i \in \pi^n[0,t]} a_{t_i} \Delta_i S,$$

where $\pi^n[0, t]$ is a sequence of partitions of $[0, t]$ with mesh tending to 0 as $n \to \infty$; $\Delta_i S = S_{t_{i+1}} - S_{t_i}$; and with convergence in u.c.p. (uniform in time on compacts and converging in probability). Thus inspired by (1) we let

$$G_t = \int_{0+}^t a_s \,\mathrm{d}S_s$$

and G is called the (*financial*) *gain process generated by* a. A trading strategy in the money market account, $b = (b_t)_{t\geqslant 0}$, is defined in an analogous fashion except that we only require that b is optional and $b \in L(R)$. We will call the pair (a, b), as defined above, a trading strategy.

Definition 6. A trading strategy (a, b) is called *self-financing* if

$$a_t S_t + b_t R_t = a_0 S_0 + b_0 R_0 + \int_0^t a_s \,\mathrm{d}S_s + \int_0^t b_s \,\mathrm{d}R_s \tag{7}$$

for all $t \geqslant 0$.

Note that the equality (7) above implies that $a_t S_t + b_t R_t$ is càdlàg.

To justify this definition heuristically, let us assume the spot interest rate is constant and zero: that is, $r = 0$ which implies that $R_t = 1$ for all $t \geqslant 0$, a.s. We can do this by the principle of numéraire invariance; see Section 3.6, later in this article. We then have

$$a_t S_t + b_t R_t = a_t S_t + b_t.$$

Assume for the moment that a and b are semimartingales, and as such let us denote them X and Y, respectively.[11] If at time t we change our position in the risky asset, to be self-financing we must change also the amount in our money market account; thus we need to have the equality:

$$(X_{t+\mathrm{d}t} - X_t)S_{t+\mathrm{d}t} = -(Y_{t+\mathrm{d}t} - Y_t),$$

which is algebraically equivalent to

$$(S_{t+\mathrm{d}t} - S_t)(X_{t+\mathrm{d}t} - X_t) + (X_{t+\mathrm{d}t})S_t = -(Y_{t+\mathrm{d}t} - Y_t),$$

which implies in continuous time:

$$S_{t-}\,\mathrm{d}X_t + \mathrm{d}[S, X]_t = -\mathrm{d}Y_t.$$

Using integration by parts, we get

$$X_t S_t - X_{t-}\,\mathrm{d}S_t = -\mathrm{d}Y_t,$$

and integrating yields the desired equality

$$X_t S_t + Y_t = \int_0^t X_{s-}\,\mathrm{d}S_s + X_0 S_0 + Y_0. \tag{8}$$

Finally we drop the assumption that X and Y are semimartingales, and replacing X_- with a and Y with b, respectively, Eq. (8) becomes

$$a_t S_t + b_t R_t = a_0 S_0 + b_0 + \int_0^t a_s\,\mathrm{d}S_s + (b_t - b_0),$$

as we have in Eq. (7).

The next concept is of fundamental importance. An *arbitrage opportunity* is the chance to make a profit *without risk*. The standard way of modeling this mathematically is as follows:

Definition 7. A model is *arbitrage free* if there does not exist a self-financing trading strategy (a, b) such that $V_0(a, b) = 0$, $V_T(a, b) \geqslant 0$, and $P(V_T(a, b) > 0) > 0$.

[11] Since X is assumed to be a semimartingale, it is right continuous, and thus is not in general predictable; hence when it is the integrand of a stochastic integral we need to replace X_s with X_{s-}, which of course denotes the left continuous version of X.

3.5 The fundamental theorem of asset pricing

In Section 2 we saw that with the "no arbitrage" assumption, at least in the case of a very simple example, a reasonable price of a derivative was obtained by taking expectations and changing from the "true" underlying probability measure, P, to an equivalent one, P^*. More formally, under the assumption that $r = 0$, or equivalently that $R_t = 1$ for all t, the price of a derivative C was not $E\{C\}$ as one might expect, but rather $E^*\{C\}$. (If the process R_t is not constant and equal to one, then we consider the expectation of the discounted claim $E^*\{C/R_T\}$.)

The idea underlying the equivalent change of measure was to find a probability P^* that gave the price process S a constant expectation. This simple insight readily generalizes to more complex stochastic processes. In continuous time, a sufficient condition for the price process $S = (S_t)_{t\geqslant 0}$ to have constant expectation is that it be a martingale. That is, if S is a martingale then the function $t \to E\{S_t\}$ is constant. Actually this property is not far from characterizing martingales. A classic theorem from martingale theory is the following (cf., e.g., Protter, 2005):

Theorem 2. *Let $S = (S_t)_{t\geqslant 0}$ be càdlàg and suppose $E\{S_\tau\} = E\{S_0\}$ for any bounded stopping time τ (and of course $E\{|S_\tau|\} < \infty$). Then S is a martingale.*

That is, if we require constant expectation at stopping times (instead of only at fixed times), then S is a martingale.

Based on this simple pricing example and the preceding theorem, one is lead naturally to the following conjecture.

Conjecture. Let S be a price process on a given space $(\Omega, \mathcal{F}, \mathbb{F}, P)$. Then there are no arbitrage opportunities if and only if there exists a probability P^*, equivalent to P, such that S is a martingale under P^*.

The origins of the preceding conjecture can be traced back to Harrison and Kreps (1979) for the case where $\mathcal{F}_T$ is finite, and later to Dalang et al. (1990) for the case where $\mathcal{F}_T$ is infinite, but time is discrete. Before stating a more rigorous theorem [our version is due to Delbaen and Schachermeyer (1994); see also Delbaen and Schachermayer (1998)], let us examine a needed hypothesis.

We need to avoid problems that arise from the classical doubling strategy in gambling. Here a player bets \$1 at a fair bet. If he wins, he stops. If he loses he next bets \$2. Whenever he wins, he stops, and his profit is \$1. If he continues to lose, he continues to play, each time doubling his bet. This strategy leads to a certain gain of \$1 without risk. However, the player needs to be able to tolerate arbitrarily large losses before he gains his certain profit. Of course, no one has such infinite resources to play such a game. Mathematically one can eliminate this type of problem by requiring trading strategies to give martingales that are bounded below by a constant. Thus the player's resources, while they can be huge, are nevertheless finite and bounded by a nonrandom constant. This leads to the next definition.

Definition 8. Let $\alpha > 0$, and let S be a semimartingale. A predictable trading strategy θ is *α-admissible* if $\theta_0 = 0$, $\int_0^t \theta_s \, \mathrm{d}S_s \geqslant -\alpha$, all $t \geqslant 0$. θ is called *admissible* if there exists $\alpha > 0$ such that θ is α-admissible.

Before we make more definitions, let us recall the basic approach. Suppose θ is an admissible, self-financing trading strategy with $\theta_0 S_0 = 0$ and $\theta_T S_T \geqslant 0$. In the next section we will see that without loss of generality we can neglect the bond or "numéraire" process by a "change of numéraire," so that the self-financing condition reduces to

$$\theta_T S_T = \theta_0 S_0 + \int_0^T \theta_s \, \mathrm{d}S_s.$$

Then if P^* exists such that $\int \theta_s \, \mathrm{d}S_s$ is a martingale, we have

$$E^*\{\theta_T S_T\} = 0 + E^*\left\{ \int_0^T \theta_s \, \mathrm{d}S_s \right\}.$$

In general, if S is continuous then $\int_0^t \theta_s \, \mathrm{d}S_s$ is only a *local* martingale.[12] If S is merely assumed to be a càdlàg semimartingale, then $\int_0^t \theta_s \, \mathrm{d}S_s$ need only be a σ martingale.[13] However if for some reason we do know that it is a true martingale then $E^*\{\int_0^T \theta_s \, \mathrm{d}S_s\} = 0$, whence $E^*\{\theta_T S_T\} = 0$, and since $\theta_T S_T \geqslant 0$ we deduce $\theta_T S_T = 0$, P^* a.s., and since P^* is equivalent to P, we have $\theta_T S_T = 0$ a.s. ($\mathrm{d}P$) as well. This implies no arbitrage exists. The technical part of this argument is to show $\int_0^t \theta_s \, \mathrm{d}S_s$ is a P^* true martingale, and not just a local martingale (see the proof of the Fundamental Theorem that follows). The converse is typically harder: that is, that no arbitrage implies P^* exists. The converse is proved using a version of the Hahn–Banach theorem.

[12] A process M is a *local martingale* if there exists a sequence of stopping times $(T_n)_{n \geqslant 1}$ increasing to ∞ a.s. such that $(M_{t \wedge T_n})_{t \geqslant 0}$ is a martingale for each $n \geqslant 1$.

[13] A process X is a σ martingale if there exists an $\mathbb{R}^d$ valued martingale M and a predictable $\mathbb{R}_+$ valued M-integrable process H such that X is the stochastic integral of H with respect to M. See Protter (2005, pp. 237–239) for more about σ martingales.

Following Delbaen and Schachermayer, we make a sequence of definitions:

$$K_0 = \left\{ \int_0^\infty \theta_s \, \mathrm{d}S_s \;\middle|\; \theta \text{ is admissible and } \lim_{t\to\infty} \int_0^t \theta_s \, \mathrm{d}S_s \text{ exists a.s.} \right\} C_0$$

$$= \{\text{all functions dominated by elements of } K_0\}$$

$$= K_0 - L^0_+, \text{ where } L^0_+ \text{ are positive, finite random variables,}$$

$$K = K_0 \cap L^\infty,$$

$$C = C_0 \cap L^\infty,$$

$$\overline{C} = \text{the closure of } C \text{ under } L^\infty.$$

Definition 9. A semimartingale price process S satisfies

(i) the *no arbitrage* condition if $\mathbb{C} \cap L^\infty_+ = \{0\}$ (this corresponds to no chance of making a profit without risk);

(ii) the *no free lunch with vanishing risk* condition (NFLVR) if $\overline{\mathbb{C}} \cap L^\infty_+ = \{0\}$, where $\overline{\mathbb{C}}$ is the closure of $\mathbb{C}$ in L^∞.

Definition 10. A probability measure P^* is called an *equivalent martingale measure*, or alternatively a *risk neutral probability*, if P^* is equivalent to P, and if under P^* the price process S is a σ martingale.

Clearly condition (ii) implies condition (i). Condition (i) is slightly too restrictive to imply the existence of an equivalent martingale measure P^*. (One can construct a trading strategy of $H_t(\omega) = 1_{\{[0,1]\setminus\mathbb{Q}\times\Omega\}}(t,\omega)$, which means one sells before each rational time and buys back immediately after it; combining H with a specially constructed càdlàg semimartingale shows that (i) does not imply the existence of P^* – see Delbaen and Schachermayer, 1994, p. 511.)

Let us examine then condition (ii). If NFLVR is not satisfied then there exists an $f_0 \in L^\infty_+$, $f_0 \not\equiv 0$, and also a sequence $f_n \in \mathbb{C}$ such that $\lim_{n\to\infty} f_n = f_0$ a.s., such that for each n, $f_n \geqslant f_0 - \frac{1}{n}$. In particular $f_n \geqslant -\frac{1}{n}$. This is almost the same as an arbitrage opportunity, since any element of $f \in \overline{\mathbb{C}}$ is the limit in the L^∞ norm of a sequence $(f_n)_{n\geqslant 1}$ in $\mathbb{C}$. This means that if $f \geqslant 0$ then the sequence of possible losses $(f_n^-)_{n\geqslant 1}$ tends to zero uniformly as $n \to \infty$, which means that the risk vanishes in the limit.

Theorem 3 *(Fundamental Theorem; Bounded Case). Let S be a bounded semimartingale. There exists an equivalent martingale measure P^* for S if and only if S satisfies NFLVR.*

Proof. Let us assume we have NFLVR. Since S satisfies the no arbitrage property we have $\mathbb{C} \cap L^\infty_+ = \{0\}$. However one can use the property NFLVR to show $\mathbb{C}$ is weak* closed in L^∞ (that is, it is closed in $\sigma(L^1, L^\infty)$), and hence there will exist a probability P^* equivalent to P with $E^*\{f\} \leqslant 0$, all f in $\mathbb{C}$. (This is the Kreps–Yan separation theorem – essentially the Hahn–Banach theorem; see, e.g., Yan, 1980). For each $s < t$, $B \in \mathcal{F}_s$, $\alpha \in \mathbb{R}$, we deduce $\alpha(S_t - S_s)1_B \in \mathbb{C}$, since S is bounded. Therefore $E^*\{(S_t - S_s)1_B\} = 0$, and S is a martingale under P^*.

For the converse, note that NFLVR remains unchanged with an equivalent probability, so without loss of generality we may assume S is a martingale under P itself. If θ is admissible, then $(\int_0^t \theta_s \, \mathrm{d}S_s)_{t\geqslant 0}$ is a local martingale, hence it is a supermartingale. Since $E\{\theta_0 S_0\} = 0$, we have as well $E\{\int_0^\infty \theta_s \, \mathrm{d}S_s\} \leqslant E\{\theta_s S_0\} = 0$. This implies that for any $f \in \mathbb{C}$, we have $E\{f\} \leqslant 0$. Therefore it is true as well for $f \in \overline{\mathbb{C}}$, the closure of $\mathbb{C}$ in L^∞. Thus we conclude $\overline{\mathbb{C}} \cap L^\infty_+ = \{0\}$. $\square$

Theorem 4 *(Corollary). Let S be a locally bounded semimartingale. There is an equivalent probability measure P^* under which S is a local martingale if and only if S satisfies NFLVR.*

The measure P^* in the corollary is known as a *local martingale measure*. We refer to Delbaen and Schachermayer (1994, p. 479) for the proof of the corollary. Examples show that in general P^* can make S only a local martingale, not a martingale. We also note that any semimartingale with continuous paths is locally bounded. However in the continuous case there is a considerable simplification: the no arbitrage property alone, properly interpreted, implies the existence of an equivalent local martingale measure P^* (see Delbaen, 1995). Indeed using the Girsanov theorem this implies that under the No Arbitrage assumption the semimartingale must have the form

$$S_t = M_t + \int_0^t H_s \, \mathrm{d}[M, M]_s,$$

where M is a local martingale under P, and with restrictions on the predictable process H. Indeed, if one has $\int_0^\epsilon H_s^2 \, \mathrm{d}[M, M]_s = \infty$ for some $\epsilon > 0$, then S admits "immediate arbitrage," a fascinating concept introduced by Delbaen and Schachermayer (1995).

For the general case, we have the impressive theorem of Delbaen and Schachermayer (1995, see for a proof), as follows:

Theorem 5 *(Fundamental Theorem; General Case). Let S be a semimartingale. There exists an equivalent probability measure P^* such that S is a sigma martingale under P^* if and only if S satisfies NFLVR.*[14]

[14] See Protter (2005, Section 9 of Chapter IV, pp. 237ff), for a treatment of sigma martingales; alternatively, see Jacod and Shiryaev (2002, Section 6e of Chapter III, pp. 214ff).

Caveat. *In the remainder of the paper we will abuse language by referring to the equivalent probability measure P^* which makes S into a sigma martingale, as an* equivalent martingale measure. *For clarity let us repeat:* if P^* is an equivalent martingale measure, then we can *a priori* conclude no more than that S is a sigma martingale (*or local martingale, if S has continuous paths*).

3.6 Numéraire invariance

Our portfolio as described in Section 3.4 consists of

$$V_t(a,b) = a_t S_t + b_t R_t,$$

where (a,b) are trading strategies, S is the risky security price, and $R_t = \exp(\int_0^t r_s \, \mathrm{d}s)$ is the price of a money market account. The process R is often called a *numéraire*. One can then deflate future monetary values by multiplying by $\frac{1}{R_t} = \exp(-\int_0^t r_s \, \mathrm{d}s)$. Let us write $Y_t = \frac{1}{R_t}$ and we shall refer to the process Y_t as a *deflator*. By multiplying S and R by $Y = \frac{1}{R}$, we can effectively reduce the situation to the case where the price of the money market account is constant and equal to one. The next theorem allows us to do just that.

Theorem 6 *(Numéraire Invariance). Let (a,b) be a trading strategy for (S,R). Let $Y = \frac{1}{R}$. Then (a,b) is self-financing for (S,R) if and only if (a,b) is self-financing for $(YS,1)$.*

Proof. Let $Z = \int_0^t a_s \, \mathrm{d}S_s + \int_0^t b_s \, \mathrm{d}R_s$. Then using integration by parts we have (since Y is continuous and of finite variation)

$$\begin{aligned}
\mathrm{d}(Y_t Z_t) &= Y_t \, \mathrm{d}Z_t + Z_t \, \mathrm{d}Y_t \\
&= Y_t a_t \, \mathrm{d}S_t + Y_t b_t \, \mathrm{d}R_t + \left(\int_0^t a_s \, \mathrm{d}S_s + \int_0^t b_s \, \mathrm{d}R_s \right) \mathrm{d}Y_t \\
&= a_t (Y_t \, \mathrm{d}S_t + S_t \, \mathrm{d}Y_t) + b_t (Y_t \, \mathrm{d}R_t + R_t \, \mathrm{d}Y_t) \\
&= a_t \, \mathrm{d}(YS)_t + b_t \, \mathrm{d}(YR)_t
\end{aligned}$$

and since $YR = \frac{1}{R} R = 1$, this is

$$= a_t \, \mathrm{d}(YS)_t$$

since $\mathrm{d}YR = 0$ because YR is constant. Therefore

$$a_t S_t + b_t R_t = a_0 S_0 + b_0 + \int_0^t a_s \, \mathrm{d}S_s + \int_0^t b_s \, \mathrm{d}R_s$$

if and only if

$$a_t \frac{1}{R_t} S_t + b_t = a_0 S_0 + b_0 + \int_0^t a_s \, \mathrm{d}\left(\frac{1}{R} S\right)_s .$$

$\square$

Theorem 6 allows us to assume $R \equiv 1$ without loss of generality. Note that one can easily check that there is no arbitrage for (a, b) with (S, R) if and only if there is no arbitrage for (a, b) with $(\frac{1}{R} S, 1)$. By renormalizing, we no longer write $(\frac{1}{R} S, 1)$, but simply S.

The preceding theorem is the standard version, but in many applications (for example those arising in the modeling of stochastic interest rates), one wants to assume that the numéraire is a strictly positive semimartingale (instead of only a continuous finite variation process as in the previous theorem). We consider here the general case, where the numéraire is a (not necessarily continuous) semimartingale. For examples of how such a change of numéraire theorem can be used (albeit for the case where the deflator is assumed continuous), see for example (Geman et al., 1995). A reference to the literature for a result such as the following theorem is (Huang, 1985, p. 223).

Theorem 7 *(Numéraire Invariance; General Case). Let S, R be semimartingales, and assume R is strictly positive. Then the deflator $Y = \frac{1}{R}$ is a semimartingale and (a, b) is self-financing for (S, R) if and only if (a, b) is self-financing for $(\frac{S}{R}, 1)$.*

Proof. Since $f(x) = \frac{1}{x}$ is $\mathcal{C}^2$ on $(0, \infty)$, we have that Y is a (strictly positive) semimartingale by Itô's formula. By the self-financing hypothesis we have

$$\begin{aligned} V_t(a, b) &= a_t S_t + b_t R_t \\ &= a_0 S_0 + b_0 R_0 + \int_0^t a_s \, \mathrm{d}S_s + \int_0^t b_s \, \mathrm{d}R_s . \end{aligned}$$

Let us assume $S_0 = 0$, and $R_0 = 1$. The integration by parts formula for semimartingales gives

$$\mathrm{d}(S_t Y_t) = \mathrm{d}\left(\frac{S_t}{R_t}\right) = S_{t-} \, \mathrm{d}\left(\frac{1}{R_t}\right) + \frac{1}{R_{t-}} \, \mathrm{d}S_t + \mathrm{d}\left[S, \frac{1}{R}\right]_t$$

and

$$\mathrm{d}\left(\frac{V_t}{R_t}\right) = V_{t-} \, \mathrm{d}\left(\frac{1}{R_t}\right) + \frac{1}{R_{t-}} \, \mathrm{d}V_t + \mathrm{d}\left[V, \frac{1}{R}\right]_t .$$

We can next use the self-financing assumption to write:

$$\begin{aligned}
\mathrm{d}\left(\frac{V_t}{R_t}\right) &= a_t S_{t-}\,\mathrm{d}\left(\frac{1}{R_t}\right) + b_t R_{t-}\,\mathrm{d}\left(\frac{1}{R_t}\right) + \frac{1}{R_{t-}} a_t\,\mathrm{d}S_t + \frac{1}{R_{t-}} b_t\,\mathrm{d}R_t \\
&\quad + a_t\,\mathrm{d}\left[S, \frac{1}{R}\right]_t + b_t\,\mathrm{d}\left[R, \frac{1}{R}\right]_t \\
&= a_t\left(S_{t-}\,\mathrm{d}\left(\frac{1}{R}\right) + \frac{1}{R_{t-}}\,\mathrm{d}S + \mathrm{d}\left[S, \frac{1}{R}\right]\right) \\
&\quad + b_t\left(R_{t-}\,\mathrm{d}\left(\frac{1}{R}\right) + \frac{1}{R_{t-}}\,\mathrm{d}R + \mathrm{d}\left[R, \frac{1}{R}\right]\right) \\
&= a_t\,\mathrm{d}\left(S\frac{1}{R}\right) + b_t\,\mathrm{d}\left(R\frac{1}{R}\right).
\end{aligned}$$

Of course $R_t\frac{1}{R_t} = 1$, and $\mathrm{d}(1) = 0$; hence

$$\mathrm{d}\left(\frac{V_t}{R_t}\right) = a_t\,\mathrm{d}\left(S_t\frac{1}{R_t}\right).$$

In conclusion we have

$$V_t = a_t S_t + b_t R_t = b_0 + \int_0^t a_s\,\mathrm{d}S_s + \int_0^t b_s\,\mathrm{d}R_s,$$

and

$$a_t\left(\frac{S_t}{R_t}\right) + b_t = \frac{V_t}{R_t} = b_0 + \int_0^t a_s\,\mathrm{d}\left(\frac{S_s}{R_s}\right).$$

□

3.7 Redundant derivatives

Let us assume given a security price process S, and by the results in Section 3.6 we take $R_t \equiv 1$. Let $\mathcal{F}_t^0 = \sigma(S_r; r \leqslant t)$ and let $\mathcal{F}_t^{\sim} = \mathcal{F}_t^0 \vee \mathcal{N}$ where $\mathcal{N}$ are the null sets of $\mathcal{F}$ and $\mathcal{F} = \bigvee_t \mathcal{F}_t^0$, under P, defined on $(\Omega, \mathcal{F}, P)$. Finally we take $\mathcal{F}_t = \bigcap_{u>t} \mathcal{F}_u^{\sim}$. A *derivative* on S is then a random variable $C \in \mathcal{F}_T$, for some fixed time T. Note that we pay a small price here for the simplification of taking $R_t \equiv 1$, since if R_t were to be a nonconstant stochastic process, it might well change the minimal filtration we are taking, because then the processes of interest would be (R_t, S_t), in place of just S_t/R_t.

One goal of Finance Theory is to show there exists a self financing trading strategy (a, b) that one can use either to obtain C at time T, or to come as close as possible – in an appropriate sense – to obtaining C. This is the issue we discuss in this section.

Definition 11. Let S be the price process of a risky security and let R be the price process of a money market account (numéraire), which we setting equal to the constant process 1.[15] A derivative $C \in \mathcal{F}_T$ is said to be *redundant* if there exists an admissible self-financing trading strategy (a, b) such that

$$C = a_0 S_0 + b_0 R_0 + \int_0^T a_s \, \mathrm{d}S_s + \int_0^T b_s \, \mathrm{d}R_s.$$

Let us normalize S by writing $M = \frac{1}{R}S$; then C will still be redundant under M and hence we have (taking $R_t = 1$, all t):

$$C = a_0 M_0 + b_0 + \int_0^T a_s \, \mathrm{d}M_s.$$

Next note that if P^* is any equivalent martingale measure making M a martingale, and if C has finite expectation under P^*, we then have

$$E^*\{C\} = E^*\{a_0 M_0 + b_0\} + E^*\left\{ \int_0^T a_s \, \mathrm{d}M_s \right\}$$

provided all expectations exist,

$$= E^*\{a_0 M_0 + b_0\} + 0.$$

Theorem 8. *Let C be a redundant derivative such that there exists an equivalent martingale measure P^* with $C \in \mathcal{L}^*(M)$. (See the second definition following for a definition of $\mathcal{L}^*(M)$.) Then there exists a unique no arbitrage price of C and it is $E^*\{C\}$.*

Proof. First we note that the quantity $E^*\{C\}$ *is the same* for every equivalent martingale measure. Indeed if Q_1 and Q_2 are both equivalent martingale measures, then

$$E_{Q_i}\{C\} = E_{Q_i}\{a_0 M_0 + b_0\} + E_{Q_i}\left\{ \int_0^T a_s \, \mathrm{d}M_s \right\}.$$

But $E_{Q_i}\{\int_0^T a_s \, \mathrm{d}M_s\} = 0$, and $E_{Q_i}\{a_0 M_0 + b_0\} = a_0 M_0 + b_0$, since we assume a_0, M_0, and b_0 are known at time 0 and thus without loss of generality are taken to be constants.

[15] Although R is taken to be constant and equal to 1, we include it initially in the definition to illustrate the role played by being able to take it a constant process.

Next suppose one offers a price $v > E^*\{C\} = a_0 M_0 + b_0$. Then one follows the strategy $a = (a_s)_{s\geqslant 0}$ and (we are ignoring transaction costs) at time T one has C to present to the purchaser of the option. One thus has a sure profit (that is, risk free) of $v - (a_0 M_0 + b_0) > 0$. This is an arbitrage opportunity. On the other hand, if one can buy the claim C at a price $v < a_0 M_0 + b_0$, analogously at time T one will have achieved a risk-free profit of $(a_0 M_0 + b_0) - v$. □

Definition 12. If C is a derivative, and there exists an admissible self-financing trading strategy (a, b) such that

$$C = a_0 M_0 + b_0 + \int_0^T a_s \, \mathrm{d}M_s;$$

then the strategy a is said to *replicate* the derivative C.

Theorem 9 *(Corollary). If C is a redundant derivative, then one can replicate C in a self-financing manner with initial capital equal to $E^*\{C\}$, where P^* is any equivalent martingale measure for the normalized price process M.*

At this point we return to the issue of *put–call parity* mentioned in the introduction (Section 2). Recall that we had the trivial relation

$$M_T - K = (M_T - K)^+ - (K - M_T)^+,$$

which, by taking expectations under P^*, shows that the price of a call at time 0 equals the price of a put plus the stock price minus K. More generally at time t, $E^*\{(M_T - K)^+ \mid \mathcal{F}_t\}$ equals the value of a put at time t plus the stock price at time t minus K, by the P^* martingale property of M.

It is tempting to consider markets where all derivatives are redundant. Unfortunately, this is too large a space of random variables; we wish to restrict ourselves to derivatives that have good integrability properties as well.

Let us fix an equivalent martingale measure P^*, so that M is a martingale (or even a local martingale) under P^*. We consider all self-financing trading strategies (a, b) such that the process $(\int_0^t a_s^2 \, \mathrm{d}[M, M]_s)^{1/2}$ is locally integrable: that means that there exists a sequence of stopping times $(T_n)_{n\geqslant 1}$ which can be taken $T_n \leqslant T_{n+1}$, a.s., such that $\lim_{n\to\infty} T_n \geqslant T$ a.s. and $E^*\{(\int_0^{T_n} a_s^2 \, \mathrm{d}[M, M]_s)^{1/2}\} < \infty$, each T_n. Let $\mathcal{L}^*(M)$ denote the class of such strategies, under P^*. We remark that we are cheating a little here: we are letting our definition of a complete market (which follows) depend on the measure P^*, and it would be preferable to define it in terms of the objective probability P. How to go about doing this is a nontrivial issue. In the happy case where the price process is already a local martingale under the objective probability measure, this issue of course disappears.

Definition 13. A market model $(M, \mathcal{L}^*(M), P^*)$ is *complete* if every derivative $C \in L^1(\mathcal{F}_T, \mathrm{d}P^*)$ is redundant for $\mathcal{L}^*(M)$. That is, for any $C \in L^1(\mathcal{F}_T, \mathrm{d}P^*)$, there exists an admissible self-financing trading strategy (a, b) with $a \in \mathcal{L}^*(M)$ such that

$$C = a_0 M_0 + b_0 + \int_0^T a_s \, \mathrm{d}M_s,$$

and such that $(\int_0^t a_s \, \mathrm{d}M_s)_{t \geqslant 0}$ is uniformly integrable. In essence, then, a complete market is one for which every derivative is redundant.

We point out that the above definition is one of many possible definitions of a complete market. For example, one could limit attention to nonnegative random payoffs and/or payoffs that are in $L^2(\mathcal{F}_T, \mathrm{d}P^*)$.

We note that in probability theory a martingale M is said to have the *predictable representation property* if for any $C \in L^2(\mathcal{F}_T)$ one has

$$C = E\{C\} + \int_0^T a_s \, \mathrm{d}M_s$$

for some predictable $a \in \mathcal{L}(M)$. This is, of course, essentially the property of market completeness. Martingales with predictable representation are well studied and this theory can usefully be applied to Finance. For example, suppose we have a model (S, R) where by a change of numéraire we take $R = 1$. Suppose further there is an equivalent martingale measure P^* such that S is a Brownian motion under P^*. Then the model is complete for all claims C in $L^1(\mathcal{F}_T, P^*)$ such that $C \geqslant -\alpha$, for some $\alpha \geqslant 0$. (α can depend on C.) To see this, we use martingale representation (see, e.g., Protter, 2005) to find a predictable process a such that for $0 \leqslant t \leqslant T$:

$$E^*\{C \mid \mathcal{F}_t\} = E^*\{C\} + \int_0^t a_s \, \mathrm{d}S_s.$$

Let

$$V_t(a, b) = a_0 S_0 + b_0 + \int_0^t a_s \, \mathrm{d}S_s + \int_0^t b_s \, \mathrm{d}R_s;$$

we need to find b such that (a, b) is an admissible, self-financing trading strategy. Since $R_t = 1$, we have $\mathrm{d}R_t = 0$, hence we need

$$a_t S_t + b_t R_t = b_0 + \int_0^t a_s \, \mathrm{d}S_s,$$

and taking $b_0 = E^*\{C\}$, we have

$$b_t = b_0 + \int_0^t a_s \, \mathrm{d}S_s - a_t S_t$$

provides such a strategy. It is admissible since $\int_0^t a_s \, \mathrm{d}S_s \geqslant -\alpha$ for some α which depends on C.

Unfortunately, having the predictable representation property is rather delicate, and few martingales possess this property. Examples include Brownian motion, the compensated Poisson process (but *not* mixtures of the two nor even the difference of two Poisson processes) (although see Jeanblanc and Privault, 2002 for sufficient conditions when one can mix the two and have completeness), and the Azéma martingales. (One can consult Protter, 2005 for background, and Dritschel and Protter, 1999 for more on the Azéma martingales.) One can mimic a complete market in the case (for example) of two independent noises, each of which is complete alone. Several authors have done this with Brownian noise together with compensated Poisson noise, by proposing hedging strategies for each noise separately. A recent example of this is Kusuoka (1999) (where the Poisson intensity can depend on the Brownian motion) in the context of default risk models. A more traditional example is Jeanblanc-Piqué and Pontier (1990).

Most models are therefore *not* complete, and most practitioners believe the a financial world being modeled is at best only approximately complete. We will return again to the notion of an incomplete market later on in this section. First, we need to characterize complete markets. In this regard, we have the following result:

Theorem 10. *Suppose there is an equivalent martingale measure P^* such that M is a local martingale. Then P^* is the unique equivalent martingale measure only if the market is complete.*

This theorem is a trivial consequence of Dellacherie's approach to martingale representation: if there is a unique probability making a process M a local martingale, then M must have the martingale representation property. The theory has been completely resolved in the work of Jacod and Yor. [See for example Protter (2005, Chapter IV, Section 4), for a pedagogic approach to the theory.]

To give an example of what can happen, let $\mathcal{M}^2$ be the set of equivalent probabilities making M an L^2-martingale. Then M has the predictable representation property (and hence market completeness) for every extremal element of the convex set $\mathcal{M}^2$. If $\mathcal{M}^2 = \{P^*\}$, only one element, then of course P^* is extremal. (See Protter, 2005, Theorem 40, p. 186.) Indeed P^* is in fact unique in the proto-typical example of Brownian motion; since many diffusions can be constructed as pathwise functionals of Brownian motion they inherit the completeness of the Brownian model. But there are examples where one has complete markets without the uniqueness of the equivalent martingale measure (see Artzner and Heath, 1995 in this regard, as well as Jarrow et al., 1999). Nevertheless the situation is simpler when we assume our models have continuous paths.

The next theorem is a version of what is known as *the second fundamental theorem of asset pricing*. We state and prove it for the case of L^2 derivatives only. We note that this theorem has a long and illustrious history, going back to the fundamental paper of Harrison and Kreps (1979, p. 392) for the discrete case, and to Harrison and Pliska (1981, p. 241) for the continuous case, although in Harrison and Pliska (1981) the theorem below is stated only for the "only if" direction.

Theorem 11. *Let M have continuous paths. There is a unique P^* such that M is an L^2 P^*-martingale if and only if the market is complete.*

Proof. The theorem follows easily from Theorems 38, 39, and 340 of Protter (2005, pp. 185–186); we will assume those results and prove the theorem. Theorem 39 shows that if P^* is unique then the market model is complete. If P^* is not unique but the model is nevertheless complete, then by Theorem 37 P^* is nevertheless extremal in the space of probability measures making M an L^2 martingale. Let Q be another such extremal probability, and let $L_\infty = \frac{\mathrm{d}Q}{\mathrm{d}P^*}$ and $L_t = E_P\{L_\infty \mid \mathcal{F}_t\}$, with $L_0 = 1$. Let $T_n = \inf\{t > 0\colon |L_t| \geqslant n\}$. L will be continuous by Theorem 40 of Protter (2005, p. 186), hence $L_t^n = L_{t\wedge T_n}$ is bounded. We then have, for bounded $C \in \mathcal{F}_s$:

$$E_Q\{M_{t\wedge T_n} C\} = E^*\big\{M_{t\wedge T_n} L_t^n C\big\},$$
$$E_Q\{M_{s\wedge T_n} C\} = E^*\big\{M_{s\wedge T_n} L_s^n C\big\}.$$

The two left sides of the above equalities are equal and this implies that ML^n is a martingale, and thus L^n is a bounded P^*-martingale orthogonal to M. It is hence constant by Theorem 39 of Protter (2005, p. 185). We conclude $L_\infty \equiv 1$ and thus $Q = P^*$. □

Note that if C is a redundant derivative, then the no arbitrage price of C is $E^*\{C\}$, for any equivalent martingale measure P^*. (If C is redundant then we have seen the quantity $E^*\{C\}$ is the same under every P^*.) However, if a market model is not complete, then

- there will arise nonredundant claims, and
- there will be more than one equivalent martingale measure P^*.

We now have the conundrum: if C is nonredundant, what is the no arbitrage price of C? We can no longer argue that it is $E^*\{C\}$, because there are many such values! The absence of this conundrum is a large part of the appeal of complete markets. One resolution of this conundrum is to use an