

Monte Carlo Simulation and Finance

DON L. McLEISH



WILEY

John Wiley & Sons, Inc.

Monte Carlo Simulation and Finance

Founded in 1807, John Wiley & Sons is the oldest independent publishing company in the United States. With offices in North America, Europe, Australia, and Asia, Wiley is globally committed to developing and marketing print and electronic products and services for our customers' professional and personal knowledge and understanding.

The Wiley Finance series contains books written specifically for finance and investment professionals as well as sophisticated individual investors and their financial advisors. Book topics range from portfolio management to e-commerce, risk management, financial engineering, valuation and financial instrument analysis, as well as much more.

For a list of available titles, please visit our Web site at: www.WileyFinance.com.

Monte Carlo Simulation and Finance

DON L. McLEISH



WILEY

John Wiley & Sons, Inc.

Copyright © 2005 by Don L. McLeish. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400, fax 978-646-8600, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, 201-748-6011, fax 201-748-6008, or online at <http://www.wiley.com/go/permissions>.

Limit of Liability/Disclaimer of Warranty: While the publisher and the author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor the author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information about our other products and services, please contact our Customer Care Department within the United States at 800-762-2974, outside the United States at 317-572-3993 or fax 317-572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Cataloging-in-Publication Data:

McLeish, Don L.

Monte Carlo simulation and finance / Don L. Mcleish.

p. cm. – (Wiley finance series)

Includes bibliographical references and index.

ISBN-13 978-0-471-67778-9 (cloth/website)

ISBN-10 0-471-67778-7 (cloth/website)

1. Financial futures. 2. Monte Carlo method. 3. Options (Finance)

I. Title. II. Series.

HG6024.3.M357 2005

332.64'5'0151828–dc22

2004025812

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

Dedication

To those who make us laugh **and** think. To Charles William McLeish, James Allen McLeish, Michael Lewis. Their laughter has since dissolved, but the smiles and thoughts they generated live on. To my family, Cynthia, Erin, Jen, Marianne, William, Grace, Shirley, Joanne and David, Rob and Laurie, Bill and Christina. To the many others family and friends, whose unfailing support, encouragement and love is so clearly unconstrained by my due.

Contents

Acknowledgments	xi
CHAPTER 1	
Introduction	1
CHAPTER 2	
Some Basic Theory of Finance	10
Introduction to Pricing: Single-Period Models	10
Multiperiod Models	17
Determining the Process B_t	23
Minimum-Variance Portfolios and the Capital Asset Pricing Model	27
Entropy: Choosing a Q Measure	45
Models in Continuous Time	53
Problems	73
CHAPTER 3	
Basic Monte Carlo Methods	77
Uniform Random Number Generation	78
Apparent Randomness of Pseudo-Random Number Generators	86
Generating Random Numbers from Nonuniform Continuous Distributions	92
Generating Random Numbers from Discrete Distributions	132
Random Samples Associated with Markov Chains	141
Simulating Stochastic Partial Differential Equations	149
Problems	158
CHAPTER 4	
Variance Reduction Techniques	163
Introduction	163
Variance Reduction for One-Dimensional Monte Carlo Integration	166
Problems	203

CHAPTER 5**Simulating the Value of Options 206**

Asian Options	206
Pricing a Call Option under Stochastic Interest Rates	215
Simulating Barrier and Lookback Options	218
Survivorship Bias	236
Problems	242

CHAPTER 6**Quasi-Monte Carlo Multiple Integration 244**

Introduction	244
Errors in Numerical Integration	247
Theory of Low-Discrepancy Sequences	249
Examples of Low-Discrepancy Sequences	252
Problems	262

CHAPTER 7**Estimation and Calibration 264**

Introduction	264
Finding a Root	267
Maximization of Functions	272
Maximum-Likelihood Estimation	280
Using Historical Data to Estimate the Parameters in Diffusion Models	300
Estimating Volatility	304
Estimating Hedge Ratios and Correlation Coefficients	317
Problems	322

CHAPTER 8**Sensitivity Analysis, Estimating Derivatives, and the Greeks 325**

Estimating Derivatives	333
The Score Function Estimator	335
Infinitesimal Perturbation Analysis (IPA): Pathwise Differentiation	347
Calibrating a Model Using Simulations	355
Problems	362

CHAPTER 9	
Other Methods and Conclusions	364
Alternative Models	364
ARCH and GARCH	365
Conclusions	370
Notes	372
References	375
Index	382

Acknowledgments

I am grateful to all of the past students of Statistics 906 and the Master's of Finance program at the University of Waterloo for their patient reading and suggestions to improve this material, especially Keldon Drudge and Hristo Sendov. I am also indebted to my colleagues, Adam Kolkiewicz and Phelim Boyle, for their contributions to my understanding of this material.

Introduction

Experience, how much and of what, is a valuable commodity. It represents a major difference between an airline pilot and a New York cab driver, a surgeon and a butcher, and a successful financier and a cashier at your local grocer's. Experience with data and its analysis, experience constructing portfolios, experience in trading, and even experience losing money (one experience we could all do without) are part of the education of the financially literate. Of course, few of us have the courage to approach the manager of the local bank and ask for a few million so that we can acquire this experience, and fewer bank managers have the courage to accede to our request. The "joy of simulation" is that you do not need to have a Boeing 767 to fly one, and you don't need millions of dollars to acquire considerable experience in valuing financial products, constructing portfolios, and testing trading rules. Of course, if your trading rule is to buy condos in Florida because you expect that all boomers will wish to retire there, a computer simulation will do little to help you since the ingredients of your decision are based largely on psychology (yours and theirs). But if your rule is that you should hedge your current investment in condos using financial derivatives written on real estate companies, the methods of computer simulation become relevant.

This book concerns the simulation and analysis of models for financial markets, particularly traded assets such as stocks and bonds. We pay particular attention to financial derivatives such as options. These are financial instruments that derive their value from some associated asset. For example, a call option is written on a particular stock, and its value depends on the price of the stock at expiry. But there are many other types of financial derivatives, traded on assets such as bonds, currency markets or foreign exchange markets, and commodities. Indeed, there is a growing interest in so-called real options, those written on some real-world physical process such as the temperature or the amount of rainfall.

In general, an option gives the holder a right, not an obligation, to sell or buy a prescribed asset (the underlying asset) at a price determined by the contract (the exercise or strike price). For example, if you own a call option on shares of IBM with expiry date October 20, 2005, and exercise price \$120, then on October 20, 2005, you have the right to purchase a fixed number, say 100, of shares of IBM at the price of \$120. If IBM is selling for \$130 on that date, then your option is worth \$10 per share on expiry. If IBM is selling for \$120 or less, then your option is worthless. We need to know what a fair value would be for this option when it is sold, say, on February 1, 2005. Determining this fair value relies on sophisticated models both for the movements in the underlying asset and the relationship of this asset with the derivative, and this is the subject of a large part of this book. You may have bought an IBM option for one of two reasons, either because you are speculating on an increase in the stock price or to hedge a promise that you have made to deliver IBM stocks to someone in the future against possible increases in the stock price. The second use of derivatives is similar to the use of an insurance policy against movements in an asset price that could damage or bankrupt the holder of a portfolio. It is this second use of derivatives that has fueled most of the phenomenal growth in their trading. With the globalization of economies, industries are subject to more and more economic forces that they are unable to control but nevertheless wish some form of insurance against. This requires hedges against a whole litany of disadvantageous moves of the market, such as increases in the cost of borrowing, decreases in the value of assets held, and changes in foreign currency exchange rates.

The advanced theory of finance, like many other areas in which complex mathematics plays an important part, is undergoing a revolution aided by the computer and the proliferation of powerful simulation and symbolic mathematical tools. This is the mathematical equivalent of the invention of the printing press. The numerical and computational power once reserved for the most highly trained mathematicians, scientists, and engineers is now available to any competent programmer.

One of the first hurdles faced in adopting stochastic or random models in finance is the recognition that, for all practical purposes, the prices of equities in an efficient market are *random variables*; that is, although they may show some dependence on fiscal and economic processes and policies, they have a component of randomness that makes them unpredictable. This appears on the surface to be contrary to the training we all receive that every effect has a cause, and every change in the price of a stock must be driven by some factor in the company or the economy. But we should remember that random models are often applied to systems that are essentially causal when measuring and analyzing the various factors influencing the process

and their effects is too monumental a task. Even in the simple toss of a fair coin, the result is determined by the forces applied to the coin during and after it is tossed. In spite of this, we model it as a random variable because we have insufficient information on these forces to make a more accurate prediction of the outcome. Most financial processes in an advanced economy are of a similar nature. Exchange rates, interest rates, and equity prices are subject to the pressures of a large number of traders, government agencies, and speculators, as well as the forces applied by international trade and the flow of information. In the aggregate there are many forces and much information that influence the process. Although we might hope to predict some features of the process such as the average change in price or the volatility, a precise estimate of the price of an asset one year from today is clearly impossible. This is the basic argument necessitating stochastic models in finance. Adoption of a stochastic model implies neither that the process is pure noise nor that we are unable to develop a forecast. Such a model is adopted whenever we acknowledge that a process is not *perfectly* predictable and the *nonpredictable* component of the process is of sufficient importance to warrant modeling.

Now, if we accept that the price of a stock is a random variable, what are the constants in our model? Is a dollar of constant value, and if so, the dollar of which nation? Or should we accept one unit of an index that in some sense represents a share of the global economy as the constant? This question concerns our choice of what is called the *numeraire* in deference to the French influence on the theory of probability, or the process against which the value of our assets will be measured. There is no unique answer to this question, nor does that matter for most purposes. We can use a bond denominated in Canadian dollars as the numeraire, or one in U.S. dollars. Provided we account for the variability in the exchange rate, the price of an asset will be the same. Since to some extent our choice of numeraire is arbitrary, we may pick whatever is most convenient for the problem at hand.

One of the most important modern tools for analyzing a stochastic system is simulation. Simulation is the imitation of a real-world process or system. It is essentially a model, often a mathematical model of a process. In finance, a basic model for the evolution of stock prices, interest rates, exchange rates, and other factors would be necessary to determine a fair price of a derivative security. Simulations, like purely mathematical models, usually make assumptions about the behavior of the system being modeled. This model requires inputs, often called the *parameters* of the model, and outputs, a result that might measure the performance of a system, the price of a given financial instrument, or the weights on a portfolio chosen to have some desirable property. We usually construct the model in such a way that

inputs are easily changed over a given set of values, as this allows for a more complete picture of the possible outcomes.

Why use simulation? The simple answer is that it transfers work to the computer. Models can be handled that involve greater complexity and fewer assumptions, and a more faithful representation of the real world is possible. By changing parameters we can examine interactions and sensitivities of the system to various factors. Experimenters may use a simulation to provide a numerical answer to a question, assign a price to a given asset, identify optimal settings for controllable parameters, examine the effects of exogenous variables, or identify which of several schemes is more efficient or more profitable. The variables that have the greatest effect on a system can be isolated. We can also use simulation to verify the results obtained from an analytic solution. For example, many of the tractable models used in finance to select portfolios and price derivatives are wrong. They put too little weight on the extreme observations, the large positive and negative movements (crashes), which have the most dramatic effect on the results. Is this lack of fit of major concern when we use a standard model such as the Black-Scholes model to price a derivative? Questions such as this can be answered in part by examining simulations that accord more closely with the real world but are intractable to mathematical analysis.

Simulation is also used to answer questions starting with “what if.” For example, what would be the result if interest rates rose 3 percentage points over the next 12 months? In engineering, determining what would happen under extreme circumstances is often referred to as *stress testing*, and simulation is a particularly valuable tool here since the scenarios of concern are those that we observe too rarely to have substantial experience with. Simulations are used, for example, to determine the effect on an aircraft of extreme conditions and to analyze flight data information in the event of an accident. Simulation often provides experience at a lower cost compared with the alternatives.

But these advantages are not without some sacrifice. Two individuals may choose to model the same phenomenon in different ways and, as a result, may derive quite different simulation results. Because the output from a simulation is random, it is sometimes hard to analyze; statistical experience and tools are valuable assets. Building models and writing simulation code is not always easy. Time is required to construct the simulation, validate it, and analyze the results. And simulation does not render mathematical analysis unnecessary. If a reasonably simple analytic expression for a solution exists, it is always preferable to a simulation. A simulation may provide an approximate numerical answer for one or more possible parameter values, but only an expression for the solution provides insight into the way it responds to the individual parameters, the sensitivities of the solution.

In constructing a simulation, there are a number of distinct steps:

1. Formulate the problem at hand. Why do we need to use simulation?
2. Set the objectives as specifically as possible. This should include what measures on the process are of most interest.
3. Suggest candidate models. Which of these are closest to the real world? Which are fairly easy to write computer code for? What parameter values are of interest?
4. If possible, collect real data and identify which of the models in step 3 is most appropriate. Which does the best job of generating the general characteristics of the real data?
5. Implement the model. Write computer code to run simulations.
6. Verify (debug) the model. Using simple special cases, ensure that the code is doing what you think it is doing.
7. Validate the model. Ensure that it generates data with the characteristics of the real data.
8. Determine simulation design parameters. How many simulations are to be run, and what alternatives are to be simulated?
9. Run the simulation. Collect and analyze the output.
10. Are there surprises? Do we need to change the model or the parameters? Do we need more runs?
11. Finally, document the results and conclusions in the light of the simulation results. Tables of numbers are to be avoided. Well-chosen graphs are often better ways of gleaning qualitative information from a simulation.

In this book, we will not always follow our own advice, leaving some of the steps for the reader to fill in. Nevertheless, the importance of model validation, for example, cannot be overstated. Particularly in finance, where data can be plentiful, highly complex mathematical models are too often applied without any evidence that they fit the observed data adequately. The reader is advised to consult and address each of the steps above with each new simulation (and many of the examples in this text). Helpful information is provided in the appendixes, which may be found on the Web at www.wiley.com/go/mcleish.

Example

Let us consider an example illustrating a simple use for a simulation model. We are considering a buyout bid for the shares of a company. Although the company's stock is presently valued at around \$11.50 per share, a careful analysis has determined that it fits sufficiently well with our current assets that if the buyout were successful, it would be worth approximately \$14.00 per share in our hands. We are considering only three alternatives, an

immediate cash offer of \$12.00, \$13.00, or \$14.00 per share for outstanding shares of the company. Naturally, we would like to bid as low as possible, but we expect a competitor virtually simultaneously to make a bid for the company, and the competitor values the shares differently. The competitor has three bidding strategies, which we will simply identify as I, II, and III. There are costs associated with any pair of strategies (our bid combined with the competitor's bidding strategy), including costs associated with losing a given bid to the competitor or paying too much for the company. In other words, the payoff to our firm depends on the amount bid by the competitor, and the possible scenarios are given in Table 1.1.

The payoffs to the competitor are somewhat different and are given in Table 1.2. For example, the combination of Our bid = \$13 per share and competitor strategy II results in a loss of 4 units (for example, four dollars per share) to us and a gain of 4 units to our competitor. However, it is not always the case that our loss is the same as our competitor's gain. A game with the property that, under all possible scenarios, the gains add to a constant is called a *zero-sum game*, and these are much easier to analyze analytically. Define the 3×3 matrix of payoffs to your company by A and the payoff matrix to our competitor by B :

$$A = \begin{pmatrix} 3 & 2 & -2 \\ 1 & -4 & 4 \\ 0 & -5 & 5 \end{pmatrix}, \quad B = \begin{pmatrix} -1 & -2 & 3 \\ 0 & 4 & -6 \\ 0 & 5 & -5 \end{pmatrix}$$

Suppose we play strategy $i = 1, 2, 3$ (i.e., bid \$12, \$13, \$14) with probabilities p_1, p_2, p_3 , respectively, and the probabilities of the competitor's strategies are q_1, q_2, q_3 . Then if we denote

$$p = \begin{pmatrix} p_1 \\ p_2 \\ p_3 \end{pmatrix}, \quad q = \begin{pmatrix} q_1 \\ q_2 \\ q_3 \end{pmatrix}$$

we can write our expected payoff in the form $\sum_{i=1}^3 \sum_{j=1}^3 p_i A_{ij} q_j$. Written as a vector-matrix product, this takes the form $p^T A q$. This might be thought of

TABLE 1.1 Payoffs to Our Firm

Our Bid	Competitor's Strategy		
	I	II	III
12	3	2	-2
13	1	-4	4
14	0	-5	5

TABLE 1.2 Payoffs to Competitor

Our Bid	Competitor's Strategy		
	I	II	III
12	-1	-2	3
13	0	4	-6
14	0	5	-5

as the average return to our firm in the long run if this game were repeated many times, although in the real world the game is played only once. If the vector q were known, we would clearly choose $p_i = 1$ for the row i corresponding to the maximum component of Aq since this maximizes our payoff. Similarly, if our competitor knew p , they would choose $q_j = 1$ for the column j corresponding to the maximum component of $p^T B$. Over the long haul, if this game were indeed repeated many times, we would likely keep track of our opponent's frequencies and replace the unknown probabilities by the frequencies. However, we assume that both the actual move made by our opponent and the probabilities that they use in selecting their move are unknown at the time we commit to our strategy. However, if the game is repeated many times, each player obtains information about the opponent's taste in moves, and this would seem to be a reasonable approach to building a simulation model for this game. Suppose the game is played repeatedly, with each of the two players updating their estimated probabilities using information gathered about their opponent's historical use of their available strategies. We may record the number of times each strategy is used by each player and hope that the relative frequencies approach a sensible limit. This is carried out by the following Matlab function.

```
function [p,q]=nonzerosum(A,B,nsim)
% A and B are payoff matrices to the two participants
in a game.
Outputs
%mixed strategies p and q determined by simulation
conducted nsim times
n=size(A);
% A and B have the same size
p=ones(1,n(1)); q=ones(n(2),1);
% initialize with positive weights on all strategies
for i=1:nsim
% runs the simulation nsim times
```

```

[m,s]=max(A*q);
% s=index of optimal strategy for us
[m,t]=max(p*B);
% =index of optimal strategy for competitor
p(s)=p(s)+1;      % augment counts for us
q(t)=q(t)+1;
% augment counts for competitor
end
p=p-ones(1,n(1)); p=p/sum(p);
% remove initial weights from counts and then
q=q-ones(n(2),1); q=q/sum(q);
% convert counts to relative frequencies

```

The following output results from running this function for 50,000 simulations.

```
[p,q]=nonzerosum(A,B,50000)
```

This results in approximately $p' = [\frac{2}{3} \ 0 \ \frac{1}{3}]$ and $q' = [0 \ \frac{1}{2} \ \frac{1}{2}]$ with an average payoff to us of 0 and to the competitor of $1/3$. This seems to indicate that the strategies should be “mixed” or random. We should choose a bid of \$12.00 with probability around $2/3$, and \$14.00 with probability $1/3$. It appears that the competitor need only toss a fair coin and select between II and III based on its outcome. Why randomize our choice? The average value of the game to us is 0 if we use the probabilities given (in fact, if our competitor chooses probabilities $q' = [0 \ \frac{1}{2} \ \frac{1}{2}]$, it doesn’t matter what our frequencies are, because the average is 0). If we were to believe that a single fixed strategy is always our “best,” then our competitor could presumably determine what our “best” strategy is and act to reduce our return (i.e., to substantially less than 0) while increasing theirs. Only randomization provides the necessary insurance that neither player can guess the strategy to be employed by the other. This is a rather simple example of a two-person game with a nonconstant sum (in the sense that $A + B$ is not a constant matrix). The mathematical analysis of such games can be quite complex. In such a case, provided we can ensure cooperation, participants may cooperate for a greater total return.

There is no assurance that the solution given here is optimal. In fact, the solution is worth an average per game of 0 to us and $1/3$ to our competitor. If we revise our strategy to $p' = [\frac{2}{3} \ \frac{2}{9} \ \frac{1}{9}]$, for example, our average return is still 0 but we have succeeded in reducing that of our opponent to $1/9$, though it is unclear what our motivation for this would be. The solution we arrived at in this case seems to be a sensible solution, achieved with little

TABLE 1.3

If	$U < 2/3$	$2/3 \leq U < 1$
Bid	12	13 14

effort. Evidently, in a game such as this, there is no clear definition of what an optimal strategy would be, since one might plan one's play based on the worst-case or the best-case scenario, or something in between, such as an average. Do we attempt to collaborate with our competitor for greater total return and then subsequently divide this in some fashion? This simulation has modeled a simple form of competitor behavior and arrived at a reasonable solution—the best we can hope for without further assumptions.

There remains the question of how we actually select a bid with probabilities $2/3$, 0 , and $1/3$, respectively. First, let us assume that we are able to choose a “random number” U in the interval $[0, 1]$ such that the probability that it falls in any given subinterval is proportional to the length of that subinterval. This means that the random number has a uniform distribution on the interval $[0, 1]$. Then we could determine our bid based on the value of this random number from Table 1.3.

The way in which U is generated on a computer will be discussed in more detail in Chapter 2, but for the present note that each of the three alternative bids has the correct probability.

Some Basic Theory of Finance

INTRODUCTION TO PRICING: SINGLE-PERIOD MODELS

Let us begin with a very simple example designed to illustrate the no-arbitrage approach to pricing derivatives. Consider a stock whose price at present is \$ s . Over a given period, the stock may move either up or down—up to a value su , where $u > 1$ with probability p , or down to a value sd , where $d < 1$ with probability $1 - p$. In this model, these are the only moves possible for the stock in a single period. Over a longer period, of course, many other values are possible. In this market, we also assume that there is a so-called risk-free bond available returning a guaranteed rate of $r\%$ per period. Such a bond cannot default; there is no random mechanism governing its price, and its return is known upon purchase. An investment of \$1 at the beginning of the period returns a guaranteed $\$(1 + r)$ at the end. Then a portfolio purchased at the beginning of a period consisting of y stocks and x bonds will return at the end of the period an amount $\$x(1 + r) + ysZ$, where Z is a random variable taking values u and d with probabilities p and $1 - p$, respectively. We permit owning a negative amount of a stock or bond, corresponding to shorting or borrowing the corresponding asset for immediate sale.

An ambitious investor might seek a portfolio whose initial cost is zero (i.e., $x + ys = 0$) such that the return is greater than or equal to zero with positive probability. Such a strategy is called an *arbitrage*. This means that the investor is able to achieve a positive probability of future profits with no down-side risk with a net investment of \$0. In mathematical terms, the investor seeks a point (x, y) such that $x + ys = 0$ (the net cost of the portfolio is zero) and

$$\begin{aligned}x(1 + r) + ysu &\geq 0 \\x(1 + r) + ysd &\geq 0\end{aligned}$$

with at least one of the two inequalities strict (so there is never a loss and a nonzero chance of a positive return). Alternatively, is there a point on the line $y = -\frac{1}{s}x$ that lies on or above both of the two lines

$$y = -\frac{1+r}{su}x$$

$$y = -\frac{1+r}{sd}x$$

and strictly above one of them? Since all three lines pass through the origin, we need only compare the slopes: an arbitrage will not be possible if

$$-\frac{1+r}{sd} \leq -\frac{1}{s} \leq -\frac{1+r}{su} \quad (2.1)$$

and otherwise there is a point (x, y) permitting an arbitrage. The condition for no arbitrage (2.1) reduces to

$$\frac{d}{1+r} \leq 1 \leq \frac{u}{1+r} \quad (2.2)$$

So the condition for no arbitrage demands that $(1+r-u)$ and $(1+r-d)$ have opposite sign, or $d \leq (1+r) \leq u$. Unless this occurs, the stock *always* has either a better or a worse return than the bond, which makes no sense in a free market where both are traded without compulsion. Under a no-arbitrage assumption, since $d \leq (1+r) \leq u$, the bond payoff is a *convex combination* or a weighted average of the two possible stock payoffs; that is, there are probabilities $0 \leq q \leq 1$ and $1-q$ such that $(1+r) = qu + (1-q)d$. In fact, it is easy to solve this equation to determine the values of q and $1-q$.

$$q = \frac{(1+r) - d}{u - d} \quad \text{and} \quad 1 - q = \frac{u - (1+r)}{u - d}$$

Denote by \mathcal{Q} the probability distribution that puts probabilities q and $1-q$ on the points su and sd . Then if S_1 is the value of the stock at the end of the period, note that

$$\frac{1}{1+r} E_{\mathcal{Q}}(S_1) = \frac{1}{1+r} (qsu + (1-q)sd) = \frac{1}{1+r} s(1+r) = s$$

where $E_{\mathcal{Q}}$ denotes the expectation assuming that \mathcal{Q} describes the probabilities of the two outcomes.

In other words, *if there is to be no arbitrage, there exists a probability measure \mathcal{Q} such that the expected price of the future value of the stock S_1*

discounted to the present using the return from a risk-free bond is exactly the present value of the stock. The measure \underline{Q} is called the *risk-neutral* measure, and the probabilities that it assigns to the possible outcomes of S are not necessarily those that determine the future behavior of the stock. The risk-neutral measure embodies both the current consensus beliefs in the future value of the stock and the consensus investors' attitude to risk avoidance. It is not usually true that $\frac{1}{1+r}E_P(S_1) = s$, with P denoting the actual probability distribution describing the future probabilities of the stock. Indeed, it is highly unlikely that an investor would wish to purchase a risky stock if he or she could achieve exactly the same expected return with no risk at all using a bond. We generally expect that for a risky investment to be attractive, its expected return should be greater than that of a risk-free investment. Notice in this example that the risk-neutral measure \underline{Q} did not use the probabilities p and $1 - p$ that the stock will go up or down, and this seems contrary to intuition. Surely if a stock is more likely to go up, then a call option on the stock should be valued higher!

Let us suppose, for example, that we have a friend willing, in a private transaction, to buy or sell a stock at a price determined from his subjectively assigned distribution P , different from \underline{Q} . The friend believes that the stock is presently worth

$$\frac{1}{1+r}E_P S_1 = \frac{psu + (1-p)sd}{1+r} \neq s \quad \text{since } p \neq q$$

The friend offers his assets as a sacrifice to the gods of arbitrage. If the friend's assessed price is greater than the current market price, we can buy on the open market and sell to the friend. Otherwise, we can do the reverse. Either way one is enriched monetarily (and perhaps impoverished socially!).

So why should we use the \underline{Q} measure to determine the price of a given asset in a market (assuming, of course, there is a risk-neutral \underline{Q} measure and we are able to determine it)? Not because it precisely describes the future behavior of the stock, *but because if we use any other distribution, we offer an intelligent investor (and there are many!) an arbitrage opportunity, or an opportunity to make money at no risk and at our expense.*

Derivatives are investments that derive their value from that of a corresponding asset, such as a stock. A *European call option* is an option that permits you (but does not compel you) to purchase the stock at a fixed future date (the *maturity date*) for a given predetermined price (the *exercise price* of the option). For example, a call option with an exercise price of \$10 on a stock whose future value is denoted S_1 , is worth on expiry $S_1 - 10$ if $S_1 > 10$ but nothing at all if $S_1 < 10$. The difference $S_1 - 10$ between the value of the stock on expiry and the exercise price of the option is your profit if you exercise the option, purchasing the stock for \$10 and selling it on the

open market at $\$S_1$. However, if $S_1 < 10$, there is no point in exercising your option, as you are not compelled to do so and your return is $\$0$. In general, your payoff from purchasing the option is a simple function of the future price of the stock, such as $V(S_1) = \max(S_1 - 10, 0)$. We denote this by $(S_1 - 10)^+$. The future value of the option is a random variable, but it derives its value from that of the stock; hence it is called a *derivative* and the stock is the *underlying*.

A function of the stock price $V(S_1)$, which may represent the return from a portfolio of stocks and derivatives, is called a *contingent claim*. $V(S_1)$ represents the payoff to an investor from a certain financial instrument or derivative when the stock price at the end of the period is S_1 . In our simple binomial example above, the random variable takes only two possible values $V(su)$ and $V(sd)$. We will show that there is a portfolio, called a *replicating* portfolio, consisting of an investment solely in the above stock and bond that reproduces these values, $V(su)$ and $V(sd)$, exactly. We can determine the corresponding weights on the bond and stocks (x, y) simply by solving the two equations in two unknowns

$$x(1+r) + ysu = V(su)$$

$$x(1+r) + ysd = V(sd)$$

Solving, $y^* = \frac{V(su) - V(sd)}{su - sd}$ and $x^* = \frac{V(su) - y^*su}{1+r}$. By buying y^* units of the stock and x^* units of the bond, we are able to replicate the contingent claim $V(S_1)$ exactly—that is, produce a portfolio of stocks and bonds with exactly the same return as the contingent claim. So, in this case at least, there can be only one possible present value for the contingent claim, and that is the present value of the replicating portfolio, $x^* + y^*s$. If the market placed any other value on the contingent claim, then a trader could guarantee a positive return by a simple trade, shorting the contingent claim and buying the equivalent portfolio or buying the contingent claim and shorting the replicating portfolio. Thus this is the only price that precludes an arbitrage opportunity. There is a simpler expression for the current price of the contingent claim in this case:

$$\begin{aligned} \frac{1}{1+r} E_Q V(S_1) &= \frac{1}{1+r} (qV(su) + (1-q)V(sd)) \\ &= \frac{1}{1+r} \left(\frac{1+r-d}{u-d} V(su) + \frac{u-(1+r)}{u-d} V(sd) \right) \\ &= x^* + y^*s \end{aligned}$$

In words, *the discounted expected value of the contingent claim is equal to the no-arbitrage price of the derivative where the expectation is taken using*

the \mathcal{Q} measure. Indeed, any contingent claim that is attainable must have its price determined in this way. Although we have developed this for an extremely simple case, it extends much more generally.

Suppose we have a total of N risky assets whose prices at times $t = 0, 1$ are given by $(S_0^j, S_1^j), j = 1, 2, \dots, N$. We denote by S_0 and S_1 the column vectors of initial and final prices:

$$S_0 = \begin{pmatrix} S_0^1 \\ S_0^2 \\ \vdots \\ S_0^N \end{pmatrix}, \quad S_1 = \begin{pmatrix} S_1^1 \\ S_1^2 \\ \vdots \\ S_1^N \end{pmatrix}$$

where at time 0, S_0 is known and S_1 is random. Assume also that there is a riskless asset (a bond) paying interest rate r over one unit of time. Suppose we borrow money (this is the same as shorting bonds) at the risk-free rate to buy w_j units of stock j at time 0 for a total cost of $\sum w_j S_0^j$. The value of this portfolio at time $t = 1$ is $T(w) = \sum w_j (S_1^j - (1+r)S_0^j)$. If there are weights w_j so that this sum is always nonnegative, and $P(T(w) > 0) > 0$, then this is an arbitrage opportunity. Similarly, by replacing the weights w_j by their negative $-w_j$, there is an arbitrage opportunity if for some weights the sum is nonpositive and negative with positive probability. In summary, there are *no arbitrage opportunities* if for all weights w_j , $P(T(w) > 0) > 0$ and $P(T(w) < 0) > 0$, so $T(w)$ takes both positive and negative values. We assume that the moment-generating function $M(w) = E[\exp(\sum w_j \times (S_1^j - (1+r)S_0^j))]$ exists and is an analytic function of w . Roughly, the condition that the moment-generating function is analytic ensures that we can expand the function in a series expansion in w . This is the case, for example, if the values of S_1 and S_0 are bounded. The following theorem is a special case of Rogers (1994) and provides the equivalence of the no-arbitrage condition and the existence of an equivalent measure \mathcal{Q} .

Theorem 1 *A necessary and sufficient condition that there be no arbitrage opportunities is that there exists a measure \mathcal{Q} equivalent to P such that $E_{\mathcal{Q}}(S_1^j) = \frac{1}{1+r}S_0^j$ for all $j = 1, \dots, N$.*

Proof Define $M(w) = E \exp(T(w)) = E[\exp(\sum w_j (S_1^j - (1+r)S_0^j))]$ and consider the problem

$$\min_w \ln(M(w))$$

The no-arbitrage condition implies that for each j there exists $\varepsilon > 0$ such that

$$P[S_1^j - (1+r)S_0^j > \varepsilon] > 0$$

and therefore as $w_j \rightarrow \infty$ while the other weights $w_k, k \neq j$, remain fixed,

$$\begin{aligned} M(w) &= E \left[\exp \left(\sum w_j (S_1^j - (1+r)S_0^j) \right) \right] > C \exp(w_j \varepsilon) \\ &\times P[S_1^j - (1+r)S_0^j > \varepsilon] \rightarrow \infty \quad \text{as } w_j \rightarrow \infty \end{aligned}$$

Similarly, $M(w) \rightarrow \infty$ as $w_j \rightarrow -\infty$. From the properties of a moment-generating function, $M(w)$ is convex, continuous, and analytic and $M(0) = 1$. Therefore, the function $M(w)$ has a minimum w^* satisfying

$$\frac{\partial M(w)}{\partial w_j} = 0 \quad \text{for all } j \tag{2.3}$$

or

$$E[S_1^j \exp(T(w))] = (1+r)S_0^j E[\exp(T(w))]$$

or

$$S_0^j = \frac{E[\exp(T(w))S_1^j]}{(1+r)E[\exp(T(w))]}$$

Define a distribution or probability measure \mathcal{Q} as follows: For any event A and $w = w^*$,

$$\mathcal{Q}(A) = \frac{E_P[I_A \exp(w^*S_1)]}{E_P[\exp(w^*S_1)]} = \int_A \left(\frac{d\mathcal{Q}}{dP} \right) dP$$

where the Radon-Nikodym derivative is

$$\frac{d\mathcal{Q}}{dP} = \frac{\exp(w^*S_1)}{E_P[\exp(w^*S_1)]}$$

Since $\infty > \frac{d\mathcal{Q}}{dP} > 0$, the measure \mathcal{Q} is equivalent to the original probability measure P (in the intuitive sense that it has the same support). When we calculate expected values under this new measure, note that for each j ,

$$\begin{aligned} E_{\mathcal{Q}}(S_1^j) &= E_P \left[\frac{d\mathcal{Q}}{dP} S_1^j \right] \\ &= \frac{E_P[S_1^j \exp(w^*S_1)]}{E_P[\exp(w^*S_1)]} \\ &= (1+r)S_0^j \end{aligned}$$

or

$$S_0^j = \frac{1}{1+r} E_{\mathbb{Q}}(S_1^j)$$

Therefore, the current price of each stock is the discounted expected value of the future price under this “risk-neutral” measure \mathbb{Q} .

Conversely, if

$$E_{\mathbb{Q}}(S_1^j) = \frac{1}{1+r} S_0^j \quad \text{for all } j \quad (2.4)$$

holds for some measure \mathbb{Q} , then $E_{\mathbb{Q}}[T(w)] = 0$ for all w , and this implies that the random variable $T(w)$ is either identically 0 or admits both positive and negative values. Therefore, the existence of the measure \mathbb{Q} satisfying (2.4) implies that there are no arbitrage opportunities. ■

The so-called risk-neutral measure \mathbb{Q} above is constructed to minimize the cross-entropy between \mathbb{Q} and P subject to the constraints $E(S_1 - (1+r) \times S_0) = 0$ (cross-entropy is defined later). If there are J possible values of the random variables S_1 and S_0 , then (2.3) consists of J equations in J unknowns, and so it is reasonable to expect a unique solution. In this case, the \mathbb{Q} measure is unique and we call the market *complete*.

The theory of pricing derivatives in a complete market is rooted in a rather trivial observation because in a complete market, the derivative can be replicated with a portfolio of other marketable securities. If we can reproduce exactly the same (random) returns as the derivative provides using a portfolio that combines other marketable securities (which have prices assigned by the market), then the derivative must have the same price as the portfolio. Any other price would provide arbitrage opportunities.

Of course, in the real world there are costs associated with trading; these costs are usually related to a bid-ask spread. There essentially are different prices for buying a security and for selling it. The argument above assumes a frictionless market with no trading costs, with borrowing any amount at the risk-free bond rate possible, and a completely liquid market in which any amount of any security can be bought or sold. Moreover, it is usually assumed that the market is complete, although it is doubtful that complete markets exist. If a derivative security can be perfectly replicated using other marketable instruments, then its only purpose in the market is packaging. All models, excepting those on *Fashion File*, have deficiencies and critics. The merit of the frictionless trading assumption is that it provides an accurate approximation to increasingly liquid real-world markets. Like all useful models, this permits tentative conclusions that are subject to constant study and improvement.