Springer Finance

Springer
New York
Berlin
Heidelberg
Hong Kong
London
Milan
Paris
Tokyo

Springer Finance

Springer Finance is a programme of books aimed at students, academics, and practitioners working on increasingly technical approaches to the analysis of financial markets. It aims to cover a variety of topics, not only mathematical finance but foreign exchanges, term structure, risk management, portfolio theory, equity derivatives, and financial economics.

M. Ammann, Credit Risk Valuation: Methods, Models, and Applications (2001)
E. Barucci, Financial Markets Theory: Equilibrium, Efficiency and Information (2003)
N.H. Bingham and R. Kiesel, Risk-Neutral Valuation: Pricing and Hedging of Financial Derivatives, 2nd Edition (2004)
T.R. Bielecki and M. Rutkowski, Credit Risk: Modeling, Valuation and Hedging (2001)
D. Brigo amd F. Mercurio, Interest Rate Models: Theory and Practice (2001)
R. Buff, Uncertain Volatility Models – Theory and Application (2002)
R.-A. Dana and M. Jeanblanc, Financial Markets in Continuous Time (2003)
G. Deboeck and T. Kohonen (Editors), Visual Explorations in Finance with Self-Organizing Maps (1998)
R.J. Elliott and P.E. Kopp, Mathematics of Financial Markets (1999)
H. Geman, D. Madan, S.R. Pliska and T. Vorst (Editors), Mathematical Finance – Bachelier Congress 2000 (2001)
M. Gundlach and F. Lehrbass (Editors), CreditRisk+ in the Banking Industry (2004)
Y.-K. Kwok, Mathematical Models of Financial Derivatives (1998)
M. Külpmann, Irrational Exuberance Reconsidered: The Cross Section of Stock Returns, 2nd Edition (2004)
A. Pelsser, Efficient Methods for Valuing Interest Rate Derivatives (2000)
J.-L. Prigent, Weak Convergence of Financial Markets (2003)
B. Schmid, Credit Risk Pricing Models: Theory and Practice, 2nd Edition (2004)
S.E. Shreve, Stochastic Calculus for Finance I: The Binomial Asset Pricing Model (2004)
S.E. Shreve, Stochastic Calculus for Finance II: Continuous-Time Models (2004)
M. Yor, Exponential Functionals of Brownian Motion and Related Processes (2001)
R. Zagst, Interest-Rate Management (2002)
Y.-l. Zhu and I.-L Chern, Derivative Securities and Difference Methods (2004)
A. Ziegler, Incomplete Information and Heterogeneous Beliefs in Continuous-Time Finance (2003)
A. Ziegler, A Game Theory Analysis of Options: Corporate Finance and Financial Intermediation in Continuous Time, 2nd Edition (2004)

Steven E. Shreve

Stochastic Calculus for Finance I

The Binomial Asset Pricing Model

With 33 Figures

Springer

To my students

Preface

Origin of This Text

This text has evolved from mathematics courses in the Master of Science in Computational Finance (MSCF) program at Carnegie Mellon University. The content of this book has been used successfully with students whose mathematics background consists of calculus and calculus-based probability. The text gives precise statements of results, plausibility arguments, and even some proofs, but more importantly, intuitive explanations developed and refined through classroom experience with this material are provided. Exercises conclude every chapter. Some of these extend the theory and others are drawn from practical problems in quantitative finance.

The first three chapters of Volume I have been used in a half-semester course in the MSCF program. The full Volume I has been used in a full-semester course in the Carnegie Mellon Bachelor's program in Computational Finance. Volume II was developed to support three half-semester courses in the MSCF program.

Dedication

Since its inception in 1994, the Carnegie Mellon Master's program in Computational Finance has graduated hundreds of students. These people, who have come from a variety of educational and professional backgrounds, have been a joy to teach. They have been eager to learn, asking questions that stimulated thinking, working hard to understand the material both theoretically and practically, and often requesting the inclusion of additional topics. Many came from the finance industry, and were gracious in sharing their knowledge in ways that enhanced the classroom experience for all.

This text and my own store of knowledge have benefited greatly from interactions with the MSCF students, and I continue to learn from the MSCF alumni. I take this opportunity to express gratitude to these students and former students by dedicating this work to them.

Acknowledgments

Conversations with several people, including my colleagues David Heath and Dmitry Kramkov, have influenced this text. Lukasz Kruk read much of the manuscript and provided numerous comments and corrections. Other students and faculty have pointed out errors in and suggested improvements of earlier drafts of this work. Some of these are Jonathan Anderson, Bogdan Doytchinov, Steven Gillispie, Sean Jones, Anatoli Karolik, Andrzej Krause, Petr Luksan, Sergey Myagchilov, Nicki Rasmussen, Isaac Sonin, Massimo Tassan-Solet, David Whitaker and Uwe Wystup. In some cases, users of these earlier drafts have suggested exercises or examples, and their contributions are acknowledged at appropriate points in the text. To all those who aided in the development of this text, I am most grateful.

During the creation of this text, the author was partially supported by the National Science Foundation under grants DMS-9802464, DMS-0103814, and DMS-0139911. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

Pittsburgh, Pennsylvania, USA
December 2003

Steven E. Shreve

Contents

Introduction

Background

By awarding Harry Markowitz, William Sharpe, and Merton Miller the 1990 Nobel Prize in Economics, the Nobel Prize Committee brought to worldwide attention the fact that the previous forty years had seen the emergence of a new scientific discipline, the "theory of finance." This theory attempts to understand how financial markets work, how to make them more efficient, and how they should be regulated. It explains and enhances the important role these markets play in capital allocation and risk reduction to facilitate economic activity. Without losing its application to practical aspects of trading and regulation, the theory of finance has become increasingly mathematical, to the point that problems in finance are now driving research in mathematics.

Harry Markowitz's 1952 Ph.D. thesis *Portfolio Selection* laid the groundwork for the mathematical theory of finance. Markowitz developed a notion of mean return and covariances for common stocks that allowed him to quantify the concept of "diversification" in a market. He showed how to compute the mean return and variance for a given portfolio and argued that investors should hold only those portfolios whose variance is minimal among all portfolios with a given mean return. Although the language of finance now involves stochastic (Itô) calculus, management of risk in a quantifiable manner is the underlying theme of the modern theory and practice of quantitative finance.

In 1969, Robert Merton introduced stochastic calculus into the study of finance. Merton was motivated by the desire to understand how prices are set in financial markets, which is the classical economics question of "equilibrium," and in later papers he used the machinery of stochastic calculus to begin investigation of this issue.

At the same time as Merton's work and with Merton's assistance, Fischer Black and Myron Scholes were developing their celebrated option pricing formula. This work won the 1997 Nobel Prize in Economics. It provided a satisfying solution to an important practical problem, that of finding a fair price for a European call option (i.e., the right to buy one share of a given stock at a specified price and time). In the period 1979–1983, Harrison, Kreps, and Pliska used the general theory of continuous-time stochastic processes to put the Black-Scholes option-pricing formula on a solid theoretical basis, and, as a result, showed how to price numerous other "derivative" securities.

Many of the theoretical developments in finance have found immediate application in financial markets. To understand how they are applied, we digress for a moment on the role of financial institutions. A principal function of a nation's financial institutions is to act as a risk-reducing intermediary among customers engaged in production. For example, the insurance industry pools premiums of many customers and must pay off only the few who actually incur losses. But risk arises in situations for which pooled-premium insurance is unavailable. For instance, as a hedge against higher fuel costs, an airline may want to buy a security whose value will rise if oil prices rise. But who wants to sell such a security? The role of a financial institution is to design such a security, determine a "fair" price for it, and sell it to airlines. The security thus sold is usually "derivative" (i.e., its value is based on the value of other, identified securities). "Fair" in this context means that the financial institution earns just enough from selling the security to enable it to trade in other securities whose relation with oil prices is such that, if oil prices do indeed rise, the firm can pay off its increased obligation to the airlines. An "efficient" market is one in which risk-hedging securities are widely available at "fair" prices.

The Black-Scholes option pricing formula provided, for the first time, a theoretical method of fairly pricing a risk-hedging security. If an investment bank offers a derivative security at a price that is higher than "fair," it may be underbid. If it offers the security at less than the "fair" price, it runs the risk of substantial loss. This makes the bank reluctant to offer many of the derivative securities that would contribute to market efficiency. In particular, the bank only wants to offer derivative securities whose "fair" price can be determined in advance. Furthermore, if the bank sells such a security, it must then address the hedging problem: how should it manage the risk associated with its new position? The mathematical theory growing out of the Black-Scholes option pricing formula provides solutions for both the pricing and hedging problems. It thus has enabled the creation of a host of specialized derivative securities. This theory is the subject of this text.

Relationship between Volumes I and II

Volume II treats the continuous-time theory of stochastic calculus within the context of finance applications. The presentation of this theory is the raison d'être of this work. Volume II includes a self-contained treatment of the probability theory needed for stochastic calculus, including Brownian motion and its properties.

Volume I presents many of the same finance applications, but within the simpler context of the discrete-time binomial model. It prepares the reader for Volume II by treating several fundamental concepts, including martingales, Markov processes, change of measure and risk-neutral pricing in this less technical setting. However, Volume II has a self-contained treatment of these topics, and strictly speaking, it is not necessary to read Volume I before reading Volume II. It is helpful in that the difficult concepts of Volume II are first seen in a simpler context in Volume I.

In the Carnegie Mellon Master's program in Computational Finance, the course based on Volume I is a prerequisite for the courses based on Volume II. However, graduate students in computer science, finance, mathematics, physics and statistics frequently take the courses based on Volume II without first taking the course based on Volume I.

The reader who begins with Volume II may use Volume I as a reference. As several concepts are presented in Volume II, reference is made to the analogous concepts in Volume I. The reader can at that point choose to read only Volume II or to refer to Volume I for a discussion of the concept at hand in a more transparent setting.

Summary of Volume I

Volume I presents the binomial asset pricing model. Although this model is interesting in its own right, and is often the paradigm of practice, here it is used primarily as a vehicle for introducing in a simple setting the concepts needed for the continuous-time theory of Volume II.

Chapter 1, *The Binomial No-Arbitrage Pricing Model*, presents the no-arbitrage method of option pricing in a binomial model. The mathematics is simple, but the profound concept of risk-neutral pricing introduced here is not. Chapter 2, *Probability Theory on Coin Toss Space*, formalizes the results of Chapter 1, using the notions of martingales and Markov processes. This chapter culminates with the risk-neutral pricing formula for European derivative securities. The tools used to derive this formula are not really required for the derivation in the binomial model, but we need these concepts in Volume II and therefore develop them in the simpler discrete-time setting of Volume I. Chapter 3, *State Prices*, discusses the change of measure associated with risk-neutral pricing of European derivative securities, again as a warm-up exercise for change of measure in continuous-time models. An interesting application developed here is to solve the problem of optimal (in the sense of expected utility maximization) investment in a binomial model. The ideas of Chapters 1 to 3 are essential to understanding the methodology of modern quantitative finance. They are developed again in Chapters 4 and 5 of Volume II.

The remaining three chapters of Volume I treat more specialized concepts. Chapter 4, *American Derivative Securities*, considers derivative securities whose owner can choose the exercise time. This topic is revisited in a continuous-time context in Chapter 8 of Volume II. Chapter 5, *Random Walk*, explains the reflection principle for random walk. The analogous reflection principle for Brownian motion plays a prominent role in the derivation of pricing formulas for exotic options in Chapter 7 of Volume II. Finally, Chapter 6, *Interest-Rate-Dependent Assets*, considers models with random interest rates, examining the difference between forward and futures prices and introducing the concept of a forward measure. Forward and futures prices reappear at the end of Chapter 5 of Volume II. Forward measures for continuous-time models are developed in Chapter 9 of Volume II and used to create forward LIBOR models for interest rate movements in Chapter 10 of Volume II.

Summary of Volume II

Chapter 1, *General Probability Theory*, and Chapter 2, *Information and Conditioning*, of Volume II lay the measure-theoretic foundation for probability theory required for a treatment of continuous-time models. Chapter 1 presents probability spaces, Lebesgue integrals, and change of measure. Independence, conditional expectations, and properties of conditional expectations are introduced in Chapter 2. These chapters are used extensively throughout the text, but some readers, especially those with exposure to probability theory, may choose to skip this material at the outset, referring to it as needed.

Chapter 3, *Brownian Motion*, introduces Brownian motion and its properties. The most important of these for stochastic calculus is quadratic variation, presented in Section 3.4. All of this material is needed in order to proceed, except Sections 3.6 and 3.7, which are used only in Chapter 7, *Exotic Options* and Chapter 8, *Early Exercise*.

The core of Volume II is Chapter 4, *Stochastic Calculus*. Here the Itô integral is constructed and Itô's formula (called the Itô-Doeblin formula in this text) is developed. Several consequences of the Itô-Doeblin formula are worked out. One of these is the characterization of Brownian motion in terms of its quadratic variation (Lévy's theorem) and another is the Black-Scholes equation for a European call price (called the Black-Scholes-Merton equation in this text). The only material which the reader may omit is Section 4.7, *Brownian Bridge*. This topic is included because of its importance in Monte Carlo simulation, but it is not used elsewhere in the text.

Chapter 5, *Risk-Neutral Pricing*, states and proves Girsanov's Theorem, which underlies change of measure. This permits a systematic treatment of risk-neutral pricing and the Fundamental Theorems of Asset Pricing (Section 5.4). Section 5.5, *Dividend-Paying Stocks*, is not used elsewhere in the text. Section 5.6, *Forwards and Futures*, appears later in Section 9.4 and in some exercises.

Chapter 6, *Connections with Partial Differential Equations*, develops the connection between stochastic calculus and partial differential equations. This is used frequently in later chapters.

With the exceptions noted above, the material in Chapters 1–6 is fundamental for quantitative finance is essential for reading the later chapters. After Chapter 6, the reader has choices.

Chapter 7, *Exotic Options*, is not used in subsequent chapters, nor is Chapter 8, *Early Exercise*. Chapter 9, *Change of Numéraire*, plays an important role in Section 10.4, *Forward LIBOR model*, but is not otherwise used. Chapter 10, *Term Structure Models*, and Chapter 11, *Introduction to Jump Processes*, are not used elsewhere in the text.

1

The Binomial No-Arbitrage Pricing Model

1.1 One-Period Binomial Model

The ***binomial asset-pricing model*** provides a powerful tool to understand ***arbitrage pricing theory*** and probability. In this chapter, we introduce this tool for the first purpose, and we take up the second in Chapter 2. In this section, we consider the simplest binomial model, the one with only one period. This is generalized to the more realistic multiperiod binomial model in the next section.

For the general one-period model of Figure 1.1.1, we call the beginning of the period *time zero* and the end of the period *time one*. At time zero, we have a stock whose price per share we denote by S_0, a positive quantity known at time zero. At time one, the price per share of this stock will be one of two positive values, which we denote $S_1(H)$ and $S_1(T)$, the H and T standing for *head* and *tail*, respectively. Thus, we are imagining that a coin is tossed, and the outcome of the coin toss determines the price at time one. We do not assume this coin is fair (i.e., the probability of head need not be one-half). We assume only that the probability of head, which we call p, is positive, and the probability of tail, which is $q = 1 - p$, is also positive.

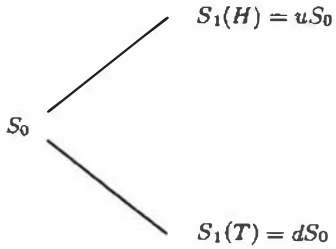

Fig. 1.1.1. General one-period binomial model.

The outcome of the coin toss, and hence the value which the stock price will take at time one, is known at time one but not at time zero. We shall refer to any quantity not known at time zero as *random* because it depends on the random experiment of tossing a coin.

We introduce the two positive numbers

$$u = \frac{S_1(H)}{S_0}, \quad d = \frac{S_1(T)}{S_0}. \tag{1.1.1}$$

We assume that $d < u$: if we instead had $d > u$, we may achieve $d < u$ by relabeling the sides of our coin. If $d = u$, the stock price at time one is not really random and the model is uninteresting. We refer to u as the *up factor* and d as the *down factor*. It is intuitively helpful to think of u as greater than one and to think of d as less than one, and hence the names *up factor* and *down factor*, but the mathematics we develop here does not require that these inequalities hold.

We introduce also an *interest rate* r. One dollar invested in the money market at time zero will yield $1+r$ dollars at time one. Conversely, one dollar borrowed from the money market at time zero will result in a debt of $1+r$ at time one. In particular, the interest rate for borrowing is the same as the interest rate for investing. It is almost always true that $r \geq 0$, and this is the case to keep in mind. However, the mathematics we develop requires only that $r > -1$.

An essential feature of an efficient market is that if a trading strategy can turn nothing into something, then it must also run the risk of loss. Otherwise, there would be an *arbitrage*. More specifically, we define *arbitrage* as a trading strategy that begins with no money, has zero probability of losing money, and has a positive probability of making money. A mathematical model that admits arbitrage cannot be used for analysis. Wealth can be generated from nothing in such a model, and the questions one would want the model to illuminate are provided with paradoxical answers by the model. Real markets sometimes exhibit arbitrage, but this is necessarily fleeting; as soon as someone discovers it, trading takes places that removes it.

In the one-period binomial model, to rule out arbitrage we must assume

$$0 < d < 1 + r < u. \tag{1.1.2}$$

The inequality $d > 0$ follows from the positivity of the stock prices and was already assumed. The two other inequalities in (1.1.2) follow from the absence of arbitrage, as we now explain. If $d \geq 1+r$, one could begin with zero wealth and at time zero borrow from the money market in order to buy stock. Even in the worst case of a tail on the coin toss, the stock at time one will be worth enough to pay off the money market debt and has a positive probability of being worth strictly more since $u > d \geq 1 + r$. This provides an arbitrage. On the other hand, if $u \leq 1 + r$, one could sell the stock short and invest the proceeds in the money market. Even in the best case for the stock, the cost of replacing it at time one will be less than or equal to the value of the money market investment, and since $d < u \leq 1 + r$, there is a positive probability that the cost of replacing the stock will be strictly less than the value of the money market investment. This again provides an arbitrage.

We have argued in the preceding paragraph that if there is to be no arbitrage in the market with the stock and the money market account, then we must have (1.1.2). The converse of this is also true. If (1.1.2) holds, then there is no arbitrage. See Exercise 1.1.

It is common to have $d = \frac{1}{u}$, and this will be the case in many of our examples. However, for the binomial asset-pricing model to make sense, we only need to assume (1.1.2).

Of course, stock price movements are much more complicated than indicated by the binomial asset-pricing model. We consider this simple model for three reasons. First of all, within this model, the concept of arbitrage pricing and its relation to risk-neutral pricing is clearly illuminated. Secondly, the model is used in practice because, with a sufficient number of periods, it provides a reasonably good, computationally tractable approximation to continuous-time models. Finally, within the binomial asset-pricing model, we can develop the theory of conditional expectations and martingales, which lies at the heart of continuous-time models.

Let us now consider a *European call option*, which confers on its owner the right but not the obligation to buy one share of the stock at time one for the *strike price* K. The interesting case, which we shall assume here, is that $S_1(T) < K < S_1(H)$. If we get a tail on the toss, the option expires worthless. If we get a head on the coin toss, the option can be *exercised* and yields a profit of $S_1(H) - K$. We summarize this situation by saying that the option at time one is worth $(S_1 - K)^+$, where the notation $(\cdots)^+$ indicates that we take the maximum of the expression in parentheses and zero. Here we follow the usual custom in probability of omitting the argument of the random variable S_1. The fundamental question of option pricing is how much the option is worth at time zero before we know whether the coin toss results in head or tail.

The *arbitrage pricing theory* approach to the option-pricing problem is to replicate the option by trading in the stock and money markets. We illustrate this with an example, and then we return to the general one-period binomial model.

Example 1.1.1. For the particular one-period model of Figure 1.1.2, let $S(0) = 4$, $u = 2$, $d = \frac{1}{2}$, and $r = \frac{1}{4}$. Then $S_1(H) = 8$ and $S_1(T) = 2$. Suppose the strike price of the European call option is $K = 5$. Suppose further that we begin with an initial wealth $X_0 = 1.20$ and buy $\Delta_0 = \frac{1}{2}$ shares of stock at time zero. Since stock costs 4 per share at time zero, we must use our initial wealth $X_0 = 1.20$ and borrow an additional 0.80 to do this. This leaves us with a cash position $X_0 - \Delta_0 S_0 = -0.80$ (i.e., a debt of 0.80 to the money market). At time one, our cash position will be $(1 + r)(X_0 - \Delta_0 S_0) = -1$

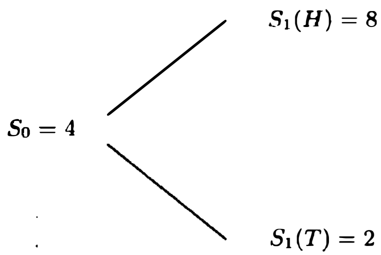

Fig. 1.1.2. Particular one-period binomial model.

(i.e., we will have a debt of 1 to the money market). On the other hand, at time one we will have stock valued at either $\frac{1}{2}S_1(H) = 4$ or $\frac{1}{2}S_1(T) = 1$. In particular, if the coin toss results in a head, the value of our portfolio of stock and money market account at time one will be

$$X_1(H) = \frac{1}{2}S_1(H) + (1+r)(X_0 - \Delta_0 S_0) = 3;$$

if the coin toss results in a tail, the value of our portfolio of stock and money market account at time one will be

$$X_1(T) = \frac{1}{2}S_1(T) + (1+r)(X_0 - \Delta_0 S_0) = 0.$$

In either case, the value of the portfolio agrees with the value of the option at time one, which is either $(S_1(H) - 5)^+ = 3$ or $(S_1(T) - 5)^+ = 0$. We have *replicated* the option by trading in the stock and money markets.

The initial wealth 1.20 needed to set up the replicating portfolio described above is the *no-arbitrage price of the option at time zero.* If one could sell the option for more than this, say, for 1.21, then the seller could invest the excess 0.01 in the money market and use the remaining 1.20 to replicate the option. At time one, the seller would be able to pay off the option, regardless of how the coin tossing turned out, and still have the 0.0125 resulting from the money market investment of the excess 0.01. This is an arbitrage because the seller of the option needs no money initially, and without risk of loss has 0.0125 at time one. On the other hand, if one could buy the option above for less than 1.20, say, for 1.19, then one should buy the option and set up the reverse of the replicating trading strategy described above. In particular, sell short one-half share of stock, which generates income 2. Use 1.19 to buy the option, put 0.80 in the money market, and in a separate money market account put the remaining 0.01. At time one, if there is a head, one needs 4 to replace the half-share of stock. The option bought at time zero is worth 3, and the 0.80 invested in the money market at time zero has grown to 1. At time one, if there is a tail, one needs 1 to replace the half-share of stock.

The option is worthless, but the 0.80 invested in the money market at time zero has grown to 1. In either case, the buyer of the option has a net zero position at time one, plus the separate money market account in which 0.01 was invested at time zero. Again, there is an arbitrage.

We have shown that in the market with the stock, the money market, and the option, there is an arbitrage unless the time-zero price of the option is 1.20. If the time-zero price of the option is 1.20, then there is no arbitrage (see Exercise 1.2). □

The argument in the example above depends on several assumptions. The principal ones are:

- shares of stock can be subdivided for sale or purchase,
- the interest rate for investing is the same as the interest rate for borrowing,
- the purchase price of stock is the same as the selling price (i.e., there is zero *bid–ask spread*),
- at any time, the stock can take only two possible values in the next period.

All these assumptions except the last also underlie the Black-Scholes-Merton option-pricing formula. The first of these assumptions is essentially satisfied in practice because option pricing and hedging (replication) typically involve lots of options. If we had considered 100 options rather than one option in Example 1.1.1, we would have hedged the short position by buying $\Delta_0 = 50$ shares of stock rather than $\Delta_0 = \frac{1}{2}$ of a share. The second assumption is close to being true for large institutions. The third assumption is not satisfied in practice. Sometimes the bid–ask spread can be ignored because not too much trading is taking place. In other situations, this departure of the model from reality becomes a serious issue. In the Black-Scholes-Merton model, the fourth assumption is replaced by the assumption that the stock price is a geometric Brownian motion. Empirical studies of stock price returns have consistently shown this not to be the case. Once again, the departure of the model from reality can be significant in some situations, but in other situations the model works remarkably well. We shall develop a modeling framework that extends far beyond the geometric Brownian motion assumption, a framework that includes many of the more sophisticated models that are not tied to this assumption.

In the general one-period model, we define a *derivative security* to be a security that pays some amount $V_1(H)$ at time one if the coin toss results in head and pays a possibly different amount $V_1(T)$ at time one if the coin toss results in tail. A European call option is a particular kind of derivative security. Another is the *European put option*, which pays off $(K - S_1)^+$ at time one, where K is a constant. A third is a *forward contract*, whose value at time one is $S_1 - K$.

To determine the price V_0 at time zero for a derivative security, we replicate it as in Example 1.1.1. Suppose we begin with wealth X_0 and buy Δ_0 shares of stock at time zero, leaving us with a cash position $X_0 - \Delta_0 S_0$. The value of our portfolio of stock and money market account at time one is

$$X_1 = \Delta_0 S_1 + (1+r)(X_0 - \Delta_0 S_0) = (1+r)X_0 + \Delta_0(S_1 - (1+r)S_0).$$

We want to choose X_0 and Δ_0 so that $X_1(H) = V_1(H)$ and $X_1(T) = V_1(T)$. (Note here that $V_1(H)$ and $V_1(T)$ are given quantities, the amounts the derivative security will pay off depending on the outcome of the coin tosses. At time zero, we know what the two possibilities $V_1(H)$ and $V_1(T)$ are; we do not know which of these two possibilities will be realized.) Replication of the derivative security thus requires that

$$X_0 + \Delta_0\left(\frac{1}{1+r}S_1(H) - S_0\right) = \frac{1}{1+r}V_1(H), \tag{1.1.3}$$

$$X_0 + \Delta_0\left(\frac{1}{1+r}S_1(T) - S_0\right) = \frac{1}{1+r}V_1(T). \tag{1.1.4}$$

One way to solve these two equations in two unknowns is to multiply the first by a number $\tilde{p}$ and the second by $\tilde{q} = 1 - \tilde{p}$ and then add them to get

$$X_0 + \Delta_0\left(\frac{1}{1+r}\left[\tilde{p}S_1(H) + \tilde{q}S_1(T)\right] - S_0\right) = \frac{1}{1+r}\left[\tilde{p}V_1(H) + \tilde{q}V_1(T)\right]. \tag{1.1.5}$$

If we choose $\tilde{p}$ so that

$$S_0 = \frac{1}{1+r}\left[\tilde{p}S_1(H) + \tilde{q}S_1(T)\right], \tag{1.1.6}$$

then the term multiplying Δ_0 in (1.1.5) is zero, and we have the simple formula for X_0

$$X_0 = \frac{1}{1+r}\left[\tilde{p}V_1(H) + \tilde{q}V_1(T)\right]. \tag{1.1.7}$$

We can solve for $\tilde{p}$ directly from (1.1.6) in the form

$$S_0 = \frac{1}{1+r}\left[\tilde{p}uS_0 + (1-\tilde{p})dS_0\right] = \frac{S_0}{1+r}\left[(u-d)\tilde{p} + d\right].$$

This leads to the formulas

$$\tilde{p} = \frac{1+r-d}{u-d}, \quad \tilde{q} = \frac{u-1-r}{u-d}. \tag{1.1.8}$$

We can solve for Δ_0 by simply subtracting (1.1.4) from (1.1.3) to get the *delta-hedging formula*

$$\Delta_0 = \frac{V_1(H) - V_1(T)}{S_1(H) - S_1(T)}. \tag{1.1.9}$$

In conclusion, if an agent begins with wealth X_0 given by (1.1.7) and at time zero buys Δ_0 shares of stock, given by (1.1.9), then at time one, if the coin toss results in head, the agent will have a portfolio worth $V_1(H)$, and if the coin toss results in tail, the portfolio will be worth $V_1(T)$. The agent has *hedged a* *short position in the derivative security.* The derivative security that pays V_1 at time one should be priced at

$$V_0 = \frac{1}{1+r}\left[\tilde{p}V_1(H) + \tilde{q}V_1(T)\right] \tag{1.1.10}$$

at time zero. This price permits the seller to hedge the short position in the claim. This price does not introduce an arbitrage when the derivative security is added to the market comprising the stock and money market account; any other time-zero price would introduce an arbitrage.

Although we have determined the no-arbitrage price of a derivative security by setting up a hedge for a short position in the security, one could just as well consider the hedge for a long position. An agent with a long position owns an asset having a certain value, and the agent may wish to set up a hedge to protect against loss of that value. This is how practitioners think about hedging. The number of shares of the underlying stock held by a long position hedge is the negative of the number determined by (1.1.9). Exercises 1.6 and 1.7 consider this is more detail.

The numbers $\tilde{p}$ and $\tilde{q}$ given by (1.1.8) are both positive because of the no-arbitrage condition (1.1.2), and they sum to one. For this reason, we can regard them as probabilities of head and tail, respectively. They are not the actual probabilities, which we call p and q, but rather the so-called *risk-neutral probabilities.* Under the actual probabilities, the average rate of growth of the stock is typically strictly greater than the rate of growth of an investment in the money market; otherwise, no one would want to incur the risk associated with investing in the stock. Thus, p and $q = 1 - p$ should satisfy

$$S_0 < \frac{1}{1+r}\left[pS_1(H) + qS_1(T)\right],$$

whereas $\tilde{p}$ and $\tilde{q}$ satisfy (1.1.6). If the average rate of growth of the stock were exactly the same as the rate of growth of the money market investment, then investors must be neutral about risk—they do not require compensation for assuming it, nor are they willing to pay extra for it. This is simply not the case, and hence $\tilde{p}$ and $\tilde{q}$ cannot be the actual probabilities. They are only numbers that assist us in the solution of the two equations (1.1.3) and (1.1.4) in the two unknowns X_0 and Δ_0. They assist us by making the term multiplying the unknown Δ_0 in (1.1.5) drop out. In fact, because they are chosen to make the mean rate of growth of the stock appear to equal the rate of growth of the money market account, they make the mean rate of growth of any portfolio of stock and money market account appear to equal the rate of growth of the money market asset. If we want to construct a portfolio whose value at time one is V_1, then its value at time zero must be given by (1.1.7), so that its mean rate of growth under the risk-neutral probabilities is the rate of growth of the money market investment.

The concluding equation (1.1.10) for the time-zero price V_0 of the derivative security V_1 is called the *risk-neutral pricing formula* for the one-period binomial model. One should not be concerned that the actual probabilities do not appear in this equation. We have constructed a hedge for a short position in the derivative security, and this hedge works regardless of whether the stock goes up or down. The probabilities of the up and down moves are irrelevant. What matters is the size of the two possible moves (the values of u and d). In the binomial model, the prices of derivative securities depend on the set of possible stock price paths but not on how probable these paths are. As we shall see in Chapters 4 and 5 of Volume II, the analogous fact for continuous-time models is that prices of derivative securities depend on the volatility of stock prices but not on their mean rates of growth.

1.2 Multiperiod Binomial Model

We now extend the ideas in Section 1.1 to multiple periods. We toss a coin repeatedly, and whenever we get a head the stock price moves "up" by the factor u, whereas whenever we get a tail, the stock price moves "down" by the factor d. In addition to this stock, there is a money market asset with a constant interest rate r. The only assumption we make on these parameters is the no-arbitrage condition (1.1.2).

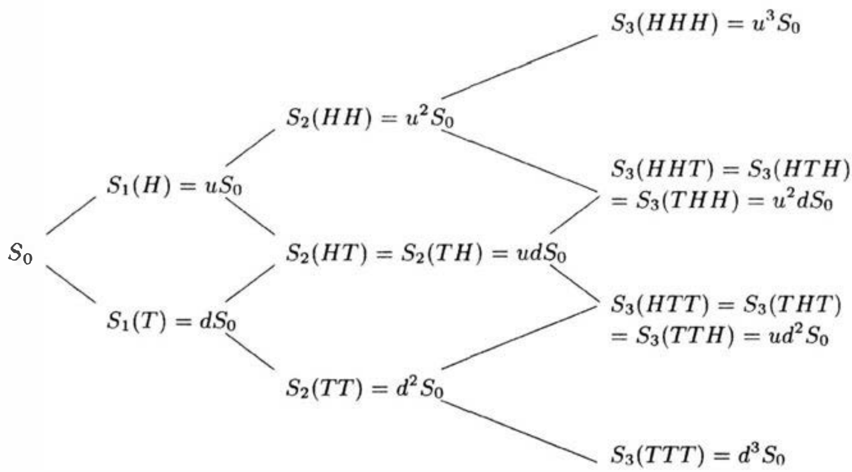

Fig. 1.2.1. General three-period model.

We denote the initial stock price by S_0, which is positive. We denote the price at time one by $S_1(H) = uS_0$ if the first toss results in head and by $S_1(T) = dS_0$ if the first toss results in tail. After the second toss, the price will be one of:

$$S_2(HH) = uS_1(H) = u^2S_0, \; S_2(HT) = dS_1(H) = duS_0,$$
$$S_2(TH) = uS_1(T) = udS_0, \; S_2(TT) = dS_1(T) = d^2S_0.$$

After three tosses, there are eight possible coin sequences, although not all of them result in different stock prices at time 3. See Figure 1.2.1.

Example 1.2.1. Consider the particular three-period model with $S_0 = 4$, $u = 2$, and $d = \frac{1}{2}$. We have the binomial "tree" of possible stock prices shown in Figure 1.2.2. □

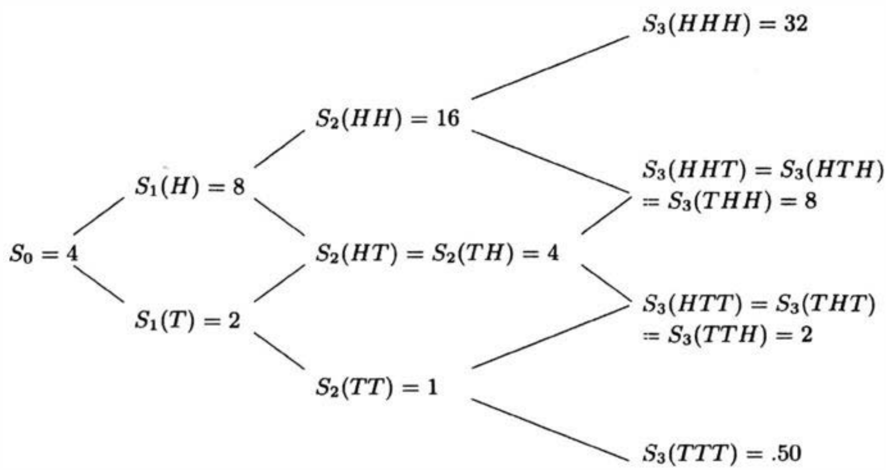

Fig. 1.2.2. A particular three-period model.

Let us return to the general three-period binomial model of Figure 1.2.1 and consider a European call that confers the right to buy one share of stock for K dollars at time two. After the discussion of this option, we extend the analysis to an arbitrary European derivative security that expires at time $N \geq 2$.

At expiration, the payoff of a call option with strike price K and expiration time two is $V_2 = (S_2 - K)^+$, where V_2 and S_2 depend on the first and second coin tosses. We want to determine the no-arbitrage price for this option at time zero. Suppose an agent sells the option at time zero for V_0 dollars, where V_0 is still to be determined. She then buys Δ_0 shares of stock, investing $V_0 - \Delta_0 S_0$ dollars in the money market to finance this. (The quantity $V_0 - \Delta_0 S_0$ will turn out to be negative, so the agent is actually borrowing $\Delta_0 S_0 - V_0$ dollars from the money market.) At time one, the agent has a portfolio (excluding the short position in the option) valued at

$$X_1 = \Delta_0 S_1 + (1 + r)(V_0 - \Delta_0 S_0). \tag{1.2.1}$$

Although we do not indicate it in the notation, S_1 and therefore X_1 depend on the outcome of the first coin toss. Thus, there are really two equations implicit in (1.2.1):

$$X_1(H) = \Delta_0 S_1(H) + (1+r)(V_0 - \Delta_0 S_0), \tag{1.2.2}$$
$$X_1(T) = \Delta_0 S_1(T) + (1+r)(V_0 - \Delta_0 S_0). \tag{1.2.3}$$

After the first coin toss, the agent has a portfolio valued at X_1 dollars and can readjust her hedge. Suppose she decides now to hold Δ_1 shares of stock, where Δ_1 is allowed to depend on the first coin toss because the agent knows the result of this toss at time one when she chooses Δ_1. She invests the remainder of her wealth, $X_1 - \Delta_1 S_1$, in the money market. In the next period, her wealth will be given by the right-hand side of the following equation, and she wants it to be V_2. Therefore, she wants to have

$$V_2 = \Delta_1 S_2 + (1+r)(X_1 - \Delta_1 S_1). \tag{1.2.4}$$

Although we do not indicate it in the notation, S_2 and V_2 depend on the outcomes of the first two coin tosses. Considering all four possible outcomes, we can write (1.2.4) as four equations:

$$V_2(HH) = \Delta_1(H)S_2(HH) + (1+r)(X_1(H) - \Delta_1(H)S_1(H)), \tag{1.2.5}$$
$$V_2(HT) = \Delta_1(H)S_2(HT) + (1+r)(X_1(H) - \Delta_1(H)S_1(H)), \tag{1.2.6}$$
$$V_2(TH) = \Delta_1(T)S_2(TH) + (1+r)(X_1(T) - \Delta_1(T)S_1(T)), \tag{1.2.7}$$
$$V_2(TT) = \Delta_1(T)S_2(TT) + (1+r)(X_1(T) - \Delta_1(T)S_1(T)). \tag{1.2.8}$$

We now have six equations, the two represented by (1.2.1) and the four represented by (1.2.4), in the six unknowns V_0, Δ_0, $\Delta_1(H)$, $\Delta_1(T)$, $X_1(H)$, and $X_1(T)$.

To solve these equations, and thereby determine the no-arbitrage price V_0 at time zero of the option and the replicating portfolio Δ_0, $\Delta_1(H)$, and $\Delta_1(T)$, we begin with the last two equations, (1.2.7) and (1.2.8). Subtracting (1.2.8) from (1.2.7) and solving for $\Delta_1(T)$, we obtain the *delta-hedging formula*

$$\Delta_1(T) = \frac{V_2(TH) - V_2(TT)}{S_2(TH) - S_2(TT)}, \tag{1.2.9}$$

and substituting this into either (1.2.7) or (1.2.8), we can solve for

$$X_1(T) = \frac{1}{1+r}[\tilde{p}V_2(TH) + \tilde{q}V_2(TT)], \tag{1.2.10}$$

where $\tilde{p}$ and $\tilde{q}$ are the risk-neutral probabilities given by (1.1.8). We can also obtain (1.2.10) by multiplying (1.2.7) by $\tilde{p}$ and (1.2.8) by $\tilde{q}$ and adding them together. Since

$$\tilde{p}S_2(TH) + \tilde{q}S_2(TT) = (1+r)S_1(T),$$

this causes all the terms involving $\Delta_1(T)$ to drop out. Equation (1.2.10) gives the value the replicating portfolio should have at time one if the stock goes down between times zero and one. We define this quantity to be the *price of the option at time one if the first coin toss results in tail*, and we denote it by $V_1(T)$. We have just shown that

$$V_1(T) = \frac{1}{1+r}[\tilde{p}V_2(TH) + \tilde{q}V_2(TT)], \tag{1.2.11}$$

which is another instance of the *risk-neutral pricing formula*. This formula is analogous to formula (1.1.10) but postponed by one period. The first two equations, (1.2.5) and (1.2.6), lead in a similar way to the formulas

$$\Delta_1(H) = \frac{V_2(HH) - V_2(HT)}{S_2(HH) - S_2(HT)} \tag{1.2.12}$$

and $X_1(H) = V_1(H)$, where $V_1(H)$ is the *price of the option at time one if the first toss results in head*, defined by

$$V_1(H) = \frac{1}{1+r}[\tilde{p}V_2(HH) + \tilde{q}V_2(HT)]. \tag{1.2.13}$$

This is again analogous to formula (1.1.10), postponed by one period. Finally, we plug the values $X_1(H) = V_1(H)$ and $X_1(T) = V_1(T)$ into the two equations implicit in (1.2.1). The solution of these equations for Δ_0 and V_0 is the same as the solution of (1.1.3) and (1.1.4) and results again in (1.1.9) and (1.1.10).

To recap, we have three *stochastic processes*, (Δ_0, Δ_1), (X_0, X_1, X_2), and (V_0, V_1, V_2). By *stochastic process*, we mean a sequence of random variables indexed by time. These quantities are random because they depend on the coin tosses; indeed, the subscript on each variable indicates the number of coin tosses on which it depends. If we begin with any initial wealth X_0 and specify values for Δ_0, $\Delta_1(H)$, and $\Delta_1(T)$, then we can compute the value of the portfolio that holds the number of shares of stock indicated by these specifications and finances this by borrowing or investing in the money market as necessary. Indeed, the value of this portfolio is defined recursively, beginning with X_0, via the *wealth equation*

$$X_{n+1} = \Delta_n S_{n+1} + (1+r)(X_n - \Delta_n S_n). \tag{1.2.14}$$

One might regard this as a contingent equation; it defines *random variables*, and actual values of these random variables are not resolved until the outcomes of the coin tossing are revealed. Nonetheless, already at time zero this equation permits us to compute what the value of the portfolio will be at every subsequent time under every coin-toss scenario.

For a derivative security expiring at time two, the random variable V_2 is contractually specified in a way that is contingent upon the coin tossing (e.g., if the coin tossing results in $\omega_1\omega_2$, so the stock price at time two is $S_2(\omega_1\omega_2)$, then for the European call we have $V_2(\omega_1\omega_2) = (S_2(\omega_1\omega_2) - K)^+)$. We want to determine a value of X_0 and values for Δ_0, $\Delta_1(H)$, and $\Delta_1(T)$ so that X_2 given by applying (1.2.14) recursively satisfies $X_2(\omega_1\omega_2) = V_2(\omega_1\omega_2)$, regardless of the values of ω_1 and ω_2. The formulas above tell us how to do this. We call V_0 the value of X_0 that allows us to accomplish this, and we define $V_1(H)$ and $V_1(T)$ to be the values of $X_1(H)$ and $X_1(T)$ given by (1.2.14) when X_0 and Δ_0 are chosen by the prescriptions above. In general, we use the symbols Δ_n and X_n to represent the number of shares of stock held by the portfolio and the corresponding portfolio values, respectively, regardless of how the initial wealth X_0 and the Δ_n are chosen. When X_0 and the Δ_n are chosen to replicate a derivative security, we use the symbol V_n in place of X_n and call this the *(no-arbitrage) price of the derivative security at time n.*

The pattern that emerged with the European call expiring at time two persists, regardless of the number of periods and the definition of the final payoff of the derivative security. (At this point, however, we are considering only payoffs that come at a specified time; there is no possibility of early exercise.)

Theorem 1.2.2 (Replication in the multiperiod binomial model). *Consider an N-period binomial asset-pricing model, with* $0 < d < 1 + r < u$, *and with*

$$\tilde{p} = \frac{1+r-d}{u-d}, \quad \tilde{q} = \frac{u-1-r}{u-d}. \tag{1.2.15}$$

Let V_N be a random variable (a derivative security paying off at time N) depending on the first N coin tosses $\omega_1\omega_2\ldots\omega_N$. Define recursively backward in time the sequence of random variables $V_{N-1}, V_{N-2}, \ldots, V_0$ by

$$V_n(\omega_1\omega_2\ldots\omega_n) = \frac{1}{1+r}[\tilde{p}V_{n+1}(\omega_1\omega_2\ldots\omega_n H) + \tilde{q}V_{n+1}(\omega_1\omega_2\ldots\omega_n T)], \tag{1.2.16}$$

so that each V_n depends on the first n coin tosses $\omega_1\omega_2\ldots\omega_n$, where n ranges between $N-1$ and 0. Next define

$$\Delta_n(\omega_1\ldots\omega_n) = \frac{V_{n+1}(\omega_1\ldots\omega_n H) - V_{n+1}(\omega_1\ldots\omega_n T)}{S_{n+1}(\omega_1\ldots\omega_n H) - S_{n+1}(\omega_1\ldots\omega_n T)}, \tag{1.2.17}$$

where again n ranges between 0 and $N-1$. If we set $X_0 = V_0$ and define recursively forward in time the portfolio values $X_1, X_2, \ldots, X_N$ by (1.2.14), then we will have

$$X_N(\omega_1\omega_2\ldots\omega_N) = V_N(\omega_1\omega_2\ldots\omega_n) \text{ for all } \omega_1\omega_2\ldots\omega_N. \tag{1.2.18}$$

Definition 1.2.3. *For $n = 1, 2, \ldots, N$, the random variable $V_n(\omega_1\ldots\omega_n)$ in Theorem 1.2.2 is defined to be the* price of the derivative security at time *n if the outcomes of the first n tosses are $\omega_1\ldots\omega_n$. The* price of the derivative security at time zero *is defined to be V_0.*

PROOF OF THEOREM 1.2.2: We prove by forward induction on n that

$$X_n(\omega_1\omega_2\ldots\omega_n) = V_n(\omega_1\omega_2\ldots\omega_n) \text{ for all } \omega_1\omega_2\ldots\omega_n, \tag{1.2.19}$$

where n ranges between 0 and N. The case of $n = 0$ is given by the definition of X_0 as V_0. The case of $n = N$ is what we want to show.

For the induction step, we assume that (1.2.19) holds for some value of n less than N and show that it holds for $n + 1$. We thus let $\omega_1\omega_2\ldots\omega_n\omega_{n+1}$ be fixed but arbitrary and assume as the induction hypothesis that (1.2.19) holds for the particular $\omega_1\omega_2\ldots\omega_n$ we have fixed. We don't know whether $\omega_{n+1} = H$ or $\omega_{n+1} = T$, so we consider both cases. We first use (1.2.14) to compute $X_{n+1}(\omega_1\omega_2\ldots\omega_n H)$, to wit

$$\begin{aligned}
&X_{n+1}(\omega_1\omega_2\ldots\omega_n H)\\
&= \Delta_n(\omega_1\omega_2\ldots\omega_n)uS_n(\omega_1\omega_2\ldots\omega_n)\\
&\qquad +(1+r)\Big(X_n(\omega_1\omega_2\ldots\omega_n) - \Delta_n(\omega_1\omega_2\ldots\omega_n)S_n(\omega_1\omega_2\ldots\omega_n)\Big).
\end{aligned}$$

To simplify the notation, we suppress $\omega_1\omega_2\ldots\omega_n$ and write this equation simply as

$$X_{n+1}(H) = \Delta_n uS_n + (1+r)(X_n - \Delta_n S_n). \tag{1.2.20}$$

With $\omega_1\omega_2\ldots\omega_n$ similarly suppressed, we have from (1.2.17) that

$$\Delta_n = \frac{V_{n+1}(H) - V_{n+1}(T)}{S_{n+1}(H) - S_{n+1}(T)} = \frac{V_{n+1}(H) - V_{n+1}(T)}{(u-d)S_n}.$$

Substituting this into (1.2.20) and using the induction hypothesis (1.2.19) and the definition (1.2.16) of V_n, we see that

$$\begin{aligned}
X_{n+1}(H) &= (1+r)X_n + \Delta_n S_n(u-(1+r))\\
&= (1+r)V_n + \frac{(V_{n+1}(H) - V_{n+1}(T))(u-(1+r))}{u-d}\\
&= (1+r)V_n + \tilde{q}V_{n+1}(H) - \tilde{q}V_{n+1}(T)\\
&= \tilde{p}V_{n+1}(H) + \tilde{q}V_{n+1}(T) + \tilde{q}V_{n+1}(H) - \tilde{q}V_{n+1}(T)\\
&= V_{n+1}(H).
\end{aligned}$$

Reinstating the suppressed coin tosses $\omega_1\omega_2\ldots\omega_n$, we may write this as

$$X_{n+1}(\omega_1\omega_2\ldots\omega_n H) = V_{n+1}(\omega_1\omega_2\ldots\omega_n H).$$

A similar argument (see Exercise 1.4) shows that

$$X_{n+1}(\omega_1\omega_2\ldots\omega_n T) = V_{n+1}(\omega_1\omega_2\ldots\omega_n T).$$

Consequently, regardless of whether $\omega_{n+1} = H$ or $\omega_{n+1} = T$, we have

$$X_{n+1}(\omega_1\omega_2\ldots\omega_n\omega_{n+1}) = V_{n+1}(\omega_1\omega_2\ldots\omega_n\omega_{n+1}).$$

Since $\omega_1\omega_2\ldots\omega_n\omega_{n+1}$ is arbitrary, the induction step is complete. □

The multiperiod binomial model of this section is said to be *complete* because every derivative security can be replicated by trading in the underlying stock and the money market. In a complete market, every derivative security has a unique price that precludes arbitrage, and this is the price of Definition 1.2.3.

Theorem 1.2.2 applies to so-called *path-dependent* options as well as to derivative securities whose payoff depends only on the final stock price. We illustrate this point with the following example.

Example 1.2.4. Suppose as in Figure 1.2.2 that $S_0 = 4$, $u = 2$, and $d = \frac{1}{2}$. Assume the interest rate is $r = \frac{1}{4}$. Then $\tilde{p} = \tilde{q} = \frac{1}{2}$. Consider a *lookback option* that pays off

$$V_3 = \max_{0\leq n\leq 3} S_n - S_3$$

at time three. Then

$$\begin{aligned}
V_3(HHH) &= S_3(HHH) - S_3(HHH) = 32 - 32 = 0,\\
V_3(HHT) &= S_2(HH) - S_3(HHT) = 16 - 8 = 8,\\
V_3(HTH) &= S_1(H) - S_3(HTH) = 8 - 8 = 0,\\
V_3(HTT) &= S_1(H) - S_3(HTT) = 8 - 2 = 6,\\
V_3(THH) &= S_3(THH) - S_3(THH) = 8 - 8 = 0,\\
V_3(THT) &= S_2(TH) - S_3(THT) = 4 - 2 = 2,\\
V_3(TTH) &= S_0 - S_3(TTH) = 4 - 2 = 2,\\
V_3(TTT) &= S_0 - S_3(TTT) = 4 - 0.50 = 3.50.
\end{aligned}$$

We compute the price of the option at other times using the backward recursion (1.2.16). This gives

$$V_2(HH) = \frac{4}{5}\left[\frac{1}{2}V_3(HHH) + \frac{1}{2}V_3(HHT)\right] = 3.20,$$

$$V_2(HT) = \frac{4}{5}\left[\frac{1}{2}V_3(HTH) + \frac{1}{2}V_3(HTT)\right] = 2.40,$$

$$V_2(TH) = \frac{4}{5}\left[\frac{1}{2}V_3(THH) + \frac{1}{2}V_3(THT)\right] = 0.80,$$

$$V_2(TT) = \frac{4}{5}\left[\frac{1}{2}V_3(TTH) + \frac{1}{2}V_3(TTT)\right] = 2.20,$$

and then

$$V_1(H) = \frac{4}{5}\left[\frac{1}{2}V_2(HH) + \frac{1}{2}V_2(HT)\right] = 2.24,$$

$$V_1(T) = \frac{4}{5}\left[\frac{1}{2}V_2(TH) + \frac{1}{2}V_2(TT)\right] = 1.20,$$

and finally

$$V_0 = \frac{4}{5}\left[\frac{1}{2}V_1(H) + \frac{1}{2}V_1(T)\right] = 1.376.$$

If an agent sells the lookback option at time zero for 1.376, she can hedge her short position in the option by buying

$$\Delta_0 = \frac{V_1(H) - V_1(T)}{S_1(H) - S_1(T)} = \frac{2.24 - 1.20}{8 - 2} = 0.1733$$

shares of stock. This costs 0.6933 dollars, which leaves her with $1.376 - 0.6933 = 0.6827$ to invest in the money market at 25% interest. At time one, she will have 0.8533 in the money market. If the stock goes up in price to 8, then at time one her stock is worth 1.3867, and so her total portfolio value is 2.24, which is $V_1(H)$. If the stock goes down in price to 2, then at time one her stock is worth 0.3467 and so her total portfolio value is 1.20, which is $V_1(T)$. Continuing this process, the agent can be sure to have a portfolio worth V_3 at time three, no matter how the coin tossing turns out. □

1.3 Computational Considerations

The amount of computation required by a naive implementation of the derivative security pricing algorithm given in Theorem 1.2.2 grows exponentially with the number of periods. The binomial models used in practice often have 100 or more periods, and there are $2^{100} \approx 10^{30}$ possible outcomes for a sequence of 100 coin tosses. An algorithm that begins by tabulating 2^{100} values for V_{100} is not computationally practical.

Fortunately, the algorithm given in Theorem 1.2.2 can usually be organized in a computationally efficient manner. We illustrate this with two examples.

Example 1.3.1. In the model with $S_0 = 4$, $u = 2$, $d = \frac{1}{2}$ and $r = \frac{1}{4}$, consider the problem of pricing a European put with strike price $K = 5$, expiring at time three. The risk-neutral probabilities are $\tilde{p} = \frac{1}{2}$, $\tilde{q} = \frac{1}{2}$. The stock process is shown in Figure 1.2.2. The payoff of the option, given by $V_3 = (5 - S_3)^+$, can be tabulated as

$$V_3(HHH) = 0, \quad V_3(HHT) = V_3(HTH) = V_3(THH) = 0$$
$$V_3(HTT) = V_3(THT) = V_3(TTH) = 3, \quad V_3(TTT) = 4.50.$$

There are $2^3 = 8$ entries in this table, but an obvious simplification is possible. Let us denote by $v_3(s)$ the payoff of the option at time three when the stock price at time three is s. Whereas V_3 has the sequence of three coin tosses as its argument, the argument of v_3 is a stock price. At time three there are only four possible stock prices, and we can tabulate the relevant values of v_3 as

$$v_3(32) = 0, \;\; v_3(8) = 0, \;\; v_3(2) = 3, \;\; v_3(.50) = 4.50.$$

If the put expired after 100 periods, the argument of V_{100} would range over the 2^{100} possible outcomes of the coin tosses whereas the argument of v_{100} would range over the 101 possible stock prices at time 100. This is a tremendous reduction in computational complexity.

According to Theorem 1.2.2, we compute V_2 by the formula

$$V_2(\omega_1\omega_2) = \frac{2}{5}\Big[V_3(\omega_1\omega_2 H) + V_3(\omega_1\omega_2 T)\Big]. \tag{1.3.1}$$

Equation (1.3.1) represents four equations, one for each possible choice of $\omega_1\omega_2$. We let $v_2(s)$ denote the price of the put at time two if the stock price at time two is s. In terms of this function, (1.3.1) takes the form

$$v_2(s) = \frac{2}{5}\Big[v_3(2s) + v_3(\tfrac{1}{2}s)\Big],$$

and this represents only three equations, one for each possible value of the stock price at time two. Indeed, we may compute

$$\begin{aligned} v_2(16) &= \frac{2}{5}\Big[v_3(32) + v_3(8)\Big] = 0, \\ v_2(4) &= \frac{2}{5}\Big[v_3(8) + v_3(2)\Big] = 1.20, \\ v_2(1) &= \frac{2}{5}\Big[v_3(2) + v_3(.50)\Big] = 3. \end{aligned}$$

Similarly,

$$\begin{aligned} v_1(8) &= \frac{2}{5}\Big[v_2(16) + v_2(4)\Big] = 0.48, \\ v_1(2) &= \frac{2}{5}\Big[v_2(4) + v_2(1)\Big] = 1.68, \end{aligned}$$

where $v_1(s)$ denotes the price of the put at time one if the stock price at time one is s. The price of the put at time zero is

$$v_0(4) = \frac{2}{5}\Big[v_1(8) + v_1(2)\Big] = 0.864.$$

At each time $n = 0, 1, 2$, if the stock price is s, the number of shares of stock that should be held by the replicating portfolio is

$$\delta_n(s) = \frac{v_{n+1}(2s) - v_{n+1}(\frac{1}{2}s)}{2s - \frac{1}{2}s}.$$

This is the analogue of formula (1.2.17). □

In Example 1.3.1, the price of the option at any time n was a function of the stock price S_n at that time and did not otherwise depend on the coin tosses. This permitted the introduction of the functions v_n related to the random variables V_n by the formula $V_n = v_n(S_n)$. A similar reduction is often possible when the price of the option does depend on the stock price path rather than just the current stock price. We illustrate this with a second example.

Example 1.3.2. Consider the lookback option of Example 1.2.4. At each time n, the price of the option can be written as a function of the stock price S_n and the maximum stock price $M_n = \max_{0 \le k \le n} S_k$ to date. At time three, there are six possible pairs of values for (S_3, M_3), namely

$$(32,32),\ (8,16),\ (8,8),\ (2,8),\ (2,4),\ (.50,4).$$

We define $v_3(s,m)$ to be the payoff of the option at time three if $S_3 = s$ and $M_3 = m$. We have

$$v_3(32,32) = 0,\ v_3(8,16) = 8,\ v_3(8,8) = 0,$$
$$v_3(2,8) = 6,\ v_3(2,4) = 2,\ v_3(.50,4) = 3.50.$$

In general, let $v_n(s,m)$ denote the value of the option at time n if $S_n = s$ and $M_n = m$. The algorithm of Theorem 1.2.2 can be rewritten in terms of the functions v_n as

$$v_n(s,m) = \frac{2}{5}\Big[v_{n+1}(2s, m \vee (2s)) + v_{n+1}(\tfrac{1}{2}s, m)\Big],$$

where $m \vee (2s)$ denotes the maximum of m and $2s$. Using this algorithm, we compute

$$v_2(16,16) = \frac{2}{5}\Big[v_3(32,32) + v_2(8,16)\Big] = 3.20,$$
$$v_2(4,8) = \frac{2}{5}\Big[v_3(8,8) + v_3(2,8)\Big] = 2.40,$$
$$v_2(4,4) = \frac{2}{5}\Big[v_3(8,8) + v_3(2,4)\Big] = 0.80,$$
$$v_2(1,4) = \frac{2}{5}\Big[v_3(2,4) + v_3(.50,4)\Big] = 2.20,$$

then compute

$$v_1(8,8) = \frac{2}{5}\Big[v_2(16,16) + v_2(4,8)\Big] = 2.24,$$
$$v_1(2,4) = \frac{2}{5}\Big[v_1(4,4) + v_1(1,4)\Big] = 1.20,$$

and finally obtain the time-zero price

$$v_0(4,4) = \frac{2}{5}\Big[v_1(8,8) + v_1(2,4)\Big] = 1.376.$$

At each time $n = 0, 1, 2$, if the stock price is s and the maximum stock price to date is m, the number of shares of stock that should be held by the replicating portfolio is

$$\delta_n(s, m) = \frac{v_{n+1}(2s, m \vee (2s)) - v_{n+1}(\frac{1}{2}s, m)}{2s - \frac{1}{2}s}.$$

This is the analogue of formula (1.2.17). □

1.4 Summary

This chapter considers a multiperiod binomial model. At each period in this model, we toss a coin whose outcome determines whether the stock price changes by a factor of u or a factor of d, where $0 < d < u$. In addition to the stock, there is a money market account with per-period rate of interest r. This is the rate of interest applied to both investing and borrowing.

Arbitrage is a trading strategy that begins with zero capital and trades in the stock and money markets in order to make money with positive probability without any possibility of losing money. The multiperiod binomial model admits no arbitrage if and only if

$$0 < d < 1 + r < u. \tag{1.1.2}$$

We shall always impose this condition.

A derivative security pays off at some expiration time N contingent upon the coin tosses in the first N periods. The *arbitrage pricing theory* method of assigning a price to a derivative security prior to expiration can be understood in two ways. First, one can ask how to assign a price so that one cannot form an arbitrage by trading in the derivative security, the underlying stock, and the money market. This no-arbitrage condition uniquely determines the price at all times of the derivative security. Secondly, at any time n prior to the expiration time N, one can imagine selling the derivative security for a price and using the income from this sale to form a portfolio, dynamically trading the stock and money market asset from time n until the expiration time N. This portfolio hedges the short position in the derivative security if its value at time N agrees with the payoff of the derivative security, regardless of the outcome of the coin tossing between times n and N. The amount for which the derivative security must be sold at time n in order to construct this hedge of the short position is the same no-arbitrage price obtained by the first pricing method.

The no-arbitrage price of the derivative security that pays V_N at time N can be computed recursively, backward in time, by the formula

$$V_n(\omega_1\omega_2\ldots\omega_n) = \frac{1}{1+r}[\tilde{p}V_{n+1}(\omega_1\omega_2\ldots\omega_n H) + \tilde{q}V_{n+1}(\omega_1\omega_2\ldots\omega_n T)]. \tag{1.2.16}$$

The number of shares of the stock that should be held by a portfolio hedging a short position in the derivative security is given by

$$\Delta_n(\omega_1 \ldots \omega_n) = \frac{V_{n+1}(\omega_1 \ldots \omega_n H) - V_{n+1}(\omega_1 \ldots \omega_n T)}{S_{n+1}(\omega_1 \ldots \omega_n H) - S_{n+1}(\omega_1 \ldots \omega_n T)}. \tag{1.2.17}$$

The numbers $\tilde{p}$ and $\tilde{q}$ appearing in (1.2.16) are the *risk-neutral probabilities* given by

$$\tilde{p} = \frac{1+r-d}{u-d}, \quad \tilde{q} = \frac{u-1-r}{u-d}. \tag{1.2.15}$$

These risk-neutral probabilities are positive because of (1.1.2) and sum to 1. They have the property that, at any time, the price of the stock is the discounted risk-neutral average of its two possible prices at the next time:

$$S_n(\omega_1 \ldots \omega_n) = \frac{1}{1+r}\left[\tilde{p}S_{n+1}(\omega_1 \ldots \omega_n H) + \tilde{q}S_{n+1}(\omega_1 \ldots \omega_n T)\right].$$

In other words, under the risk-neutral probabilities, the mean rate of return for the stock is r, the same as the rate of return for the money market. Therefore, if these probabilities actually governed the coin tossing (in fact, they do not), then an agent trading in the money market account and stock would have before him two opportunities, both of which provide the same mean rate of return. Consequently, no matter how he invests, the mean rate of return for his portfolio would also be r. In particular, if it is time $N-1$ and he wants his portfolio value to be $V_N(\omega_1 \ldots \omega_N)$ at time N, then at time $N-1$ his portfolio value must be

$$\frac{1}{1+r}\left[\tilde{p}V_N(\omega_1 \ldots \omega_{N-1} H) + \tilde{q}V_N(\omega_1 \ldots \omega_{N-1} T)\right].$$

This is the right-hand side of (1.2.16) with $n = N-1$, and repeated application of this argument yields (1.2.16) for all values of n.

The explanation of (1.2.16) above was given under a condition contrary to fact, namely that $\tilde{p}$ and $\tilde{q}$ govern the coin tossing. One can ask whether such an argument can result in a valid conclusion. It does result in a valid conclusion for the following reason. When hedging a short position in a derivative security, we want the hedge to give us a portfolio that agrees with the payoff of the derivative security *regardless of the coin tossing.* In other words, the hedge must work *on all stock price paths.* If a path is possible (i.e., has positive probability), we want the hedge to work along that path. The actual value of the probability is irrelevant. We find these hedges by solving a system of equations along the paths, a system of the form (1.2.2)–(1.2.3), (1.2.5)–(1.2.8). There are no probabilities in this system. Introducing the risk-neutral probabilities allows us to argue as above and find a solution to the system. Introducing any other probabilities would not allow such an argument because only the risk-neutral probabilities allow us to state that no matter how the agent invests, the mean rate of return for his portfolio is r. The risk-neutral probabilities provide a shortcut to solving the system of equations. The actual probabilities are no help in solving this system. Under the actual probabilities, the mean rate of return for a portfolio depends on the portfolio, and when we are trying to solve the system of equations, we do not know which portfolio we should use.

Alternatively, one can explain (1.2.16) without recourse to any discussion of probability. This was the approach taken in the proof of Theorem 1.2.2. The numbers $\tilde{p}$ and $\tilde{q}$ were used in that proof, but they were not regarded as probabilities, just numbers defined by the formula (1.2.15).

1.5 Notes

No-arbitrage pricing is implicit in the work of Black and Scholes [5], but its first explicit development is provided by Merton [34], who began with the axiom of no-arbitrage and obtained a surprising number of conclusions. No arbitrage pricing was fully developed in continuous-time models by Harrison and Kreps [17] and Harrison and Pliska [18]. These authors introduced martingales (Sections 2.4 in this text and Section 2.3 in Volume II) and risk-neutral pricing. The binomial model is due to Cox, Ross, Rubinstein [11]; a good reference is [12]. The binomial model is useful in its own right, and as Cox et al. showed, one can rederive the Black-Scholes-Merton formula as a limit of the binomial model (see Theorem 3.2.2 in Chapter 3 of Volume II for the log-normality of the stock price obtained in the limit of the binomial model.)

1.6 Exercises

Exercise 1.1. Assume in the one-period binomial market of Section 1.1 that both H and T have positive probability of occurring. Show that condition (1.1.2) precludes arbitrage. In other words, show that if $X_0 = 0$ and

$$X_1 = \Delta_0 S_1 + (1+r)(X_0 - \Delta_0 S_0),$$

then we cannot have X_1 strictly positive with positive probability unless X_1 is strictly negative with positive probability as well, and this is the case regardless of the choice of the number Δ_0.

Exercise 1.2. Suppose in the situation of Example 1.1.1 that the option sells for 1.20 at time zero. Consider an agent who begins with wealth $X_0 = 0$ and at time zero buys Δ_0 shares of stock and Γ_0 options. The numbers Δ_0 and Γ_0 can be either positive or negative or zero. This leaves the agent with a cash position of $-4\Delta_0 - 1.20\Gamma_0$. If this is positive, it is invested in the money market; if it is negative, it represents money borrowed from the money market. At time one, the value of the agent's portfolio of stock, option, and money market assets is

$$X_1 = \Delta_0 S_1 + \Gamma_0 (S_1 - 5)^+ - \frac{5}{4}(4\Delta_0 + 1.20\Gamma_0).$$

Assume that both H and T have positive probability of occurring. Show that if there is a positive probability that X_1 is positive, then there is a positive probability that X_1 is negative. In other words, one cannot find an arbitrage when the time-zero price of the option is 1.20.

Exercise 1.3. In the one-period binomial model of Section 1.1, suppose we want to determine the price at time zero of the derivative security $V_1 = S_1$ (i.e., the derivative security pays off the stock price.) (This can be regarded as a European call with strike price $K = 0$). What is the time-zero price V_0 given by the risk-neutral pricing formula (1.1.10)?

Exercise 1.4. In the proof of Theorem 1.2.2, show under the induction hypothesis that

$$X_{n+1}(\omega_1\omega_2\ldots\omega_n T) = V_{n+1}(\omega_1\omega_2\ldots\omega_n T).$$

Exercise 1.5. In Example 1.2.4, we considered an agent who sold the lookback option for $V_0 = 1.376$ and bought $\Delta_0 = 0.1733$ shares of stock at time zero. At time one, if the stock goes up, she has a portfolio valued at $V_1(H) = 2.24$. Assume that she now takes a position of $\Delta_1(H) = \frac{V_2(HH)-V_2(HT)}{S_2(HH)-S_2(HT)}$ in the stock. Show that, at time two, if the stock goes up again, she will have a portfolio valued at $V_2(HH) = 3.20$, whereas if the stock goes down, her portfolio will be worth $V_2(HT) = 2.40$. Finally, under the assumption that the stock goes up in the first period and down in the second period, assume the agent takes a position of $\Delta_2(HT) = \frac{V_3(HTH)-V_3(HTT)}{S_3(HTH)-S_3(HTT)}$ in the stock. Show that, at time three, if the stock goes up in the third period, she will have a portfolio valued at $V_3(HTH) = 0$, whereas if the stock goes down, her portfolio will be worth $V_3(HTT) = 6$. In other words, she has hedged her short position in the option.

Exercise 1.6 (Hedging a long position-one period). Consider a bank that has a long position in the European call written on the stock price in Figure 1.1.2. The call expires at time one and has strike price $K = 5$. In Section 1.1, we determined the time-zero price of this call to be $V_0 = 1.20$. At time zero, the bank owns this option, which ties up capital $V_0 = 1.20$. The bank wants to earn the interest rate 25% on this capital until time one (i.e., without investing any more money, and regardless of how the coin tossing turns out, the bank wants to have

$$\frac{5}{4} \cdot 1.20 = 1.50$$

at time one, after collecting the payoff from the option (if any) at time one). Specify how the bank's trader should invest in the stock and money markets to accomplish this.

Exercise 1.7 (Hedging a long position-multiple periods). Consider a bank that has a long position in the lookback option of Example 1.2.4. The bank intends to hold this option until expiration and receive the payoff V_3. At time zero, the bank has capital $V_0 = 1.376$ tied up in the option and wants to earn the interest rate of 25% on this capital until time three (i.e., without investing any more money, and regardless of how the coin tossing turns out, the bank wants to have

$$\left(\frac{5}{4}\right)^3 \cdot 1.376 = 2.6875$$

at time three, after collecting the payoff from the lookback option at time three). Specify how the bank's trader should invest in the stock and the money market account to accomplish this.

Exercise 1.8 (Asian option). Consider the three-period model of Example 1.2.1, with $S_0 = 4$, $u = 2$, $d = \frac{1}{2}$, and take the interest rate $r = \frac{1}{4}$, so that $\tilde{p} = \tilde{q} = \frac{1}{2}$. For $n = 0, 1, 2, 3$, define $Y_n = \sum_{k=0}^{n} S_k$ to be the sum of the stock prices between times zero and n. Consider an *Asian call option* that expires at time three and has strike $K = 4$ (i.e., whose payoff at time three is $\left(\frac{1}{4}Y_3 - 4\right)^+$). This is like a European call, except the payoff of the option is based on the average stock price rather than the final stock price. Let $v_n(s, y)$ denote the price of this option at time n if $S_n = s$ and $Y_n = y$. In particular, $v_3(s, y) = \left(\frac{1}{4}y - 4\right)^+$.

(i) Develop an algorithm for computing v_n recursively. In particular, write a formula for v_n in terms of v_{n+1}.
(ii) Apply the algorithm developed in (i) to compute $v_0(4, 4)$, the price of the Asian option at time zero.
(iii) Provide a formula for $\delta_n(s, y)$, the number of shares of stock that should be held by the replicating portfolio at time n if $S_n = s$ and $Y_n = y$.

Exercise 1.9 (Stochastic volatility, random interest rate). Consider a binomial pricing model, but at each time $n \geq 1$, the "up factor" $u_n(\omega_1\omega_2\ldots\omega_n)$, the "down factor" $d_n(\omega_1\omega_2\ldots\omega_n)$, and the interest rate $r_n(\omega_1\omega_2\ldots\omega_n)$ are allowed to depend on n and on the first n coin tosses $\omega_1\omega_2\ldots\omega_n$. The initial up factor u_0, the initial down factor d_0, and the initial interest rate r_0 are not random. More specifically, the stock price at time one is given by

$$S_1(\omega_1) = \begin{cases} u_0 S_0 & \text{if } \omega_1 = H, \\ d_0 S_0 & \text{if } \omega_1 = T, \end{cases}$$

and, for $n \geq 1$, the stock price at time $n + 1$ is given by

$$S_{n+1}(\omega_1\omega_2\ldots\omega_n\omega_{n+1}) = \begin{cases} u_n(\omega_1\omega_2\ldots\omega_n)S_n(\omega_1\omega_2\ldots\omega_n) & \text{if } \omega_{n+1} = H, \\ d_n(\omega_1\omega_2\ldots\omega_n)S_n(\omega_1\omega_2\ldots\omega_n) & \text{if } \omega_{n+1} = T. \end{cases}$$

One dollar invested in or borrowed from the money market at time zero grows to an investment or debt of $1 + r_0$ at time one, and, for $n \geq 1$, one dollar invested in or borrowed from the money market at time n grows to an investment or debt of $1 + r_n(\omega_1\omega_2 \ldots \omega_n)$ at time $n + 1$. We assume that for each n and for all $\omega_1\omega_2 \ldots \omega_n$, the no-arbitrage condition

$$0 < d_n(\omega_1\omega_2 \ldots \omega_n) < 1 + r_n(\omega_1\omega_2 \ldots \omega_n) < u_n(\omega_1\omega_2 \ldots \omega_n)$$

holds. We also assume that $0 < d_0 < 1 + r_0 < u_0$.

(i) Let N be a positive integer. In the model just described, provide an algorithm for determining the price at time zero for a derivative security that at time N pays off a random amount V_N depending on the result of the first N coin tosses.

(ii) Provide a formula for the number of shares of stock that should be held at each time n $(0 \leq n \leq N - 1)$ by a portfolio that replicates the derivative security V_N.

(iii) Suppose the initial stock price is $S_0 = 80$, with each head the stock price increases by 10, and with each tail the stock price decreases by 10. In other words, $S_1(H) = 90$, $S_1(T) = 70$, $S_2(HH) = 100$, etc. Assume the interest rate is always zero. Consider a European call with strike price 80, expiring at time five. What is the price of this call at time zero?