

Option Pricing: A Simplified Approach[†]

John C. Cox
Massachusetts Institute of Technology and Stanford University

Stephen A. Ross
Yale University

Mark Rubinstein
University of California, Berkeley

March 1979 (revised July 1979)
(published under the same title in *Journal of Financial Economics* (September 1979))

[1978 winner of the Pomeranze Prize of the Chicago Board Options Exchange]

[reprinted in *Dynamic Hedging: A Guide to Portfolio Insurance*, edited by Don Luskin (John Wiley and Sons 1988)]

[reprinted in *The Handbook of Financial Engineering*, edited by Cliff Smith and Charles Smithson (Harper and Row 1990)]

[reprinted in *Readings in Futures Markets* published by the Chicago Board of Trade, Vol. VI (1991)]

[reprinted in *Vasicek and Beyond: Approaches to Building and Applying Interest Rate Models*, edited by Risk Publications, Alan Brace (1996)]

[reprinted in *The Debt Market*, edited by Stephen Ross and Franco Modigliani (Edward Lear Publishing 2000)]

[reprinted in *The International Library of Critical Writings in Financial Economics: Options Markets* edited by G.M. Constantinides and A.G. Malliaris (Edward Lear Publishing 2000)]

Abstract

This paper presents a simple discrete-time model for valuing options. The fundamental economic principles of option pricing by arbitrage methods are particularly clear in this setting. Its development requires only elementary mathematics, yet it contains as a special limiting case the celebrated Black-Scholes model, which has previously been derived only by much more difficult methods. The basic model readily lends itself to generalization in many ways. Moreover, by its very construction, it gives rise to a simple and efficient numerical procedure for valuing options for which premature exercise may be optimal.

[†] Our best thanks go to William Sharpe, who first suggested to us the advantages of the discrete-time approach to option pricing developed here. We are also grateful to our students over the past several years. Their favorable reactions to this way of presenting things encouraged us to write this article. We have received support from the National Science Foundation under Grants Nos. SOC-77-18087 and SOC-77-22301.

1. Introduction

An option is a security that gives its owner the right to trade in a fixed number of shares of a specified common stock at a fixed price at any time on or before a given date. The act of making this transaction is referred to as exercising the option. The fixed price is termed the strike price, and the given date, the expiration date. A call option gives the right to buy the shares; a put option gives the right to sell the shares.

Options have been traded for centuries, but they remained relatively obscure financial instruments until the introduction of a listed options exchange in 1973. Since then, options trading has enjoyed an expansion unprecedented in American securities markets.

Option pricing theory has a long and illustrious history, but it also underwent a revolutionary change in 1973. At that time, Fischer Black and Myron Scholes presented the first completely satisfactory equilibrium option pricing model. In the same year, Robert Merton extended their model in several important ways. These path-breaking articles have formed the basis for many subsequent academic studies.

As these studies have shown, option pricing theory is relevant to almost every area of finance. For example, virtually all corporate securities can be interpreted as portfolios of puts and calls on the assets of the firm.¹ Indeed, the theory applies to a very general class of economic problems — the valuation of contracts where the outcome to each party depends on a quantifiable uncertain future event.

Unfortunately, the mathematical tools employed in the Black-Scholes and Merton articles are quite advanced and have tended to obscure the underlying economics. However, thanks to a suggestion by William Sharpe, it is possible to derive the same results using only elementary mathematics.²

In this article we will present a simple discrete-time option pricing formula. The fundamental economic principles of option valuation by arbitrage methods are particularly clear in this setting. Sections 2 and 3 illustrate and develop this model for a call option on a stock that pays no dividends. Section 4 shows exactly how the model can be used to lock in pure arbitrage profits if the market price of an option differs from the value given by the model. In section 5, we will show that our approach includes the Black-Scholes model as a special limiting case. By taking the limits in a different way, we will also obtain the Cox-Ross (1975) jump process model as another special case.

¹ To take an elementary case, consider a firm with a single liability of a homogeneous class of pure discount bonds. The stockholders then have a “call” on the assets of the firm which they can choose to exercise at the maturity date of the debt by paying its principal to the bondholders. In turn, the bonds can be interpreted as a portfolio containing a default-free loan with the same face value as the bonds and a short position in a put on the assets of the firm.

² Sharpe (1978) has partially developed this approach to option pricing in his excellent new book, *Investments*. Rendleman and Bartter (1978) have recently independently discovered a similar formulation of the option pricing problem.

Other more general option pricing problems often seem immune to reduction to a simple formula. Instead, numerical procedures must be employed to value these more complex options. Michael Brennan and Eduardo Schwartz (1977) have provided many interesting results along these lines. However, their techniques are rather complicated and are not directly related to the economic structure of the problem. Our formulation, by its very construction, leads to an alternative numerical procedure that is both simpler, and for many purposes, computationally more efficient.

Section 6 introduces these numerical procedures and extends the model to include puts and calls on stocks that pay dividends. Section 7 concludes the paper by showing how the model can be generalized in other important ways and discussing its essential role in valuation by arbitrage methods.

2. The Basic Idea

Suppose the current price of a stock is $S = \$50$, and at the end of a period of time, its price must be either $S^* = \$25$ or $S^* = \$100$. A call on the stock is available with a strike price of $K = \$50$, expiring at the end of the period.³ It is also possible to borrow and lend at a 25% rate of interest. The one piece of information left unfurnished is the current value of the call, C . However, if riskless profitable arbitrage is not possible, we can deduce from the given information *alone* what the value of the call *must* be!

Consider the following levered hedge:

- (1) write 3 calls at C each,
- (2) buy 2 shares at \$50 each, and
- (3) borrow \$40 at 25%, to be paid back at the end of the period.

Table 1 gives the return from this hedge for each possible level of the stock price at expiration. Regardless of the outcome, the hedge exactly breaks even on the expiration date. Therefore, to prevent profitable riskless arbitrage, its current cost must be zero; that is,

$$3C - 100 + 40 = 0$$

The current value of the call must then be $C = \$20$.

³ To keep matters simple, assume for now that the stock will pay no cash dividends during the life of the call. We also ignore transaction costs, margin requirements and taxes.

Table 1
Arbitrage Table Illustrating the Formation of a Riskless Hedge

	present date	expiration date	
		$S^* = \$25$	$S^* = \$100$
write 3 calls	$3C$	—	-150
buy 2 shares	-100	50	200
borrow	40	-50	-50
total		—	—

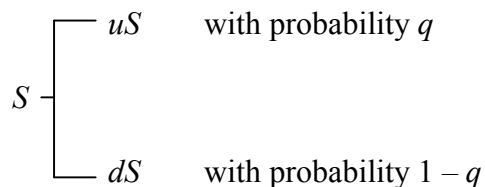
If the call were not priced at \$20, a sure profit would be possible. In particular, if $C = \$25$, the above hedge would yield a current cash inflow of \$15 and would experience no further gain or loss in the future. On the other hand, if $C = \$15$, then the same thing could be accomplished by buying 3 calls, selling short 2 shares, and lending \$40.

Table 1 can be interpreted as demonstrating that *an appropriately levered position in stock will replicate the future returns of a call*. That is, if we buy shares and borrow against them in the right proportion, we can, in effect, duplicate a pure position in calls. In view of this, it should seem less surprising that all we needed to determine the *exact* value of the call was its *strike price, underlying stock price, range of movement in the underlying stock price, and the rate of interest*. What may seem more incredible is what we do not need to know: among other things, *we do not need to know the probability that the stock price will rise or fall*. Bulls and bears must agree on the value of the call, relative to its underlying stock price!

This example is very simple, but it shows several essential features of option pricing. And we will soon see that it is not as unrealistic as it seems.

3. The Binomial Option Pricing Formula

In this section, we will develop the framework illustrated in the example into a complete valuation method. We begin by assuming that the stock price follows a multiplicative binomial process over discrete periods. The rate of return on the stock over each period can have two possible values: $u - 1$ with probability q , or $d - 1$ with probability $1 - q$. Thus, if the current stock price is S , the stock price at the end of the period will be either uS or dS . We can represent this movement with the following diagram:



We also assume that the interest rate is constant. Individuals may borrow or lend as much as they wish at this rate. To focus on the basic issues, we will continue to assume that there are no

taxes, transaction costs, or margin requirements. Hence, individuals are allowed to sell short any security and receive full use of the proceeds.⁴

Letting r denote one plus the riskless interest rate over one period, we require $u > r > d$. If these inequalities did not hold, there would be profitable riskless arbitrage opportunities involving only the stock and riskless borrowing and lending.⁵

To see how to value a call on this stock, we start with the simplest situation: the expiration date is just one period away. Let C be the current value of the call, C_u be its value at the end of the period if the stock price goes to uS and C_d be its value at the end of the period if the stock price goes to dS . Since there is now only one period remaining in the life of the call, we know that the terms of its contract and a rational exercise policy imply that $C_u = \max[0, uS - K]$ and $C_d = \max[0, dS - K]$. Therefore,

$$C \begin{cases} C_u = \max[0, uS - K] & \text{with probability } q \\ C_d = \max[0, dS - K] & \text{with probability } 1 - q \end{cases}$$

Suppose we form a portfolio containing Δ shares of stock and the dollar amount B in riskless bonds.⁶ This will cost $\Delta S + B$. At the end of the period, the value of this portfolio will be

$$\Delta S + B \begin{cases} \Delta uS + rB & \text{with probability } q \\ \Delta dS + rB & \text{with probability } 1 - q \end{cases}$$

Since we can select Δ and B in any way we wish, suppose we choose them to equate the end-of-period values of the portfolio and the call for each possible outcome. This requires that

$$\begin{aligned} \Delta uS + rB &= C_u \\ \Delta dS + rB &= C_d \end{aligned}$$

Solving these equations, we find

$$\Delta = \frac{C_u - C_d}{(u - d)S}, \quad B = \frac{uC_d - dC_u}{(u - d)r} \quad (1)$$

⁴ Of course, restitution is required for payouts made to securities held short.

⁵ We will ignore the uninteresting special case where q is zero or one and $u = d = r$.

⁶ Buying bonds is the same as lending; selling them is the same as borrowing.

With Δ and B chosen in this way, we will call this the hedging portfolio.

If there are to be no riskless arbitrage opportunities, the current value of the call, C , cannot be less than the current value of the hedging portfolio, $\Delta S + B$. If it were, we could make a riskless profit with no net investment by buying the call and selling the portfolio. It is tempting to say that it also cannot be worth more, since then we would have a riskless arbitrage opportunity by reversing our procedure and selling the call and buying the portfolio. But this overlooks the fact that the person who bought the call we sold has the right to exercise it immediately.

Suppose that $\Delta S + B < S - K$. If we try to make an arbitrage profit by selling calls for more than $\Delta S + B$, but less than $S - K$, then we will soon find that we are the source of arbitrage profits rather than the recipient. Anyone could make an arbitrage profit by buying our calls and exercising them immediately.

We might hope that we will be spared this embarrassment because everyone will somehow find it advantageous to hold the calls for one more period as an investment rather than take a quick profit by exercising them immediately. But each person will reason in the following way. If I do not exercise now, I will receive the same payoff as a portfolio with ΔS in stock and B in bonds. If I do exercise now, I can take the proceeds, $S - K$, buy this same portfolio and some extra bonds as well, and have a higher payoff in every possible circumstance. Consequently, no one would be willing to hold the calls for one more period.

Summing up all of this, we conclude that if there are to be no riskless arbitrage opportunities, it must be true that

$$C = \Delta S + B = \frac{C_u - C_d}{u - d} + \frac{uC_d - dC_u}{(u - d)r} = \left[\left(\frac{r - d}{u - d} \right) C_u + \left(\frac{u - r}{u - d} \right) C_d \right] / r \quad (2)$$

if this value is greater than $S - K$, and if not, $C = S - K$.⁷

Equation (2) can be simplified by defining

$$p \equiv \frac{r - d}{u - d} \quad \text{and} \quad 1 - p \equiv \frac{u - r}{u - d}$$

so that we can write

$$C = [pC_u + (1 - p)C_d] / r \quad (3)$$

It is easy to see that in the present case, with no dividends, this will always be greater than $S - K$ as long as the interest rate is positive. To avoid spending time on the unimportant situations where the interest rate is less than or equal to zero, we will now assume that r is always greater

⁷ In some applications of the theory to other areas, it is useful to consider options that can be exercised only on the expiration date. These are usually termed European options. Those that can be exercised at any earlier time as well, such as we have been examining here, are then referred to as American options. Our discussion could be easily modified to include European calls. Since immediate exercise is then precluded, their values would always be given by (2), even if this is less than $S - K$.

than one. Hence, (3) is the exact formula for the value of a call one period prior to the expiration in terms of S , K , u , d , and r .

To confirm this, note that if $uS \leq K$, then $S < K$ and $C = 0$, so $C > S - K$. Also, if $dS \geq K$, then $C = S - (K/r) > S - K$. The remaining possibility is $uS > K > dS$. In this case, $C = p(uS - K)/r$. This is greater than $S - K$ if $(1 - p)dS > (p - r)K$, which is certainly true as long as $r > 1$.

This formula has a number of notable features. First, the probability q does not appear in the formula. This means, surprisingly, that even if different investors have different subjective probabilities about an upward or downward movement in the stock, they could still agree on the relationship of C to S , u , d , and r .

Second, the value of the call does not depend on investors' attitudes toward risk. In constructing the formula, the only assumption we made about an individual's behavior was that he prefers more wealth to less wealth and therefore has an incentive to take advantage of profitable riskless arbitrage opportunities. We would obtain the same formula whether investors are risk-averse or risk-preferring.

Third, the only random variable on which the call value depends is the stock price itself. In particular, it does not depend on the random prices of other securities or portfolios, such as the market portfolio containing all securities in the economy. If another pricing formula involving other variables was submitted as giving equilibrium market prices, we could immediately show that it was incorrect by using our formula to make riskless arbitrage profits while trading at those prices.

It is easier to understand these features if it is remembered that the formula is only a relative pricing relationship giving C in terms of S , u , d , and r . Investors' attitudes toward risk and the characteristics of other assets may indeed influence call values indirectly, through their effect on these variables, but they will not be separate determinants of call value.

Finally, observe that $p \equiv (r - d)/(u - d)$ is always greater than zero and less than one, so it has the properties of a probability. In fact, p is the value q would have in equilibrium if investors were risk-neutral. To see this, note that the expected rate of return on the stock would then be the riskless interest rate, so

$$q(uS) + (1 - q)(dS) = rS$$

and

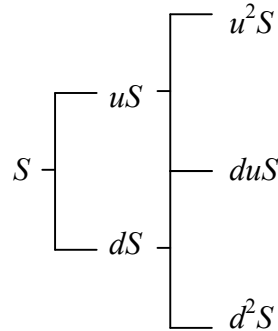
$$q = (r - d)/(u - d) = p$$

Hence, the value of the call can be interpreted as the expectation of its discounted future value in a risk-neutral world. In light of our earlier observations, this is not surprising. Since the formula does not involve q or any measure of attitudes toward risk, then it must be the same for any set of preferences, including risk neutrality.

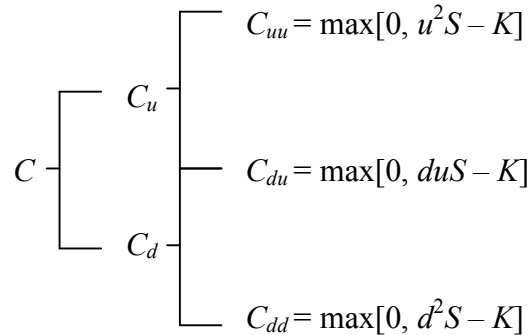
It is important to note that this does not imply that the equilibrium expected rate of return on the call is the riskless interest rate. Indeed, our argument has shown that, in equilibrium, holding the call over the period is exactly equivalent to holding the hedging portfolio. Consequently, the risk

and expected rate of return of the call must be the same as that of the hedging portfolio. It can be shown that $\Delta \geq 0$ and $B \leq 0$, so the hedging portfolio is equivalent to a particular levered long position in the stock. In equilibrium, the same is true for the call. Of course, if the call is currently mispriced, its risk and expected return over the period will differ from that of the hedging portfolio.

Now we can consider the next simplest situation: a call with two periods remaining before its expiration date. In keeping with the binomial process, the stock can take on three possible values after two periods,



Similarly, for the call,



C_{uu} stands for the value of a call two periods from the current time if the stock price moves upward each period; C_{du} and C_{dd} have analogous definitions.

At the end of the current period there will be one period left in the life of the call, and we will be faced with a problem identical to the one we just solved. Thus, from our previous analysis, we know that when there are two periods left,

$$C_u = [pC_{uu} + (1 - p)C_{du}]/r \tag{4}$$

and

$$C_d = [pC_{du} + (1 - p)C_{dd}]/r$$

Again, we can select a portfolio with ΔS in stock and B in bonds whose end-of-period value will be C_u if the stock price goes to uS and C_d if the stock price goes to dS . Indeed, the

functional form of Δ and B remains unchanged. To get the new values of Δ and B , we simply use equation (1) with the new values of C_u and C_d .

Can we now say, as before, that an opportunity for profitable riskless arbitrage will be available if the current price of the call is not equal to the new value of this portfolio or $S - K$, whichever is greater? Yes, but there is an important difference. With one period to go, we could plan to lock in a riskless profit by selling an overpriced call and using part of the proceeds to buy the hedging portfolio. At the end of the period, we knew that the market price of the call must be equal to the value of the portfolio, so the entire position could be safely liquidated at that point. But this was true only because the end of the period was the expiration date. Now we have no such guarantee. At the end of the current period, when there is still one period left, the market price of the call could still be in disequilibrium and be greater than the value of the hedging portfolio. If we closed out the position then, selling the portfolio and repurchasing the call, we could suffer a loss that would more than offset our original profit. However, we could always avoid this loss by maintaining the portfolio for one more period. The value of the portfolio at the end of the current period will always be exactly sufficient to purchase the portfolio we would want to hold over the last period. In effect, we would have to readjust the proportions in the hedging portfolio, but we would not have to put up any more money.

Consequently, we conclude that even with two periods to go, there is a strategy we could follow which would guarantee riskless profits with no net investment if the current market price of a call differs from the maximum of $\Delta S + B$ and $S - K$. Hence, the larger of these is the current value of the call.

Since Δ and B have the same functional form in each period, the current value of the call in terms of C_u and C_d will again be $C = [pC_u + (1 - p)C_d]/r$ if this is greater than $S - K$, and $C = S - K$ otherwise. By substituting from equation (4) into the former expression, and noting that $C_{du} = C_{ud}$, we obtain

$$C = [p^2 C_{uu} + 2p(1 - p)C_{ud} + (1 - p)^2 C_{dd}]/r^2 \tag{5}$$

$$= [p^2 \max[0, u^2 S - K] + 2p(1 - p) \max[0, duS - K] + (1 - p)^2 \max[0, d^2 S - K]]/r^2$$

A little algebra shows that this is always greater than $S - K$ if, as assumed, r is always greater than one, so this expression gives the exact value of the call.⁸

All of the observations made about formula (3) also apply to formula (5), except that the number of periods remaining until expiration, n , now emerges clearly as an additional determinant of the call value. For formula (5), $n = 2$. That is, the full list of variables determining C is S, K, n, u, d , and r .

⁸ In the current situation, with no dividends, we can show by a simple direct argument that if there are no arbitrage opportunities, then the call value must always be greater than $S - K$ before the expiration date. Suppose that the call is selling for $S - K$. Then there would be an easy arbitrage strategy that would require no initial investment and would always have a positive return. All we would have to do is buy the call, short the stock, and invest K dollars in bonds. See Merton (1973). In the general case, with dividends, such an argument is no longer valid, and we must use the procedure of checking every period.

We now have a recursive procedure for finding the value of a call with any number of periods to go. By starting at the expiration date and working backwards, we can write down the general valuation formula for any n :

$$C = \left[\sum_{j=0}^n \left(\frac{n!}{j!(n-j)!} \right) p^j (1-p)^{n-j} \max[0, u^j d^{n-j} S - K] \right] / r^n \quad (6)$$

This gives us the complete formula, but with a little additional effort we can express it in a more convenient way.

Let a stand for the minimum number of upward moves that the stock must make over the next n periods for the call to finish in-the-money. Thus a will be the smallest non-negative integer such that $u^a d^{n-a} S > K$. By taking the natural logarithm of both sides of this inequality, we could write a as the smallest non-negative integer greater than $\log(K/Sd^n)/\log(u/d)$.

For all $j < a$,

$$\max[0, u^j d^{n-j} S - K] = 0$$

and for all $j \geq a$,

$$\max[0, u^j d^{n-j} S - K] = u^j d^{n-j} S - K$$

Therefore,

$$C = \left[\sum_{j=a}^n \left(\frac{n!}{j!(n-j)!} \right) p^j (1-p)^{n-j} [u^j d^{n-j} S - K] \right] / r^n$$

Of course, if $a > n$, the call will finish out-of-the-money even if the stock moves upward every period, so its current value must be zero.

By breaking up C into two terms, we can write

$$C = S \left[\sum_{j=a}^n \left(\frac{n!}{j!(n-j)!} \right) p^j (1-p)^{n-j} \left(\frac{u^j d^{n-j}}{r^n} \right) \right] - Kr^{-n} \left[\sum_{j=a}^n \left(\frac{n!}{j!(n-j)!} \right) p^j (1-p)^{n-j} \right]$$

Now, the latter bracketed expression is the complementary binomial distribution function $\phi[a; n, p]$. The first bracketed expression can also be interpreted as a complementary binomial distribution function $\phi[a; n, p']$, where

$$p' \equiv (u/r)p \text{ and } 1-p' \equiv (d/r)(1-p)$$

p' is a probability, since $0 < p' < 1$. To see this, note that $p < (r/u)$ and

$$p^j (1-p)^{n-j} \left(\frac{u^j d^{n-j}}{r^n} \right) = \left[\frac{u}{r} p \right]^j \left[\frac{d}{r} (1-p) \right]^{n-j} = p'^j (1-p')^{n-j}$$

In summary:

Binomial Option Pricing Formula

$$C = S\phi[a; n, p] - Kr^{-n}\phi[a; n, p]$$

where

$$p \equiv (r - d)/(u - d) \text{ and } p' \equiv (u/r)p$$

$a \equiv$ the smallest non-negative integer
greater than $\log(K/Sd^n)/\log(u/d)$

If $a > n$, then $C = 0$.

It is now clear that all of the comments we made about the one period valuation formula are valid for any number of periods. In particular, the value of a call should be the expectation, in a risk-neutral world, of the discounted value of the payoff it will receive. In fact, that is exactly what equation (6) says. Why, then, should we waste time with the recursive procedure when we can write down the answer in one direct step? The reason is that while this one-step approach is always technically correct, it is really useful only if we know in advance the circumstances in which a rational individual would prefer to exercise the call before the expiration date. If we do not know this, we have no way to compute the required expectation. In the present example, a call on a stock paying no dividends, it happens that we can determine this information from other sources: the call should never be exercised before the expiration date. As we will see in section 6, with puts or with calls on stocks that pay dividends, we will not be so lucky. Finding the optimal exercise strategy will be an integral part of the valuation problem. The full recursive procedure will then be necessary.

For some readers, an alternative “complete markets” interpretation of our binomial approach may be instructive. Suppose that π_u and π_d represent the state-contingent discount rates to states u and d , respectively. Therefore, π_u would be the current price of one dollar received at the end of the period, if and only if state u occurs. Each security — a riskless bond, the stock, and the option — must all have returns discounted to the present by π_u and π_d if no riskless arbitrage opportunities are available. Therefore,

$$\begin{aligned} 1 &= \pi_u r + \pi_d r \\ S &= \pi_u (uS) + \pi_d (dS) \\ C &= \pi_u C_u + \pi_d C_d \end{aligned}$$

The first two equations, for the bond and the stock, imply

$$\pi_u = \left(\frac{r-d}{u-d} \right) \frac{1}{r} \quad \text{and} \quad \pi_d = \left(\frac{u-r}{u-d} \right) \frac{1}{r}$$

Substituting these equalities for the state-contingent prices in the last equation for the option yields equation (3).

It is important to realize that we are not assuming that the riskless bond and the stock and the option are the only three securities in the economy, or that other securities must follow a binomial process. Rather, however these securities are priced in relation to others in equilibrium, among themselves they must conform to the above relationships.

From either the hedging or complete markets approaches, it should be clear that three-state or trinomial stock price movements will not lead to an option pricing formula based solely on arbitrage considerations. Suppose, for example, that over each period the stock price could move to uS or dS or remain the same at S . A choice of Δ and B that would equate the returns in two states could not in the third. That is, a riskless arbitrage position could not be taken. Under the complete markets interpretation, with three equations in now three unknown state-contingent prices, we would lack the redundant equation necessary to price one security in terms of the other two.

4. Riskless Trading Strategies

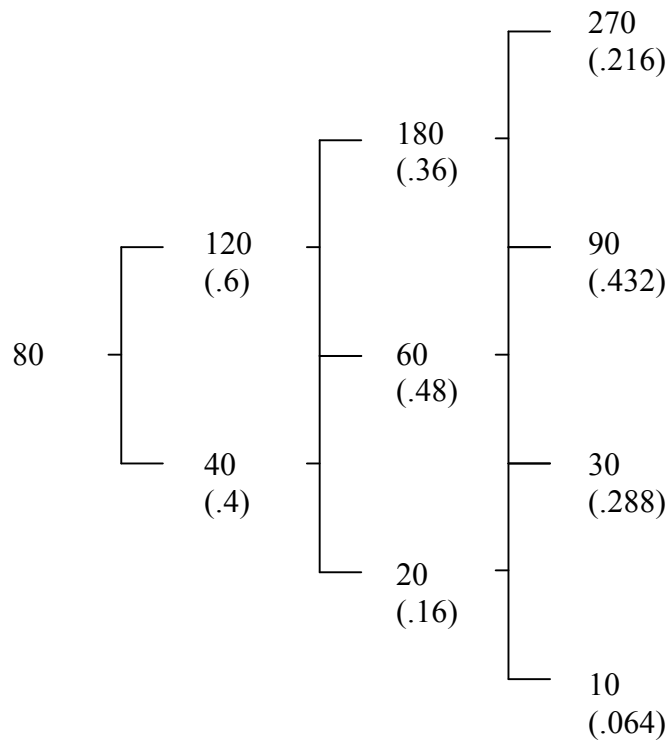
The following numerical example illustrates how we could use the formula if the current *market price* M ever diverged from its *formula value* C . If $M > C$, we would hedge, and if $M < C$, “reverse hedge”, to try and lock in a profit. Suppose the values of the underlying variables are

$$S = 80, \quad n = 3, \quad K = 80, \quad u = 1.5, \quad d = 0.5, \quad r = 1.1$$

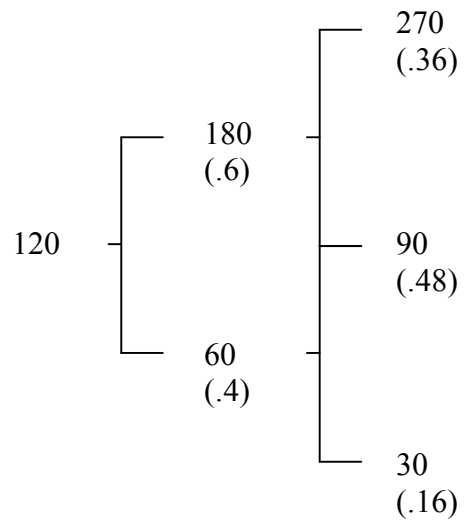
In this case, $p = (r-d)/(u-d) = 0.6$. The relevant values of the discount factor are

$$r^{-1} = 0.909, \quad r^{-2} = 0.826, \quad r^{-3} = 0.751$$

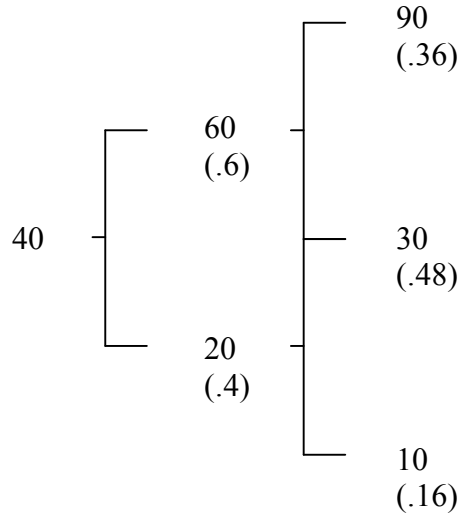
The paths the stock price may follow and their corresponding probabilities (using probability p) are, when $n = 3$, with $S = 80$,



when $n = 2$, if $S = 120$,



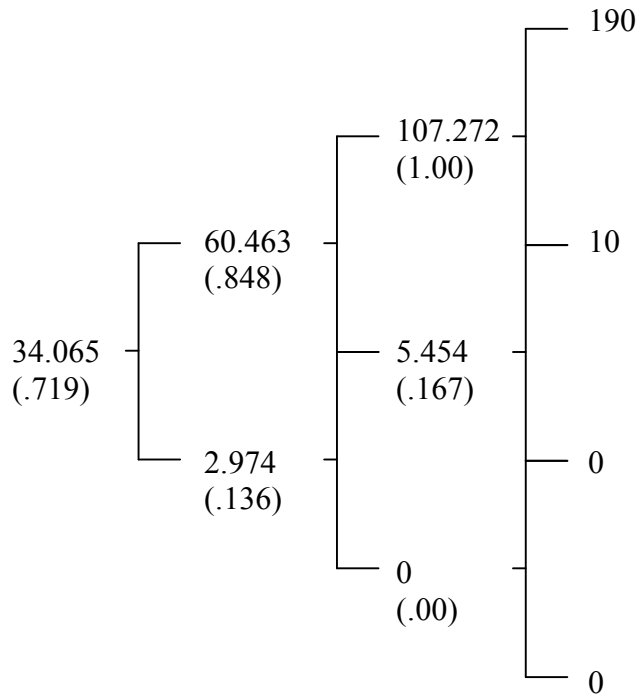
when $n = 2$, if $S = 40$,



Using the formula, the current value of the call would be

$$C = 0.751[0.064(0) + 0.288(0) + 0.432(90 - 80) + 0.216(270 - 80)] = 34.065.$$

Recall that to form a riskless hedge, for each call we sell, we buy and subsequently keep adjusted a portfolio with ΔS in stock and B in bonds, where $\Delta = (C_u - C_d)/(u - d)S$. The following tree diagram gives the paths the call value may follow and the corresponding values of Δ :



With this preliminary analysis, we are prepared to use the formula to take advantage of mispricing in the market. Suppose that when $n = 3$, the market price of the call is 36. Our formula tells us the call should be worth 34.065. The option is overpriced, so we could plan to sell it and assure ourselves of a profit equal to the mispricing differential. Here are the steps you could take for a typical path the stock might follow.

Step 1 ($n = 3$): Sell the call for 36. Take 34.065 of this and invest it in a portfolio containing $\Delta = 0.719$ shares of stock by borrowing $0.719(80) - 34.065 = 23.455$. Take the remainder, $36 - 34.065 = 1.935$, and put it in the bank.

Step 2 ($n = 2$): Suppose the stock goes to 120 so that the new Δ is 0.848. Buy $0.848 - 0.719 = 0.129$ more shares of stock at 120 per share for a total expenditure of 15.480. Borrow to pay the bill. With an interest rate of 0.1, you already owe $23.455(1.1) = 25.801$. Thus, your total current indebtedness is $25.801 + 15.480 = 41.281$.

Step 3 ($n = 1$): Suppose the stock price now goes to 60. The new Δ is 0.167. Sell $0.848 - 0.167 = 0.681$ shares at 60 per share, taking in $0.681(60) = 40.860$. Use this to pay back part of your borrowing. Since you now owe $41.281(1.1) = 45.409$, the repayment will reduce this to $45.409 - 40.860 = 4.549$.

Step 4d ($n = 0$): Suppose the stock price now goes to 30. The call you sold has expired worthless. You own 0.167 shares of stock selling at 30 per share, for a total value of $0.167(30) = 5$. Sell the stock and repay the $4.549(1.1) = 5$ that you now owe on the borrowing. Go back to the bank and withdraw your original deposit, which has now grown to $1.935(1.1)^3 = 2.341$.

Step 4u ($n = 0$): Suppose, instead, the stock price goes to 90. The call you sold is in the money at the expiration date. Buy back the call, or buy one share of stock and let it be exercised, incurring a loss of $90 - 80 = 10$ either way. Borrow to cover this, bringing your current indebtedness to $5 + 10 = 15$. You own 0.167 shares of stock selling at 90 per share, for a total value of $0.167(90) = 15$. Sell the stock and repay the borrowing. Go back to the bank and withdraw your original deposit, which has now grown to $1.935(1.1)^3 = 2.341$.

In summary, if we were correct in our original analysis about stock price movements (which did not involve the unenviable task of predicting whether the stock price would go up or down), and if we faithfully adjust our portfolio as prescribed by the formula, then we can be assured of walking away in the clear at the expiration date, while still keeping the original differential and the interest it has accumulated. It is true that closing out the position before the expiration date, which involves buying back the option at its then current market price, might produce a loss which would more than offset our profit, but this loss could always be avoided by waiting until the expiration date. Moreover, if the market price comes into line with the formula value before the expiration date, we can close out the position then with no loss and be rid of the concern of keeping the portfolio adjusted.

It still might seem that we are depending on rational behavior by the person who bought the call we sold. If instead he behaves foolishly and exercises at the wrong time, could he make things worse for us as well as for himself? Fortunately, the answer is no. Mistakes on his part can only

mean greater profits for us. Suppose that he exercises too soon. In that circumstance, the hedging portfolio will always be worth more than $S - K$, so we could close out the position then with an extra profit.

Suppose, instead, that he fails to exercise when it would be optimal to do so. Again there is no problem. Since exercise is now optimal, our hedging portfolio will be worth $S - K$.⁹ If he had exercised, this would be exactly sufficient to meet the obligation and close out the position. Since he did not, the call will be held at least one more period, so we calculate the new values of C_u and C_d and revise our hedging portfolio accordingly. But now the amount required for the portfolio, $\Delta S + B$, is less than the amount we have available, $S - K$. We can withdraw these extra profits now and still maintain the hedging portfolio. The longer the holder of the call goes on making mistakes, the better off we will be.

Consequently, we can be confident that things will eventually work out right no matter what the other party does. The return on our total position, when evaluated at prevailing market prices at intermediate times, may be negative. But over a period ending no later than the expiration date, it will be positive.

In conducting the hedging operation, the essential thing was to maintain the proper proportional relationship: for each call we are short, we hold Δ shares of stock and the dollar amount B in bonds in the hedging portfolio. To emphasize this, we will refer to the number of shares held for each call as the hedge ratio. In our example, we kept the number of calls constant and made adjustments by buying or selling stock and bonds. As a result, our profit was independent of the market price of the call between the time we initiated the hedge and the expiration date. If things got worse before they got better, it did not matter to us.

Instead, we could have made the adjustments by keeping the number of shares of stock constant and buying or selling calls and bonds. However, this could be dangerous. Suppose that after initiating the position, we needed to increase the hedge ratio to maintain the proper proportions. This can be achieved in two ways:

- (a) buy more stock, or
- (b) buy back some of the calls.

If we adjust through the stock, there is no problem. If we insist on adjusting through the calls, not only is the hedge no longer riskless, but it could even end up losing money! This can happen if the call has become even more overpriced. We would then be closing out part of our position in calls at a loss. To remain hedged, the number of calls we would need to buy back depends on their value, not their price. Therefore, since we are uncertain about their price, we then become uncertain about the return from the hedge. Worse yet, if the call price gets high enough, the loss on the closed portion of our position could throw the hedge operation into an overall loss.

⁹ If we were reverse hedging by buying an undervalued call and selling the hedging portfolio, then we would ourselves want to exercise at this point. Since we will receive $S - K$ from exercising, this will be exactly enough money to buy back the hedging portfolio.

To see how this could happen, let us rerun the hedging operation, where we adjust the hedge ratio by buying and selling calls.

Step 1 ($n = 3$): Same as before.

Step 2 ($n = 2$): Suppose the stock goes to 120, so that the new $\Delta = 0.848$. The call price has gotten further out of line and is now selling for 75. Since its value is 60.463, it is now overpriced by 14.537. With 0.719 shares, you must buy back $1 - 0.848 = 0.152$ calls to produce a hedge ratio of $0.848 = 0.719/0.848$. This costs $75(0.152) = 11.40$. Borrow to pay the bill. With the interest rate of 0.1, you already owe $23.455(1.1) = 25.801$. Thus, your total current indebtedness is $25.801 + 11.40 = 37.201$.

Step 3 ($n = 1$): Suppose the stock goes to 60 and the call is selling for 5.454. Since the call is now fairly valued, no further excess profits can be made by continuing to hold the position. Therefore, liquidate by selling your 0.719 shares for $0.719(60) = 43.14$ and close out the call position by buying back 0.848 calls for $0.848(5.454) = 4.625$. This nets $43.14 - 4.625 = 38.515$. Use this to pay back part of your borrowing. Since you now owe $37.20(1.1) = 40.921$, after repayment you owe 2.406. Go back to the bank and withdraw your original deposit, which has now grown to $1.935(1.1)^2 = 2.341$. Unfortunately, after using this to repay your remaining borrowing, you still owe 0.065.

Since we adjusted our position at Step 2 by buying overpriced calls, our profit is reduced. Indeed, since the calls were considerably overpriced, we actually lost money despite apparent profitability of the position at Step 1. We can draw the following adjustment rule from our experiment: *To adjust a hedged position, never buy an overpriced option or sell an underpriced option.* As a corollary, whenever we can adjust a hedged position by buying more of an underpriced option or selling more of an overpriced option, our profit will be enhanced if we do so. For example, at Step 3 in the original hedging illustration, had the call still been overpriced, it would have been better to adjust the position by selling more calls rather than selling stock. In summary, by choosing the right side of the position to adjust at intermediate dates, *at a minimum* we can be assured of earning the original differential and its accumulated interest, and we may earn considerably more.

5. Limiting Cases

In reading the previous sections, there is a natural tendency to associate with each period some particular length of calendar time, perhaps a day. With this in mind, you may have had two objections. In the first place, prices a day from now may take on many more than just two possible values. Furthermore, the market is not open for trading only once a day, but, instead, trading takes place almost continuously.

These objections are certainly valid. Fortunately, our option pricing approach has the flexibility to meet them. Although it might have been natural to think of a period as one day, there was nothing that forced us to do so. We could have taken it to be a much shorter interval — say an hour — or even a minute. By doing so, we have met both objections simultaneously. Trading

would take place far more frequently, and the stock price could take on hundreds of values by the end of the day.

However, if we do this, we have to make some other adjustments to keep the probability small that the stock price will change by a large amount over a minute. We do not want the stock to have the same percentage up and down moves for one minute as it did before for one day. But again there is no need for us to have to use the same values. We could, for example, think of the price as making only a very small percentage change over each minute.

To make this more precise, suppose that h represents the elapsed time between successive stock price changes. That is, if t is the fixed length of calendar time to expiration, and n is the number of periods of length h prior to expiration, then

$$h \equiv t/n$$

As trading takes place more and more frequently, h gets closer and closer to zero. We must then adjust the interval-dependent variables r , u , and d in such a way that we obtain empirically realistic results as h becomes smaller, or, equivalently, as $n \rightarrow \infty$.

When we were thinking of the periods as having a fixed length, r represented both the interest rate over a fixed length of calendar time and the interest rate over one period. Now we need to make a distinction between these two meanings. We will let r continue to mean one plus the interest rate over a fixed length of calendar time. When we have occasion to refer to one plus the interest rate over a period (trading interval) of length h , we will use the symbol \hat{r} .

Clearly, the size of \hat{r} depends on the number of subintervals, n , into which t is divided. Over the n periods until expiration, the total return is \hat{r}^n , where $n = t/h$. Now not only do we want \hat{r} to depend on n , but we want it to depend on n in a particular way — so that as n changes the total return \hat{r}^n over the fixed time t remains the same. This is because the interest rate obtainable over some fixed length of calendar time should have nothing to do with how we choose to think of the length of the time interval h .

If r (without the “hat”) denotes one plus the rate of interest over a *fixed* unit of calendar time, then over elapsed time t , r^t is the total return.¹⁰ Observe that this measure of total return does not depend on n . As we have argued, we want to choose the dependence of \hat{r} on n , so that

$$\hat{r}^n = r^t$$

for any choice of n . Therefore, $\hat{r} = r^{t/n}$. This last equation shows how \hat{r} must depend on n for the total return over elapsed time t to be independent of n .

We also need to define u and d in terms of n . At this point, there are two significantly different paths we can take. Depending on the definitions we choose, as $n \rightarrow \infty$ (or, equivalently, as $h \rightarrow 0$), we can have either a continuous or a jump stochastic process. In the first situation, very

¹⁰ The scale of this unit (perhaps a day, or a year) is unimportant as long as r and t are expressed in the same scale.

small random changes in the stock price will be occurring in each very small time interval. The stock price will fluctuate incessantly, but its path can be drawn without lifting pen from paper. In contrast, in the second case, the stock price will usually move in a smooth deterministic way, but will occasionally experience sudden discontinuous changes. Both can be derived from our binomial process simply by choosing how u and d depend on n . We examine in detail only the continuous process that leads to the option pricing formula originally derived by Fischer Black and Myron Scholes. Subsequently, we indicate how to develop the jump process formula originally derived by John Cox and Stephen Ross.

Recall that we supposed that over each period the stock price would experience a one plus rate of return of u with probability q and d with probability $1 - q$. It will be easier and clearer to work, instead, with the natural logarithm of the one plus rate of return, $\log u$ or $\log d$. This gives the continuously compounded rate of return on the stock over each period. It is a random variable which, in each period, will be equal to $\log u$ with probability q and $\log d$ with probability $1 - q$.

Consider a typical sequence of five moves, say u, d, u, u, d . Then the final stock price will be $S^* = uduudS$; $S^*/S = u^3d^2$, and $\log(S^*/S) = 3 \log u + 2 \log d$. More generally, over n periods,

$$\log(S^*/S) = j \log u + (n - j) \log d = j \log(u/d) + n \log d$$

where j is the (random) number of upward moves occurring during the n periods to expiration. Therefore, the expected value of $\log(S^*/S)$ is

$$E[\log(S^*/S)] = \log(u/d) \cdot E(j) + n \log d$$

and its variance is

$$Var[\log(S^*/S)] = [\log(u/d)]^2 \cdot Var(j)$$

Each of the n possible upward moves has probability q . Thus, $E(j) = nq$. Also since the variance each period is $q(1 - q)^2 + (1 - q)(0 - q)^2 = q(1 - q)$, then $Var(j) = nq(1 - q)$. Combining all of this, we have

$$\begin{aligned} E[\log(S^*/S)] &= [q \log(u/d) + \log d]n \equiv \hat{\mu}n \\ Var[\log(S^*/S)] &= q(1 - q)[\log(u/d)]^2 n \equiv \hat{\sigma}^2 n \end{aligned}$$

Let us go back to our discussion. We were considering dividing up our original longer time period (a day) into many shorter periods (a minute or even less). Our procedure calls for, over fixed length of calendar time t , making n larger and larger. Now if we held everything else constant while we let n become large, we would be faced with the problem we talked about earlier. In fact, we would certainly not reach a reasonable conclusion if either $\hat{\mu}n$ or $\hat{\sigma}^2 n$ went to zero or infinity as n became large. Since t is a fixed length of time, in searching for a realistic result, we must make the appropriate adjustments in u , d , and q . In doing that, we would at least want the mean and variance of the continuously compounded rate of return of the assumed stock price movement to coincide with that of the actual stock price as $n \rightarrow \infty$.

Suppose we label the actual empirical values of $\hat{\mu}n$ and $\hat{\sigma}^2n$ as μt and $\sigma^2 t$, respectively. Then we would want to choose u , d , and q so that

$$\begin{aligned} [q \log(u/d) + \log d]n &\rightarrow \mu t \\ &\text{as } n \rightarrow \infty \\ q(1-q)[\log(u/d)]^2 n &\rightarrow \sigma^2 t \end{aligned}$$

A little algebra shows we can accomplish this by letting

$$u = e^{\sigma\sqrt{t/n}}, \quad d = e^{-\sigma\sqrt{t/n}}, \quad q = \frac{1}{2} + \frac{1}{2}(\mu/\sigma)\sqrt{t/n}$$

In this case, for any n ,

$$\hat{\mu}n = \mu t \quad \text{and} \quad \hat{\sigma}^2n = [\sigma^2 - \mu^2(t/n)]t$$

Clearly, as $n \rightarrow \infty$, $\hat{\sigma}^2n \rightarrow \sigma^2 t$ while $\hat{\mu}n = \mu t$ for all values of n .

Alternatively, we could have chosen u , d , and q so that the mean and variance of the future stock price for the discrete binomial process approach the prespecified mean and variance of the actual stock price as $n \rightarrow \infty$. However, just as we would expect, the same values will accomplish this as well. Since this would not change our conclusions, and it is computationally more convenient to work with the continuously compounded rates of return, we will proceed in that way.

This satisfies our initial requirement that the limiting means and variances coincide, but we still need to verify that we are arriving at a sensible limiting probability distribution of the continuously compounded rate of return. The mean and variance only describe certain aspects of that distribution.

For our model, the random continuously compounded rate of return over a period of length t is the sum of n independent random variables, each of which can take the value $\log u$ with probability q and $\log d$ with probability $1 - q$. We wish to know about the distribution of this sum as n becomes large and q , u , and d are chosen in the way described. We need to remember that as we change n , we are not simply adding one more random variable to the previous sum, but instead are changing the probabilities and possible outcomes for every member of the sum. At this point, we can rely on a form of the central limit theorem which, when applied to our problem, says that, as $n \rightarrow \infty$, if

$$\frac{q|\log u - \hat{\mu}|^3 + (1-q)|\log d - \hat{\mu}|^3}{\hat{\sigma}^3 \sqrt{n}} \rightarrow 0$$

then

$$\text{Prob} \left[\left(\frac{\log(S^*/S) - \hat{\mu}n}{\hat{\sigma}\sqrt{n}} \right) \leq z \right] \rightarrow N(z)$$

where $N(z)$ is the standard normal distribution function. Putting this into words, as the number of periods into which the fixed length of time to expiration is divided approaches infinity, the probability that the standardized continuously compounded rate of return of the stock through the expiration date is not greater than the number z approaches the probability under a standard normal distribution.

The initial condition says roughly that higher-order properties of the distribution, such as how it is skewed, become less and less important, relative to its standard deviation, as $n \rightarrow \infty$. We can verify that the condition is satisfied by making the appropriate substitutions and finding

$$\frac{q|\log u - \hat{\mu}|^3 + (1-q)|\log d - \hat{\mu}|^3}{\hat{\sigma}^3 \sqrt{n}} = \frac{(1-q)^2 + q^2}{\sqrt{nq(1-q)}}$$

which goes to zero as $n \rightarrow \infty$ since $q = \frac{1}{2} + \frac{1}{2}(\mu/\sigma)\sqrt{t/n}$. Thus, the multiplicative binomial model for stock prices includes the lognormal distribution as a limiting case.

Black and Scholes began directly with continuous trading and the assumption of a lognormal distribution for stock prices. Their approach relied on some quite advanced mathematics. However, since our approach contains continuous trading and the lognormal distribution as a limiting case, the two resulting formulas should then coincide. We will see shortly that this is indeed true, and we will have the advantage of using a much simpler method. It is important to remember, however, that the economic arguments we used to link the option value and the stock price are exactly the same as those advanced by Black and Scholes (1973) and Merton (1973, 1977).

The formula derived by Black and Scholes, rewritten in terms of our notation, is

Black-Scholes Option Pricing Formula

$$C = SN(x) - Kr^{-t}N(x - \sigma\sqrt{t})$$

where

$$x = \frac{\log(S/Kr^{-t})}{\sigma\sqrt{t}} + \frac{1}{2}\sigma\sqrt{t}$$

We now wish to confirm that our binomial formula converges to the Black-Scholes formula when t is divided into more and more subintervals, and \hat{r} , u , d , and q are chosen in the way we described — that is, in a way such that the multiplicative binomial probability distribution of stock prices goes to the lognormal distribution.

For easy reference, let us recall our binomial option pricing formula:

$$C = S\phi[a; n, p'] - K\hat{r}^{-n}\phi[a; n, p]$$

The similarities are readily apparent. \hat{r}^{-n} is, of course, always equal to r^{-t} . Therefore, to show the two formulas converge, we need only show that as $n \rightarrow \infty$

$$\phi[a; n, p'] \rightarrow N(x) \quad \text{and} \quad \phi[a; n, p] \rightarrow N(x - \sigma\sqrt{t})$$

We will consider only $\phi[a; n, p]$, since the argument is exactly the same for $\phi[a; n, p']$.

The complementary binomial distribution function $\phi[a; n, p]$ is the probability that the sum of n random variables, each of which can take on the value 1 with the probability p and 0 with the probability $1 - p$, will be greater than or equal to a . We know that the random value of this sum, j , has mean np and standard deviation $\sqrt{np(1-p)}$. Therefore,

$$1 - \phi[a; n, p] = \text{Prob}[j \leq a - 1] = \text{Prob}\left[\frac{j - np}{\sqrt{np(1-p)}} \leq \frac{a - 1 - np}{\sqrt{np(1-p)}}\right]$$

Now we can make an analogy with our earlier discussion. If we consider a stock which in each period will move to uS with probability p and dS with probability $1 - p$, then $\log(S^*/S) = j \log(u/d) + n \log d$. The mean and variance of the continuously compounded rate of return of this stock are

$$\hat{\mu}_p = p \log(u/d) + \log d \quad \text{and} \quad \hat{\sigma}_p^2 = p(1-p)[\log(u/d)]^2$$

Using these equalities, we find that

$$\frac{j - np}{\sqrt{np(1-p)}} = \frac{\log(S^*/S) - \hat{\mu}_p n}{\hat{\sigma}_p \sqrt{n}}$$

Recall from the binomial formula that

$$a - 1 = \log(K / Sd^n) / \log(u/d) - \varepsilon = [\log(K/S) - n \log d] / \log(u/d) - \varepsilon,$$

where ε is a number between zero and one. Using this and the definitions of $\hat{\mu}_p$ and $\hat{\sigma}_p^2$, with a little algebra, we have

$$\frac{a - 1 - np}{\sqrt{np(1-p)}} = \frac{\log(K/S) - \hat{\mu}_p n - \varepsilon \log(u/d)}{\hat{\sigma}_p \sqrt{n}}$$

Putting these results together,

$$1 - \phi[a; n, p] = \text{Prob} \left[\frac{\log(S^*/S) - \hat{\mu}_p n}{\hat{\sigma}_p \sqrt{n}} \leq \frac{\log(K/S) - \hat{\mu}_p n - \varepsilon \log(u/d)}{\hat{\sigma}_p \sqrt{n}} \right]$$

We are now in a position to apply the central limit theorem. First, we must check if the initial condition,

$$\frac{p |\log u - \hat{\mu}_p|^3 + (1-p) |\log d - \hat{\mu}_p|^3}{\hat{\sigma}_p \sqrt{n}} = \frac{(1-p)^2 + p^2}{\sqrt{np(1-p)}} \rightarrow 0$$

as $n \rightarrow \infty$, is satisfied. By first recalling that $p \equiv (\hat{r} - d)/(u - d)$, and then $\hat{r} = r^{t/n}$, $u = e^{\sigma\sqrt{t/n}}$, and $d = e^{-\sigma\sqrt{t/n}}$, it is possible to show that as $n \rightarrow \infty$,

$$p \rightarrow \frac{1}{2} + \frac{1}{2} \left(\frac{\log r - \frac{1}{2}\sigma^2}{\sigma} \right) \sqrt{\frac{t}{n}}$$

As a result, the initial condition holds, and we are justified in applying the central limit theorem.

To do so, we need only evaluate $\hat{\mu}_p n$, $\hat{\sigma}_p^2 n$ and $\log(u/d)$ as $n \rightarrow \infty$.¹¹ Examination of our discussion for parameterizing q shows that as $n \rightarrow \infty$

¹¹ A surprising feature of this evaluation is that although $p \neq q$ and thus $\hat{\mu}_p \neq \hat{\mu}$ and $\hat{\sigma}_p \neq \hat{\sigma}$, nonetheless $\hat{\sigma}_p \sqrt{n}$ and $\hat{\sigma} \sqrt{n}$ have the same limiting value as $n \rightarrow \infty$. By contrast, since $\mu \neq \log r - \left(\frac{1}{2}\sigma^2\right)$, $\hat{\mu}_p n$ and $\hat{\mu} n$ do not. This results from the way we needed to specify u and d to obtain convergence to a lognormal distribution. Rewriting this as $\sigma\sqrt{t} = (\log u)\sqrt{n}$, it is clear that the limiting value σ of the standard deviation does not depend on p or q , and hence must be the same for either. However, at any point before the limit, since

$$\hat{\sigma}^2 n = \left(\sigma^2 - \mu^2 \frac{t}{n} \right) t \quad \text{and} \quad \hat{\sigma}_p^2 n = \left[\sigma^2 - \left(\log r - \frac{1}{2}\sigma^2 \right)^2 \frac{t}{n} \right] t$$

$\hat{\sigma}$ and $\hat{\sigma}_p$ will generally have different values.

The fact that $\hat{\mu}_p n \rightarrow \left(\log r - \frac{1}{2}\sigma^2 \right) t$ can also be derived from the property of the lognormal distribution that

$$\log E[S^*/S] = \mu_p t + \frac{1}{2}\sigma^2 t$$

where E and μ_p are measured with respect to probability p . Since $p = (\hat{r} - d)/(u - d)$, it follows that $\hat{r} = pu + (1-p)d$. For independently distributed random variables, the expectation of a product equals the product of their expectations. Therefore,

$$E[S^*/S] = [pu + (1-p)d]^n = \hat{r}^n = r^t$$

Substituting r^t for $E[S^*/S]$ in the previous equation, we have

$$\hat{\mu}_p n \rightarrow \left(\log r - \frac{1}{2} \sigma^2 \right) t \quad \text{and} \quad \hat{\sigma}_p \sqrt{n} \rightarrow \sigma \sqrt{t}$$

Furthermore, $\log(u/d) \rightarrow 0$ as $n \rightarrow \infty$.

For this application of the central limit theorem, then, since

$$\frac{\log(K/S) - \hat{\mu}_p n - \varepsilon \log(u/d)}{\hat{\sigma}_p \sqrt{n}} \rightarrow z = \frac{\log(K/S) - \left(\log r - \frac{1}{2} \sigma^2 \right) t}{\sigma \sqrt{t}}$$

we have

$$1 - \phi[a; n, p] \rightarrow N(z) = N \left[\frac{\log(Kr^{-t}/S)}{\sigma \sqrt{t}} + \frac{1}{2} \sigma \sqrt{t} \right]$$

The final step in the argument is to use the symmetry property of the standard normal deviation distribution that $1 - N(z) = N(-z)$. Therefore, as $n \rightarrow \infty$

$$\phi[a; n, p] \rightarrow N(-z) = N \left[\frac{\log(S/Kr^{-t})}{\sigma \sqrt{t}} - \frac{1}{2} \sigma \sqrt{t} \right] = N(x - \sigma \sqrt{t})$$

Since a similar argument holds for $\phi[a; n, p]$, this completes our demonstration that the binomial option pricing formula contains the Black-Scholes formula as a limiting case.^{12,13}

$$\mu_p = \log r - \frac{1}{2} \sigma^2$$

¹² The only difference is that, as $n \rightarrow \infty$, $p' \rightarrow \frac{1}{2} + \frac{1}{2} \left[\left(\log r + \frac{1}{2} \sigma^2 \right) / \sigma \right] \sqrt{t/n}$. Further, it can be shown that as $n \rightarrow \infty$, $\Delta \rightarrow N(x)$. Therefore, for the Black-Scholes model, $\Delta S = SN(x)$ and $B = -Kr^{-t} N(x - \sigma \sqrt{t})$.

¹³ In our original development, we obtained the following equation (somewhat rewritten) relating the call prices in successive periods:

$$\left(\frac{\hat{r} - d}{u - d} \right) C_u + \left(\frac{u - \hat{r}}{u - d} \right) C_d - \hat{r} C = 0$$

By their more difficult methods, Black and Scholes obtained directly a partial differential equation analogous to our discrete-time difference equation. Their equation is

$$\frac{1}{2} \sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} + (\log r) S \frac{\partial C}{\partial S} - \frac{\partial C}{\partial t} - (\log r) C = 0.$$

The value of the call, C , was then derived by solving this equation subject to the boundary condition $C^* = \max[0, S^* - K]$.

As we have remarked, the seeds of both the Black-Scholes formula and a continuous-time jump process formula are both contained within the binomial formulation. At which end point we arrive depends on how we take limits. Suppose, in place of our former correspondence for u , d , and q , we instead set

$$u = u, \quad d = e^{\zeta(t/n)}, \quad q = \lambda(t/n).$$

This correspondence captures the essence of a pure jump process in which each successive stock price is almost always close to the previous price ($S \rightarrow dS$), but occasionally, with low but continuing probability, significantly different ($S \rightarrow uS$). Observe that, as $n \rightarrow \infty$, the probability of a change by d becomes larger and larger, while the probability of a change by u approaches zero.

With these specifications, the initial condition of the central limit theorem we used is no longer satisfied, and it can be shown the stock price movements converge to a log-Poisson rather than a lognormal distribution as $n \rightarrow \infty$. Let us define

$$\Psi[x; y] \equiv \sum_{i=x}^{\infty} \frac{e^{-y} y^i}{i!}$$

as the complementary Poisson distribution function. The limiting option pricing formula for the above specifications of u , d and q is then

Jump Process Option Pricing Formula

$$C = S\Psi[x; y] - Kr^{-t}\Psi[x; y/u],$$

where

$$y \equiv (\log r - \zeta)ut / (u - 1),$$

and

$x \equiv$ the smallest non-negative integer greater than $(\log(K/S) - \zeta t) / \log u$.

A very similar formula holds if we let $u = e^{\zeta(t/n)}$, $d = d$, and $1 - q = \lambda(t/n)$.

Based on our previous analysis, we would now suspect that, as $n \rightarrow \infty$, our difference equation would approach the Black-Scholes partial differential equation. This can be confirmed by substituting our definitions of \hat{r} , u , d in terms of n in the way described earlier, expanding C_u , C_d in a Taylor series around $(e^{\sigma\sqrt{h}}S, t-h)$ and $(e^{-\sigma\sqrt{h}}S, t-h)$, respectively, and then expanding $e^{\sigma\sqrt{h}}$, $e^{-\sigma\sqrt{h}}$, and r^h in a Taylor series, substituting these in the equation and collecting terms. If we then divide by h and let $h \rightarrow 0$, all terms of higher order than h go to zero. This yields the Black-Scholes equation.

6. Dividends and Put Pricing

So far we have been assuming that the stock pays no dividends. It is easy to do away with this restriction. We will illustrate this with a specific dividend policy: the stock maintains a constant yield, δ , on each ex-dividend date. Suppose there is one period remaining before expiration and the current stock price is S . If the end of the period is an ex-dividend date, then an individual who owned the stock during the period will receive at that time a dividend of either δuS or δdS . Hence, the stock price at the end of the period will be either $u(1 - \delta)^v S$ or $d(1 - \delta)^v S$, where $v = 1$ if the end of the period is an ex-dividend date and $v = 0$ otherwise. Both δ and v are assumed to be known with certainty.

When the call expires, its contract and a rational exercise policy imply that its value must be either

$$C_u = \max[0, u(1 - \delta)^v S - K]$$

or

$$C_d = \max[0, d(1 - \delta)^v S - K]$$

Therefore,

$$C \begin{cases} C_u = \max[0, u(1 - \delta)^v S - K] \\ C_d = \max[0, d(1 - \delta)^v S - K] \end{cases}$$

Now we can proceed exactly as before. Again, we can select a portfolio of Δ shares of stock and the dollar amount B in bonds that will have the same end-of-period value as the call.¹⁴ By retracting our previous steps, we can show that

$$C = [pC_u + (1 - p)C_d] / \hat{r}$$

if this is greater than $S - K$ and $C = S - K$ otherwise. Here, once again, $p = (\hat{r} - d)/(u - d)$ and $\Delta = (C_u - C_d)/(u - d)S$.

Thus far the only change is that $(1 - \delta)^v S$ has replaced S in the values for C_u and C_d . Now we come to the major difference: early exercise may be optimal. To see this, suppose that $v = 1$ and $d(1 - \delta)S > K$. Since $u > d$, then, also, $u(1 - \delta)S > K$. In this case, $C_u = u(1 - \delta)S - K$ and $C_d = d(1 - \delta)S - K$. Therefore, since $(u/\hat{r})p + (d/\hat{r})(1 - p) = 1$, then

$$[pC_u + (1 - p)C_d] / \hat{r} = (1 - \delta)S - (K/\hat{r})$$

¹⁴ Remember that if we are long the portfolio, we will receive the dividend at the end of the period; if we are short, we will have to make restitution for the dividend.

For sufficiently high stock prices, this can obviously be less than $S - K$. Hence, there are definitely some circumstances in which no one would be willing to hold the call for one more period.

In fact, there will always be a critical stock price, \hat{S} , such that if $S > \hat{S}$, the call should be exercised immediately. \hat{S} will be the stock price at which $[pC_u + (1-p)C_d]/\hat{r} = S - K$.¹⁵ That is, it is the lowest stock price at which the value of the hedging portfolio exactly equals $S - K$. This means \hat{S} will, other things equal, be lower the higher the dividend yield, the lower the interest rate, and the lower the strike price.

We can extend the analysis to an arbitrary number of periods in the same way as before. There is only one additional difference, a minor modification in the hedging operation. Now the funds in the hedging portfolio will be increased by any dividends received, or decreased by the restitution required for dividends paid while the stock is held short.

Although the possibility of optimal exercise before the expiration date causes no conceptual difficulties, it does seem to prohibit a simple closed-form solution for the value of a call with many periods to go. However, our analysis suggests a sequential numerical procedure that will allow us to calculate the continuous-time value to any desired degree of accuracy.

Let C be the current value of a call with n periods remaining. Define

$$\bar{v}(n, i) \equiv \sum_{k=1}^{n-i} v_k$$

so that $\bar{v}(n, i)$ is the number of ex-dividend dates occurring during the next $n - i$ periods. Let $C(n, i, j)$ be the value of the call $n - i$ periods from now, given that the current stock price S has changed to $u^j d^{n-i-j} (1 - \delta)^{\bar{v}(n, i)} S$, where $j = 0, 1, 2, \dots, n - i$.

With this notation, we are prepared to solve for the current value of the call by working backward in time from the expiration date. At expiration, $i = 0$, so that

$$C(n, 0, j) = \max[0, u^j d^{n-j} (1 - \delta)^{\bar{v}(n, 0)} S - K] \quad \text{for } j = 0, 1, 2, \dots, n$$

One period before the expiration date, $i = 1$ so that

$$C(n, 1, j) = \max \left[u^j d^{n-1-j} (1 - \delta)^{\bar{v}(n, 1)} S - K, [pC(n, 0, j+1) + (1-p)C(n, 0, j)]/\hat{r} \right] \\ \text{for } j = 0, 1, 2, \dots, n - 1$$

More generally, i periods before expiration

¹⁵ Actually solving for \hat{S} explicitly is straightforward but rather tedious, so we will omit it.

$$C(n, i, j) = \max \left[u^j d^{n-i-j} (1 - \delta)^{\bar{v}(n,i)} S - K, [pC(n, i - 1, j + 1) + (1 - p)C(n, i - 1, j)] / \hat{r} \right] \\ \text{for } j = 0, 1, 2, \dots, n - i$$

Observe that each prior step provides the inputs needed to evaluate the right-hand arguments of each succeeding step. The number of calculations decreases as we move backward in time. Finally, with n periods before expiration, since $i = n$,

$$C = C(n, n, 0) = \max [S - K, [pC(n, n - 1, 1) + (1 - p)C(n, n - 1, 0)] / \hat{r}]$$

and the hedge ratio is

$$\Delta = \frac{C(n, n - 1, 1) - C(n, n - 1, 0)}{(u - d)S}$$

We could easily expand the analysis to include dividend policies in which the amount paid on any ex-dividend date depends on the stock price at that time in a more general way.¹⁶ However, this will cause some minor complications. In our present example with a constant dividend yield, the possible stock prices $n - i$ periods from now are completely determined by the total number of upward moves (and ex-dividend dates) occurring during that interval. With other types of dividend policies, the enumeration will be more complicated, since then the terminal stock price will be affected by the timing of the upward moves as well as their total number. But the basic principle remains the same. We go to the expiration date and calculate the call value for all of the possible prices that the stock could have then. Using this information, we step back one period and calculate the call values for all possible stock prices at that time, and so forth.

We will now illustrate the use of the binomial numerical procedure in approximating continuous-time call values. In order to have an exact continuous-time formula to use for comparison, we will consider the case with no dividends. Suppose that we are given the inputs required for the Black-Scholes option pricing formula: S , K , t , σ , and r . To convert this information into the inputs d , u , and \hat{r} required for the binomial numerical procedure, we use the relationships:

$$d = 1/u, \quad u = e^{\sigma\sqrt{t/n}}, \quad \hat{r} = r^{t/n}$$

Table 2 gives us a feeling for how rapidly option values approximated by the binomial method approach the corresponding limiting Black-Scholes values given by $n = \infty$. At $n = 5$, the values differ by at most \$0.25, and at $n = 20$, they differ by at most \$0.07. Although not shown, at $n = 50$, the greatest difference is less than \$0.03, and at $n = 150$, the values are identical to the penny.

To derive a method for valuing puts, we again use the binomial formulation. Although it has been convenient to express the argument in terms of a particular security, a call, this is not essential in any way. The same basic analysis can be applied to puts.

¹⁶ We could also allow the amount to depend on previous stock prices.

Letting P denote the current price of a put, with one period remaining before expiration, we have

$$P \begin{cases} P_u = \max[0, K - u(1 - \delta)^v S] \\ P_d = \max[0, K - d(1 - \delta)^v S] \end{cases}$$

Once again, we can choose a portfolio with ΔS in stock and B in bonds which will have the same end-of-period values as the put. By a series of steps that are formally equivalent to the ones we followed in section 3, we can show that

$$P = [pP_u + (1 - p)P_d] / \hat{r}$$

if this is greater than $K - S$, and $P = K - S$ otherwise. As before, $p = (\hat{r} - d)/(u - d)$ and $\Delta = (P_u - P_d)/(u - d)S$. Note that for puts, since $P_u \leq P_d$, then $\Delta \leq 0$. This means that if we sell an overvalued put, the hedging portfolio that we buy will involve a short position in the stock.

We might hope that with puts we will be spared the complications caused by optimal exercise before the expiration date. Unfortunately, this is not the case. In fact, the situation is even worse in this regard. Now there are always some possible circumstances in which no one would be willing to hold the put for one more period.

To see this, suppose $K > u(1 - \delta)^v S$. Since $u > d$, then, also, $K > d(1 - \delta)^v S$. In this case, $P_u = K - u(1 - \delta)^v S$ and $P_d = K - d(1 - \delta)^v S$. Therefore, since $(u/\hat{r})p + (d/\hat{r})(1 - p) = 1$, then

$$[pP_u + (1 - p)P_d] / \hat{r} = (K/\hat{r}) - (1 - \delta)^v S$$

If there are no dividends (that is, $v = 0$), then this is certainly less than $K - S$. Even with $v = 1$, it will be less for a sufficiently low stock price.

Thus, there will now be a critical stock price, \hat{S} , such that if $S < \hat{S}$, the put should be exercised immediately. By analogy with our discussion for the call, we can see that this is the stock price at which $[pP_u + (1 - p)P_d] / \hat{r} = K - S$. Other things equal, \hat{S} will be higher the lower the dividend yield, the higher the interest rate, and the higher the strike price. Optimal early exercise thus becomes more likely if the put is deep-in-the-money and the interest rate is high. The effect of dividends yet to be paid diminishes the advantages of immediate exercise, since the put buyer will be reluctant to sacrifice the forced declines in the stock price on future ex-dividend dates.

This argument can be extended in the same way as before to value puts with any number of periods to go. However, the chance for optimal exercise before the expiration date once again seems to preclude the possibility of expressing this value in a simple form. But our analysis also indicates that, with slight modification, we can value puts with the same numerical techniques

Table 2
Binomial Approximation of Continuous-time Call Values ($S = 40$ and $r = 1.05$)[†]

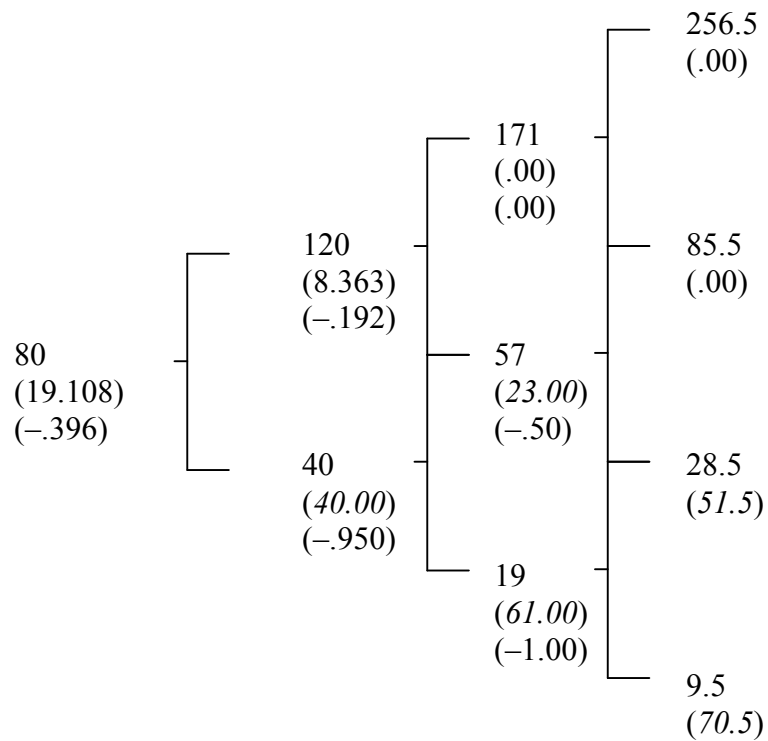
		$n = 5$			$n = 20$			$n = \infty$		
σ	K	JAN	APR	JUL	JAN	APR	JUL	JAN	APR	JUL
0.2	35	5.14	5.77	6.45	5.15	5.77	6.39	5.15	5.76	6.40
	40	1.05	2.26	3.12	0.99	2.14	2.97	1.00	2.17	3.00
	45	0.02	0.54	1.15	0.02	0.51	1.11	0.02	0.51	1.10
0.3	35	5.21	6.30	7.15	5.22	6.26	7.19	5.22	6.25	7.17
	40	1.53	3.21	4.36	1.44	3.04	4.14	1.46	3.07	4.19
	45	0.11	1.28	2.12	0.15	1.28	2.23	0.16	1.25	2.24
0.4	35	5.40	6.87	7.92	5.39	6.91	8.05	5.39	6.89	8.09
	40	2.01	4.16	5.61	1.90	3.93	5.31	1.92	3.98	5.37
	45	0.46	1.99	3.30	0.42	2.09	3.42	0.42	2.10	3.43

[†] The January options have one month to expiration, the Aprils, four months, and the Julys, seven months; r and σ are expressed in annual terms.

we use for calls. Reversing the difference between the stock price and the strike price at each stage is the only change.¹⁷

The diagram presented in table 3 shows the stock prices, put values, and values of Δ obtained in this way for the example given in section 4. The values used there were $S = 80$, $K = 80$, $n = 3$, $u = 1.5$, $d = 0.5$, and $\hat{r} = 1.1$. To include dividends as well, we assumed that a cash dividend of five percent ($\delta = 0.05$) will be paid at the end of the last period before the expiration date. Thus, $(1 - \delta)^{\bar{v}(n,0)} = 0.95$, $(1 - \delta)^{\bar{v}(n,1)} = 0.95$, and $(1 - \delta)^{\bar{v}(n,2)} = 1.0$. Put values in *italics* indicate that immediate exercise is optimal.

Table 3
Three-period Binomial Tree for an American Put



¹⁷ Michael Parkinson (1977) has suggested a similar numerical procedure based on a trinomial process, where the stock price can increase, decrease, or remain unchanged. In fact, given the theoretical basis for the binomial numerical procedure provided, the numerical method can be generalized to permit $k + 1 \leq n$ jumps to new stock prices in each period. We can consider exercise only every k periods, using the binomial formula to leap across intermediate periods. In effect, this means permitting $k + 1$ possible new stock prices before exercise is again considered. That is, instead of considering exercise n times, we would only consider it about n/k times. For fixed t and k , as $n \rightarrow \infty$, option values will approach their continuous-time values.

This alternative procedure is interesting, since it may enhance computer efficiency. At one extreme, for calls on stocks which do not pay dividends, setting $k + 1 = n$ gives the most efficient results. However, when the effect of potential early exercise is important and greater accuracy is required, the most efficient results are achieved by setting $k = 1$, as in our description above.

7. Conclusion

It should now be clear that whenever stock price movements conform to a discrete binomial process, or to a limiting form of such a process, options can be priced solely on the basis of arbitrage considerations. Indeed, we could have significantly complicated the simple binomial process while still retaining this property.

The probabilities of an upward or downward move did not enter into the valuation formula. Hence, we would obtain the same result if q depended on the current or past stock prices or on other random variables. In addition, u and d could have been deterministic functions of time. More significantly, the size of the percentage changes in the stock price over each period could have depended on the stock price at the beginning of each period or on previous stock prices.¹⁸ However, if the size of the changes were to depend on any other random variable, not itself perfectly correlated with the stock price, then our argument will no longer hold. If any arbitrage result is then still possible, it will require the use of additional assets in the hedging portfolio.

We could also incorporate certain types of imperfections into the binomial option pricing approach, such as differential borrowing and lending rates and margin requirements. These can be shown to produce upper and lower bounds on option prices, outside of which riskless profitable arbitrage would be possible.

Since all existing preference-free option pricing results can be derived as limiting forms of a discrete two-state process, we might suspect that two-state stock price movements, with the qualifications mentioned above, must be in some sense necessary, as well as sufficient, to derive option pricing formulas based solely on arbitrage considerations. To price an option by arbitrage methods, there must exist a portfolio of other assets that exactly replicates in every state of nature the payoff received by an optimally exercised option. Our basic proposition is the following. Suppose, as we have, that markets are perfect, that changes in the interest rate are never random, and that changes in the stock price are always random. In a discrete time model, a necessary and sufficient condition for options of all maturities and strike prices to be priced by arbitrage using only the stock and bonds in the portfolio is that in each period,

- (a) the stock price can change from its beginning-of-period value to only two ex-dividend values at the end of the period, and
- (b) the dividends and the size of each of the two possible changes are presently known functions depending at most on: (i) current and past stock prices, (ii) current and past values of random variables whose changes in each period are perfectly correlated with the change in the stock price, and (iii) calendar time.

¹⁸ Of course, different option pricing formulas would result from these more complex stochastic processes. See Cox and Ross (1976) and Geske (1979). Nonetheless, all option pricing formulas in these papers can be derived as limiting forms of a properly specified discrete two-state process.

The sufficiency of the condition can be established by a straightforward application of the methods we have presented. Its necessity is implied by the discussion at the end of section 3.^{19,20,21}

This rounds out the principal conclusion of this paper: the simple two-state process is really the essential ingredient of option pricing by arbitrage methods. This is surprising, perhaps, given the mathematical complexities of some of the current models in this field. But it is reassuring to find such simple economic arguments at the heart of this powerful theory.

¹⁹ Note that option values need not depend on the present stock price alone. In some cases, formal dependence on the entire series of past values of the stock price and other variables can be summarized in a small number of state variables.

²⁰ In some circumstances, it will be possible to value options by arbitrage when this condition does not hold by using additional assets in the hedging portfolio. The value of the option will then in general depend on the values of these other assets, although in certain cases only parameters describing their movement will be required.

²¹ Merton's (1976) model, with both continuous and jump components, is a good example of a stock price process for which no exact option pricing formula is obtainable purely from arbitrage considerations. To obtain an exact formula, it is necessary to impose restrictions on the stochastic movements of other securities, as Merton did, or on investor preferences. For example, Rubinstein (1976) has been able to derive the Black-Scholes option pricing formula, under circumstances that do not admit arbitrage, by suitably restricting investor preferences. Additional problems arise when interest rates are stochastic, although Merton (1973) has shown that some arbitrage results may still be obtained.

References

- Black, F. and M. Scholes, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy* 81, No. 3 (May-June 1973), pp. 637-654.
- Brennan, M.J. and E.S. Schwartz, "The Valuation of American Put Options," *Journal of Finance* 32, (1977), pp. 449-462.
- Cox, J.C. and S.A. Ross, "The Pricing of Options for Jump Processes," unpublished working paper #2-75, University of Pennsylvania, (April 1975).
- Cox, J.C. and S.A. Ross, "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics* 3, No. 1 (January-March 1976) pp. 145-166.
- Geske, R., "The Valuation of Compound Options," *Journal of Financial Economics* 7, No. 1 (March 1979), pp. 63-81.
- Harrison, J.M. and D.M. Kreps, "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory* 20, No. 3 (July 1979), pp. 381-408.
- Merton, R.C., "The Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science* 4, No. 1(Spring 1973), pp. 141-183.
- Merton, R.C., "Option Pricing When Underlying Stock Returns are Discontinuous," *Journal of Financial Economics* 3, No. 1 (January-March 1976), pp. 125-144.
- Merton, R.C., "On the Pricing of Contingent Claims and the Modigliani-Miller Theorem," *Journal of Financial Economics* 5, No. 2 (November 1977), pp. 241-250.
- Parkinson, M., "Option Pricing: The American Put," *Journal of Business* 50, (1977), pp. 21-36.
- Rendleman, R.J. and B.J. Bartter, "Two-State Option Pricing," unpublished working paper, Northwestern University (1978).
- Rubinstein, M., "The Valuation of Uncertain Income Streams and the Pricing of Options," *Bell Journal of Economics* 7, No. 2 (Autumn 1976), pp. 407-425.
- Sharpe, W.F., *Investments*, Prentice-Hall (1978).