Binomial lattice model for stock prices

Here we model the price of a stock in discrete time by a Markov chain of the recursive form $S_{n+1} = S_n Y_{n+1},\ n \geq 0$, where the $\{Y_i\}$ are iid with distribution $P(Y = u) = p,\ P(Y = d) = 1 - p$. Here $0 < d < 1 + r < u$ are constants with r the risk-free interest rate ($(1 + r)x$ is the payoff you would receive one unit of time later if you bought \$$x$ worth of the risk-free asset (a bond for example, or placed money in a savings account at that fixed rate) at time $n = 0$). Given the value of S_n,

$$S_{n+1} = \begin{cases} uS_n, & \text{w.p. } p; \\ dS_n, & \text{w.p. } 1-p, \end{cases} \quad n \geq 0,$$

independent of the past. Thus the stock either goes up ("u") or down ("d") in each time period, and the randomness is due to iid Bernoulli (p) rvs (flips of a coin so to speak) where we can view "up=success", and "down=failure".

Expanding the recursion yields

$$S_n = S_0 \times Y_1 \times \cdots \times Y_n,\ n \geq 1, \tag{1}$$

where S_0 is the initial price per share and S_n is the price per share at time n. [1]

It follows from (1) that for a given n, $S_n = u^i d^{n-i} S_0$ for some $i \in \{0, \dots n\}$, meaning that the stock went up i times and down $n - i$ times during the first n time periods (i "successes" and $n - i$ "failures" out of n independent Bernoulli (p) trials). The corresponding probabilities are thus determined by the binomial(n, p) distribution;

$$P(S_n = u^i d^{n-i} S_0) = \binom{n}{i} p^i (1-p)^{n-i},\ 0 \leq i \leq n,$$

which is why we refer to this model as the *binomial lattice model (BLM)*. The lattice is the set of points $\{u^i d^{n-i} S_0 : 0 \leq i \leq n < \infty\}$, which is the state space for this Markov chain. Note that this lattice depends on the initial price S_0 and the values of u, d.

Portfolios of stock and a risk-free asset

In addition to our stock there is a *risk-free* asset (money) with fixed interest rate $0 < r < 1$ that costs \$1.00 per share; x shares bought now (at time $t = 0$) would be worth the deterministic amount $x(1 + r)^n$ at time $t = n,\ n \geq 1$ (interest is compounded each time unit). Buying this asset is lending money. Selling this asset is borrowing money (shorting this asset).

We must have $1 + r < u$ for otherwise there would be no reason to invest in the stock: you could instead obtain a riskless payoff of $S_0(1+r) \geq S_0 u$ at time $t = 1$ by buying S_0 shares of the

[1] This model is meant to approximate the continuous-time geometric Brownian motion (GBM) $S(t) = S_0 e^{X(t)}$ model for stock, where $X(t) = \sigma B(t) + \mu t$ is Brownian motion (BM) with drift μ and variance term σ^2. The idea is to break up the time interval $(0, t]$ into n small subintervals of length $h = t/n$, $(0, h], (h, 2h],\ \dots ((n-1)h, nh = t]$, and re-write

$$S(t) = S(0) \times H_1 \times \cdots \times H_n,$$

where $H_i = S(ih)/S((i-1)h),\ i \geq 1$ are the succesive price ratios, and are in fact iid (due to the stationary and independent increments of the BM $X(t)$). Then we find an appropriate p, u, d so that the distribution of H is well approximated by the two-point distribution of Y (typically done by fitting the first two moments of H with those of Y). As $h \downarrow 0$ the approximation becomes exact.

risk-free asset at time $t = 0$ and selling them at time $t = 1$, thus earning at least as much, with certainty, than is ever possible from the stock. Similarly we have $d < 1 + r$ for otherwise there would be no reason to invest in the risk-free asset. (Inherent in our argument is the economic assumption of *non-arbitrage*, meaning that it is not possible, with certainty, for people to make a profit from nothing.)

A portfolio of stock and risk-free asset is a pair (α, β) describing our total investment at a given time; α shares of stock and β shares of the risk-free asset. We allow the values of α and β to be positive or negative or zero and they do not have to be integers. Negative values refer to shorting (borrowing). For example $(2.3, -7.4)$ means that we bought 2.3 shares of stock, and shorted 7.4 shares of the risk free asset (meaning we borrowed 7 dollars and 40 cents at interest rate r.)

Observe that a portfolio of stock and risk-free asset always has a well-defined price: a portfolio's price (cost) at time $t = 0$ is its worth, $\alpha S_0 + \beta$, and its price at time $t = n$, $n \geq 0$ is its worth at that time, $\alpha S_n + \beta(1 + r)^n$. For example at time $t = 1$ our $(2.3, -7.4)$ portfolio is worth $2.3S_1 - 7.4(1 + r)$ meaning that we now have $2.3S_1$ dollars worth of stock and owe $7.4(1 + r)$ dollars.

0.1 Pricing the European call option when the expiration date is $t = 1$

Now consider a European call option for one share of the stock, with strike price K, and expiration date $t = 1$. The payoff to the holder of this option at time $t = 1$ is a random variable given by $C_1 = (S_1 - K)^+$; the buyer of such an option is thus betting that the stock price will be above K at the expiration date. This random payoff has only two possible values: $C_1 = C_u = (uS_0 - K)^+$ if the stock goes up, and $C_1 = C_d = (dS_0 - K)^+$ if the stock moves down. Both of these values are known since they depend only on the known values u, d, S_0, and K. We next proceed to determine what a fair price should be for this option and denote this price by C_0. Clearly $C_0 \leq S_0$ because the payoff is less: $C_1 = (S_1 - K)^+ \leq S_1$. That is why people buy options, they are cheaper than the stock itself, but potentially can yield high payoffs. Unlike a portfolio of stock and risk-free asset, however, it is not immediate what this price should be, but we can use a portfolio to figure it out. To this end we will construct a portfolio (α, β) of stock and risk free asset, which if bought at time $t = 0$, then goes on to *replicate*, at time $t = 1$, the option payoff C_1: a portfolio that at time $t = 1$ yields payoff C_u if the stock goes up and C_d if it goes down. But the payoff of the portfolio at time $t = 1$ is $\alpha S_1 + \beta(1 + r)$, so we simply need to find the values α and β such that $\alpha S_1 + \beta(1 + r) = C_1$: find α and β such that $\alpha u S_0 + \beta(1+r) = C_u$ and $\alpha d S_0 + \beta(1+r) = C_d$. Once we do this, since the two investments yield the same payoff at time $t = 1$ they must have the same price at time $t = 0$:

$$C_0 = \textit{ the price of the replicating portfolio} = \alpha S_0 + \beta. \tag{2}$$

The point is that, in effect, they are the same investment, and thus must cost the same.

The solution to the two equations with two unknowns, $\alpha u S_0 + \beta(1 + r) = C_u$ and $\alpha d S_0 + \beta(1 + r) = C_d$, is

$$\alpha = \frac{C_u - C_d}{S_0(u - d)} \tag{3}$$

$$\beta = \frac{uC_d - dC_u}{(1 + r)(u - d)}. \tag{4}$$

Plugging this solution into (2) yields

$$C_0 = \frac{C_u - C_d}{(u - d)} + \frac{uC_d - dC_u}{(1 + r)(u - d)},$$

which when algabraically simplified (details left to the reader) yields:

$$\begin{aligned}
C_0 &= \frac{1}{1+r}(p^*C_u + (1-p^*)C_d) & (5)\\
&= \frac{1}{1+r}(p^*(uS_0-K)^+ + (1-p^*)(dS_0-K)^+), \text{ where} & (6)\\
p^* &\stackrel{\text{def}}{=} \frac{1+r-d}{u-d} & (7)\\
1-p^* &= \frac{u-(1+r)}{u-d}. & (8)
\end{aligned}$$

Since $1+r < u$ (by assumption), we see that $0 < p^* < 1$ is a probability, and C_0 as given in (5) is expressed elegantly as the discounted expected payoff of the option if $p = p^*$ for the underlying "up" probability p for the stock;

$$C_0 = \frac{1}{1+r}E^*(C_1), \tag{9}$$

where E^* denotes expected value when $p = p^*$ for the stock price. p^* is called the *risk-neutral* probability, for reasons we shall take up in the next section.

The point here is that the real expected payoff is given by

$$E(C_1) = pC_u + (1-p)C_d,$$

where p is the underlying up probability for the stock. But when pricing the option, it is not the real p that ends up being used in the pricing formula, it is the risk-neutral p^* instead. Noticing that p^* in (7) only depends on r, u and d, we conclude that *the price of the option does not depend at all on* p, only on S_0, u, d and r. So to price the option we never need to know the real p.[2].

This irrelevancy of p will later, when we study stock models in continous time, express itself in the famous Black-Scholes pricing formula which does not depend on the mean μ of the underlying Brownian motion, but only on the variance σ^2.

0.2 Pricing other European call options when the expiration date is $t = 1$

The option pricing method of the previous section goes thru for any option in which the payoff (denoted by C_1) occurs at the expiration date $t = 1$;

$$C_0 = \frac{1}{1+r}E^*(C_1) = \frac{1}{1+r}(p^*C_u + (1-p^*)C_d). \tag{10}$$

It is only required that the payoff values C_u and C_d, whatever they are, are known; $C_1 = C_u$ if the stock goes up, $C_1 = C_d$ if the stock goes down. Examples include a *put* option, with payoff $(K - S_1)^+$, in which the holder has the option to sell a share of the stock at price K at the expiration date $t = 1$; $C_u = (K - uS_0)^+,\ C_d = (K - dS_0)^+$.

In words, the option pricing formula says

> *The price of the option is equal to the present value of the expected payoff of the option under the risk-neutral probability.*

[2]But hidden in here is the economic fact that a stock with a higher p would have a higher S_0

0.2.1 Risk-neutral measure

We saw that the price of the European call option can be expressed as an expected value (9) if we use the risk-neutral probability p^* defined in (7). Moreover, p^* only depends on r, u and d, but not on the real value of p underlying the stock's randomness. We conclude that we never need to know what the real p is to compute C_0. We need to know the values of the payoff outcomes, C_u and C_d, but not their probabilities of ocurrence.

p^* has a nice interpretation as the unique probability p making the stock price move in a "fair" way, meaning that given the initial price S_0, the present value of the expected price at time $t = 1$ is yet again S_0: on average, the stock (when discounted) neither goes up nor down in price, it is risk-neutral;

$$(1 + r)^{-1}E(S_1|S_0) = S_0, \ \textit{if } p = p^*. \tag{11}$$

To see that this is so, expanding the expected value in (11) yields the equation

$$(1 + r)^{-1}(puS_0 + (1 - p)dS_0) = S_0,$$

or simply

$$(1 + r)^{-1}(pu + (1 - p)d) = 1,$$

with unique solution

$$p = p^* = \frac{1 + r - d}{u - d}.$$

Thus by imagining that the stock price evolves "fairly" (that is, $p = p^*$), the price of the option can be realized as the expected discounted payoff of the option at time $t = 1$. Changing from p to p^* is sometimes referred to as a *change of measure*, since we have changed the way that the probabilities of stock outcomes are measured; $P(S_1 = uS_0)$ has been changed from p to p^*, and $P(S_1 = dS_0)$ has been changed from $1 - p$ to $1 - p^*$. We thus sometimes say that the stock pricing is being considered under the "risk-neutral measure", meaning that we are using p^*.

Since $\{S_n : n \geq 0\}$ is a Markov process, (11) and the analysis that followed imply that

$$(1 + r)^{-(n+1)}E^*(S_{n+1}|S_n, S_{n-1}, \ldots, S_0) = (1 + r)^{-n}S_n, \ n \geq 0, \tag{12}$$

which means that under the risk-neutral measure, the stochastic process $\{(1 + r)^{-n}S_n : n \geq 0\}$ of discounted prices is a *martingale*[3]. Thus p^* is the unique probability making the discounted stock prices form a martingale. In particular, $(1 + r)^{-n}E^*(S_n) = E^*(S_0), \ n \geq 0$, so if we buy the stock now at time $t = 0$ at price S_0, then $(1 + r)^{-n}E^*(S_n) = S_0$, meaning that under p^* the PV of the expected value of the stock at any time is the same as the initial price we paid.

0.3 Pricing the European options when the expiration date is $t \geq 2$

If the expiration date of a European style option is $t = T$, then we denote the payoff at time T by the random variable C_T. For example, $C_T = (S_T - K)^+$ for the European call option. The various payoff values at time T depend on the outcomes over the T time units. For example

[3] A martingale is a stochastic process $\{X_n : n \geq 0\}$ with the fundamental property that $E(X_{n+1}|X_0, \ldots, X_n) = X_n, \ n \geq 0$. Martingales capture the notion of a fair game in the context of gambling: Letting X_n denote your total fortune right after your n^{th} gamble, the martingale property states that regardless of your past gambles, the next gamble will, on average, neither give you a gain or a loss; each gamble is fair. It immediately follows that $E(X_n) = E(X_0), \ n \geq 0$: when you finish gambling, your expected total fortune is the same as what you started with.

if $T = 2$, then there are the four values $C_{2,uu},\ C_{2,ud},\ C_{2,du},\ C_{2,dd}$ reflecting the up and down outcomes over the two time periods. The probabilities of these are $p^2,\ p(1-p), (1-p)p, (1-p)^2$ respectively, and so the expected payoff is given by

$$E(C_2) = p^2 C_{2,uu} + p(1-p)C_{2,ud} + (1-p)pC_{2,du} + (1-p)^2 C_{2,dd}.$$

For the European call option, order does not matter; $C_{2,ud} = C_{2,du} = (udS_0 - K)^+$, but in general, order will matter. More generally, for the European call option, the payoff at time T is always of the form $(u^i d^{T-i} S_0 - K)^+$ and does not depend on the order in which the ups and downs occurred; for other options order may matter; they are called *path-dependent* options. Examples include an Asian call option with payoff $(\frac{1}{T}\sum_{n=1}^{T} S_n - K)^+$.)

The following is a beautiful generalization of (10):

Theorem 0.1 *Under the Binomial lattice model for stock pricing, the price of a European style option with expiration date $t = T$ is given by*

$$C_0 = \frac{1}{(1+r)^T} E^*(C_T). \tag{13}$$

E^* denotes expected value under the risk-neutral probability p^* for stock price (defined in (7)). In words: "the price of the option is equal to the present value of the expected payoff of the option under the risk-neutral measure".

Applying Theorem 0.1 to a European call option where the order of the ups and downs is irrelevant yields the discrete-time analog of the famous *Black-Scholes-Merton* pricing formula (for European call options):

Corollary 0.1 *Under the Binomial lattice model for stock pricing, the price of a European call option with strike price K and expiration date $t = T$ is given by*

$$\begin{aligned} C_0 &= \frac{1}{(1+r)^T} E^*(C_T) && (14)\\ &= \frac{1}{(1+r)^T} E^*(S_T - K)^+ && (15)\\ &= \frac{1}{(1+r)^T} \sum_{i=0}^{T} \binom{T}{i} (p^*)^i (1-p^*)^{T-i} (u^i d^{T-i} S_0 - K)^+. && (16) \end{aligned}$$

We will prove Theorem 0.1 for $T = 2$, since the $T > 2$ case is analogous. To this end we must show that

$$\begin{aligned} C_0 &= \frac{1}{(1+r)^2} E^*(C_2) && (17)\\ &= \frac{1}{(1+r)^2} \Big[C_{2,uu}(p^{*2}) + C_{2,ud}(p^*(1-p^*)) + C_{2,du}(p^*(1-p^*)) + && (18)\\ &\quad C_{2,dd}(1-p^*)^2 \Big]. && (19) \end{aligned}$$

The key idea: Although we can't exercise the option at the earlier time $t = 1$, we can sell it, so it does have a "price" at that time which we can view as a potential "payoff". At time $t = 1$, we would know what the new price of the stock is, S_1, and we thus could sell the option which then would have an expiration date of $T = 1$. For example, if $S_1 = uS_0$, then we use the $T = 1$ price formula in (10) with outcomes $C_u = C_{2,u,u}$ and $C_d = C_{2,ud}$ yielding the price (denoted by $C_{1,u}$, the price if the stock went up at $t = 1$)

$$C_{1,u} = \frac{1}{1+r}[p^* C_{2,u,u} + (1 - p^*) C_{2,ud}].$$

Similarly, if $S_1 = dS_0$, then

$$C_{1,d} = \frac{1}{1+r}[p^* C_{2,du} + (1 - p^*) C_{2,dd}].$$

But now we can go one more time step back to $t = 0$: We have these known "payoff" values at time $t = 1$ of $C_{1,u}$ and $C_{1,d}$, which we just computed, and thus we can now use them in the $T = 1$ formula (10) again to obtain

$$C_0 = \frac{1}{1+r}[p^* C_{1,u} + (1 - p^*) C_{1,d}] \tag{20}$$

$$= \frac{1}{(1+r)^2}\Big[C_{2,uu}(p^{*2}) + C_{2,ud}(p^*(1-p^*)) + C_{2,du}(p^*(1-p^*)) + \tag{21}$$

$$C_{2,dd}(1-p^*)^2\Big]. \tag{22}$$

In general, the proof proceeds by starting at time T and moving back in time step-by-step to each node on the lattice until finally reaching time $t = 0$. This procedure yields not only C_0 but all the intermediary prices as well.

0.4 Pricing options that allow early exercise

Some options (other than European) allow one to exercise early. For example, in a American put option with expiration date T, the holder has the right to exercise the option at any time $1 \leq t \leq T$. If exercised at time t, the payoff is $(K - S_t)^+$.

Although the pricing formula in Theorem 0.1 is no longer valid, the same method used in its proof yields a method for pricing here too, and also yields the optimal time at which the holder should exercise. We will illustrate here for the put option when $T = 2$.

Pricing an American put option with expiration date $T = 2$

At time $t = 1$ we need to decide wether to exercise or not. Suppose that the stock went up. Then if we exercise we receive payoff $(K - uS_0)^+$. On the other hand the pricing formula (10) yields the price if we hold on to it:

$$V_{1,u} = \frac{1}{1+r}[p^*(K - u^2 S_0)^+ + (1 - p^*)(K - udS_0)^+].$$

Thus we need to compare and choose the one that is larger,

$$C_{1,u} = \max\{V_{1,u}, (K - uS_0)^+\}.$$

Only if $(K - uS_0)^+ > V_{1,u}$ would we exercise early. $C_{1,u}$ is how much the option is worth at time $t = 1$ if the stock went up.

Similarly, if the stock goes down at time $t = 1$, we have

$$V_{1,d} = \frac{1}{1+r}[p^*(K - udS_0)^+ + (1 - p^*)(K - d^2 S_0)^+],$$

and

$$C_{1,d} = \max\{V_{1,d}, (K - dS_0)^+\}$$

is how much the option is worth. Only if $(K - dS_0)^+ > V_{1,d}$ would we exercise early. The two values $C_{1,u}$ and $C_{1,d}$ are then the prices at time $t = 1$, and finally the price C_0 is then given by applying (10) to the two values

$$C_0 = \frac{1}{1+r}[p^* C_{1,u} + (1 - p^*) C_{1,d}].$$

In general, when the expiration date is T, one computes all the prices at all the intermediary points step-by-step in an analogous way. The optimal time to exercise is then determined as follows: at time $t = 1$, when the new value of the stock is known, one checks to see if the payoff from exercising exceeds the computed V value (using (10)) at that node. If it does then exercise, otherwise wait one more unit of time and check again, and so on until (if at all) finding the first time for which it is optimal to exercise.

It is never optimal to exercise early for an American call option

For the American call option with $T = 2$,

$$V_{1,u} = \frac{1}{1+r}[p^*(Ku^2 S_0 - K)^+ + (1 - p^*)(udS_0 - K)^+],$$

and it is easily seen that $V_{1,u} \geq (uS_0 - K)^+$. Similarly

$$V_{1,d} = \frac{1}{1+r}[p^*(KudS_0 - K)^+ + (1 - p^*)(d^2 S_0 - K)^+],$$

and $V_{1,d} \geq (dS_0 - K)^+$. So it is always optimal to wait until the end at time $T = 2$. One can generalize easily to see that this is so for any expiration date T. This means that the two call options (European, American) are really identical, and have the same price.

0.5 Monte Carlo simulation for pricing options

As a motivating example, suppose we wish to compute the expected payoff of some option $E^*(C_T)$ (using the risk-free measure), so as to get the price $C_0 = (1 + r)^{-T} E^*(C_T)$. C_T is a random variable and under the binomial lattice model, $S_n = S_0 Y_1 \times \cdots \times Y_n$, it is some function h of S_0 and $Y_1, \ldots, Y_T$; $C_T = h(S_0, Y_1, \ldots, Y_T)$. For example, if $C_T = (S_T - K)^+$, the payoff for the European call, then $h(S_0, Y_1, \ldots, Y_T) = (S_0 Y_1 \times \cdots \times Y_T - K)^+$.

Note that if U is uniformly distributed over the continuous interval $(0, 1)$, then Y defined by $Y = uI\{U \leq p^*\} + dI\{U > p^*\}$ has the correct distribution of the iid Y_i; $P(Y = u) = p^*$, $P(Y = d) = 1 - p^*$.

Thus we can ask our computer to hand us T iid uniforms $U_1, \ldots, U_T$; construct the iid $Y_i = uI\{U_i \leq p^*\} + dI\{U_i > p^*\}$, and then compute a first sample $X_1 = h(S_0, Y_1, \ldots, Y_T)$. Then, independently, we do this again to obtain a second sample, X_2, and keep on doing so a total of n times where n is "large", obtaining n iid copies $X_1, \ldots, X_n$. Then we use the estimate

$$E^*(C_T) \approx \frac{1}{n} \sum_{i=1}^{n} X_i, \tag{23}$$

which is justified via the strong law of large numbers which asserts that wp1, the approximation becomes exact as $n \to \infty$.

We illustrate with some simple examples:

1. (Down and out call option)

   Suppose that we start with a European call option but we also introduce a level $0 < b < S_0$ and some pre-specified times $0 < n_1 < n_2 < \cdots < n_k < T$ at which it must hold that $S_{n_i} > b,\ i \in \{1, \ldots, k\}$. If at any one such time n_i is holds that $S_{n_i} \leq b$, then the option becomes worthless. Thus the payoff at time T is given by

   $$C_T = (S_T - K)^+ I\{S_{n_1} > b,\ S_{n_2} > b, \ldots, S_{n_k} > b\}.$$

   Computing exactly the expected value of this payoff is not possible in general, so we will estimate it via Monte Carlo. For concreteness, let's assume that $T = 10$, and there are three checking times $n_1 = 2,\ n_2 = 4,\ n_3 = 7$. First we generate Y_1 and Y_2 so as to simulate the first needed value, $S_2 = S_0 Y_1 Y_2$. If $S_2 \leq b$, stop and set $C_T = 0$; otherwise generate (independently) Y_3, Y_4, to obtain $S_4 = S_2 Y_3 Y_4$. If $S_4 \leq b$, stop and set $C_T = 0$; otherwise generate (independently) Y_5, Y_6, Y_7, to obtain $S_7 = S_4 Y_5 Y_6 Y_7$. If $S_7 \leq b$, stop and set $C_T = 0$; otherwise generate (independently) Y_8, Y_9, Y_{10}, to obtain $S_{10} = S_7 Y_8, Y_9, Y_{10}$ and set $C_T = (S_{10} - K)^+$. This would give us our first copy of C_T, denote by X_1. Now repeat this procedure independently again and again until finally obtaining n such copies $X_1, \ldots, X_n$, finally using (23) as our estimate of the desired expected value.

2. (Asian call option)

   Here the payoff involves the entire average over all T time periods instead of only S_T:

   $$C_T = (\frac{1}{T} \sum_{i=1}^{T} S_i - K)^+.$$

   In this case we need to sequentially generate all T of the Y_i and use them to construct the sum $\frac{1}{T} \sum_{i=1}^{T} S_i$, and our copy of $C_T = (\frac{1}{T} \sum_{i=1}^{T} S_i - K)^+$. Then we repeat this n times and so on.