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.

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 $h = t/n$, $(0, h], (h, 2h],\ \ldots((n-1)h, nh = t]$, 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.

It follows from (1) that for a given n, $S_n = u^i d^{n-i} S_0$ for some $i \in \{0, \ldots n\}$, 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 $S_0(1+r)\geq S_0u$ at time $t=1$ by buying S_0 shares of the 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 unit 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 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.

$C_1=C_u=(uS_0-K)^+$ if the stock goes up, and $C_1=C_d=(dS_0-K)^+$ if the stock moves down.

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 $\alpha uS_0+\beta(1+r)=C_u$ and $\alpha dS_0+\beta(1+r)=C_d$. 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, $\alpha uS_0+\beta(1+r)=C_u$ and $\alpha dS_0+\beta(1+r)=C_d$, 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 can be 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^* (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.

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.

0.1.1 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 the real p plays no role in the pricing of the option; we never need to know what it 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. For a given S_0, r, u and d, different values of p yield the same price C_0[1].

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;

$$(1+r)^{-1}E(S_1|S_0) = S_0, \ \textit{if } p = p^*. \tag{10}$$

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

To see that this is so, expanding the expected value in (10) 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, (10) and the analysis that followed imply that

$$(1+r)^{-(n+1)}E^*(S_{n+1}|S_n, S_{n-1}, \ldots, S_0) = (1+r)^{-n}S_n, \ n \geq 0, \tag{11}$$

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*[2]. 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.

0.2 Pricing the European call option when the expiration date is $t \geq 2$

The following is a beautiful generalization of (9) and is the discrete-time analog of the famous *Black-Scholes-Merton* pricing formula for European call options:

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

$$C_0 = \frac{1}{(1+r)^n}E^*(C_n) \tag{12}$$

$$= \frac{1}{(1+r)^n}E^*(S_n - K)^+ \tag{13}$$

$$= \frac{1}{(1+r)^n}\sum_{i=0}^{n}\binom{n}{i}(p^*)^i(1-p^*)^{n-i}(u^i d^{n-i}S_0 - K)^+. \tag{14}$$

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".

[2] 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.

It in fact can be shown that the above price converges exactly to the Black-Scholes-Merton price for a European call option (under geometric BM) as the approximating interval length $h \downarrow 0$.