Understanding Volatility and the Binomial Market Model

Roger E. Ison, Mantic Software Corporation

This article and the associated software are designed to give you some practical insight into the *binomial model* of stock price movements. The binomial model is presented in every finance text that covers the theory of option valuation. It's the most widely used method for computing the "fair value" of American-style options (contracts permitting the holder to exercise them before expiration).

Its theoretical foundation is the famous "random walk" hypothesis, but textbooks don't present it from that perspective. I thought it would be instructive and entertaining to see and test the real implications of this model, so I wrote some software to help visualize the workings and statistical properties of a binomial market.

It's a simple model, driven by a uniformly distributed random number generator, with straightforward statistical properties and no "market logic" or intelligent trading agents. Yet it generates remarkably market-like price behavior, and not just in broad statistics. Long- and short-term price runs and cyclical movements appeared plainly in the graphs, which are much more realistic-looking than I'd expected. As Figure 1 illustrates, even the long- and short-term moving averages have a surprisingly realistic appearance, at least to the "naked eye". Why do price series derived from random numbers exhibit so much apparent structure?

Figure 1

The answer is that the binomial model *inevitably* produces a kind of non-periodic, cyclic price motion. Prices in a binomial series are more than random points with a certain statistical distribution, even though a random number generator helps produce them. Like real market prices, each binomial price is related to other, recent prices, and carries an implicit forecast about the rate of change and size of future price movements.

Binomial cycles result from a particular *process* of price movements, and the binomial process is a reasonable description of some aspects of real markets. To the extent that it reflects reality, it must contribute to the cyclical patterns we see in real stock prices. It may even be possible to distinguish this inherent, non-periodic cyclicality from meaningful changes in a stock's price behavior.

One conclusion I personally take away from this mathematical experiment is that statistics by themselves are weak tools for characterizing market behavior. To *really* understand markets, we will have to model and understand the underlying processes which produce the price movements we observe.

A Binomial Refresher

Binomial price movements depend on two parameters: the *risk-free interest rate* (the rate earned on loans with essentially no risk of default, e.g. the money-market mutual fund rate), and the stock's *price volatility*.

Volatility, denoted by the Greek letter σ, is familiar to option traders as a measure of how much the stock price is likely to change over time. Mathematically, σ is the standard deviation of the logarithm of the daily rate of return on a stock. From a sequence of daily prices $P_1, P_2, P_3 \ldots P_n$ we calculate each daily rate of return as today's price divided by yesterday's price, $R_{i+1} = P_{i+1} / P_i$. This is just the relative or percent change in stock price from each day to the next. Then we take the logarithms of these numbers, $X_i = \log(R_i)$, and compute the standard deviation of the series X_i in the usual way:

$$s = \sqrt{\frac{\sum (X_i - \overline{X})^2}{n-1}} \qquad \text{where } \overline{X} \text{ is the average of all the } X_i$$

This daily σ can stretched mathematically to span any time period - an hour, a month, or a year - but it's conventionally stated as an annualized value. According to statistical theory, if a normally distributed (Gaussian) random variable has standard deviation σ over time *t*, then its standard deviation over time *T* is $s\sqrt{T/t}$.

Notice that standard deviation actually contains no information about the *direction* of price movements, only their *magnitude*. You can't use standard deviation to spot or characterize trends. The formula takes the squared difference $(X_i - \overline{X})^2$, where $\overline{X}$ is the average of all the X_i. Squared numbers are always positive, so the direction of each price movement is lost. A long up or down trend may have the same standard deviation as a chaotic sequence of up-and-down moves.

Binomial price movements

To create binomial price movements, we essentially reverse the process used to compute σ. Whenever we want to know the "next" price, we flip a biased coin to decide whether the next price movement will be up or down. The new price is calculated by multiplying the current price by an "up" percentage or a "down" percentage:

$$P_{i+1} = u \cdot P_i \qquad \text{if the price moves up}$$
$$P_{i+1} = d \cdot P_i \qquad \text{if the price moves down}$$

You can see that *u* and *d* are just conditional, periodic rates of return; they're special cases of the daily return P_{i+1} / P_i used to calculate σ. The trick is choosing a bias for the coin that produces the correct proportion of up and down movements, and values of *u* and *d* so that the standard deviation of the generated series $P_1, P_2, \ldots P_n$ actually converges to σ if we generate a long enough price series.

Cox, Ross and Rubinstein showed how to choose these values so that in the long run, the distribution of binomial stock price movements is almost log-normal. Then the values of $\log\left(P_{i+1} / P_i\right)$ naturally distribute into a classical bell-shaped curve, like the grades of students in a large class, and have the desired standard deviation. The C-R-R formulas are:

$u = e^{s\sqrt{t}}$	the "up" multiplier
$d = 1/u$	the "down" multiplier
$p_u = \dfrac{(r-d)}{(u-d)}$	probability of an "up" movement
$p_d = 1 - p_u$	probability of a "down" movement

Why not exactly log-normal? The model assumes traders can freely perform arbitrage among mispriced puts and calls, stock, and risk-free bonds; so the risk-free interest rate *r* over the time between consecutive binomial prices figures into the equations. Thus *r* and sigma combine to skew the distribution slightly away from true log-normal; the binomial series converges exactly upon log-normal only in the special case when $p_u = p_d$.

Physical meaning of the binomial parameters

We should take a moment to think about the physical and economic meaning of the binomial parameters.

The up multiplier *u* predicts that if the stock price makes an up movement, it will rise by an incremental amount related to time and volatility. The expression $u = e^{s\sqrt{t}}$ actually has its roots in physics of "Brownian motion", not statistics or finance. Brownian motion is the motion of gas molecules as they travel randomly through space, having their direction of motion altered by chance as they bang into each other. It's the original random walk! Einstein studied this motion and showed that on average, the distance traveled in a random walk along a single dimension is characterized by $e^{s\sqrt{t}}$. In physics, temperature is a key factor determining the value of sigma.

Students of price trends and persistence say that when the actual values of *u* and *d*, and the probabilities match this physical model, the price movements are a random walk. If the distance traveled over time is less than $e^{s\sqrt{t}}$, the price series is said to be anti-persistent. If the distance is greater, the series shows more persistence (a greater tendency to sustain trends of direction) than would be consistent with a random walk.

The risk-free rate of return *r* specifies how far the price must move *consistently* to provide the same rate as a risk-free bond. In other words, if the stock moved up by a proportion $(1+r)$ on every single move, its return at the end of the process would be exactly the risk-free discount rate. Repeated upward price motions compound the risk-free rate at each price movement, so that after *n* price movements the compounded result will be $(1+r)^n$.

The binomial parameters are chosen so that the expected rate of return for holding stock over a single price step equals the risk-free rate: $E = u \cdot p_u + d \cdot p_d = r$. You can verify this for yourself by substituting values from the Cox-Ross-Rubinstein equations.

Moreover, for the model to make economic sense, it requires that *u* > *r* > *d*. If the return for owning a stock, given that it goes up, doesn't exceed the risk-free rate, a rational investor would always prefer to own a risk-free bond that pays better, so the model requires that $u > r$. Similarly, if the return for owning the stock, given that it goes down, exceeds the risk-free rate, nobody would hold the bonds.

How realistic is this model?

A number of academic studies have measured whether actual stock price movements have log-normal distributions over long time spans. In general, these studies conclude with a qualified *yes*. The humped shape and long tails of a log-normal distribution, illustrated by the green bars at the right side of Figure 3, form a pretty good statistical picture of stock price movements for a large universe of stocks. The tails (the price extremes outside the main hump) of actual stock price distributions tend to be "fatter" than theory predicts, and the hump not quite as large, so real prices show some systematic deviation from log-normality.

In other words, the probability of extreme up or down movements is higher than we'd expect from theory. The usual interpretation of this observation is that price volatility is not "stationary" - it occasionally changes (jumps to a substantially different value) when a company's prospects are suddenly perceived to change. The basic binomial model doesn't capture these volatility jumps, and while there are several plausible ways to extend it, nobody has shown conclusively that jump processes actually produce better results. In fact, some option pricing studies have come to the opposite conclusion.

On the other hand, it's also reasonable to be skeptical about such studies. Statistical characterizations have their value, but there are other skewed, long-tailed distributions that may be plausible candidates to describe market price movements. Fat humps and long tails tell us plainly that actual price movements are either statistically persistent, or antipersistent, compared with the random walk hypothesis.

From a trader's point of view, the binomial model is a reasonable one under "calm" circumstances. When buyers and sellers arrive at market in a more-or-less random order, typically when there's no news about the stock, we'd expect roughly random sequences of up and down ticks. We expect clusters of buyers or sellers to drive the price up or down only when something noteworthy happens. Even then, market specialists are expected to use their inventory and capital to maintain smooth, incremental stock price changes.

It's fractal

The binomial model is *fractal* - it has similar behavior and statistical properties at any time or price scale. It's fractal with respect to price because, at each step, the current price is multiplied by u or d, changing it in proportion to its magnitude. It's fractal with respect to time because interest earnings and the size of up and down moves are scaled by the factor $e^{s\sqrt{t}}$. The model also preserves its internal consistency even if the stock price drifts very high or low after a long series of movements.

By implication, any conclusions we draw from the model should hold true whether it's applied over short, medium or long time spans. In the real world, that just isn't true. For instance Malkiel's popular book, *A Random Walk Down Wall Street*, presents data showing that the expected return for holding stock more than five years is much better than the return for short holding periods.

An example

We saw earlier that the standard deviation statistic contains no information about price trends. To explain why price trends appear in binomial series, we must examine the specific sequences of up-and-down movements the binomial process can generate, and estimate their probabilities. That's the only way to really understand what's going on in Figure 1. Since the model is fractal, our analysis of price movement sequences will apply at all time and price scales.

σ in the C-R-R equations is an annualized standard deviation; *t* is the time between consecutive price samples, stated as a fraction of a year; and *r* is the interest rate over time *t*. If the risk-free rate is 4% and our simulated stock has annual volatility of 0.15, then for daily price movements:

$t = 1/365 = 0.002740$	one day as a fraction of a year
$r = (1+0.04)^t = 1.000107$	daily return on risk-free bonds
$u = e^{s\sqrt{t}} = 1.007882$	magnitude of movement if price rises
$d = 1/u = 0.992179$	magnitude of movement if price falls
$p_u = 0.504880$	probability that the price will rise
$p_d = 0.495120$	probability that the price will fall

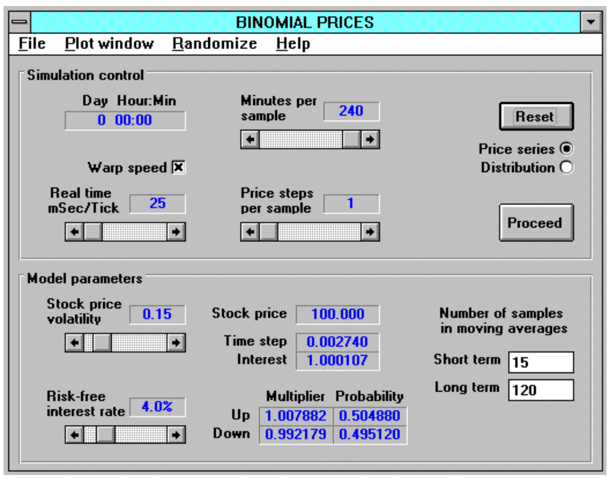

Figure 2

Suppose the stock's price is 100 dollars at time t_0. The next "moment" in the series, time t_1, is one day later, when the price will be either $P_1 = 100 \cdot u = 100.7882$ (with probability 50.4880%) or $P_1 = 100 \cdot d = 99.2179$ (with probability 49.5120%). If the price goes up at step t_1 then falls at t_2, price $P_2 = P_0 \cdot u \cdot d = P_0$ since $d = 1/u$.

Over time, this process carries the stock price through an arbitrary series of up and down movements like *UUDDUDUDUDDD*.... Figure 3 illustrates this by mapping out every possible sequence of up and down movements for 15 steps from an initial price of 100, when σ = 0.15, the risk-free rate is 4%, and the price steps are 30 days apart. Since there are two possible moves (up or down) at each time step, there are 2^n possible paths of *n* steps through the map; but the reciprocal relationship of *u* and *d* means there are only *n* possible prices the stock can have after *n* price movements.

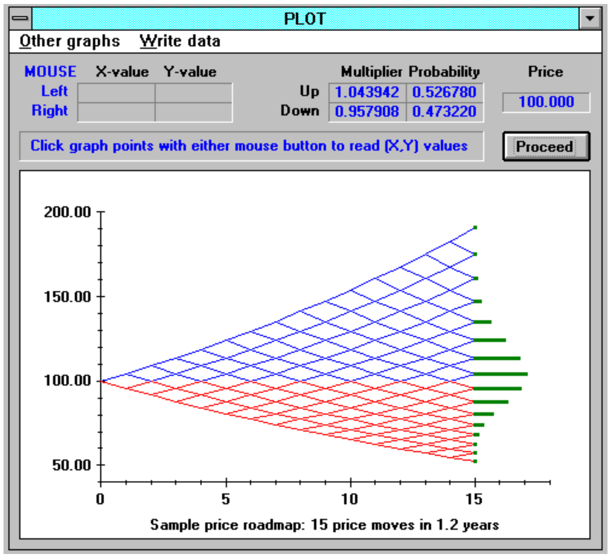

Figure 3

The green bars at the right end of the price map show the probability of each price outcome at the final step. Knowing that the probability at each step is either p_u or p_d, we can explore all the possible paths leading to each terminal node of the map and accumulate the probabilities along each path as we proceed.

A better way to find these path probabilities is to observe that if a node at the end of an *n*-move path has *j* more "up" moves than "down" moves, the number of possible paths to that node must be *C(n,j)*, the number of combinations of *n* things taken *j* at a time. This is true because the order of the moves is irrelevant: *UUUDDUD* gets to the same place as *DDDUUUU*, or any other combination of four "ups" and three "downs". Each such path to the end has *j* up moves with probability p_u, and *(n-j)* down moves with probability p_d, so the total probability of reaching that particular end node is the probability of that price outcome, times the number of ways to get there: $C(n, j)\, p_u{}^{j}\, p_d{}^{n-j}$.

The probability bars in Figure 3 have the characteristic shape of a log-normal probability distribution. This shape reflects the fact that the price can never reach zero, no matter how many down moves take place; each downward move shrinks the previous price by a fixed fraction *d*, always less than one. The upside potential is unlimited, since *u* is always greater than one.

You can see why the probability of ending up at prices not too far from the initial price is higher than the chance of reaching the extremes. There are many combinations of up/down motions leading to the middle of the price range, but the two price extremes can only be reached by relatively improbable sequences of all-up or all-down movements.

It's sometimes asserted that in the long run, the binomial process converges to a log-normal distribution. That isn't quite true. Figure 4 illustrates how the interest rate distorts the probability distribution of long binomial prices runs.

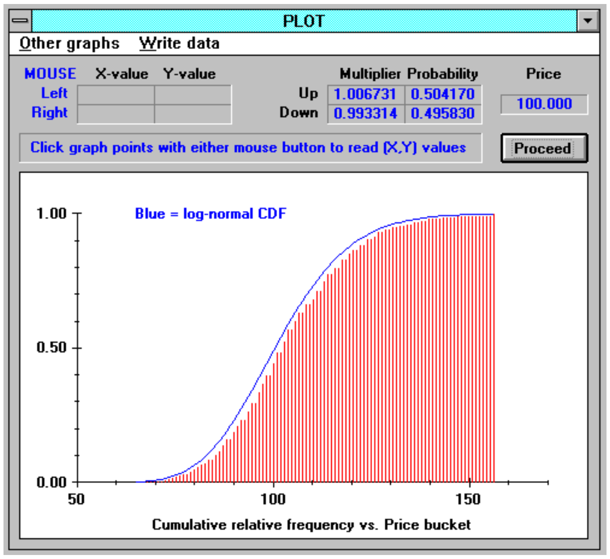

Figure 4

The data in the Figure were produced by repeatedly generating sequences of 500 binomial price moves over one simulated year, starting each sequence at an initial price of 100. The 500th (year-end) price of each sequence is taken as a sample data point, and the graph shows the cumulative probability distribution function (CDF) of more than 1000 data points. In other words, it shows the probability that the binomial stock price after 500 moves will be less than or equal to each price *P* on the x-axis. You can see that this binomial distribution deviates systematically from the log-normal CDF drawn in blue. If I pick values of σ and interest rate to make $p_d > p_u$, the binomial CDF will spill out to the left of the log-normal line. The curves match only in the special case where $p_u = p_d$. At 4% interest, this happens when σ is about 0.28.

Price runs

So, where do the prominent price cycles in Figure 1 come from?

Consider any price path that always stays within the blue region of Figure 3. In such a “price run”, the simulated stock price never falls below the initial price. It’s a simplified form of price cycle: the price heads up and stays up for a while, until it eventually falls back below the initial price. Similarly, a “down” run stays within the red area until it finally moves back above the initial price.

A run that reaches the right edge of the blue area somewhere above the initial price after *n* price steps, requires additional price steps before it can possibly fall back to the initial price. So in Figure 3, *all runs of length 15 or longer* pass through the right edge of the price map. Obviously the map can be extended to cover any number of price steps.

How likely are such price runs?

I don't think there's a simple, closed formula involving combinations or permutations that gives the answer, because the constraint that runs can't cross the middle of the price map excludes certain combinations of up-and-down movements. For example, *UDDUU* is an excluded (wandering) path while *UUDDU* is a valid run in the blue region. Still, we can count the runs. Look first at the path consisting exclusively of upward moves. This is the upper edge of the blue region in Figure 3, and there is only one way to get to any node along this path: all the moves leading to it are up moves.

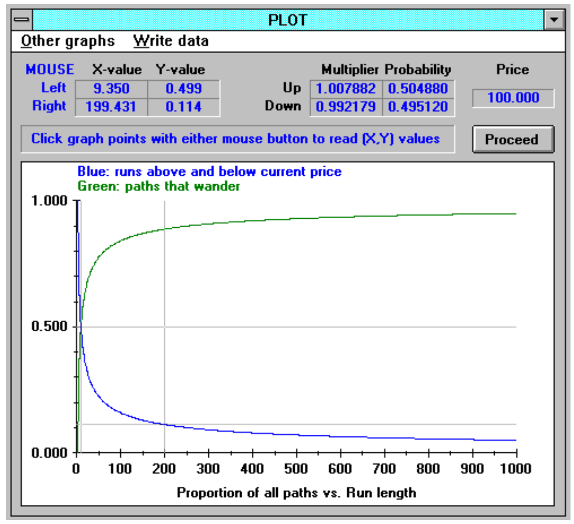

Figure 5

Now if you look at any "interior" node adjacent to this upper edge, but still within the blue area, there will be two ways to get to it (down or up from the preceding nodes), except that the bottom blue nodes are reachable from just a single blue predecessor. Obviously the total number of paths to any interior node is the sum of the number paths to its one or two direct predecessors. With this observation, you can work your way through the price map from left to right. When you reach a node at the right end, you'll have accumulated the total number of paths leading to it.

This counting procedure is easy to perform, but the bookkeeping is best done by a computer.

First, let's see what proportion of all the possible *n*-move paths are runs, and what proportion wander above and below the initial price. Figure 5 plots the answer for typical conditions (σ = 0.15, interest rate 4%, daily price movements). Each point on the blue curve indicates the proportion of all the possible paths of length *n* price steps *or longer*, that are runs above or below the initial price. The green curve counts all the other "wandering" paths.

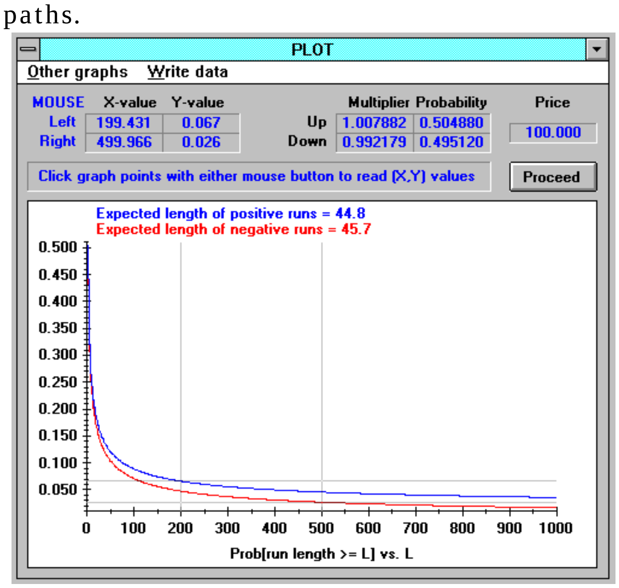

Figure 6

Using the mouse, I marked two points in Figure 5 with hairlines, and the corresponding (x,y) values are displayed in boxes above the graph. You can see that more than half of all paths of length 9 or less are runs. Even when we consider paths of at least 200 moves, more than 10% are runs.

If a price outcome after *n* steps is *j* "up moves" above the initial price, the probability of any run leading to that price is, as before, $p_u^{\,j} p_d^{\,n-j}$. And we know how to count the runs leading to each price, so we can calculate the cumulative probability of runs to each price at the right edge of the price map, and add them up to get the total probability of runs with length greater or equal to value of *n*.

Figure 6 plots the result of this computation for our familiar case when σ = 0.15 and interest is 4%. It shows that the chance of a run exceeding 200 moves is more than 10%, and likelihood of runs exceeding 1000 moves is nearly 5%. Given these probabilities for runs of various lengths, we can also estimate the probability-weighted, expected length of runs above and below the initial price - about 44 steps in the present example. But notice how the curves flatten as *n* increases. The probability of a run at least 500 steps long is not much less than the probability of a run at least 200 steps long. Intuitively, the reason is that runs which do migrate far from the original price require many moves before they can possibly return to that original price. Expect some very long runs.

10,000 steps? 20,000? Each new price move doubles the number of possible paths in the map, and every path that was a "wanderer" after *n* steps forks into two wandering paths in the next price step. However, that next step also turns some paths, which had been successful runs, into wanderers. Therefore the proportion of wanderers always increases, and the probability of an awesomely long run is vanishingly small.

All good runs come to an end; but in real markets, it ain't necessarily so. As I mentioned earlier, studies of actual market history indicate that if you hold "an average stock" long enough, the return on your investment will become deterministically positive. This is obviously contrary to our conclusion about the binomial model.

The Cox-Ross-Rubinstein option value

Having come this far, we ought in passing to discuss the binomial value of an American option. Figure 7 illustrates call values attached to a price map. Initially, the map spans the time from "now" to the terminal price when the option (in this example, a $100 call) expires. Off each price node at the right edge of the map, the program displays the value the option would have if the stock has the corresponding price at expiration. If the stock closes at $100 or below, the option expires worthless. Above $100, its value rises dollar-for-dollar with the stock.

Now let's back up one step in the price map (Figure[8]). As we move left to the predecessor node of each terminal price, we can calculate the expected value of the option at the preceding price, since we know probabilities p_u and p_d. The expected value is just the probability-weighted sum of the two terminal values. To be absolutely precise, we discount this value by the risk-free interest rate for the time between the nodes.

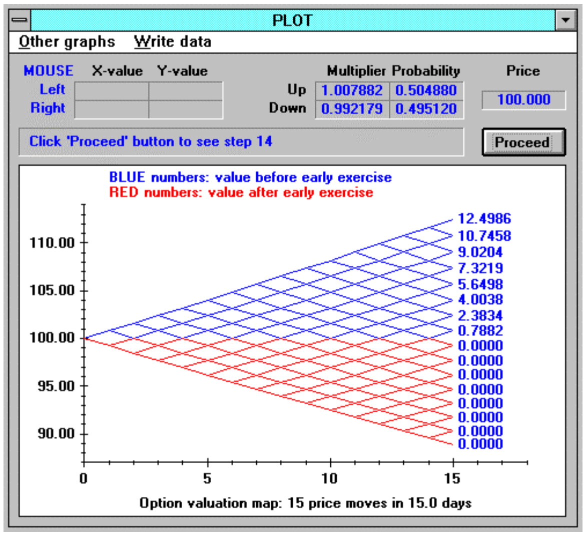

Figure 7

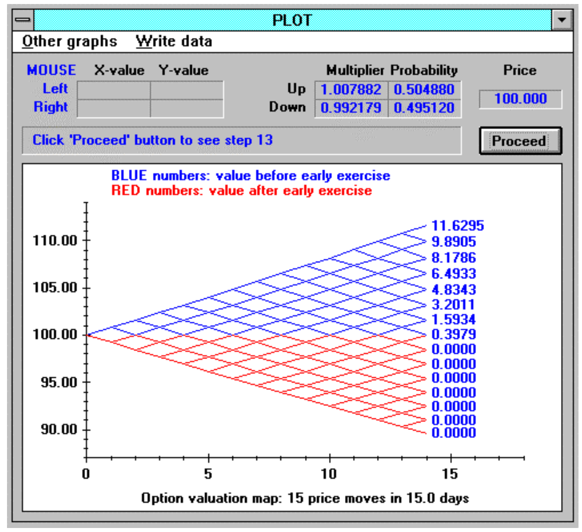

Figure 8

This procedure is simply repeated as we propagate back through each step of the price map, until the current date is reached. The result is the "fair value" of the option. As a further refinement, we can test to see whether the option should be exercised at each step. It should be exercised if its intrinsic value is less than the value that would be received by exercising it. For instance, if the option is a put, cash is received for shares when it's exercised. If the interest that can be earned on the cash is more than the value of the option, it would be economically rational to exercise it. The Binomial Market model program shows this effect when valuing an in-the-money put.

The binomial forecast

Essentially, the binomial model acts out a forecast of future price movements implied by the value of σ, the stock price volatility. It's true that any particular path produced by a simulation is selected by a random number generator; but as the analysis demonstrates and the eye confirms, *every* binomial path of any significant length exhibits runs and cycles.

So we have a statistically driven price movement model, with no trading or economic information content, that inherently produces price cycles much like those in real markets. What can it tell us about real markets? Here's an entertaining speculation: perhaps binomial cycles could be used as norms for comparison with real price series. For instance, we might try to “subtract” the inherent cyclicality implicit in a particular value of σ from actual price series. What's the motivation for this idea?

The binomial process approximates what mathematicians call a *continuous diffusion* process. Roughly speaking, this means small, incremental stock price movements in the absence of surprises. These processes assume we can accurately characterize the expected return for owning the stock, which presumptively has no “unexpected returns” when the price gaps up or down.

Of course, it's those unexpected returns that traders want to spot; and the sooner, the better. We'd like to notice when a stock starts experiencing “discontinuous diffusion”.

Under the continuous diffusion assumptions, a stock's price volatility implies certain statistical expectations about its future price movements within a specified time period: in other words, probable limits. We usually describe this in static terms. For instance, we might calculate a 99% chance (three standard deviations) that the stock price will be between 85 and 119 after six months. It could exceed those limits, but the chance of great deviations is much less than the chance of small ones.

A more fundamental interpretation of volatility refers to the underlying price change process. The relative price change P_{i+1} / P_i is an estimate of the slope or first derivative of the price line over the time period $(t_{i+1} - t_i)$; thus σ actually characterizes the *rate of change* of the stock price. The price change over a specified (long) time is the sum of all the little changes that take place during that interval, which binomial prices approximate with discrete up and down price movements. The price limit boundaries described in the previous paragraph are approximated by the maximum number of up or down movements that can take place in the specified time.

If you're the kind of person who likes to think about "technical" trading tools, the foregoing analysis of run probabilities suggests some ideas you might like to explore.

First, is the distribution of run lengths for actual stock price movements really fractally similar on different time scales? This could be tested by accumulating the actual number of "up" and "down" runs at different time scales: say, comparing the frequency of runs in hourly price samples with runs in the weekly and monthly numbers. If the relationship between the short-term and long-term distributions changes suddenly, that may be a sign of something interesting happening to the stock.

Second, is the distribution of "down" runs similar to that of "up" runs? We would expect so in a market at equilibrium. If the equilibrium shifts or breaks down, there should be a period when the up and down run distributions become sharply different. There are a number of tests that can be used to determine when probability distributions are significantly different from each other or a "standard" benchmark distribution.

I'd like to tell you I've done this with real data - but I haven't yet. I'd love to hear from anyone who tries these tests, or has other or better ideas.

Conclusion

These speculations are entertaining, but a more basic lesson is this: if you believe that most of the time, the "next" stock price is closely related to recent prices - which is the core assumption of the binomial process - then you must expect cycles and long runs (trends) to appear in price data even when psychology, expectations, momentum and other technical factors are neutral. Price trends alone, without measuring other tangible factors, can easily mislead unless you account for the inherent cyclicality described here.

A deeper observation has to do with the importance of *process* over statistics. Despite its simplicity, the binomial model is very carefully and cleverly contrived. Its process structure - a rigid assumption of fixed-magnitude up-or-down movements - directly constrains the simulated stock's time-rate-of-change.

Real markets, on the other hand, often exhibit powerful trends that move outside these statistical constraints. When this happens we can say "the volatility has changed", but that really means the model broke down. Although we adjust the parameters, sooner or later the model will fail again. Some market trends are driven by intangibles like trading psychology, but others are based on persistent value relationships among investment alternatives (like the relationship between stock and bond prices). Models that don't account for the non-linear, feedback properties of these relationships obscure reality, even if the resulting probability distributions have statistically plausible shapes. The most important thing to remember about empirical probability distributions is the enormous amount of information discarded to derive them.

Further reading

John Hull's book, *Options, Futures and Other Derivative Securities,* is a standard text that covers the binomial model.

I like the presentation in Rajna Gibson's *Option Valuation.* She includes a critique of the binomial model *vs.* real security prices, and a useful review of the work on price movement processes where the volatility occasionally "jumps".

You might also enjoy Edgar Peters's book, *Fractal Market Analysis.* He discusses a technique called *rescaled range analysis* or *R/S analysis* that can be used to measure and characterize non-periodic cycles like those generated by the binomial market model.

Author's biography

Roger Ison is cofounder and president of Mantic Software Corporation in Loveland, Colorado. Mantic makes financial and investment modeling software. Before founding Mantic, Ison was an R&D manager and business strategist at Hewlett-Packard. He studied political science and economics as an undergraduate, and earned his M.S. and Ph.D. in Computer Science at the University of Virginia.