Derivative Securities -- Spring 2007 — Section 4

Notes by Robert V. Kohn, modified by Steve Allen
Courant Institute of Mathematical Sciences

Lognormal price dynamics and passage to the continuum limit. After a brief recap of our pricing formula, this section introduces the lognormal model of stock price dynamics, and explains how it can be approximated using binomial trees. Then we use these binomial trees to price contingent claims. The Black-Scholes analysis is obtained in the limit $\delta t \to 0$. As usual, Baxter & Rennie captures the central ideas concisely yet completely (Section 2.4) Hull has a lot of information about the lognormal model scattered through Chapter 13.

$$*********************$$

Recap of the multi-period option pricing formula. Recall what we achieved at the end of the last section: if the risk-free rate is constant and the risky asset price evolution is described by a multiplicative binomial tree with $F_{up} = uF_{now}$ and $F_{down} = dF_{now}$ then the value at time 0 of a contingent claim with maturity $T = N\delta t$ and payoff $V(F_T)$ is

$$V_0 = e^{-rT} E_{RN}[V(F_T)]$$

where $E_{RN}[V(F_T)]$ is the expected final payoff, computed with respect to the risk-neutral probability:

$$E_{RN}\left[V\left(F_T\right)\right] = \sum_{j=0}^{N}\binom{N}{j}q^{j}\left(1-q\right)^{N-j}V\left(F_0 u^j d^{N-j}\right)$$

with q=(1-d) / (u-d). Let's check this assertion for consistency and gain some intuition by making a few observations:

What if the contingent claim pays the forward price itself? This is the case $V(F_T) = F_T$. It is replicated by the portfolio consisting of one unit of the forward (no trading). So the present value should be F_0, the price of the forward now. Let's verify that this is the same result we get by "working backward through the tree." It's enough to show that if $V(F) = F$ for every possible price F at a given time, then the same relation holds at the time just before. To see this let "now" refer to any possible forward price at the time just before. We are assuming $V(F_{up}) = F_{up}$ and $V(F_{down}) = F_{down}$ and we want to show $V(F_{now}) = F_{now}$. By definition,

$$V(F_{now}) = qF_{up} + (1-q)F_{down}$$

with $q = \dfrac{F_{now} - F_{down}}{F_{up} - F_{down}}$. So $V(F_{now}) = \dfrac{F_{now} - F_{down}}{F_{up} - F_{down}}F_{up} + \dfrac{F_{up} - F_{now}}{F_{up} - F_{down}}F_{down}$

$$= F_{now} \frac{F_{up} - F_{down}}{F_{up} - F_{down}}$$

$$= F_{now}$$

Hence, $E_{RN}[F_T] = F_0$.

What if the contingent claim is a forward contract with strike price K? Under our standing constant-interest-rate hypothesis we know the present value should be $e^{-rT}(F_0–K)$ if the maturity is $T=N\delta t$. Let's verify that any binomial tree gives the same result. The payoff is $V(F_T) = F_T - K$. Our formula

$$e^{-rT} E_{RN}[V(F_T)]$$

is linear in the payoff. $E_{RN}[K]=K$, i.e. the total risk-neutral probability is 1; this can be seen from the fact that $(q + [1 - q])^N = 1$. And $E_{RN}[F_T] = F_0$, as just shown above. Thus our formula for the value of a forward is

$$e^{-rT}E_{RN}[F_T - K] = e^{-rT}E_{RN}[F_T] - e^{-rT}E_{RN}[K] = e^{-rT}(F_0 - K)$$

 as expected.

What if the contingent claim is a European call with strike price K >> F_0? We expect such a call to be worthless, or nearly so. This is captured by the model, since only a few exceptional paths (involving an exceptional number of ``ups") will result in a positive payoff.

What if the contingent claim is a European call with strike price K << 0? We expect such a call to be worth about the same as a forward with strike price K. This too is captured by the model, since only a few exceptional paths (involving an exceptional number of ``downs") will result in a payoff different from that of the forward.

Analogous observations hold for European puts.

$$**************** $$

Lognormal price dynamics. Our simple model of a risk-free asset has a constant interest rate. A bond worth ψ_0 dollars at time 0 is worth $\psi(t) = \psi_0 e^{rt}$ dollars at time t. The quantity that's constant is not the growth rate $\dfrac{d\psi}{dt}$ but rather the interest rate

$$r = \frac{1}{\psi} \frac{d\psi}{dt} = \frac{d \log \psi}{dt} .$$

Our forward is risky, i.e. its evolution is unknown and appears to be random. We can still

describe its dynamics in terms of an equivalent interest rate for each time period. Breaking time up into intervals of length δt, the equivalent interest rate for $j\delta t < t < (j+1)\delta t$ is r_j if $F((j+1)\delta t = e_j^{r\delta}t\, F(j\delta t)$, i.e.

$$r_j = \frac{\log s\big((j+1)\delta t\big) - \log s\big(j\delta t\big)}{\delta t}$$

Standard terminology: r_j is the **return** of the forward over the relative time interval. Note that to calculate the price change over a longer interval you just add the exponents:

$$s(k\delta t) = e^{\left(r_j\delta t + r_{j+1}\delta t + \cdots + r_{k-1}\delta t\right)} s(j\delta t), \text{ for } j < k.$$

Since the price is random so is each r^j. The lognormal model of price dynamics specifies their statistics:

- The random variables $r_j\delta t$ are independent, identically distributed, Gaussian random variables with mean $\mu\delta t$ and variance $\sigma^2\delta t$, for some constants μ and σ.

The constant μ is called the **expected return** (though actually, the expected return over a time interval of length δt is $\mu\delta t$). The constant σ is called the "volatility of return," or more briefly just **volatility**. These constants are assumed to be the same regardless of the length of the interval δt. Thus we really mean the following slightly stronger statement:

- For any time interval (t_1, t_2), $\log s(t_2) - \log s(t_1)$ is a Gaussian random variable with mean $\mu (t_2 - t_1)$ and variance $\sigma^2 (t_2 - t_1)$.

- The Gaussian random variables associated with disjoint time intervals are independent.

In particular (for those who know what this means) $\log s(t)$ executes a Brownian motion with drift. Strictly speaking σ has units of $1/\sqrt{time}$, however it is common to call σ the "volatility per year".

Why should we believe this hypothesis about prices? Perhaps it would be more credible to suppose that the daily (or hourly or minute-by-minute) return is determined by a random event (arrival of news, perhaps) which we can model by flipping a coin. The lognormal model is the limit of such dynamics, as the time-frequency of the coin-flips tends to zero. We'll discuss this in detail presently.

The lognormal hypothesis will lead us to a formula for the present value of a derivative security -- but it's important to remember that the formula is no better than the stock price model it's based on. The formula doesn't agree perfectly with what one finds in the marketplace; the main reason is probably that the lognormal model isn't a perfect model of real stock prices. Much work has been done on improving it -- for example by letting the volatility itself be random rather than constant in time.

In fact, very few market practitioners believe the lognormal hypothesis. They use it as a convenience, mainly because it is simple and it is better than the competing hypothesis of equal simplicity, the normal hypothesis that $F_{t_2} - F_{t_1}$ is a Gaussian random variable. It is superior because the normal hypothesis allows prices to become negative while the lognormal hypothesis does not allow prices to become negative. And most financial variables – stock prices, foreign exchange rates, commodity prices, interest rates – and the forward prices based on them must remain positive (what would a negative price mean?) In section 8 we'll learn more about how market practioners adjust their option price results for the fact that they don't actually believe the lognormal hypothesis.

∎∎∎

Lognormal dynamics and the limit of multi-period binomial trees. We claim that lognormal dynamics can be approximated by dividing time into many intervals, and flipping a coin to determine the return for each interval.

The coin can be fair or biased; to keep things as simple as possible let's concentrate on the fair case first. To simulate a lognormal process with expected return μ and volatility σ the return should be

$$\mu \delta t + \sigma \sqrt{\delta t} \quad \text{if heads (probability 1/2}$$
$$\mu \delta t - \sigma \sqrt{\delta t} \quad \text{if tails (probability 1/2)}$$

In other words, given δt we wish to consider the recombinant binomial tree with

$$F_{up} = F_{now} e^{\mu \delta t + \sigma \sqrt{\delta t}}, \quad F_{down} = F_{now} e^{\mu \delta t - \sigma \sqrt{\delta t}}$$

and with each branch assigned (subjective) probability 1/2.

Consider any time t. What is the probability distribution of prices at time t? Let's assume for simplicity that t is a multiple of δt, specifically, $t = n \delta t$. If in arriving at this time you got heads j times and tails n - j times, then the price is

$$F_0 \exp\left[n\mu \delta t + j\sigma \sqrt{\delta t} - (n - j)\sigma \sqrt{\delta t} \right] = F_0 \left[\exp \mu t + (2j - n)\sigma \sqrt{\delta t} \right]$$

We should be able to understand the probability distribution (asymptotically as $\delta t \to 0$), since we surely understand the results of flipping a coin many times. Briefly: if you make a histogram of the proportion of heads, it will resemble (as n →∞) a Gaussian distribution centered at 1/2. We'll get the variance straight in a minute. (What we're really using here is the central limit theorem.)

To proceed more quantitatively, it's helpful to use the notation of probability. Recognizing that j is a random variable, let's change notation to make it look like one by calling it X_n:

X_n = number of times you get heads in n flips of a fair coin.

Since X_n is the sum of n independent random variables (one for each coin-flip), each taking values 0 and 1 with probability 1/2, one easily sees that

Expected value of $X_n = n / 2$, Variance of $X_n = n / 4$.

It's easy to see from this that

$$\frac{2X_n - n}{\sqrt{n}}$$

tends to a Gaussian with mean value 0 and variance 1. Since $\sqrt{\delta t} = \sqrt{t}/\sqrt{n}$ our formula for the final price can be expressed as

$$F_t = F_0 \exp\left[\mu t + \sigma\sqrt{t}\,\frac{2X_n - n}{\sqrt{n}}\right]$$

Thus asymptotically, as $\delta t \to 0$ and $n \to \infty$ with $t = n\,\delta t$ held fixed,

$$F_t = F_0 \exp\left[\mu t + \sigma\sqrt{t}Z\right]$$

where Z is a random variable with mean 0 and variance 1. In particular, $\log F_t$ $\log F_0$ is a Gaussian random variable with mean μt and variance $\sigma^2 t$, as expected.

Our assertion of lognormal dynamics said a little more: that $\log F_{t_2} - \log F_{t_1}$ was Gaussian with mean $\mu(t_2\ t_1)$ and variance $\sigma^2(t_2\ t_1)$ for all $t_1 < t_2$. The justification is the same as what we did above -- it wasn't really important that we started at 0.

Notice that our calculation used only the mean and variance of X_n, since it was based on the Central Limit Theorem. Our particular way of choosing the tree -- with $F_{up} = F_{now}e^{\mu\delta t + \sigma\sqrt{\delta t}}$, $F_{down} = F_{now}e^{\mu\delta t - \sigma\sqrt{\delta t}}$, and with each choice having probability 1/2, was not the only one possible. A more general approach would take $F_{up} = F_{now}u$ with probability p, $F_{down} = F_{now}d$ with probability 1-p, and choose the three parameters u, d, p to satisfy two constraints associated with the mean and variance. Evidently one degree of freedom remains. Thus once p is fixed the other parameters are determined.

****

Implication for pricing options. We attached subjective probabilities (always equal to 1/2) to our binomial tree because we wanted to recognize lognormal dynamics as the limit of a coin-flipping process. Now let us consider one of those binomial trees -- for some specific δt near 0 – and use it to price options.

The structure of the tree remains relevant (particularly the factors u and d determining F_{up} = uF_{now} and $F_{down} = d\,F_{now}$). The subjective probabilities (1/2 for every branch) are irrelevant because our pricing is based on arbitrage. But we know a formula for the price of the option with payoff $V(F_T)$ at time maturity T:

$$V_0 = e^{-rT} E_{RN}\left[V\left(F_T\right)\right]$$

where E_{RN} denotes the expected value relative to the risk-neutral probability. And using the risk-neutral probability instead of the subjective probability just means our coin is no longer fair. Instead it is biased, with probability of heads (stock goes up) q and probability of tails (stock goes down) 1-q, where

$$q = \frac{1-d}{u-d} = \frac{1 - e^{\mu\delta t - \sigma\sqrt{\delta t}}}{e^{\mu\delta t + \sigma\sqrt{\delta t}} + e^{\mu\delta t - \sigma\sqrt{\delta t}}}$$

One verifies (using the Taylor expansion of e^x near x = 0) that this is close to 1/2 when δt is small, and in fact

$$q = \frac{1}{2}\left(1 - \sqrt{\delta t}\,\frac{\mu + \frac{1}{2}\sigma^2}{\sigma}\right) + \text{terms of order } \delta t$$

and q(1-q) = ¼ + terms of order δt

For details of this derivation, see the appendix to this section..

Our task is now clear. All we have to do is find the distribution of final values F_T when one uses the q-biased coin, then take the expected value of $V(F_T)$ with respect to this distribution. We can use a lot of what we did above: writing X_n for the number of heads as before, we still have

$$F_t = F_0 \exp\left[\mu t + \sigma\sqrt{t}\,\frac{2X_n - n}{\sqrt{n}}\right]$$

But now X_n is the sum of n independent random variables with mean q and variance $q(1 - q)$, so X_n has mean nq and variance $nq(1 - q)$. So

$$\text{Mean of } \frac{2X_n - n}{\sqrt{n}} = (2q - 1)\sqrt{n}$$

$$\approx -\sqrt{t}\left(\frac{\mu + \frac{1}{2}\sigma^2}{\sigma}\right)$$

and

variance of $\dfrac{2X_n - n}{\sqrt{n}} \approx 1$

The central limit theorem tells us the limiting distribution is Gaussian, and the preceding calculation tells us its mean and variance. In summary: as $\delta t \to 0$, when using the biased coin associated with the risk-neutral probabilities,

$$F_t = F_0 \exp\left[\mu t + \sigma \sqrt{t} Z'\right]$$

where Z' is a Gaussian random variable with mean $\sqrt{t}\left(\dfrac{-\mu - \frac{1}{2}\sigma^2}{\sigma}\right)$ and variance 1.

Equivalently, writing $Z' = Z + \sqrt{t}\left(\dfrac{-\mu - \frac{1}{2}\sigma^2}{\sigma}\right)$,

$$F_t = F_0 \exp\left[-\frac{1}{2}\sigma^2 t + \sigma \sqrt{t} Z\right]$$

where Z is Gaussian with mean 0 and variance 1. Notice that the statistical distribution of F_t depends on σ but not on μ (we'll return to this point soon).

The value of the option is the e^{-rT} times the expected value of the payoff relative to this probability distribution. Using the distribution function of the Gaussian to evaluate the expected value, we get:

$$V_0 = e^{-rT} E_{RN}[V(F_T)] = e^{-rT} E\left[V\left(F_0 e^X\right)\right]$$

where X is a Gaussian random variable with mean $-\dfrac{1}{2}\sigma^2 T$ and variance $\sigma^2 T$, or equivalently

$$V_0 = e^{-rT} \int_{-\infty}^{\infty} V(F_0 e^x) \frac{1}{\sigma\sqrt{2\pi T}} \exp\left[\frac{-\left(x - [-\sigma^2/2]T\right)^2}{2\sigma^2 T}\right] dx$$

This (when specialized to puts and calls) is the famous Black-Scholes relation.

We'll talk later about evaluating the integral. For now let's be satisfied with working backward through the binomial tree obtained with a specific (small) value of δt. Reviewing what we found above: given a lognormal stock process with return μ and volatility σ, and given a choice of δt, the tree should be constructed so that $F_{up} = uF_{now}$, $F_{down} = dF_{now}$ with

$$u = e^{\mu\delta t + \sigma\sqrt{\delta t}}, \; d = e^{\mu\delta t - \sigma\sqrt{\delta t}}$$

These determine the risk-neutral probability q by the formula given above. Working backward through the tree is equivalent to finding the discounted expected value of $V(F_T)$ relative to the risk-neutral probability.

Let us return to the observation, made above, that the statistics of $F(t)$ relative to the risk-neutral probability depend on σ (volatility of the stock) but not on μ. It follows that for pricing derivative securities the value of μ isn't really needed. More precisely: in the limit $\delta t \to 0$, the lognormal models with different μ's but the same σ all assign the same values to options. So we may choose μ any way we please — there's no reason to require that it match the actual expected return of the stock under consideration. The two most common choices are

1. choose μ to be the expected return of the forward nevertheless; or
2. choose μ so that $\mu + \dfrac{1}{2}\sigma^2 = 0$; i.e. $\mu = -\dfrac{1}{2}\sigma^2$

The latter choice has the advantage that it puts q even closer to 1/2. This is the selection favored by Jarrow-Turnbull and many other authors.

It may seem strange that the value of an option doesn't depend on μ. Heuristic argument why this should be so: we are using arbitrage considerations, so it doesn't matter whether the price tends to go up or down, which is (mainly) what μ tells us.

Here's a more limited but less heuristic argument why the option pricing formula should not depend on μ. We start from the observation that in the special case $V = F$, i.e. if the payoff is just the value-at-maturity of the forward, then the value of the option at time 0 must be F_0. We discussed this at length at the beginning of this section. Of course it should be valid also in the continuous-time limit. (The payoff $V_T = F_T$ is replicated by a very simple trading strategy - namely hold one unit of forward and never trade - whether time is continuous or discrete.) Now consider the analysis we just completed, passing to the continuum limit via binomial trees. It tells us that when $V = F$ the value of the option is

$$e^{-rT} E\left[F_0 e^X\right]$$

where X is Gaussian with mean $-\dfrac{1}{2}\sigma^2 T$ and variance $\sigma^2 T$.

The two calculations are consistent only if for such X

$$e^{-rT} E\!\left[e^X\right] = 1$$

Are the two calculations consistent? The answer is *yes*. Moreover, *if* you accept the existence of a pricing formula $V_0 = e^{-rT} E\!\left[V\!\left(F_0 e^X\right)\right]$, with X a Gaussian random variable with variance $\sigma^2 T$, then this consistency test *forces* the mean of X to be $-\dfrac{1}{2}\sigma^2 T$.

It remains to justify our assertion of consistency. This is a special case of a theorem we will prove in the next section.

<u>Appendix</u>

Details of the Taylor expansion of q on page 5

$$\frac{1-\exp(\mu\delta t - \sigma\sqrt{\delta t})}{\exp(\mu\delta t + \sigma\sqrt{\delta t}) - \exp(\mu\delta t - \sigma\sqrt{\delta t})}$$

$$=\frac{1-\left(1+\mu\delta t - \sigma\sqrt{\delta t} + \frac{\sigma^2\delta t}{2} + o(\delta t^{1\frac{1}{2}})\right)}{\left(1+\mu\delta t + \sigma\sqrt{\delta t} + \frac{\sigma\delta t^2}{2} + o(\delta t^{1\frac{1}{2}})\right) - \left(1+\mu\delta t - \sigma\sqrt{\delta t} + \frac{\sigma^2\delta t}{2} + o(\delta t^{1\frac{1}{2}})\right)}$$

$$\text{[by Taylor expansion of } e^x = 1 + x + (x^2/2) + o(x^3)]$$

$$=\frac{\left(-\mu - \frac{\sigma^2}{2}\right)\delta t + \sigma\sqrt{\delta t} + o\left(\delta t^{1\frac{1}{2}}\right)}{2\sigma\sqrt{\delta t} + o\left(\delta t^{1\frac{1}{2}}\right)} \quad \text{[collecting terms]}$$

$$=\frac{\left(-\mu - \frac{\sigma^2}{2}\right)\delta t + \sigma\sqrt{\delta t} + o\left(\delta t^{1\frac{1}{2}}\right)}{2\sigma\sqrt{\delta t}(1 + o(\delta t))}$$

$$=\frac{\left(-\mu - \frac{\sigma^2}{2}\right)\delta t + \sigma\sqrt{\delta t} + o\left(\delta t^{1\frac{1}{2}}\right)}{2\sigma\sqrt{\delta t}}(1 + o(\delta t)) \quad \text{[since } \frac{a}{b(1+o(\delta t))} = \frac{a(1-o(\delta t))}{b} + o(\delta t)]$$

$$=\left(\frac{\left(-\mu - \frac{\sigma^2}{2}\right)\sqrt{\delta t}}{2\sigma} + \frac{1}{2}\right)(1 + o(\delta t))$$

$$=\frac{1}{2} + \frac{\left(-\mu - \frac{\sigma^2}{2}\right)\sqrt{\delta t}}{2\sigma} + o(\delta t)$$

$$= \frac{1}{2}\left(1 - \sqrt{\delta t}\,\frac{\left(\mu + \frac{1}{2}\sigma^2\right)}{\sigma}\right) + o(\delta t)$$

Proof that the variance of the "q-biased coin," on page 6, approaches ¼ in the limit

$$q = \frac{1}{2}\left(1 - \sqrt{\delta t}\,\frac{\mu + \frac{1}{2}\sigma^2}{\sigma}\right)$$

$$1\text{-}q = \frac{1}{2}\left(1 + \sqrt{\delta t}\,\frac{\mu + \frac{1}{2}\sigma^2}{\sigma}\right)$$

$$q(1\text{-}q) = \frac{1}{4} - \frac{1}{2}\sqrt{\delta t}\,\frac{\mu + \frac{1}{2}\sigma^2}{\sigma} + \frac{1}{2}\sqrt{\delta t}\,\frac{\mu + \frac{1}{2}\sigma^2}{\sigma} + o(\delta t)$$

$$= \frac{1}{4} + o(\delta t)$$