CS522 - Option Pricing: Building the Lattice (2)

Last time we have shown how the appropriate choices for U and D lead to the emergence (in the limit) of the normal distribution of stock price returns, and of the log-normal distribution of stock prices. The parameters μ and σ^2 represent the expected return of the stock price and the variance of the stock price return per unit of time.

We continue our examination of the multiperiod binomial model. We assume that the initial state occurs at time 0, the final states occur at time T, and that we divide the interval $[0, T]$ into n equal intervals of length $\Delta = \frac{T}{n}$.

In our previous discussion we have selected the probability p to be equal to $\frac{1}{2}$. We know, of course, that the true probabilities are not equal to $\frac{1}{2}$; more we know that the probabilities that count are the equivalent martingale probabilities q.

Let us remember the definition of q:[1]

$$q = \frac{Se^{r\Delta} - S_d}{S_u - S_d} = \frac{e^{r\Delta} - \frac{S_d}{S}}{\frac{S_u}{S} - \frac{S_d}{S}} = \frac{e^{r\Delta} - D}{U - D}$$

Last time we have chosen particular forms for U and D, as given below:

$$\begin{aligned} U &= \exp(\mu\Delta + \sigma\sqrt{\Delta}) \\ D &= \exp(\mu\Delta - \sigma\sqrt{\Delta}) \end{aligned}$$

From these definitions, we obtain the following expression for q:

$$\begin{aligned} q &= \frac{e^{r\Delta} - D}{U - D} \\ &= \frac{\exp((r-\mu)\Delta + \sigma\sqrt{\Delta}) - 1}{\exp(2\sigma\sqrt{\Delta}) - 1} \end{aligned}$$

We can develop a more intuitive understanding for the properties of q if we determine a series expansion of q as a sum of powers of Δ. We could do this "by hand," of course, but since we are in a course that also covers tools, we will now rely on Mathematica to help us out:

```
In[1]:=g[Delta_]:=(Exp[(r-mu)*Delta+sigma*Sqrt[Delta]]-1)
                  /(Exp[2*sigma*Sqrt[Delta]]-1)
In[2]:=Series[g[d],{d, 0, 1}]//OutputForm
```

We show (one possible form of) the trivial Mathematica program that gives us the series expansion of q. The output above can be rewritten as follows:

[1] Remember that when we introduced the definition of q, we have used a one-period model with time starting at 0, and ending at time t. Now we are in the context of the multiperiod model, when the length of the time interval associated with the individual steps is Δ.

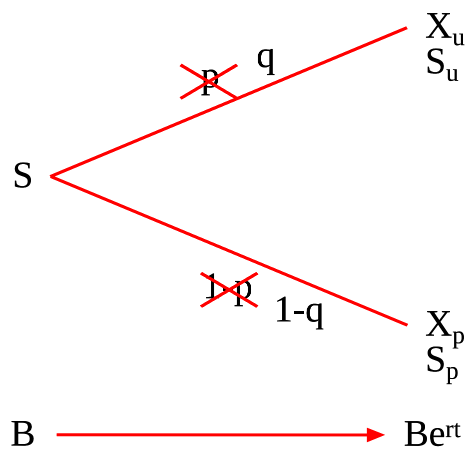

Figure 1: One-period binomial model with equivalent [martingale] probabilities.

$$\frac{1}{2} + \left(-\frac{\text{sigma}}{2} + \frac{2\,(-\text{mu} + r) + \text{sigma}^2}{4\,\text{sigma}}\right)\sqrt{d} + \left(\frac{\text{sigma}^2}{6} + \frac{1}{4}\left(-2\,(-\text{mu} + r) - \text{sigma}^2\right) + \frac{2\,(-\text{mu} + r)\,\text{sigma} + \frac{1}{2}\,\text{sigma}\,\left(2\,(-\text{mu} + r) + \text{sigma}^2\right)}{6\,\text{sigma}}\right) d + \left(\frac{1}{12}\,\text{sigma}\,\left(2\,(-\text{mu} + r) + \text{sigma}^2\right) + \frac{1}{6}\left(-2\,(-\text{mu} + r)\,\text{sigma} - \frac{1}{2}\,\text{sigma}\,\left(2\,(-\text{mu} + r) + \text{sigma}^2\right)\right) + \frac{1}{8\,\text{sigma}}\left((-\text{mu} + r)\,\left(2\,(-\text{mu} + r) + \text{sigma}^2\right) + \frac{1}{3}\,\text{sigma}\,\left(2\,(-\text{mu} + r)\,\text{sigma} + \frac{1}{2}\,\text{sigma}\,\left(2\,(-\text{mu} + r) + \text{sigma}^2\right)\right)\right)\right) d^{3/2} + O[d]^2$$

Figure 2: Screen capture of Mathematica output showing the series for q.

$$
\begin{aligned}
q &= \frac{1}{2} \\
&+ \left(\frac{-\sigma}{2} + \frac{2\,(-\mu + r) + \sigma^2}{4\,\sigma} \right) \sqrt{\Delta} \\
&+ \left(\frac{\sigma^2}{6} + \frac{-2\,(-\mu + r) - \sigma^2}{4} + \frac{2\,(-\mu + r)\,\sigma + \frac{\sigma\left(2\,(-\mu+r)+\sigma^2\right)}{2}}{6\,\sigma} \right) \Delta \\
&+ O(\Delta^{\frac{3}{2}})
\end{aligned}
$$

Retaining just the first two terms, we get the following approximation for q:

$$
q \approx \frac{1}{2} \left[1 + \sqrt{\Delta} \left(\frac{r - \mu}{\sigma} - \frac{1}{2}\sigma \right) \right]
$$

We note that when $\Delta \to 0$ (or, equivalently, when $n \to \infty$), q becomes equal to $\frac{1}{2}$.

Consider again the definitions of U and D given above. Let us introduce a binomial variable Y_n, whose value is equal to the number of "up" states that a certain stock price evolution "encounters" from time 0 to time T. The probability of an "up" state in the next interval Δ is equal to q. We then get

$$
\begin{aligned}
r_T &= \mu T + \sigma(2Y_n - n)\sqrt{\frac{T}{n}} \\
S(T) &= S(0) \exp \left[\mu T + \sigma(2Y_n - n)\sqrt{\frac{T}{n}} \right]
\end{aligned}
$$

Important note: Except for replacing variable X_n with Y_n, the expressions above are identical to those that we have derived in the preceding lecture. There is a crucial difference, however: while both X_n and Y_n are binomial, the probability of the "up" state is $\frac{1}{2}$ in the first case, and q in the second.

We can easily compute the expectation and the variance of Y_n:

$$
\begin{aligned}
E[Y_n] &= nq \\
Var[Y_n] &= nq(1 - q)
\end{aligned}
$$

We rewrite the expression for the return over $[0, T]$ to emphasize the return and variance of Y_n:

$$
\begin{aligned}
r_T &= \mu T + \sigma\sqrt{T}\frac{Y_n - \frac{n}{2}}{\sqrt{\frac{n}{4}}} \\
&= \mu T + 2\sigma\sqrt{T}\frac{(Y_n - nq) + nq - \frac{n}{2}}{\sqrt{n}} \\
&= \mu T + \sigma\sqrt{T}\left[\frac{Y_n - nq}{\sqrt{nq(1-q)}}2\sqrt{q(1-q)} + 2\left(q - \frac{1}{2}\right)\sqrt{n}\right] \\
&= \mu T + \sigma\sqrt{T}\left[\frac{Y_n - nq}{\sqrt{nq(1-q)}}2\sqrt{q(1-q)} + 2\left(q - \frac{1}{2}\right)\sqrt{n}\right]
\end{aligned}
$$

It is clear that $\lim_{n\to\infty}\sqrt{q(1-q)} = \frac{1}{2}$. Given that

$$
2\left(q - \frac{1}{2}\right)\sqrt{n} = \sqrt{T}\left(\frac{r-\mu}{\sigma} - \frac{1}{2}\sigma\right) + O(\sqrt{\Delta})
$$

we can now further rewrite r_T:

$$
\begin{aligned}
r_T &= \mu T + \sigma\sqrt{T}\left[\frac{Y_n - nq}{\sqrt{nq(1-q)}}2\sqrt{q(1-q)} + \sqrt{T}\left(\frac{r-\mu}{\sigma} - \frac{1}{2}\sigma\right) + O(\sqrt{\Delta})\right] \\
&= \mu T + \left(r - \mu - \frac{1}{2}\sigma^2\right)T + \sigma\sqrt{T}\frac{Y_n - nq}{\sqrt{nq(1-q)}}2\sqrt{q(1-q)} + O(\sqrt{\Delta}) \\
&= \left(r - \frac{1}{2}\sigma^2\right)T + \sigma\sqrt{T}\frac{Y_n - nq}{\sqrt{nq(1-q)}}2\sqrt{q(1-q)} + O(\sqrt{\Delta})
\end{aligned}
$$

By the Central Limit Theorem, the quantity $\frac{Y_n - nq}{\sqrt{nq(1-q)}}$ converges to $N(0,1)$. In the limit, we thus get that stock price returns on the interval $[0,T]$ are normally distributed, and that

$$
\begin{aligned}
E[r_T] &= \left(r - \frac{1}{2}\sigma^2\right)T \\
Var[r_T] &= \sigma^2 T
\end{aligned}
$$

Thus in the limit, when $n \to \infty$ (i.e. $\Delta \to 0$) the evolution of the stock price is still normal, but with a mean equal to $\left(r - \frac{1}{2}\sigma^2\right)T$, and the same variance as before. The fact that the variance did not change is a remarkable fact with profound implications.

Using the notation $N(0,1)$ to denote a normally distributed random variable with mean 0 and variance 1, in the limit $n \to \infty$ we can now write

$$
\begin{aligned}
r_T &= \left(r - \frac{1}{2}\sigma^2\right)T + \sigma\sqrt{T}N(0,1) \\
S(T) &= S(0)\exp\left[\left(r - \frac{1}{2}\sigma^2\right)T + \sigma\sqrt{T}N(0,1)\right]
\end{aligned}
$$

At this point, it is useful to recall that $\lim_{n\to\infty} q = \frac{1}{2}$.

Remember the definitions of U and D?

$$\begin{aligned} U &= \exp(\mu\Delta + \sigma\sqrt{\Delta}) \\ D &= \exp(\mu\Delta - \sigma\sqrt{\Delta}) \end{aligned}$$

Here μ is the expected return per unit of time of the stock price. When using the equivalent probabilities q, we have that $\mu = r - \frac{1}{2}\sigma^2$. We thus must have factors U and D defined as follows:

$$\begin{aligned} U &= \exp\left[\left(r - \frac{1}{2}\sigma^2\right)\Delta + \sigma\sqrt{\Delta}\right] \\ D &= \exp\left[\left(r - \frac{1}{2}\sigma^2\right)\Delta - \sigma\sqrt{\Delta}\right] \end{aligned}$$

We then get the following value for q:

$$q = \frac{\exp\left(\frac{1}{2}\sigma^2\Delta + \sigma\sqrt{\Delta}\right) - 1}{\exp(2\sigma\sqrt{\Delta}) - 1}$$

0.1 Estimating Model Parameters

Assume that the length of the interval of interest, T, is given. Further, assume that we have already chosen the number of equal-length subintervals into which we want to divide interval $[0, T]$ (i.e. we already chose n and Δ). In order to simulate the evolution of the stock price all we need to do is to determine the value of σ, and that of r.

Rate r might be directly observable, or - more frequently - must be inferred from the prices of various traded instruments. For our purposes, we will assume that r is equal to the continuously compoundend yield of the Treasury whose leftover maturity is closest to T. Of course, this approach has its drawbacks, as the price of the respective Treasury might be distorted significantly by some of the effects we have discussed previously, most likely by liquidity effects due to the possible off-the-run status of the respective instrument. Ideally, one would like to develop a more complex theory that incorporates the distorting effects before infering the value of r from the market price and the other relevant characteristics of the respective Treasury. On the positive side, however, it is known that r has a relatively mild effect on the value of options, so its very precise determination is not indispensable.

A more interesting problem is posed by parameter σ, the volatility of the underlying stock. In principle, one could compute this volatility from the data series of past stock prices (i.e. historical stock prices). This issues has been studied thoroughly, and it has been concluded that historical volatilities are not very useful in valuing options on stock.

Rather than trying to determine the value of σ from historical data, we can assume the price of the traded options to be correct, and we can infer what is the implied volatility

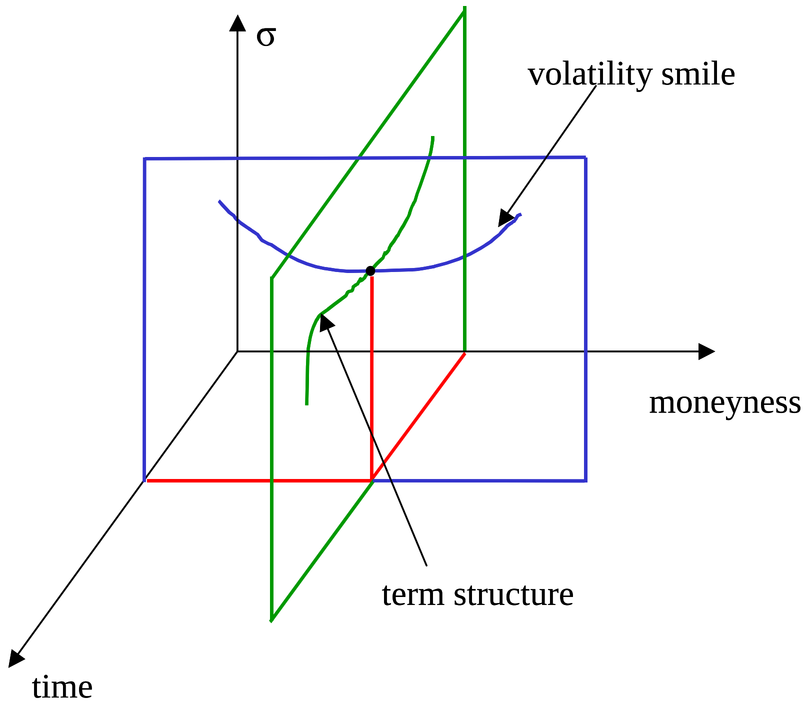

Figure 3: The volatility surface captures the dependence of the implied volatility on the time to expiration and the strike price (moneyness).

that corresponds to the respective options. For a given stock, the implied volatility should depend neither on the strike price of the underlying options, nor on the expiration date of the options; it should be constant. Such constancy of the implied volatility is not observed in practice.

We can represent the dependence of the implied volatility on the strike price and expiration date using a graph like the one in figure 3. As this figure shows, the strike price is often represented in relative terms versus the current stock price. The moneyness of the option is defined as the ratio $\frac{K}{S}$, where K is the strike price of the option and S is the current price of the underlying stock[2] is S.

The intersection between the volatility surface and a plane perpendicular to the moneyness axis is a curve that shows the term structure of the implied volatility (i.e. the volatility's dependence on the expiration date, for a given moneyness value). The intersection of the volatility surface and a plane perpendicular to the time axis shows the dependence of the volatility on the moneyness of the respective option. These curves have characteristic shapes that bear individualized names; two of these are "smile" and "skew" (see figure 4). Traders often recognize, and trade based on more sophisticated patterns.

Due to the discrete nature of the strike prices and the expiration date the volatility

[2] Or commodity, as a matter of fact.

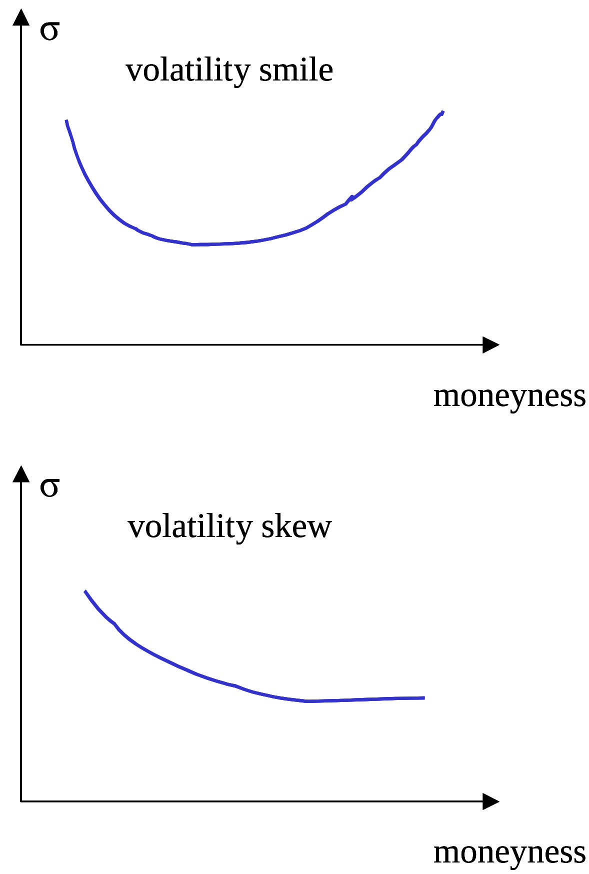

Figure 4: Smiles and skews are two typical patterns exhibited by the implied volatility of options having the same expiration date, by various strike prices.

surface can only be sampled in a relatively small number of discrete points. If we want to actually build a representation of the volatility surface, then we need to interpolate between the values we have sampled empirically. Spline approximations to the volatility surface are particularly easy to construct.

Given a volatility surface, one can use it to price options whose values have not been used in the determination of the surface. Assuming that the volatility surface does not change significantly from one trading day to the next one, one can also use the previous day's volatility surface to price the current day's options. If large enough discrepancies are noted,[3] one can trade based on the relationship between the actual price and the computed theoretical price.

[3]Small discrepancies can not be exploited due to transaction costs. Also, since models are simplified representations of the underlying reality, the model's predictions are likely to be slightly off w.r.t. the true underlying value of the option. Price discrepancies are due to model misspecification or imprecise parameter estimation.