TEACHING NOTE 96-04:

MODELING ASSET PRICES AS STOCHASTIC PROCESSES I

Version date: September 1, 2000 C:\ClassMaterial\Teaching Notes\TN96-04.WPD

The prices of assets evolve in a random manner. This means that stock prices, interest rates, foreign exchange rates, and commodity prices are largely unpredictable. Unpredictable does not mean hopelessly irrelevant. Indeed it is important that we understand the probability process driving prices because this helps us develop correct valuation models. Estimates of expected returns and volatilities and their effects on asset and derivative prices are essential in financial decision making.

A *stochastic process* is a sequence of observations from a probability distribution. Rolling dice at regular time intervals is a stochastic process. In this case the distribution is stable because the possible outcomes do not change from one roll to the next. Rolling a 6-5 three times in a row, while highly unlikely, in no way changes the probability of rolling another 6-5. A changing distribution, however, would be the case if we drew a card from a deck without replacing the previously drawn cards. Real world asset prices probably come from changing distributions though it is difficult to determine when a distribution has changed. Empirical analysis of past data can be useful in that context - not to predict the future but to know when the numbers are coming out according to different bounds of probability.

In around 1827, the Scottish scientist Robert Brown observed the random behavior of pollen particles suspended in water. This phenomenon came to be known as *Brownian Motion*. About 80 years passed before Einstein developed the mathematical properties of Brownian motion. This is not to suggest that no work was being done in the interim but scientists did not always know what other work was being done. It is not surprising that it was Einstein who received most of the credit.

Let us start by assuming that a series of numbers is coming out of a standard normal (bell-shaped) probability distribution. Let this number be denoted as ,(t) where t denotes the point in time. Since the numbers are of the standard normal type, this means they on average equal zero and have a standard deviation of 1. Numbers like this have very limited properties and in this form are not very useful for modeling asset prices. Let us transform these numbers into another process. Suppose we are currently at time t. Take any number you would like and call it W_t. This is our starting point. Now move ahead to time t+1 by drawing a number from the standard normal probability distribution. Call it ε_{t+1}. A very simple transformation of the standard normal variable into the W variable would be to add ε_{t+1} to W_t to get W_{t+1}. Another simple transformation would be to multiply ε_{t+1} by a term we call dt, which is the length of time that elapses between t and t+1. If that time interval happened to be one minute, dt would be 1/(60*24*365), or in other words, the fraction of a year that elapses between t and t+1.

One reason we like to multiply ε_{t+1} by a time factor is that we would like our model to accommodate time intervals between t and t+1 of different lengths. These statistical shocks that are the source of randomness might be larger if they were spread out over a longer time period; hence, we need to scale them by a function of time. In fact, to model asset prices evolving continuously, we need the interval between t and t+1 to be as short as possible. Mathematicians say that "in the limit" (meaning almost there but not quite), dt will approach zero. Unfortunately, the model $W_{t+1} = W_t + \varepsilon_{t+1}dt$ will give us a problem when dt is nearly zero. That comes from the fact that the variance of W_{t+1} will be nearly zero. That is because dt is very small and to obtain the variance, we have to square it, which drives it even close to zero. Thus, the variable W will have no variance, which takes away its randomness. Because it does not vary, we cannot even call it a "variable" anymore.

The problem is best solved by multiplying ε_t by the square root of dt:

$$W_{t+1} = W_t + \varepsilon_t\sqrt{dt}.$$

Then when we square it to take the variance, we obtain dt.

This model has many convenient properties. Suppose we are interested in predicting a future value of W, say at time $t + \tau$. Then the expected value of $W_{t+\tau}$ is W_t. That is because the expected random change in the process is zero. If you start off at W_t and keep incrementing it by values that average to zero, you are expected to get nowhere. The variance of $W_{t+\tau}$ is $t + \tau - t = \tau$, or in other words, however, much time elapses between now, time t, and the future point, time $t + \tau$.

This is the process called Brownian Motion. Now let us take the difference between W_{t+1} and W_t and denote it as dW_t, which will be defined as

$$dW_t = \varepsilon_t \sqrt{dt}.$$

This process, the increment to the Brownian Motion, is called a *Wiener process*, named after the American mathematician, Norbert Wiener (1894-1964), who did important work in this area. In the field of financial derivatives, we are more interested in the process dW_t than in the process W_t. We shall transform dW_t into something more useful for modeling asset prices at a later point.[1]

It is perhaps important to note that the mathematics necessary to define the expected value and variance require the mathematical technique of integration. The ordinary rules of integration, however, do not automatically apply when the terms are stochastic. Fortunately, work by the Japanese mathematician K. Itô proved that the integral, defined as a *stochastic integral*, does exist though with a slightly different definition. Consequently many of the rules of ordinary integration apply in similar or slightly different forms.

One interesting property of the Wiener process is that when you square it, it becomes perfectly predictable. This seems to be a real puzzler. How can you generate random numbers, square them and find them perfectly predictable? Suppose we draw a standard normal random variable, ε_t. We multiply it by the square root of the time interval dt. We know that this transformed value is unpredictable but we know its expected value and its variance. The expected value is obviously zero. Using the rule that the variance of a constant times a random variable is the constant squared times the variance of the random variable, we see that its variance is dt.[2]

Note, however, that if we define the variable of interest as dW_t^2, we get a different result. To determine dW_t^2 we simply draw the value of ε_t, multiply it by the square root of dt and square the entire expression. This equals $\varepsilon_t^2 dt$. The variance of this expression is found by squaring dt and multiplying it by the variance of ε_t^2. By definition, however, all values of dt^k where $k > 1$ are zero. In other words, the

[1]The terms Brownian Motion, Wiener process and Itô process are often used interchangeably.

[2]Remember that $dW_t = \varepsilon_t \sqrt{dt}$. Squaring the square root term gives dt. This is then multiplied by the variance of ε_t which is 1.0.

length of the time interval is so short that squaring it makes it shorter and effectively zero. The expected value of dW_t^2 is the expected value of $\varepsilon_t^2 dt$. This will be dt times the expected value of ε_t^2. So we must evaluate $E[\varepsilon_t^2]$. First let us use the well-known result that the variance of any random variable x is defined as $E[x^2] - E[x]^2$. Since ε_t is a standard normal variable, then $Var[\varepsilon_t] = E[\varepsilon_t^2] - E[\varepsilon_t]^2 = 1$. We know that $E[\varepsilon_t] = 0$ so $E[\varepsilon_t^2] = 1$. Thus, $E[dW_t^2] = 1*dt = dt$. Since $Var[dW_t^2] = 0$, $E[dW_t^2] = dW_t^2 = dt$. In other words, any variable that has zero variance can be expressed as its expected value. So remember this important result: $dW_t^2 = dt$! We shall see it again.

Why do these things matter? They are the foundations of the most fundamental model used to price options. Let us look at how this process can be used to model stock price movements. We know that stock price fluctuations have several important characteristics. First, over the long run, stock prices go up. They are said to "drift." This represents the return from bearing risk. The Wiener process does not drift, but as we show later, it is easy to make it drift either upward or downward. Second, stock prices are random. We know that Wiener processes are random, though we cannot use the basic Wiener process for every stock because different stocks have different volatilities. We can, however, transform the basic Wiener process to give it a different volatility. Third, it should be harder to forecast stock prices further into the future than nearby. That does not mean that stock prices are very predictable but that the margin of error, which is related to the variance of the future stock price, should be greater when predicting far into the future than when predicting into the near future. The final property is that a stock price should never be allowed to become negative. Corporate shareholders have limited liability so the minimum value of their shares is zero.

The properties of a stock's return can be described by its mean and variance.[3] Let E[R] be the stock's expected return and σ^2 be the stock's variance over a period of a year.[4] Then σ is the standard deviation, which is the square root of the variance. If we have a good model it will have an expected value and standard deviation equal to these values.

---

[3]Other properties like skewness might be important but we ignore them for this model.

[4]It is common but not required to express expected returns and variances on an annual basis.

Let $R_t h$ be the stock's return, its price change divided by its base price, over the holding period h. Now set $R_t h$ equal to $E[R]h + g_t$, where g_t represents the random component of the stock's return. We do not yet know what g_t is but we assume it is influenced by the variance. Now, let us force the expected value of our model to E[R]h. This is easily done by letting $E[g_t] = 0$. Thus, we must keep this constraint in mind when we look for a suitable form for g_t.

Now let us force the model to have the correct variance. If the variance over a year is σ^2, then the variance over the holding period h is $h\sigma^2$. The variance of our model should be such that it equals $h\sigma^2$. So far our model has two terms $E[R]h + g_t$. The term E[R]h is a constant so it has no variance. Thus, we need to make sure that $Var[g_t] = h\sigma^2$.

First let us define h to be the standard time increment in a stochastic process of dt. One model that has the appropriate variance is $g_t = \sigma\varepsilon_t\sqrt{dt}$, which is simply σdW_t. This is just a transformed Wiener process with the transformation coming from multiplying it by the stock's standard deviation. Remember that we also have to have the expected value of g_t equal to zero but that requirement is upheld because we already know that $g_t = dW_t$ and $E[dW_t] = 0$.

Now our model looks like this

$$R_t = E(R)dt + \sigma\varepsilon_t\sqrt{dt}$$

Does this model have the other properties we want it to have? If S_t is the current stock price, then the stock price after the period dt is $S_{t+dt} = S_t(1 + E[R]dt + \sigma\varepsilon_t\sqrt{dt})$. What is the variance of this future stock price? Without getting into the technical details the variance will indeed be larger if the time interval is larger.[5]

Finally, the model must not permit the stock price to ever go below zero. Let us write the model in the following form

[5]A simple proof of this point is worth noting. The variance of S_{t+dt} is completely determined by ε_t, dt and σ. The only one of these three values that automatically gets larger the further out in time you move is the time indicator, dt. The larger is dt, the larger is the variance of S_{t+dt}. A formal mathematical proof is technically required but you can be assured that the proof holds up.

$$\frac{dS_t}{S_t} = E[R]dt + \sigma\varepsilon_t\sqrt{dt}.$$

Multiplying through by S_t makes everything on the right-hand side be multiplied by S_t. If S_t ever reaches zero then dS_t is stuck at zero. Zero is said to be a natural absorbing barrier, which is equivalent to bankruptcy. It should be noted that the rules governing Brownian motion do not allow the process to go through zero to a negative value. Once it hits zero, the process stops.

One final adjustment is necessary. Most of the time the annualized expected return is written as either α or μ. We shall choose the former. Thus, the model is written as

$$\frac{dS_t}{S_t} = \alpha dt + \sigma dW_t,$$

where it is understood that $dW_t = \varepsilon_t\sqrt{dt}$. A stochastic processes of this type is called an *Itô process*. It is more generally stated in the form, $dS_t = \alpha(S,t)dt + \sigma(S,t)dW_t$ where the expected value and variance are allowed to change with S and t.

To recap, the model allows us to replicate the behavior of the stock over a short holding period. We have taken the basic Brownian Motion process and converted it into a form that models stock price movements. This model has many convenient and reasonable properties. We refer to the process as *Geometric Brownian Motion*. It is "geometric" in the sense that proportional changes, which is what we mean by the percentage change, in the stock price follow this stochastic process. Without providing the mathematical details, we can say that the returns on stocks follow a lognormal distribution. This is not the normal, or bell-shaped curve. A lognormal distribution is skewed toward positive returns in contrast to the normal distribution, which is symmetric. A lognormal distribution does imply, however, that the logarithm of the returns comes from the normal or bell-shaped distribution.[6] These properties are all desirable and fairly reasonable from an empirical standpoint.

---

[6]This type of process is also sometimes called a *lognormal diffusion process*.

Of course, no model will reproduce perfectly the process in which stock returns are generated. The real world can rarely be reduced to a set of mathematical equations. But as is nearly always the case, if a set of mathematical equations can reproduce the basic manner in which a real-world phenomenon occurs, it can have many uses. One of these uses is in pricing derivatives on assets that follow the process described by the mathematical model covered here.

References

These notes are adapted from my article

Chance, D. M. "The ABCs of Geometric Brownian Motion." *Derivatives Quarterly* 1 (Winter, 1994), 41-47.

For other treatments of this model in finance see

Dothan, M. U. *Prices in Financial Markets.* New York: Oxford University Press (1990), Chs. 7,8.

Duffie, D. *Dynamic Asset Pricing Theory*, 2nd. ed. Princeton: Princeton University Press (1996), Ch. 5.

Duffie, D. *Security Markets: Stochastic Models*. Boston: Academic Press (1988), Chs. 21-23.

Haley, C. W. and L. D. Schall. "Stochastic Calculus and Derivation of the Option Pricing Model," Appendix 10A of *The Theory of Financial Decision Making*, 2nd. ed., New York: McGraw-Hill (1979).

Hull, J. C. *Options, Futures and Other Derivative Securities*, 4th ed. Upper Saddle River, NJ: Prentice-Hall (200), Ch. 10.

Ingersoll, J. E. *Theory of Financial Decision Making*. Totowa, NJ: Rowman & Littlefield (1987), Ch. 16.

Jarrow, R. and A. Rudd. *Option Pricing*. Homewood, Illinois: Irwin (1983), Ch. 7.

Malliaris, A. G. and W. A. Brock. *Stochastic Methods in Economics and Finance*. New York: North Holland Publishing Co. (1983), Ch. 2.

Merton, R. C. "On the Mathematics and Economics Assumptions of Continuous-Time Models." *Financial Economics: Essays in Honor of Paul Cootner*, ed. by W. F. Sharpe and C. M. Cootner. Englewood Cliffs, NJ: Prentice-Hall (1982).

Neftci, S. N. *An Introduction to the Mathematics of Financial Derivatives*. San Diego: Academic Press (2000), Chs. 6,8

Nielsen, L. T. *Pricing and Hedging of Derivative Securities*. Oxford: Oxford University Press (1999), Chs. 1-3.

Shimko, D. C. *Finance in Continuous Time.* Miami: Kolb Publishing (1992), Ch. 1.

Smith, C. W. "Appendix: An Introduction to Stochastic Calculus." *The Modern Theory of Corporate Finance*, ed. by M. C. Jensen and C. W. Smith. New York: McGraw-Hill (1984).

For more of the mathematical theory, some excellent references are:

Cox, D. R. and H. D. Miller. *The Theory of Stochastic Processes*. London: Chapman & Hall (1965), Chs. 1, 2, 5.

Karatzas, I. and S. E. Shreve. *Brownian Motion and Stochastic Calculus*. New York: Springer-Verlag (1988), Ch. 2.

Karlin, S. and H. M. Taylor. *A First Course in Stochastic Processes*. New York: Academic Press (1975), Ch. 7.

Karlin, S. *A Second Course in Stochastic Processes*. New York: Academic Press (1981), Ch. 15.

Resnick, S. I. *Adventures in Stochastic Processes*. Boston: Birkhauser (1992), Ch. 6.

The first and classic applications of Brownian motion in finance were

Bachelier, L. “Theory of Speculation.” English translation by A. J. Boness, *The Random Character of Stock Market Prices*, ed. P. Cootner. Cambridge, Mass: The M.I.T. Press (1964), 17-78.

Osborne, M. F. M. “Brownian Motion in the Stock Market.” *Operations Research* 7 (March-April, 1959), 145-173.

For commentary on Bachelier’s work see

Sullivan, E. J. and T. M. Weithers. “Louis Bachelier: The Father of Modern Option Pricing Theory.” *Journal of Economic Education* 22 (Spring, 1991), 165-170.

An excellent historical treatise on Brownian motion is found in

Maiocchi, R. “The Case of Brownian Motion.” *British Journal of the History of Science* 23 (1990), 257-283.