

Asset Price Dynamics

Introduction

These notes give assumptions of asset price returns that are derived from the efficient markets hypothesis. Although a hypothesis, there is widespread empirical evidence that broadly supports the hypothesis and therefore the assumptions made on the process governing asset price changes. Continuous time stochastic processes are discussed and the geometric Brownian motion model for stock price changes is derived. We first look at rates of return as if they are known for certain and then consider the realistic case that asset price returns are unknown in advance.

Keywords: *continuously compounded rate of return, stochastic process, random walk, martingale, Markov property Wiener process, geometric Brownian motion, Ito calculus.*

Reading: You should read Hull chapter 12 and perhaps the very first part of chapter 13.

Rates of Return

The Rate of Return

The rate of return is simply the end value less the initial value as a proportion of the initial value. Thus if 100 is invested and at the end value is 120 then the rate of return is $\frac{120-100}{100} = \frac{1}{5}$ or 20%. If the the initial investment is B_0 and the end value is B_T after T periods then the rate of return is

$$r(T) = \frac{B_T - B_0}{B_0}$$

or equivalently the rate of return $r(T)$ satisfies $B_T = B_0(1 + r(T))$.

It is important to know the rate of return. However to compare rates of return on different investments with different time horizons it is also important to have a measure of the rate of return per period. One method of making this comparison is to use *continuously compounded rates of return*. To explain this we first consider compound returns and then show what happens when the compounding is continuous.

Compound Rates of Return

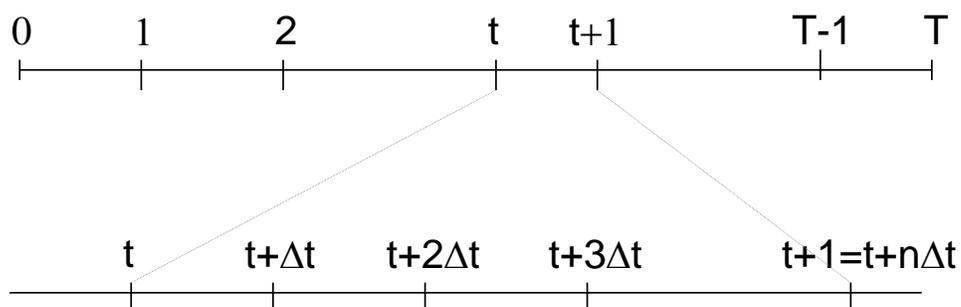
Compound interest rates are calculated by assuming that the principal (initial investment) plus interest is re-invested each period. Compounding might be done annually, semi-annually, quarterly, monthly or even daily. Assuming the re-investment is done after each period, the per-period interest rate r on the investment satisfies

$$(1 + r(T)) = (1 + r)^T.$$

Now consider dividing up each period into n sub-periods each of length Δt . This is illustrated in Figure 1. Then if the compounding is done n times per period, the compound interest rate r satisfies

$$(1 + r(T)) = \left(1 + \frac{r}{n}\right)^{nT}.$$

For example consider a time period of one-year and suppose an investment of 100 that yields 120 after two years ($T = 2$) has a rate of return $r(2) = 0.2$. If the interest rate is annualised using annual compounding ($n = 1$, $T = 2$), then $r = 0.09544$; with semi-annual compounding ($n = 2$, $T = 2$) the annualised interest rate is $r = 0.09327$; with quarterly compounding ($n = 4$, $T = 2$) the annualised interest rate is $r = 0.0922075$ etc.

Figure 1: DIVIDING A TIME INTERVAL n SUB-PERIODS

Continuous Compounding

Suppose that compounding is done n times per period and let the length of time between compounding be denoted by $\Delta t = \frac{1}{n}$. Continuous compounding occurs as $\Delta t \rightarrow 0$ or equivalently as $n \rightarrow \infty$. In this case the compounding factor r satisfies

$$\left(1 + \frac{r}{n}\right)^{nT} = (1 + r\Delta t)^{\frac{T}{\Delta t}}.$$

Let $m = \frac{1}{r\Delta t}$, then

$$(1 + r\Delta t)^{\frac{T}{\Delta t}} = \left(1 + \frac{1}{m}\right)^{mrT} = \left(\left(1 + \frac{1}{m}\right)^m\right)^{rT}.$$

compounded rates of return. This is a very convenient feature which makes using the continuously compounded rates of return especially simple. To see this let r_t denote the continuously compounded rate of return from period t to $t + 1$, that is

$$r_t = \ln \left(\frac{B_{t+1}}{B_t} \right)$$

where B_t is the value of the asset at time t . Let $r(T)$ denote the continuously compounded rate of return over the period 0 to T ,

$$r(T) = \ln \left(\frac{B_T}{B_0} \right) = \ln B_T - \ln B_0.$$

Suppose that $T = 2$ then we can write this as

$$r(2) = \ln B_2 - \ln B_0 = (\ln B_2 - \ln B_1) + (\ln B_1 - \ln B_0) = r_2 + r_1.$$

Thus the continuously compounded rate of return over two periods is simply the sum of the two period by period returns. In general for any value of T we can write

$$\begin{aligned} r(T) &= (\ln B_T - \ln B_{T-1}) + (\ln B_{T-1} - \ln B_{T-2}) + \dots + (\ln B_2 - \ln B_1) + (\ln B_1 - \ln B_0) \\ &= r_{T-1} + r_{T-2} + \dots + r_1 + r_0 = \sum_{t=0}^{T-1} r_t. \end{aligned}$$

Thus the continuously compounded rate of return over time T is simply the sum of the period by period returns. If r_t is constant over time then $r(T) = rT$.

A Differential Equation

Let r_t denote the rate of return between t and $t + 1$. Then over any sub-interval of Δt say between t and $t + \Delta t$, r_t satisfies

$$B_{t+\Delta t} = (1 + r_t \Delta t) B_t.$$

Then taking the limit as $\Delta t \rightarrow 0$ we have $B_{t+\Delta t} - B_t \rightarrow dB(t)$ where $B(t)$ is the price at time t and $\Delta t \rightarrow dt$. Hence we can write

$$dB(t) = r_t B(t) dt$$

or equivalently

$$\frac{dB(t)}{B(t)} = r_t dt.$$

This is a differential equation. If we assume that $r_t = r$ is independent of the time t , then this equation can be solved at by integration to give the asset price at time T

$$\ln B_T - \ln B_0 = \ln \left(\frac{B_T}{B_0} \right) = rT$$

or

$$B_T = B_0 e^{rT}.$$

Stochastic Processes

We have assumed so far that the rate of return was known so that we were dealing with a risk-free asset. But for most assets the rate of return is uncertain or *stochastic*. As the asset value also changes through time we say that the asset price follows a **stochastic process**. Fortunately the efficient markets hypothesis provides some strong indication of what properties this stochastic process will have.

A Coin Tossing Example

To examine the form that uncertain returns may take it is useful to think first of a very simple stochastic process. This we have already seen as the binomial model is itself a stochastic process. As an example consider the case of tossing a fair coin where one unit is won if the coin ends up Heads

and one unit is lost if the coin ends up Tails. An example of the possible payoffs for a particular sequence of Heads and Tails is illustrated in Figure 2.

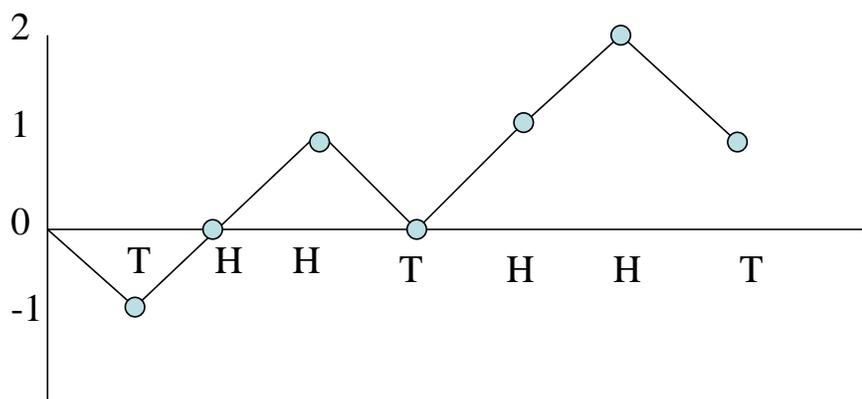


Figure 2: A COIN TOSSING STOCHASTIC PROCESS

The important properties of this example are that the distribution of returns are 1) identically distributed at each toss (there is an equal chance of a Head or a Tail); 2) independently distributed (the probability of a Head today is independent of whether there was a Head yesterday); 3) the expected return is the same each period (equal to zero); 4) the variance is constant at each period (equal to one).

There are some important implications to note about this process. First let x_t denote the winnings on the t th toss. We have $x_0 = 0$ and $E[x_1] = 0$ where $E[x_t]$ denotes the expected winnings at date t . Then we also have at any date

$$E[x_{t+1}] = x_t.$$

Any process with this property is said to be a **martingale**. Another important property is that the variance of x is increasing proportionately to the number of tosses. In particular letting σ_t^2 denote the variance of the winnings at the t th toss we have $\sigma_t^2 = t$ or in terms of the standard deviation (the square root of the variance)

$$\sigma_t = \sqrt{t}.$$

This is illustrated in Figure 3. Figure 3 shows all possible winnings through four tosses. The variance of winnings is easily calculated at each toss. For example at the second toss the expected winnings are zero so the variance is given by

$$\sigma_2^2 = \frac{1}{4}(2 - 0)^2 + \frac{1}{2}(0 - 0)^2 + \frac{1}{4}(-2 - 0)^2 = 1 + 1 = 2$$

and so the standard deviation is $\sigma_2 = \sqrt{2}$.

A Stochastic Process for Asset Prices

The efficient markets hypothesis implies that all relevant information is rapidly assimilated into asset prices. Thus asset prices will respond only to new information (news) and since news is essentially unforecastable so to are asset prices. The efficient market hypothesis also implies that it is impossible to consistently make abnormal profits by trading on publically available information and in particular the past history of asset prices. Thus only the current asset price is relevant in predicting future prices and past prices are irrelevant. This property is know as the **Markov** property for stock prices. If we add a further assumption that the variability of asset prices is roughly constant over time, then the asset price is said to follow a **random walk**. This was true of our coin tossing example above.

Let u_t denote the *random* rate of return from period t to $t + 1$. Then

$$S_{t+1} = (1 + u_t)S_t.$$

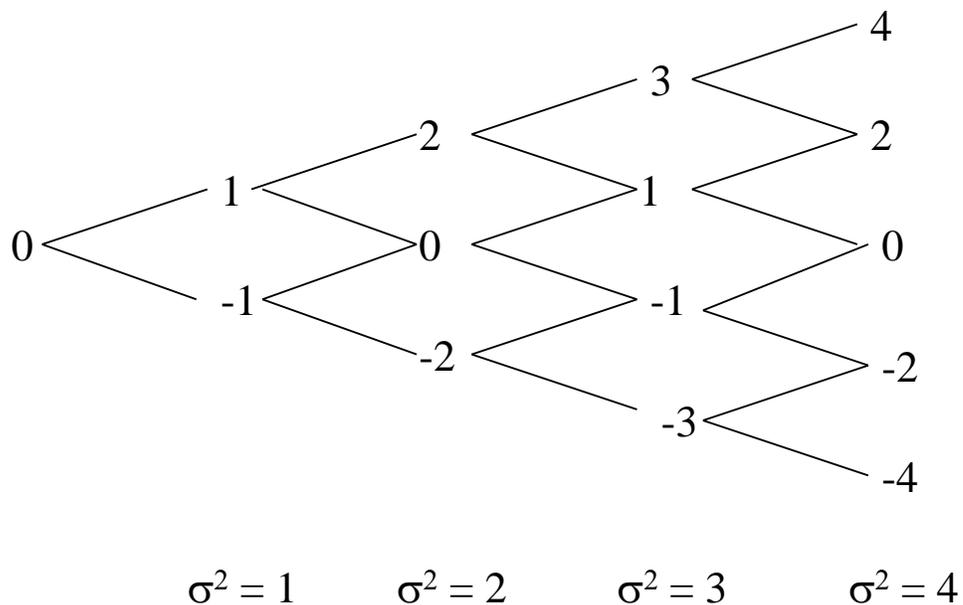


Figure 3: COIN TOSSING EXAMPLE: THE VARIANCE IS PROPORTIONAL TO TIME

The return u_t is now random because the future asset price is unknown.¹ It can be considered as a random *shock* or *disturbance*. Taking natural logarithm of both sides gives

$$\ln S_{t+1} = \ln S_t + \ln(1 + u_t).$$

¹This was the case in our binomial model where u_t takes on either the value of u in the upstate or d in the down state.

We can then see how the stochastic process for the asset price evolves. Suppose we start from a given value S_0 , then

$$\begin{aligned}\ln S_1 &= \ln S_0 + \ln(1 + u_0) \\ \ln S_2 &= \ln S_1 + \ln(1 + u_1) = \ln S_0 + \ln(1 + u_0) + \ln(1 + u_1) \\ \ln S_3 &= \ln S_2 + \ln(1 + u_2) = \ln S_0 + \ln(1 + u_0) + \ln(1 + u_1) + \ln(1 + u_2) \\ &\vdots \\ \ln S_T &= \ln S_0 + \sum_{i=0}^{T-1} \ln(1 + u_i).\end{aligned}$$

Let $\omega_t = \ln(1 + u_t)$. We can then write

$$\ln S_T = \ln S_0 + \sum_{i=0}^{T-1} \omega_i.$$

We shall assume that ω_t is a random variable which is identically and independently distributed and such that the expected value $E[\omega_t] = \nu$ and variance $\text{Var}[\omega_t] = \sigma^2$. There is a great deal of evidence to support the assumption that ω_t is independently and identically distributed and over short time horizons. It is also usually reasonable to assume that the expected value ν and variance σ^2 are independent of time for the short time horizons that we normally consider in pricing options.

We shall make a further assumption that each ω_t is normally distributed. Since the sum of randomly distributed random variables is normally distributed, and since S_0 is known the natural logarithm of the asset price will also be normally distributed. Taking expectations we can therefore show that

$$E[\ln S_T] = \ln S_0 + \nu T$$

and

$$\text{Var}[\ln S_T] = \sigma^2 T.$$

Since the logarithm of the asset price is normally distributed the asset price itself is said to be **lognormally** distributed. In practice when one looks at the empirical evidence asset prices are reasonably closely lognormally distributed.

Lognormal Random Variable

We have assumed that that $\omega_t = \ln(1 + u_t)$ is normally distributed with an expected value of ν and variance σ^2 . But $1 + u_t$ is a lognormal variable. Since $1 + u_t = e^{\omega_t}$ we might guess that the expected value of u_t is $E[u_t] = e^\nu - 1$. However this would be WRONG. The expected value of u_t is

$$E[u_t] = e^{\nu + \frac{1}{2}\sigma^2} - 1.$$

The expected value is actually higher than anticipated by half the variance. The reason why can be seen from looking at an example of the lognormal distribution which is drawn in Figure 4. The distribution is skewed and as the variance increases the lognormal distribution will spread out. It cannot spread out too much in a downward direction because the variable is always non-negative. But it can spread out upwards and this tends to increase the mean value. One can likewise show that the expected value of the asset price at time T is

$$E[S_T] = S_0 e^{(\nu + \frac{1}{2}\sigma^2)T}.$$

Letting $\mu = \nu + \frac{1}{2}\sigma^2$ we have

$$E[S_T] = s_0 e^{\mu T}$$

so that μ is the expected continuously compounded rate of return. We will explain the relationship between μ and ν a little bit further below.

Standard Normal Variable

We have seen that $\ln S_T$ is normally distributed with mean (expected value) of $\ln S_0 + \nu T$ and variance of $\sigma^2 T$. It is useful to transform this to a variable

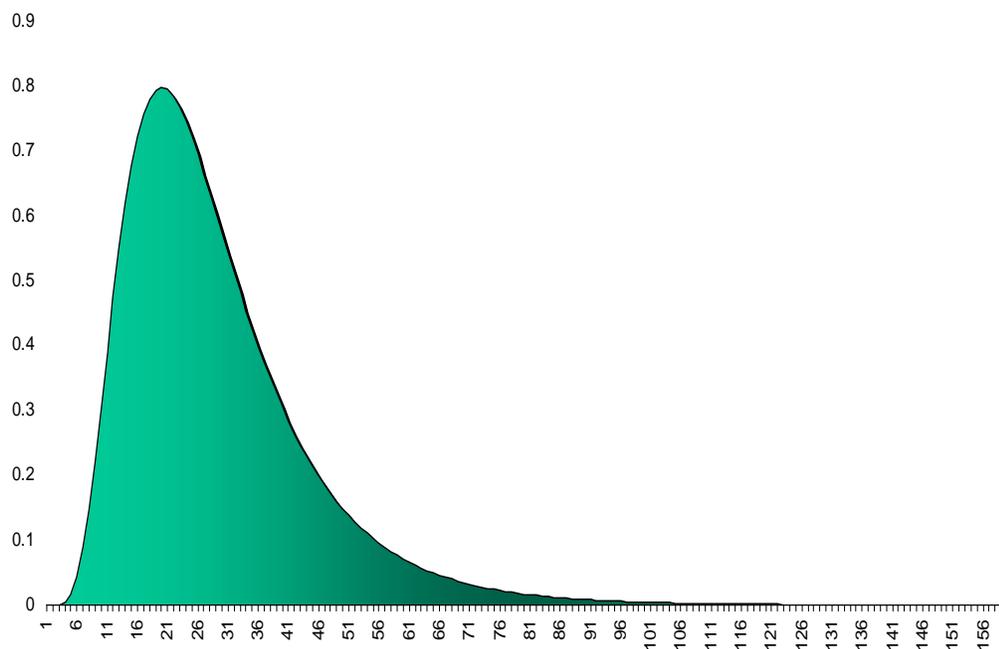


Figure 4: A LOGNORMAL DISTRIBUTION

which has a standard normal distribution with mean of zero and standard deviation of one. Such a variable is called a standard normal variable. To make this transformation, we subtract the mean and divide by the standard deviation (square root of the variance). Thus

$$\frac{\ln S_T - \ln S_0 - \nu T}{\sigma\sqrt{T}}$$

is a standard normal variable. We let $N(x)$ denote the cumulative probability that the standard normal variable is less than or equal to x . A standard normal distribution is drawn in Figure 5. It can be seen that $N(0) = 0.5$ as the normal distribution is symmetric and half the distribution is to the left of the mean value of zero. It also follows from symmetry that if $x > 0$, then $1 - N(x) = N(-x)$. We will use this property later when we look at the Black-Scholes formula.

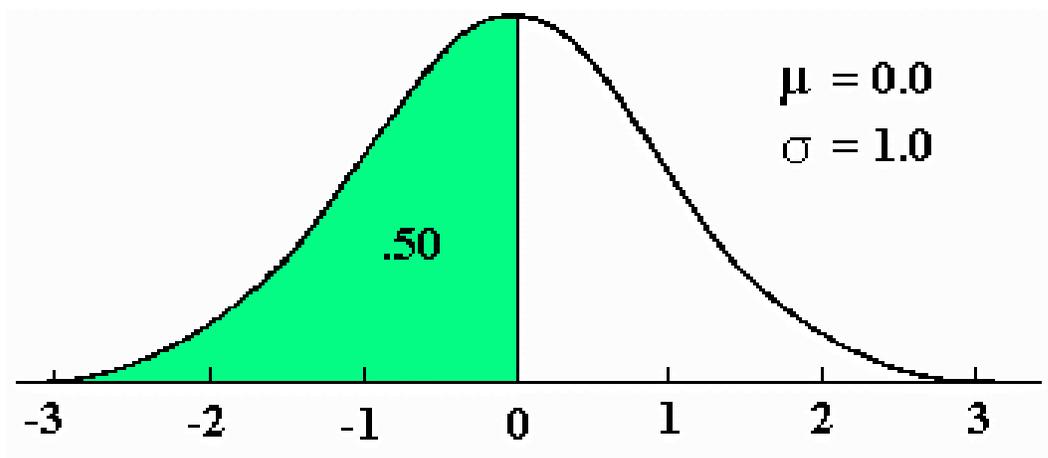


Figure 5: A STANDARD NORMAL DISTRIBUTION. $N(0) = 0.5$.

Arithmetic and geometric rates of return

We now consider μ and ν again. Suppose we have an asset worth 100 and for two successive periods it increases by 20%. Then the value at the end of the first period is 120 and the value at the end of the second period is 144.

Now suppose that instead the asset increases in the first period by 30% and in the second period by 10%. The average or **arithmetic mean** of the return is 20%. However the value of the asset is 130 at the end of first period and 143 at the end of the second period. The variability of the return has meant that the asset is worth less after two periods even though the average return is the same. We can calculate the equivalent per period return that would give the same value of 143 after two periods if there were no variance in the returns. That is the value ν that satisfies

$$143 = 100(1 + \nu)^2.$$

This value is known as the **geometric mean**. It is another measure of the average return over the two periods. Solving this equation gives the geometric

mean as $\nu = 0.195826$ or 19.58% per period² which is less than the arithmetic rate of return per period.

There is a simple relationship between the arithmetic mean return, the geometric mean return and the variance of the return. Let $\mu_1 = \mu + \sigma$ be the rate of return in the first period and let $\mu_2 = \mu - \sigma$ be the rate of return in the second period. Here the average rate of return is $\frac{1}{2}(\mu_1 + \mu_2) = \mu$ and the variance of the two rates is σ^2 . The geometric rate of return μ satisfies $(1 + \nu)^2 = (1 + \mu_1)(1 + \mu_2)$. Substituting and expanding this gives

$$1 + 2\nu + \nu^2 = (1 + \mu + \sigma)(1 + \mu - \sigma) = (1 + \mu)^2 - \sigma^2 = 1 + 2\mu + \mu^2 - \sigma^2$$

or

$$\nu = \mu - \frac{1}{2}\sigma^2 + \frac{1}{2}(\mu^2 - \nu^2).$$

Since rates of return are typically less than one, the square of the return is even smaller and hence the difference between two squared percentage terms is smaller still. Hence we have the approximation $\nu \approx \mu - \frac{1}{2}\sigma^2$ or

$$\text{geometric mean} \approx \text{arithmetic mean} - \frac{1}{2}\text{variance}.$$

This approximation will be better the smaller are the interest rates and the smaller is the variance. In the example $\mu = 0.2$ and $\sigma = 0.1$ so $\frac{1}{2}\sigma^2 = 0.005$ and $\mu - \frac{1}{2}\sigma^2 = 0.1950$ which is close to the actual geometric mean of 0.1958. Thus the difference between μ and ν is that ν is the geometric rate of return and μ is the arithmetic rate of return. It is quite usual to use the arithmetic rate and therefore to write that the expected value of the logarithm of the stock price satisfies

$$E[\ln S_T] = \ln S_0 + \left(\mu - \frac{1}{2}\sigma^2\right) T$$

and

$$\text{Var}[\ln S_T] = \sigma^2 T.$$

²The geometric mean of two numbers a and b is \sqrt{ab} . Thus strictly speaking $1 + g = 1.195826$ is the geometric mean of 1.1 and 1.3.

Continuously Compounded Rate of Return

The value ν is the continuously compounded rate of return. It is given by

$$\nu = \frac{1}{T}[\ln S_T - \ln S_0].$$

Hence taking expectations we can calculate

$$E[\nu] = \frac{1}{T}E[\ln S_T - \ln S_0] = \frac{1}{T}(E[\ln S_T] - \ln S_0) = \mu - \frac{\sigma^2}{2}.$$

Similarly the variance satisfies

$$\text{Var}[\nu] = \frac{1}{T^2}\text{Var}[\ln S_T - \ln S_0] = \frac{1}{T^2}\text{Var}[\ln S_T] = \frac{\sigma^2 T}{T^2} = \frac{\sigma^2}{T}.$$

Hence the standard deviation of ν is simply σ/\sqrt{T} .

A Wiener Process

We will now consider the stochastic process in more detail and see how to take limits as the length of the time interval goes to zero. This will produce a *continuous time stochastic process*.

Consider a variable z which takes on values at discrete points in time $t = 0, 1, \dots, T$ and suppose that z evolves according to the following rule:

$$z_{t+1} = z_t + \epsilon; \quad W_0 \text{ fixed}$$

where ϵ is a random drawing from a standardized normal distribution, that is with mean of zero and variance of one. The draws are assumed to be independently distributed. This represents a **random walk** where on average z remains unchanged each period but where the standard deviation of the realized value is one each period. At date $t = 0$, we have $E[z_T] = z_0$ and the variance $\text{Var}[z_T] = T$ as the draws are independent.

Now divide the periods into n subperiods each of length Δt . To keep the process equivalent the variance in the shock must also be reduced so that the standard deviation is $\sqrt{\Delta t}$. The resulting process is known as a **Wiener process**. The Wiener process has two important properties:

Property 1 The change in z over a small interval of time satisfies:

$$z_{t+h} = z_t + \epsilon\sqrt{\Delta t}.$$

Then as of time $t = 0$, it is still the case that $E[z_T] = z_0$ and the variance $Var[z_T] = T$. This relation may be written

$$\Delta z(t + \Delta t) = \epsilon\sqrt{\Delta t}$$

where $\Delta z(t + \Delta t) = z_{t+\Delta t} - z_t$. This has an expected value of zero and standard deviation of $\sqrt{\Delta t}$.

Property 2 The values of Δz for any two different short intervals of time are independent.

It follows from this that

$$z(T) - z(0) = \sum_i^N \epsilon_i \sqrt{\Delta t}$$

where $N = T/\Delta t$ is the number of time intervals between 0 and T . Hence we have

$$E[Z(T)] = z(0)$$

and

$$Var[z(T)] = N\Delta t = T$$

or the standard deviation of $z(T)$ is \sqrt{T} .

Now consider what happens in the limit as $\Delta t \rightarrow 0$, that is as the length of the interval becomes an infinitesimal dt . We replace $\Delta z(t + \Delta t)$ by $dz(t)$

which has a mean of zero and standard deviation of dt . This continuous time stochastic process is also known as **Brownian Motion** after its use in physics to describe the motion of particles subject to a large number of small molecular shocks.

This process is easily generalized to allow for a non-zero mean and arbitrary standard deviation. A **generalized Wiener process** for a variable x is defined in terms of $dz(t)$ as follows

$$dx = a dt + b dz$$

where a and b are constants. This formula for the change in the value of x consists of two components, a deterministic component $a dt$ and a stochastic component $b dz(t)$. The deterministic component is $dx = a dt$ or $\frac{dx}{dt} = a$ which shows that $x = x_0 + at$ so that a is simply the trend term for x . Thus the increase in the value of x over one time period is a . The stochastic component $b dz(t)$ adds noise or variability to the path for x . The amount of variability added is b times the Wiener process. Since the Wiener process has a standard deviation of one the generalized process has a standard deviation of b .

The Asset Price Process

Remember that we have

$$\ln S_{t+1} - \ln S_t = \omega_t$$

where ω_t is are independent random variables with mean ν and standard deviation of σ . The continuous time version of this equation is therefore

$$d \ln S(t) = \nu dt + \sigma dz$$

where z is a standard Wiener process. The right-hand-side of the equation is just a random variable that is evolving through time. The term ν is called the

drift parameter and the standard deviation of the continuously compounded rate of return is $\sqrt{\text{Var}[r(t)]} = \sigma\sqrt{\Delta t}$ and the term σ is referred to as the **volatility** of the asset return.

Ito Calculus

We have written the process in terms of $\ln S(t)$ rather than $S(t)$ itself. This is convenient and shows the connection to the binomial model. However it is usual to think in terms of $S(t)$ itself too. In ordinary calculus we know that

$$d \ln S(t) = \frac{dS(t)}{S(t)}.$$

Thus we might think that $dS(t)/S(t) = \nu dt + \sigma dz$. But this would be WRONG. The correct version is

$$\frac{dS(t)}{S(t)} = \left(\nu + \frac{1}{2}\sigma^2 \right) dt + \sigma dz.$$

This is a special case of **Ito's lemma**. Ito's lemma shows that for any process of the form

$$dx = a(x, t)dt + b(x, t)dz$$

then the function $G(x, t)$ follows the process

$$dG = \left(\frac{\partial G}{\partial x}a(x, t) + \frac{\partial G}{\partial t} + \frac{1}{2}\frac{\partial^2 G}{\partial x^2}b^2(x, t) \right) dt + \frac{\partial G}{\partial x}b(x, t)dz.$$

We'll see how to use Ito's lemma. We have

$$d \ln S(t) = \nu dt + \sigma dz.$$

Then let $\ln S(t) = x(t)$, so $s(T) = G(x, t) = e^x$. Then upon differentiating

$$\frac{\partial G}{\partial x} = e^x = S, \quad \frac{\partial^2 G}{\partial x^2} = e^x = S, \quad \frac{\partial G}{\partial t} = 0.$$

Hence using Ito's lemma

$$dS(t) = (\nu S(t) + 0 + \frac{1}{2}\sigma^2 S(t))dt + \sigma S(t) dz$$

or

$$dS(t) = \left(\nu + \frac{1}{2}\sigma^2\right)S(t) dt + \sigma S(t) dz$$

Since $\mu = \nu + \frac{1}{2}\sigma^2$ we can write this as

$$dS(t) = \mu S(t) dt + \sigma S(t) dz.$$

This process is known as geometric Brownian motion as it is the rate of change which is Brownian motion. Thus sometimes the above equation is written as

$$\frac{dS(t)}{S(t)} = \mu dt + \sigma dz.$$

We can also do the same calculation the other way around. Suppose that we start from the process

$$ds(t) = \mu S(t) dt + \sigma S(t) dz.$$

Now consider the function $G(S) = \ln S$. Differentiating we have

$$\frac{\partial G}{\partial S} = \frac{1}{S}, \quad \frac{\partial^2 G}{\partial S^2} = -\frac{1}{S^2}, \quad \frac{\partial G}{\partial t} = 0.$$

Hence substituting into Ito's lemma we get

$$dG = d \ln S(t) = \left(\mu - \frac{1}{2}\sigma^2\right) dt + \sigma dz.$$

The forward price

As we have seen before the forward price just depends on the current price of the underlying, the interest rate and the time to expiration. With continuous compounding we can write the forward price equation as

$$F(S(t), t) = S(t)e^{r(T-t)}.$$

This shows the forward price is a stochastic process which depends on the price of the underlying asset which itself is a stochastic process. Since we have

that the forward price is a function of a stochastic process we can use Ito's lemma. Upon differentiation we have

$$\frac{\partial F}{\partial S} = e^{r(T-t)}; \quad \frac{\partial F}{\partial t} = -rS(t)e^{r(T-t)}; \quad \frac{\partial^2 F}{\partial S^2} = 0.$$

Hence substituting into Ito's lemma

$$\begin{aligned} dF &= (e^{r(T-t)}\mu S(t) - rS(t)e^{r(T-t)}) dt + \sigma S(t)e^{r(T-t)} dz \\ &= (\mu - r)S(t)e^{r(T-t)}\sigma S(t)e^{r(T-t)} dz \\ &= (\mu - r)F(t)dt + \sigma F(t)dz. \end{aligned}$$

This shows that the forward price also follows a geometric Brownian motion process with expected return given by the risk premium on the underlying $\mu - r$ and volatility σ (the same as the underlying asset).

Summary

We have shown how returns are continuously compounded and introduced the geometric Brownian motion process for stock prices. We have shown how Ito's lemma can be used. The next thing to do will be to show how to use the assumption of geometric Brownian motion to price an option or derivative using Ito's lemma.