A Model of Stochastic Stock Prices*

Robert McDonald

April 19, 2004

Standard derivatives pricing models assume that the price of the underlying asset follows a process based on Brownian motion. In this note we discuss Brownian motion and its properties, stochastic integration, and a key result called Girsanov's theorem.

1 Some Definitions

This section defines a few terms and concepts that are needed to discuss the mathematical underpinnings of option pricing. The defintions here are intended to help you appreciate the concepts, rather than to be mathematically precise. Rather than being formal, we will try to introduce these concepts using the terminology and ideas of financial markets.

An important concept in asset pricing is that of "past information": things that an investor knows as of a given time. For example, in discussing efficient markets we might say that past information is already incorporated into prices. Let $\Phi(t)$ denote "information available at time t", i.e., past information as of time t. For example, if we are at time t, we can write the expected stock price at time $T > t$ as $E[S(T)|\Phi(t)]$. This is a conditional expectation: the calcuation is performed using all information available at time t. We assume that there is no forgetting, so that information at time $t + s$ at a minimum contains all the information available at time t, that is, $\Phi(t) \subseteq \Phi(t + s)$ ($\Phi(t)$ is a subset of $\Phi(t + s)$).

Stochastic Process A mathematical model for a random process as a function of time. For example, the set of stock prices over time fits this definition. We can express a stochastic process as a collection of random variables $\{X_t\}$, parameterized by t.

Martingale A stochastic process for which $E[X(t + s)|\Phi(t)] = X(t)$. A martingale is not expected to change over time.

Measure Another term for probability distribution.

*This is very preliminary. Comments are welcome, and I expect to produce at least one more (lengthier) version of this in the near future.

2 Brownian Motion

In chapter 11 we presented an example of a **random walk**, Z_n:

$$Z_n = \sum_{i=1}^{n} Y_i \tag{1}$$

where Y_i is a random variable that takes on the values $+1$ or -1. We refer to Y_i as the **increment** to the random walk.

2.1 Definition of Brownian Motion

This particular version of a random walk is hard to interpret economically. To create a random walk that will be useful in modeling stock prices, we need to introduce a time scale into the random walk. We can do this by modifying equation (1) in a manner consistent with the binomial tree.

Suppose the random walk covers the time from 0 to T, and we want to interpret each move, Y_i, as a return over a period of length $h = T/n$. Then we can scale each increment by $\sqrt{h}$. Also, suppose at the start of the random walk the value of Z is $Z(0)$. This gives us

$$Z_n(T) = Z(0) + \sum_{i=1}^{n} \sqrt{h}\, Y_i = Z(0) + \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \sqrt{T}\, Y_i \tag{2}$$

Brownian motion is the limit of this expression as $n \to \infty$ for fixed T. As $n \to \infty$, the contribution of each individual increment goes to 0, and by the central limit theorem, the sum of the increments is normally distributed.

We define Brownian motion as a stochastic process for which

- $Z(0) = 0$
- $Z(t + s) - Z(t)$ is normally distributed with mean 0 and variance s.
- $Z(t + s_1) - Z(t)$ is independent of $Z(t) - Z(t - s_2)$, where $s_1, s_2 > 0$. In other words, non-overlapping increments are independently distributed.
- $Z(t)$ is continuous.

You might wonder how much we assume at the outset by basing our analysis on Brownian motion, which is normally distributed. It is therefore important to note that if a process is continuous and has stationary independent increments (i.e., non-overlapping increments are uncorrelated and the variance of an increment from t to $t + s$ depends only on s), then it *is* Brownian motion.[1] The key assumptions are continuity and independence of non-overlapping increments; normality is a consequence of these assumptions. This result is intuitive: By

[1] See Karatzas and Shreve (1991, Section 3.4). They prove a stronger version of this result, which is that a general continuous stochastic process with independent increments can be expressed in terms of a Brownian motion.

assuming continuity of the path we are effectively assuming that each increment to the process has a variance that is not "too large". When each increment has a finite variance, the central limit theorem applies to the sum of the increments, creating a normally-distributed process.

2.2 Properties of Brownian Motion

We now want to understand how Brownian motion behaves. Using equation (2), we can explore properties of Brownian motion. The following derivations will be rather loose, intended to provide intuition rather than actual proofs. In particular, we will continue to construct the random walk binomially.

The **quadratic variation** of a Brownian process, $Z(t)$ is

$$\lim_{n\to\infty}\sum_{i=1}^{n}\left(Z((i+1)h)-Z(ih)\right)^2=\lim_{n\to\infty}\sum_{i=1}^{n}\left(\sqrt{h}Y_{ih}\right)^2=\lim_{n\to\infty}\sum_{i=1}^{n}hY_{ih}^2$$

Since we are treating Y_i as binomial, taking on the values ± 1, we have $Y_{ih}^2=1$ and hence

$$\lim_{n\to\infty}\sum_{i=1}^{n}hY_{ih}^2=\lim_{n\to\infty}\sum_{i=1}^{n}h=T$$

In other words,

$$\lim_{n\to\infty}\sum_{i=1}^{n}\left(Z((i+1)h)-Z(ih)\right)^2=T \tag{3}$$

We will refer to this result as the finite quadratic variation property of Brownian motion.

The **total variation** of the process is

$$\lim_{n\to\infty}\sum_{i=1}^{n}|\sqrt{h}Y_{ih}|=\lim_{n\to\infty}\sum_{i=1}^{n}\sqrt{h}|Y_{ih}|$$

Again, treating Y_i as binomial, we have $|Y_{ih}|=1$, and hence

$$\lim_{n\to\infty}\sum_{i=1}^{n}\sqrt{h}|Y_{ih}|=\lim_{n\to\infty}\sqrt{T}\sum_{i=1}^{n}\frac{1}{\sqrt{n}}=\sqrt{T}\lim_{n\to\infty}\sqrt{n}=\infty$$

This means that the absolute length of a Brownian path is infinite over any finite interval. If you think about it, in order for a path to have infinite length over a finite interval of time it must move up and down rapidly. This behavior implies the *infinite crossing* property, which states that a Brownian path will cross its starting point an infinite number of times in an interval of any finite length.

2.3 Stochastic Integrals

An integral is essentially the (infinite) sum of a function over a region. In this sense, the random walk defined by equation (2) as $n \to \infty$ is an integral. We can write

$$\int_0^T dZ(t) = Z(T) - Z(0) \tag{4}$$

In words, we are summing up the increments, $Z(t)$, from 0 to T.

A more complicated problem occurs when we want to define the integral of a *function* of $Z(t)$, for example

$$V[Z(0)] = \int_0^T f(Z(t), t) dZ(t) \tag{5}$$

To appreciate some of the issues that arise in defining this integral, consider the case where $f(Z(t), t) = Z(t)$. This particular example is widely used as a clear example of the difference between ordinary and stochastic integrals.[2]

2.3.1 The Riemann Integral

We first examine the ordinary Riemann integral, where Z is not a random variable. We want to compute

$$V(0) = \int_0^T x dx \tag{6}$$

We know from standard calculus that the value of this integral is $T^2/2$. This can be demonstrated by approximating V with a sum, and then showing that the limit of the sum is in fact $T^2/2$.

Split the interval $[0, T]$ into n subintervals, and let $h = T/n$. Then we can set $x_i = ih$. An approximation to equation (6) is

$$V(0) \approx \sum_{i=0}^{n-1} (x_i + \lambda h)[x_{i+1} - x_i]$$

where $0 \le \lambda \le 1$. The parameter λ permits us to evaluate the integral by setting x at an arbitrary point in the interval from x_i to x_{i+1}. Substituting for x_i and expanding terms gives

$$\begin{aligned} V(0) &\approx \sum_{i=0}^{n-1} (ih + \lambda h) h \\ &= \frac{1}{2} T^2 \frac{n(n+1)}{n^2} + \lambda n \frac{T^2}{n^2} \end{aligned}$$

[2]For example, see discussions of stochastic integration in Arnold (1974), Karatzas and Shreve (1991, p. 148) and Neftci (2000).

When we let $n \to \infty$, we obtain $V(0) = T^2/2$, as expected. Notice that the value of the integral does not depend upon λ.

In general

$$\int_a^b x dx = \frac{1}{2}[x^2|_{x=b} - x^2|_{x=a}] = \frac{1}{2}[b^2 - a^2] \tag{7}$$

2.3.2 Stochastic Integration

Now let $Z(t)$ be a Brownian motion, with $Z(0) = 0$. By analogy with equation (6), we wish to evalute

$$V(0) = \int_0^T Z(s)dZ(s) \tag{8}$$

We might guess, by analogy with the previous example, that this integral will equal $[Z(T)^2 - Z(0)^2]/2$. However, the point of this example will be to show the definition of this integral cannot be exactly like the Riemann integral, and moreover that with an *economically* reasonable definition of the integral, we obtain a different answer than with the Riemann integral.[3]

As before, approximate equation (8) as a sum:

$$V(0) \approx \sum_{i=0}^{n-1} Z(ih + \lambda h)[Z((i+1)h) - Z(ih)] \tag{9}$$

In rewriting this sum, we want to do so in a way that exploits two properties of Brownian motion: quadratic variation and the indepence of non-overlapping increments. Let $Z_i = Z(ih)$ and $Z_\lambda = Z(ih + \lambda h)$. Then we can rewrite equation (14) as

$$V(0) \approx \sum_{i=0}^{n-1} \{Z_i[Z_{i+1} - Z_i] + [Z_\lambda - Z_i][Z_{i+1} - Z_\lambda - (Z_i - Z_\lambda)]\} \tag{10}$$

Concentrate first on the second set of terms. We have

$$\sum_{i=0}^{n-1} [Z_\lambda - Z_i][Z_{i+1} - Z_\lambda - (Z_i - Z_\lambda)]$$
$$= \sum_{i=0}^{n-1} \{[Z_\lambda - Z_i][Z_{i+1} - Z_\lambda] + [Z_\lambda - Z_i]^2\} \tag{11}$$

Now we let $n \to \infty$. Because of quadratic variation, the second term equals λT. The first term is the infinite product of uncorrelated normal random variables. As $n \to \infty$, each term in the sum is the product of two independent normal variables with variances proportional to λh and $(1 - \lambda)h$. The result that this sum is zero in the limit is plausible: the variances of each Brownian increment

[3]For a mathematically precise discussion of stochastic integrals, see Karatzas and Shreve (1991).

shrink to 0, their product on average is zero, and hence the infinite sum of the products converges to zero.[4] Thus, we have

$$\sum_{i=0}^{n-1}[Z_\lambda - Z_i][Z_{i+1} - Z_\lambda - (Z_i - Z_\lambda)] = \lambda T \tag{12}$$

Now examine the first term. We have[5]

$$\sum_{i=0}^{n-1} Z_i[Z_{i+1} - Z_i] = \frac{1}{2}\sum_{i=0}^{n-1}[Z_{i+1}^2 - Z_i^2 - (Z_{i+1} - Z_i)^2] \tag{13}$$

If you write out the summation for the first two terms in the summation on the right-hand side of equation (13), you will see that

$$\frac{1}{2}\sum_{i=0}^{n-1}[Z_{i+1}^2 - Z_i^2] = \frac{1}{2}[Z_n^2 - Z_0^2]$$

By quadratic variation, the last term within the summation on the right-hand side of equation (13) is $T/2$. Putting together these results, we have

$$\lim_{n\to\infty}\sum_{i=0}^{n-1} Z_\lambda[Z_{i+1} - Z_i] = \frac{1}{2}[Z_T^2 - Z_0^2] + [\lambda - \frac{1}{2}]T \tag{14}$$

Unlike in the Riemann integral, the value of the stochastic integral *does* depend upon the specific choice of λ. If we set λ at any value other than zero, we induce correlation between Z_λ and $Z_{i+1} - Z_i$. Equation (14) states that this correlation does not vanish in the limit.

As it turns out, there is an economic argument for setting $\lambda = 0$. Suppose that $Z(t)$ is the stock price, and we let $\alpha(Z_t, t)$ denote the number of shares of stock we will buy ($\alpha > 0$ or sell ($\alpha < 0$) as a function of the stock price. Then the number of shares we hold at time T is the cumulated sum of our buys and sells, or

$$\int_0^T \alpha[Z(t), t]dZ(t)$$

If we define this integral with $\lambda > 0$, then the economic interpretation is that we select the trade at time t as a function of the change in the stock price *that is going to occur*! We would be basing our trade at time t on a price change that, as of time t, we had not yet observed. While the ability to see into the future would be desirable, it is not a property we want to build into a model of financial markets. The Itô integral is a stochastic integral defined with $\lambda = 0$.

[4]To prove a statement like this for an infinite sum of random variables, it is necessary to define convergence. The concept typically used for stochastic integration is convergence in quadratic mean, which essentially means that the variance of the sum goes to zero. In other words, it is highly unlikely that in the limit the sum would be anything other than zero.

[5]This uses the fact that since $(a+b)^2 = a^2+b^2+2ab$, it follows that $ab = [(a+b)^2-a^2-b^2]/2$. Set $a = Z_i$ and $b = Z_{i+1} - Z_i$.

2.4 Girsanov's Theorem

We have seen that it is possible to derive the Black-Scholes formula with the following steps:

1. Assume that the stock is lognormally distributed with expected return r instead of α
2. Compute the expected payoff of the option at maturity, time T
3. Discount this expected payoff at the riskfree rate.

While this procedure gives us a correct option price, the first step is mysterious. Replacing α with r appears to convert the stock's probability distribution from the true probability distribution to the risk-neutral distribution. However, what permits us to take this step?

Consider the two following expressions for the logarithm of the future stock price:

$$\ln(S_T) = \ln(S_0) + (\alpha - \delta - 0.5\sigma^2)T + \sigma\sqrt{T}Z(T) \tag{15}$$

$$\ln(S_T) = \ln(S_0) + (r - \delta - 0.5\sigma^2)T + \sigma\sqrt{T}Z^*(T) \tag{16}$$

In equation (15), the drift is $\alpha - \delta - 0.5\sigma^2$; in equation (16), the drift is $r - \delta - 0.5\sigma^2$. How can these two equations yield the same stock price at time T? Note that there is one other difference between the two equations: equation (15) contains the normally-distributed variable $Z(T)$, and equation (16) contains the normally-distributed variable $Z^*(T)$. The two equations can describe the same stock price at time T as long as we define Z^* to undo the change in mean return from replacing α with r.

Another name for a probability distribution is **probability measure**, and we refer to the replacement of $Z(T)$ with $Z^*(T)$ as a **change of measure**. At first glance (and perhaps at second glance) the concept of a change of measure is baffling. What permits us to redefine probabilities?

To better understand the meaning of a change of measure, consider a coin-flipping game in which you receive \$1.50 if a fair coin lands heads (this occurs with a 50% probability because the coin is fair), and you pay \$1 if the coin lands tails. You play this game repeatedly. Suppose that the coin has a 50% probability of landing heads. On any given toss of the coin, your expected payment is

$$0.5 \times \$1.50 + 0.50 \times (-\$1) = \$0.25$$

Your cash balance is not a random walk, but has a positive drift of \$0.25 per play. We could "change the measure" to make this a fair game. How? We could replace the fair coin with a coin that has a probability of landing heads, p, that satisfies

$$p \times \$1.50 + (1 - p) \times (-\$1) = 0$$

This implies a p of 0.40. While it is possible to replace the coin, it should be clear that this is a change in the physical world and is not helpful in analyzing the original game with a fair coin.

A change of measure can be helpful, however, when we want to *value* a claim. In this case the change of measure does not involve literally changing the physical coin, but rather *imagining* a hypothetical change in the physical coin. We construct hypothetical probabilities to be consistent with values of assets that we observe, and then use the hypothetical probabilities to value payoffs that are related to the original asset.

Suppose that risk-averse investors play the coin-flipping game. Because of their risk-aversion, these investors value a dollar of winnings less highly than they value a dollar of loss. Suppose, moreover, we learn that these investors are exactly indifferent about playing the coin-flipping game, in other words they would be willing to play it but only if playing were costless. In other words, investors value the game at zero.

This indifference about playing the game is powerful information, because it tells us precisely how the investors value winning as opposed to losing. Let U_H denote the utility value an investor attaches to the winnings associated with heads and U_L the utility value an investor attaches to the winnings associated with tails. Then we have

$$U_H \times \$1.50 + U_L \times (-\$1) = 0$$

or $U_L/U_H = \$1.50$. Consider using this ratio to guide us in a change of measure: since investors value wins at 50% more than losses, for the purposes of *valuing* the game, let's assign a probability of losing that is 50% greater than the probability of winning. This gives us

$$\frac{1-p}{p} = 1.50 \quad \text{or} \quad p = 0.40$$

Notice that 0.40 is also the probability that makes the coin flip game have a zero average payoff. The probability 0.40 is what we have been calling the risk-neutral probability.

Now we can see how a change of measure is useful for valuation. There are two aspects of the coin flipping game we can observe: average winnings, and investor willingness to participate in the game. When we want to understand the average winnings we will observe in the coin flipping game, we need to use the true probability of heads. However, when we observe that investors are indifferent about participation, we conclude that investors attach a zero value to the game. When we want to understand how investors *value* the game, we can use the change of measure to set $p = 0.40$, and we conclude that the game has a value of zero:

$$0.40 \times \$1.50 + (1 - 0.40) \times (-\$1) = 0$$

How does this notion of indifference apply to stocks and bonds? Investors are not in general indifferent between an investment in stocks and bonds, but in portfolio equilibrium, investors are indifferent about changes in the allocation of their wealth between stocks and bonds: this is the definition of equilibrium.

So in equilibrium, investors are indifferent *at the margin* between investments in stocks and bonds. Risk-neutral pricing effectively exploits this *marginal* indifference between risky stocks and risk-free bonds. There will be a unique risk-neutral change of measure as long as all investors are indifferent at the margin between stocks and bonds, i.e., as long as all investors are in portfolio equilibrium. At this point you should review Appendix **??**, which demonstrates a change of measure based on stocks and bonds rather than coin flips.

To this point the change of measure only verifies what we already knew: the coin flipping game has zero value. Now suppose we change the game so that with heads we receive \$3 and with tails we still pay \$1. Clearly if investors value the original game at zero, this modified game will have a positive price. There are two ways to perform the valuation.

First, we could reason that the new game is like scaling the old game by Δ and receiving the fixed amount B. Then we can find the value of the game by solving two equations in two unknowns:

$$\Delta \times \$1.50 + B = \$3$$
$$\Delta \times (-\$1.00) + B = -\$1$$

Solving, we obtain $\Delta = 1.60$ and $B = \$0.60$. Thus, the value of the new game is \$0.60, since the new game is like scaling the original game, which has a zero value, by 1.60, and adding \$0.60.

Second, we can use the change of measure. The risk-neutral probability is 0.40, so the game has the value

$$0.40 \times \$3.00 + (1 - 0.40) \times (-\$1.00) = \$0.60$$

The risk-neutral probability—once we have it—provides a simple answer, because it tells us that each extra \$1 of winnings is worth \$0.40. Hence \$1.50 of winnings is worth \$0.60.[6]

It is important to understand that the change of measure works in this case only because we observe how investors value the coin flip; there is no mechanical procedure to tell us what change of measure is appropriate. If we did not observe investor valuation of the coin flip game, there are two alternatives:

1. we would need to know something about how investors assess the risks of coin flipping; that is, we would need to know investor utility functions, or

2. we would need to be able to characterize coin flipping in terms of other risky activities we do observe. I.e., we would need *comparables*, and in the language of financial economists we would be assuming that the risk of coin flipping is *spanned*—i.e., is replicable—by other traded risks we do observe.

[6] END OF CHAPTER PROBLEM: What change of measure is appropriate if investors pay \$0.05 to pay the coin flipping game?

What if the risky payoff we wish to value is not traded, so we do not observe a price; we know nothing about investor utility; and there are no traded comparables, so the risk is not spanned? In this case there is no alternative to making assumptions, typically either about investor preferences or about comparables. If you find this frustrating, recognize that finance gains much of its power from the existence of financial markets, in which risk-averse investors trade securities. Also, recognize that financial analysts routinely assume spanning. For example, suppose an analyst must recommend that a firm either accept or reject a project. We typically think of firms as maximizing shareholder value. Thus, the analyst must assess the value of the project *from the perspective of shareholders*. Implicitly, we assume that shareholders will agree that the project's risk can be evaluated using a model like the CAPM; this amounts to assuming that the project's cash flows are spanned by the cash flows of other traded securities.