Random Walks

John Norstad

j-norstad@northwestern.edu
http://www.norstad.org

January 28, 2005
Updated: November 3, 2011

Abstract

We develop the formal mathematics of the lognormal random walk model. We start by discussing continuous compounding for risk-free investments. We introduce a random variable to model the uncertainty of a risky investment. We apply the Central Limit Theorem to argue that under three strong assumptions, the values of risky investments at any time horizon are lognormally distributed. We model the random walk using a stochastic differential equation. We define the notion of an "Ito process" and prove that it is equivalent to our formulation.

We apply the model to the S&P 500 stock market index as an example. We learn how to do parameter estimation for the model using historical time series data and how to do calculations in the model in computer programs. We discuss how uncertainty and risk increase with time horizon when investing in volatile assets like stocks, contrary to popular opinion.

We conclude by asking the all-important question of how well the simple random walk model describes how financial markets actually work. We mention known failings of the model and conclude that at best it is a rough approximation to reality and should be used for real-life financial planning with caution.

We assume that the reader is familiar with the normal and lognormal probability distributions as presented in reference [8].

Contents

List of Figures

1 Continuous Compounding

Suppose we have s_0 dollars and invest it in a savings account or other risk-free investment earning a yearly interest rate of μ. What is the value s_1 of our investment at the end of one year?

The answer depends on how often interest is compounded.

First suppose that interest is compounded yearly, and we receive a single interest payment at the end of the year. The interest payment is $s_0\mu$, and the ending value of our investment is:

$$s_1 = s_0 + s_0\mu = s_0(1+\mu)$$

This is called *simple compounding*, and the interest rate μ is called the *simply compounded rate of return* of our investment.

Now suppose that interest is compounded semi-annually. We receive one interest payment after six months using the interest rate $\mu/2$. If we denote our balance after six months by $s_{0.5}$, we have:

$$s_{0.5} = s_0 + s_0\frac{\mu}{2} = s_0\left(1+\frac{\mu}{2}\right)$$

At the end of the year we receive a second interest payment, again using the interest rate $\mu/2$, applied to our balance $s_{0.5}$:

$$\begin{aligned} s_1 &= s_{0.5} + s_{0.5}\frac{\mu}{2} \\ &= s_{0.5}\left(1+\frac{\mu}{2}\right) \\ &= s_0\left(1+\frac{\mu}{2}\right)\left(1+\frac{\mu}{2}\right) \\ &= s_0\left(1+\frac{\mu}{2}\right)^2 \end{aligned}$$

We can compute the equivalent simply compounded rate of return r:

$$\begin{aligned} s_0(1+r) &= s_0\left(1+\frac{\mu}{2}\right)^2 \\ 1+r &= \left(1+\frac{\mu}{2}\right)^2 \\ r &= \left(1+\frac{\mu}{2}\right)^2 - 1 \end{aligned}$$

Now suppose that interest is compounded monthly. We receive twelve interest payments, one at the end of each month, using the interest rate $\mu/12$. Each interest payment multiplies our balance by $1+\mu/12$. At the end of the year our balance is:

$$s_1 = s_0 \left(1 + \frac{\mu}{12}\right)^{12}$$

The equivalent simply compounded rate of return r is:

$$r = \left(1 + \frac{\mu}{12}\right)^{12} - 1$$

As one more example, suppose interest is compounded daily, and it is not a leap year, so we have 365 days in the year. In this case we have:

$$\begin{aligned} s_1 &= s_0 \left(1 + \frac{\mu}{365}\right)^{365} \\ r &= \left(1 + \frac{\mu}{365}\right)^{365} - 1 \end{aligned}$$

In general, if interest is compounded n times per year, we have:

$$s_1 = s_0 \left(1 + \frac{\mu}{n}\right)^{n} \tag{1}$$

$$r = \left(1 + \frac{\mu}{n}\right)^{n} - 1 \tag{2}$$

What happens if interest is compounded more and more frequently? In other words, what happens as n gets larger and larger in equations (1) and (2)? In the limit, we have:

$$s_1 = s_0 \lim_{n\to\infty} \left(1 + \frac{\mu}{n}\right)^{n} \tag{3}$$

$$r = \lim_{n\to\infty} \left(1 + \frac{\mu}{n}\right)^{n} - 1 \tag{4}$$

To evaluate the limits in equations (3) and (4) we use L'Hôpital's rule. Let $x = 1/n$. Then:

$$\begin{aligned} \lim_{n\to\infty} \left(1 + \frac{\mu}{n}\right)^{n} &= \lim_{x\to 0} (1+\mu x)^{1/x} \\ &= \lim_{x\to 0} e^{\log[(1+\mu x)^{1/x}]} \\ &= e^{\lim_{x\to 0} \log[(1+\mu x)^{1/x}]} \\ \lim_{x\to 0} \log[(1+\mu x)^{1/x}] &= \lim_{x\to 0} \frac{\log(1+\mu x)}{x} \\ &= \lim_{x\to 0} \frac{\frac{d}{dx}\log(1+\mu x)}{\frac{d}{dx}x} \\ &= \lim_{x\to 0} \frac{\frac{\mu}{1+\mu x}}{1} \\ &= \mu \\ \lim_{n\to\infty} \left(1 + \frac{\mu}{n}\right)^{n} &= e^{\mu} \end{aligned}$$

Thus our equations (3) and (4) become:

$$
\begin{aligned}
s_1 &= s_0 e^{\mu} && (5)\\
r &= e^{\mu} - 1 && (6)\\
\mu &= \log(r+1) && (7)
\end{aligned}
$$

These equations tell us what happens if interest is compounded continuously, at every instant of time over the year. This is called *continuous compounding* and μ is called the *continuously compounded rate of return* of our investment.

As a concrete example, suppose the simply compounded rate of return $r = 5\%$. The equivalent continuously compounded rate of return is $\mu = \log(r+1) = \log(1.05) = 4.8790\%$. If we invest $s_0 = \$100$, after one year our investment grows to:

$$
\begin{aligned}
s_1 &= s_0(1+r) = 100(1+5\%) = 100(1.05) = \$105.00\\
s_1 &= s_0 e^{\mu} = 100e^{4.8790\%} = 100e^{0.048790} = \$105.00
\end{aligned}
$$

Suppose we have a risk-free investment s with an initial value of s_0 that earns a continuously compounded rate of return μ. Let:

$$s(t) = \text{the value of the investment at time } t$$

Then:

$$s(t) = s_0 e^{\mu t}$$

Consider the value of the investment a short time later, at time $t + dt$:

$$s(t+dt) = s_0 e^{\mu(t+dt)} = s_0 e^{\mu t} e^{\mu dt} = s(t) e^{\mu dt}$$

Let:

$$ds(t) = \text{the growth of the investment over the time interval } [t, t+dt]$$

Then:

$$
\begin{aligned}
ds(t) &= s(t+dt) - s(t)\\
&= s(t)e^{\mu dt} - s(t)\\
&= s(t)(e^{\mu dt} - 1)\\
\frac{ds(t)}{s(t)} &= e^{\mu dt} - 1
\end{aligned}
$$

This equation holds at all times t, so we have the following differential equation which describes the behavior of our risk-free investment with continuous compounding:

$$\frac{ds}{s} = e^{\mu dt} - 1 \tag{8}$$

2 Uncertainty

In the previous section we examined risk-free investments that earn interest continuously over time.

In this section we turn our attention to risky investments where the change in value of the investment over time is uncertain.

Let s be a risky investment with intial value s_0.

Consider a small time interval dt and let s_1 be the value of our initial investment s_0 after dt time has passed. Over this short time interval the rate of return of our investment is some random variable Y_1, and the value s_1 of our investment at the end of the time interval is:

$$s_1 = s_0(1 + Y_1)$$

Now consider a second small time interval dt. Let s_2 be the value of our investment at the end of the second time interval, and let Y_2 be the random variable for the rate of return of our investment over the second time interval. Then:

$$s_2 = s_1(1 + Y_2) = s_0(1 + Y_1)(1 + Y_2)$$

As time goes on, over each small time interval dt the value of our investment changes by some small random amount. Let s_n be the value of our investment at the end of n time intervals, and let Y_i be the random variable for the rate of return of our investment over the time interval i. Then:

$$s_n = s_0 \prod_{i=1}^{n}(1 + Y_i)$$

Take the logarithm of both sides of this equation:

$$\begin{aligned}
\log(s_n) &= \log\left(s_0 \prod_{i=1}^{n}(1 + Y_i)\right) \\
&= \log(s_0) + \log\left(\prod_{i=1}^{n}(1 + Y_i)\right) \\
&= \log(s_0) + \sum_{i=1}^{n}\log(1 + Y_i) \\
\log(s_n/s_0) &= \log(s_n) - \log(s_0) \\
&= \sum_{i=1}^{n}\log(1 + Y_i)
\end{aligned}$$

For each i, let Z_i be the random variable $\log(1+Y_i)$. Then our equation becomes:

$$\log(s_n/s_0) = \sum_{i=1}^{n} Z_i \tag{9}$$

3 Lognormal Random Walks

The equation (9) which we derived in the previous section is not very useful without additional information about the distribution of the random variables Y_i which give the rate of return of the investment over time interval i.

We now make three strong assumptions about these random variables:[1]

1. The random variables Y_i are *independent*. What happens at one time interval does not affect what happens at subsequent time intervals. The market "has no memory."
2. The random variables Y_i are *identically distributed*. The means, standard deviations, and other attributes of the probability distributions do not change over time.
3. The random variables Y_i have *finite variance*.

Recall equation (9) from the previous section:

$$\log(s_n/s_0) = \sum_{i=1}^{n} Z_i \quad \text{where } Z_i = \log(1+Y_i)$$

If the random variables Y_i are independent, identically distributed, and have finite variance, then so do the random variables Z_i.

The Central Limit Theorem of Probability Theory says that in the limit, as $n \to \infty$, the average of n independent identically distributed random variables with finite variance is normally distributed. Thus under our three assumptions, we have:

$$\lim_{n\to\infty} \frac{1}{n} \log(s_n/s_0) \text{ is normally distributed}$$

Our time interval dt is very short. For example, if dt is one second, and each Y_i is the random rate of return of our investment over one second, then over a single seven hour trading day, we have $n = 7 * 60 * 60 = 25,200$ seconds.

It is reasonable to assume at this point that $\log(s_n/s_0)$ is normally distributed after even just one day. Indeed, it is reasonable to assume that $\log(s_n/s_0)$ is normally distributed for all n, with each Z_i normally distributed with identical means and variances.

Let:

$$\begin{aligned} \mu &= \mathrm{E}(Z_i)/dt \\ \sigma^2 &= \mathrm{Var}(Z_i)/dt \end{aligned}$$

[1]All three of these assumptions turn out to be suspect. We discuss this in our conclusion in section 7.

Then Z_i is $N[\mu dt, \sigma^2 dt]$. Define:

$$dX_i = (Z_i - \mu dt)/\sigma$$

Then:

$$Z_i = \mu dt + \sigma dX_i \quad \text{where } dX_i \text{ is } N[0, dt]$$

Define:

$$s(t) = \text{the value of investment } s \text{ at time } t$$

and let:

$$n = t/dt$$

Then:

$$\begin{aligned}
\log(s(t)/s(0)) &= \log(s_n/s_0) \\
&= \sum_{i=1}^{n} Z_i \\
&= \sum_{i=1}^{n} (\mu dt + \sigma dX_i) \quad \text{where } dX_i \text{ is } N[0, dt] \\
&= \sum_{i=1}^{n} \mu dt + \sum_{i=1}^{n} \sigma dX_i \\
&= n\mu dt + \sigma \sum_{i=1}^{n} dX_i \\
&= \mu t + \sigma \sum_{i=1}^{n} dX_i
\end{aligned}$$

The variables dX_i are independent normally distributed random variables and each has mean 0 and variance dt. By Corollary 1 in reference [8], $\sum_{i=1}^{n} dX_i$ is also normally distributed and has mean 0 and variance $n \times dt = t$. Thus we have:

$$\begin{aligned}
\log(s(t)/s(0)) &= \mu t + \sigma X \quad \text{where } X \text{ is } N[0, t] && (10) \\
s(t)/s(0) &= e^{\mu t + \sigma X} && (11) \\
s(t) &= s(0) e^{\mu t + \sigma X} && (12)
\end{aligned}$$

Note that $\mu t + \sigma X$ is normally distributed $N[\mu t, \sigma^2 t]$, so $s(t)/s(0)$ is lognormally distributed $LN[\mu t, \sigma^2 t]$.

When $t = 1$ (one year) we have:

$$\begin{aligned}
s(1) &= s(0) e^{\mu + \sigma X} \quad \text{where } X \text{ is } N[0, 1] \\
s(1) &= s(0) e^{\hat{X}} \quad \text{where } \hat{X} \text{ is } N[\mu, \sigma^2] && (13)
\end{aligned}$$

Finally, consider the change in value $ds(t)$ of the investment s over a short time interval $[t, t+dt]$. We have:

$$\begin{aligned}
s(t+dt) &= s(t)e^{\mu dt+\sigma dX} \quad \text{where } dX \text{ is } N[0, dt] \\
ds(t) &= s(t+dt) - s(t) \\
&= s(t)\left(e^{\mu dt+\sigma dX} - 1\right) \\
\frac{ds(t)}{s(t)} &= e^{\mu dt+\sigma dX} - 1
\end{aligned}$$

This equation holds at all times t, so we have the following *stochastic* differential equation which describes the behavior of our risky investment over time:

$$\frac{ds}{s} = e^{\mu dt+\sigma dX} - 1 \quad \text{where } dX \text{ is } N[0, dt] \tag{14}$$

Compare this equation (14) to the ordinary differential equation (8) we derived in section 1 for the behavior of a risk-free investment over time with continuous compounding:

$$\frac{ds}{s} = e^{\mu dt} - 1 \tag{15}$$

The difference between these equations is that equation (14) has the additional random term σdX to account for the uncertainty (riskiness) of our investment. Equation (15) is the special case of equation (14) when $\sigma = 0$.

Equation (14) is one formulation of the *lognormal random walk model*. μ is the *continuously compounded expected rate of return* of the investment, and σ is the *standard deviation of the continuously compounded returns*.

4 Ito Processes

In the formal mathematics of continuous-time finance,[2] the lognormal random walk model is usually formulated in terms of the following stochastic differential equation:

$$\frac{ds}{s} = \alpha dt + \sigma dX \quad \text{where } dX \text{ is } N[0, dt] \tag{16}$$

A random variable s which satisfies this equation is called an *Ito process*. The number α is called the *instantaneous rate of return*.

Compare (16) to our formulation (14) in the previous section:

$$\frac{ds}{s} = e^{\mu dt + \sigma dX} - 1 \quad \text{where } dX \text{ is } N[0, dt]$$

These formulations are clearly not the same. In fact, they are quite different. We can, however, show that they are equivalent, under a suitable interpretation of the variable α that appears in equation (16). To do this we need a fundamental result from the theory of stochastic calculus, which we state here without proof.

Lemma 4.1 *(Ito's Lemma) Suppose G is a random variable satisfying the stochastic differential equation*

$$dG = A(G, t)dX + B(G, t)dt$$

where dX is $N[0, dt]$, A and B are two functions of G and t, and $f(G, t)$ is a twice differentiable function of G and a once differentiable function of t. Then:

$$df = A\frac{\partial f}{\partial G}dX + \left(B\frac{\partial f}{\partial G} + \frac{1}{2}A^2\frac{\partial^2 f}{\partial G^2} + \frac{\partial f}{\partial t}\right)dt$$

If f is a function only of G and not of t, as will be the case in our application, we can state Ito's Lemma in the following form:

$$df = Af'dX + (Bf' + \frac{1}{2}A^2 f'')dt$$

Note the funny looking term $\frac{1}{2}A^2 f''$ in this equation. If X were a simple deterministic function of t instead of a random variable this term would not present, and the Lemma without the term would be a trivial application of the chain rule. This extra term illustrates how the rules of ordinary calculus to which we are accustomed change rather radically when we introduce random elements.

[2]See for example Merton [7].

We are now ready to prove the main result of this paper.

Theorem 4.1 *For a random variable s,*

$$\frac{ds}{s} = \alpha dt + \sigma dX \quad \textit{iff} \quad \frac{ds}{s} = e^{\mu dt + \sigma dX} - 1$$

where dX is $N[0, dt]$ and $\alpha = \mu + \frac{1}{2}\sigma^2$.

In this case, s follows a lognormal random walk. The logarithm of $s(1)/s(0)$ is normally distributed with mean μ and standard deviation σ. α is the yearly instantaneous expected return, μ is the yearly continuously compounded expected return and σ is the standard deviation of those returns. Over any time horizon t, $s(t)$ is lognormally distributed with:

$$s(t) = s(0)e^{\mu t + \sigma X} \quad \textit{where } X \textit{ is } N[0, t] \tag{17}$$

Proof:

First suppose $\frac{ds}{s} = \alpha dt + \sigma dX$. Apply Ito's Lemma with:

$$\begin{aligned} f(s) &= \log(s) \\ A &= \sigma s \\ B &= \alpha s \end{aligned}$$

Ito's Lemma becomes:

$$df = \sigma s f' dX + (\alpha s f' + \frac{1}{2}\sigma^2 s^2 f'')dt$$

Substitute $f' = 1/s$ and $f'' = -1/s^2$ to get:

$$\begin{aligned} df &= \sigma dX + (\alpha - \frac{1}{2}\sigma^2)dt \\ &= \sigma dX + \mu dt \end{aligned}$$

Note that we also have:

$$\begin{aligned} df &= f(s(t+dt)) - f(s(t)) \\ &= \log(s(t+dt)) - \log(s(t)) \\ &= \log\left(\frac{s(t+dt)}{s(t)}\right) \\ &= \log(\frac{ds}{s} + 1) \end{aligned}$$

Thus we have:

$$\begin{aligned} \log(\frac{ds}{s} + 1) &= \sigma dX + \mu dt \\ \frac{ds}{s} &= e^{\mu dt + \sigma dX} - 1 \end{aligned}$$

This completes one half of the proof.

For the other direction, suppose $\frac{ds}{s} = e^{\mu dt + \sigma dX} - 1$. Again let $f(s) = \log(s)$. As above,

$$\begin{aligned} df &= \log\left(\frac{ds}{s} + 1\right) \\ &= \sigma dX + \mu dt \end{aligned}$$

Apply Ito's Lemma with:

$$\begin{aligned} s(f) &= e^f = e^{\log(s)} = s \\ A &= \sigma \\ B &= \mu \end{aligned}$$

Note that $s = s' = s'' = e^f$. Ito's Lemma says that:

$$\begin{aligned} ds &= \sigma s' dX + (\mu s' + \frac{1}{2}\sigma^2 s'')dt \\ &= \sigma s dX + (\mu s + \frac{1}{2}\sigma^2 s)dt \\ &= \sigma s dX + \alpha s dt \end{aligned}$$

Divide by s to get:

$$\frac{ds}{s} = \alpha dt + \sigma dX$$

This completes the other half of the proof.

The equation (17) in the theorem is the equation (12) which we derived in section 3 on page 7.

5 Measuring Returns

In sections 3 and 4 we developed two formulations of the random walk model. The formulations used two different ways to measure the expected return of our investment. Section 3 used μ = the expected continuously compounded return, and section 4 used α = the expected instantaneous return. We showed that these two ways to measure expected returns are related by the equation $\alpha = \mu + \frac{1}{2}\sigma^2$.

In the popular literature returns are more often measured using simple compounding. We now compute the expected simply compounded return.

On page 7 we derived equation (13) for the value of a risky investment s after one year:

$$s(1) = s(0)e^{\hat{X}} \qquad \text{where } \hat{X} \text{ is } N[\mu, \sigma^2]$$

The yearly simply compounded rate of return on our investment is the random variable R given by:

$$R = s(1)/s(0) - 1 = e^{\hat{X}} - 1$$

$\hat{X}$ is $N[\mu, \sigma^2]$, so $e^{\hat{X}}$ is $LN[\mu, \sigma^2]$. By proposition 5 in reference [8], the expected value of R is:

$$\begin{aligned} \mathrm{E}(R) &= \mathrm{E}(e^{\hat{X}} - 1) \\ &= \mathrm{E}(e^{\hat{X}}) - 1 \\ &= e^{\mu + \frac{1}{2}\sigma^2} - 1 \\ &= e^{\alpha} - 1 \end{aligned}$$

We can also easily compute the median value of R. The random variable $\hat{X}$ is normally distributed with mean value = median value = μ. Thus 50% of the time the value of $\hat{X}$ is less than μ and 50% of the time the value is greater than μ. The median value of $R = e^{\hat{X}} - 1$ is therefore $e^{\mu} - 1$. This value is also called the *geometric mean return* or the *annualized return*.

To summarize, we have four different ways to measure returns:

α = instantaneous return
μ = continuously compounded return
$r1$ = simply compounded arithmetic mean (average) return
$r2$ = simply compounded geometric mean (annualized or median) return

These ways to measure returns are related by the following equations:

$$\alpha = \mu + \frac{1}{2}\sigma^2 \tag{18}$$
$$r1 = e^{\alpha} - 1 \tag{19}$$
$$r2 = e^{\mu} - 1 \tag{20}$$

It is important to distinguish between these ways to measure returns and to use the proper measures in the proper contexts.

6 Example – The S&P 500

As an extended example, we will apply what we have learned to build a random walk model for the S&P 500 stock market index. This index measures the performance of large US company stocks. It is often used in both academic finance and the popular press as a proxy for the entire US stock market, or at least the part of the stock market representing large companies.

6.1 Parameter Estimation

We have yearly return data for the S&P 500 all the way back to 1926. The data measures "total returns," which includes both capital gains and dividends and assumes that all dividends are reinvested.

Let $R =$ the time series of yearly total return data for the S&P 500 index from 1926 through 1994.[3]

The first task is to use R to get estimates of the parameters μ and σ for our model.

The time series R measures yearly returns using simple compounding. We first convert to continuous compounding by taking the natural logarithm of 1 plus each return to get a new time series $\log(1+R)$. We then set μ and σ to be the mean and standard deviation of this time series respectively:

$$\mu = \mathrm{E}(\log(1+R)) = 0.097070 = 9.7070\%$$
$$\sigma = \mathrm{Stdev}(\log(1+R)) = 0.194756 = 19.4756\%$$

When estimating variances from discrete sample data, statisticians tell us that we should divide by the number of samples minus 1 rather than by the number of samples to get a more accurate estimate. Thus we compute the variance of a sample $x = x_1 \ldots x_n$ as $\frac{1}{n-1}\sum(x_i - \mathrm{E}(x))^2$ rather than $\frac{1}{n}\sum(x_i - \mathrm{E}(x))^2$. We did this in our computations above. Note that if you use Microsoft Excel to do these calculations its built-in functions for variance and standard deviation make this adjustment for you.

[3] The data is from Table 2-4 in reference [3].

6.2 Simulating Random Walks

Now that we have estimated μ and σ, what can we do with our model?

One thing we can do is simulate random walks using a computer program.[4] We show the graphs of two such simulations in Figure 1. The program uses our random walk model with the following additional parameters:

s_0 = starting value = \$100
t = time period = 1 year
dt = .004 = 1/250 = one trading day

The smooth curve represents the median and is generated by setting $\sigma = 0$. If we run a large number of simulations, about half of them end below the median and about half end above. Note that the median ending value is \$110.19, which represents an annual return of 10.19% (using simple compounding).

The graphs look remarkably like stock market charts, don't they? Perhaps this math stuff isn't as useless as we thought!

How does one write a program to do such a simulation? It's quite easy. Here's some skeleton Java code for the algorithm:

```
double mu = .097070;
double sigma = .194756;
double s = 100.0;
double t = 1.0;
double dt = 0.004;
double sqrtdt = Math.sqrt(dt);
Random random = new Random();
for (int i = 0; i < t/dt; i++) {
   double dx = sqrtdt * random.nextGaussian();
   s *= Math.exp(mu*dt + sigma*dx);
   // Draw one new little segment of the graph by moving the
   // graphics pen right dt and up or down to the new s value.
}
```

The call to the method `nextGaussian` generates and returns a single normally distributed pseudo-random number with mean 0 and variance 1.

[4] We use the "Random Walker" program [11], available at the author's web site. All the graphs in this paper were drawn by Random Walker.

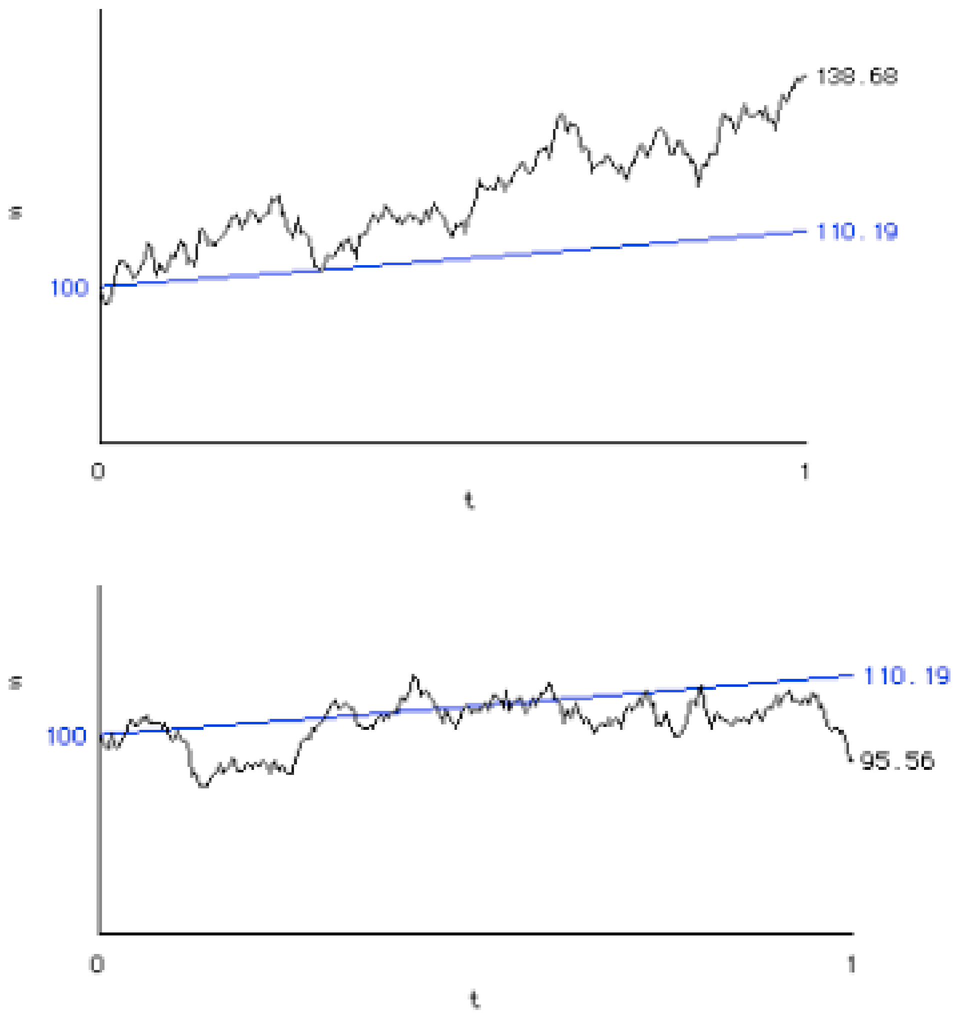

Figure 1: S&P 500 Random Walks

6.3 Density Functions

A more useful thing we can do with our model is use it to graph the density function of the ending values.

As an example, in Figure 2, we start with $100 invested in our S&P 500 model. The graphs show the density functions of the ending values of our investment after 1, 5, 10, and 20 years.

In each density graph we have also drawn two vertical lines, one at the median (on the left) and one at the mean (expected value, on the right). These values are computed from the equations we derived:

$$\text{Median} = s_0 e^{\mu t}$$
$$\text{Mean} = s_0 e^{(\mu + \frac{1}{2}\sigma^2)t}$$

	1 year	5 years	10 years	20 years
Median	\$110.19	\$162.47	\$263.98	\$696.85
Mean	\$112.30	\$178.64	\$319.10	\$1,108.28
Standard deviation	\$22.98	\$81.63	\$216.72	\$1,084.97

Note that we always have mode $<$ median $<$ mean, unlike in a normal distribution, where these three values are all the same. As the time period increases, these three values spread farther apart.

Also note how the distinctive shape of the lognormal distribution density function becomes more pronounced as the time period increases. The density function becomes increasingly skewed. To the right of the mode it grows a longer tail.

As the time period increases, the spread of the likely ending values increases rapidly, as does their standard deviation. Investing in volatile assets like stocks is clearly an uncertain proposition under this model, and the uncertainty increases with time. Look at the X axis on the graphs and see how we had to change the scale with each increase in time to accommodate this wider spread of ending values. In each of the graphs we scaled the X axis from the 1st percentile to the 99th percentile. In other words, the probability of an ending value being less than the X axis minimum is 1%, as is the probability of it being larger than the X axis maximum.

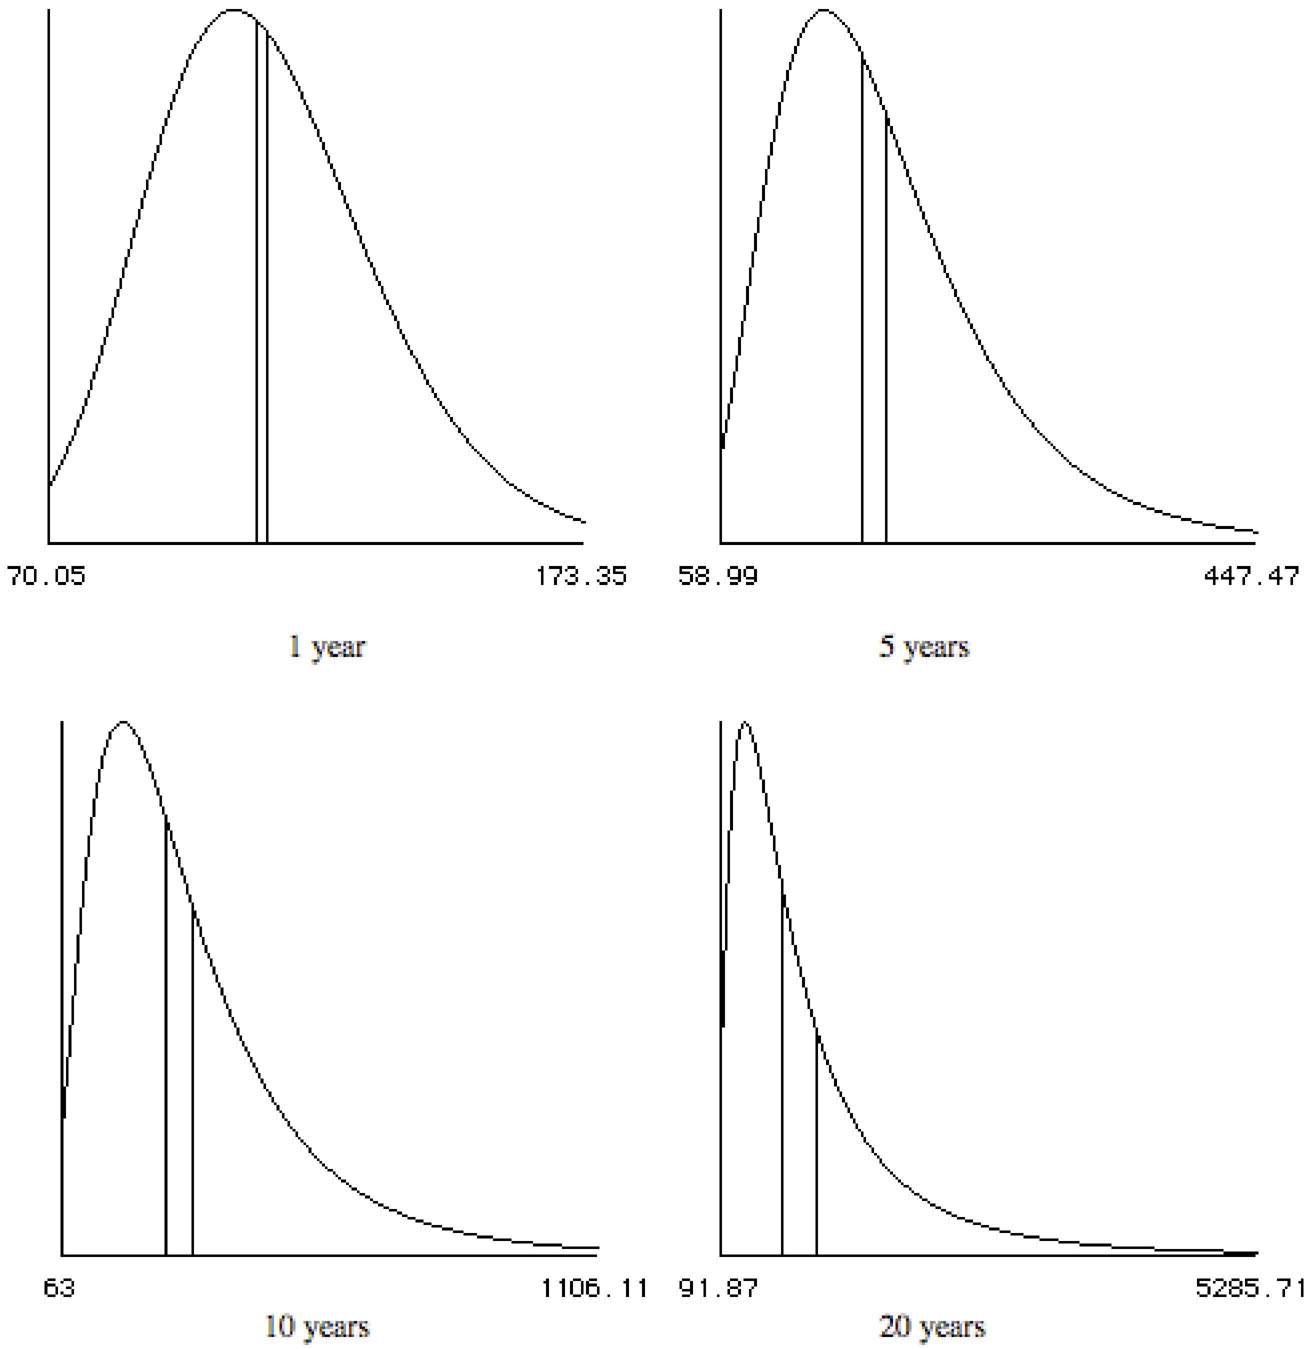

Figure 2: S&P 500 Density Functions

6.4 Cumulative Density Functions

The most useful thing we can do with our model is graph the cumulative density function. Figure 3 shows the cumulative density function for our S&P 500 model for 5 years.

For example, what is the probability that a $100 investment in the S&P 500 model will grow to at least $170 after 5 years? An estimate which is accurate to a percent or two can be read from the graph – it's about 46%.

Similarly, we can see that the probability that we'll lose money in our model over 5 years is about 13%, and that the median ending value is about $162.

As a final example, we can say that the probability of the ending value falling between $92.98 and $283.90 is 80%. Note that our graph is cut off at the 10th and 90th percentiles. We have a 1 in 10 chance of losing more than $7.02, and we have a 1 in 10 chance of making a profit of more than $183.90 (remember that our initial investment was $100, so this is a handsome profit).

A nicely drawn cumulative density graph is a useful tool for financial analysis and planning. With one simple graph we can get a good feeling for how an investment might perform over a given period of time and how volatile it is. We can use the graph to answer a variety of specific questions about the likelihood of various outcomes. Indeed, in a very real sense, the cumulative density function for an investment completely defines the investment for all practical purposes. Thus, the ability to use a computer program to quickly graph and visualize this function is a very practical and useful tool.

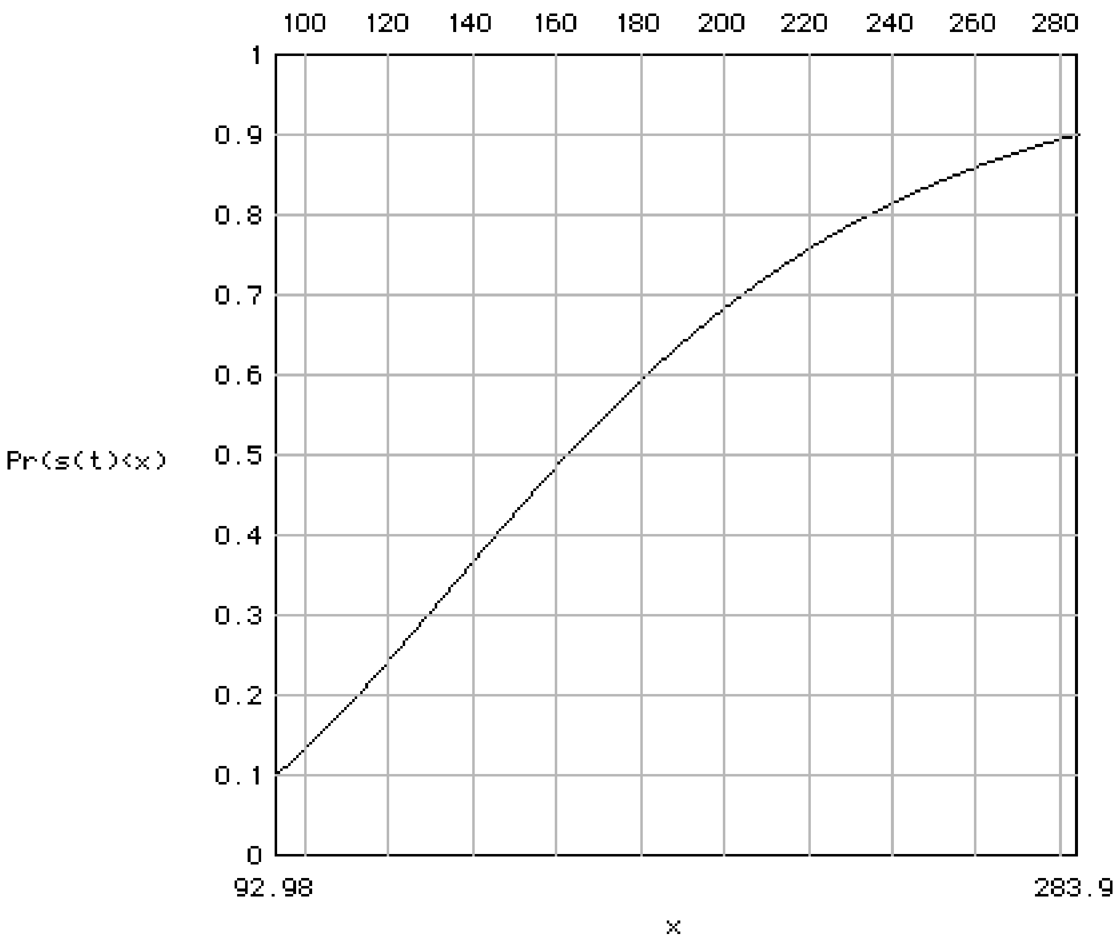

Figure 3: S&P 500 5 Year Cumulative Density Function

6.5 Computing the Density and Cumulative Density Functions

It is quite easy to use computer programs to graph the density and cumulative density functions.

The ending value $s(t)$ after t years is $LN[\log s_0 + \mu t, \sigma^2 t]$. By Proposition 8 in reference [8], the density function is:

$$\frac{1}{x\sigma\sqrt{2\pi t}} e^{-(\log(x/s_0)-\mu t)^2/2\sigma^2 t}$$

To graph the cumulative density function, we first note that:

$$\begin{aligned}
\text{Prob}(s(t) < k) &= \text{Prob}\left(s_0 e^{\sigma X + \mu t} < k\right) \\
&= \text{Prob}\left(X < \frac{\log(k/s_0) - \mu t}{\sigma}\right) \quad \text{where } X \text{ is } N[0,t] \\
&= \text{Prob}\left(\sqrt{t}Y < \frac{\log(k/s_0) - \mu t}{\sigma}\right) \quad \text{where } Y \text{ is } N[0,1] \\
&= \text{Prob}\left(Y < \frac{\log(k/s_0) - \mu t}{\sigma\sqrt{t}}\right)
\end{aligned}$$

Computer code to compute the cumulative density function $N(x)$ for the normal distribution $N[0,1]$ is readily available. With the help of this function, the cumulative density function for the ending price $s(t)$ is computed as:

$$N\left(\frac{\log(k/s_0) - \mu t}{\sigma\sqrt{t}}\right)$$

It is also quite easy the graph these functions in the Microsoft Excel spreadsheet program, using its own built-in NORMSDIST function.

Indeed, all the computation we've done in this section is trivial, and can easily be done in a good spreadsheet program like Excel.

6.6 The Notion of "Average Return"

The phrase "average return" appears frequently in the financial literature. "Average return" can mean several different things, as we discussed in section 5. This can be confusing, so it merits even more discussion.

Note once again the median and mean ending values for 1 year in our S&P 500 model, \$110.19 and \$112.30 respectively. These are equivalent to median and mean simply compounded yearly rates of return of 10.19% and 12.30%. These are numbers one hears quoted frequently in the financial literature, both academic and popular, often stated in terms such as "12.30% is the average annual return for the S&P 500" and "10.19% is the average annually compounded return for the S&P 500" or "10.19% is the average annualized return for the S&P 500."

12.30% is the simple average of the annual returns, or the arithmetic mean of the returns. 10.19% is the geometric mean of the returns.

With a discrete time series $R = r_1 \ldots r_n$ of simply compounded yearly return data, these values are computed as follows:

$$\text{Arithmetic mean} = \frac{1}{n}\sum_{i=1}^{n} r_i$$

$$\text{Geometric mean} = \sqrt[n]{\prod_{i=1}^{n}(1+r_i)} - 1 = e^{\frac{1}{n}\sum_{i=1}^{n}\log(1+r_i)} - 1 = e^{\mu} - 1$$

where μ is the average continuously compounded rate of return.

Why are the mean and the median different? The reason is the volatility of the S&P 500. Whenever the variance of a time series is greater than 0, the geometric mean is less than the arithmetic mean, and the larger the variance, the greater the difference.

To get a better feeling for why there's a difference, simplify the problem and assume for the moment that $\mu = 0$ and $\sigma > 0$. Consider what happens to \$100 with continuously compounded rates of return of +10% and -10%. In the first case our \$100 grows to $\$100e^{0.1} = \$110.52$, for a profit of \$10.52 and a simply compounded rate of return of 10.52%. In the second case our \$100 shrinks to $\$100e^{-0.1} = \$90.48$, for a loss of \$9.52 and a simply compounded rate of return of -9.52%. The profit is larger than the loss. We are assuming that the continuously compounded rates of return are normally distributed about a mean of 0. So it makes sense that the average dollar profit we enjoy when returns are positive exceeds the average dollar loss we suffer when returns are negative. Thus, even though the average rate of return is 0, and the median profit/loss is 0, the average profit is positive, and the mean is larger than the median.

Most people are more comfortable using simple compounding to measure yearly returns. Mathematicians find it more convenient to use continuous compounding, and for the S&P 500 model this is the parameter μ. When rates of return are measured using continuous compounding, the arithmetic and geometric means are the same.

To summarize, in our S&P 500 model we are currently using four notions of the "average annual return:"

- Continuously compounded average annual return $= \mu = 9.7070\%$
- Instantaneous rate of return $= \alpha = \mu + \frac{1}{2}\sigma^2 = 11.6035\%$
- Geometric mean annual return = median annual return = average annualized return $= e^{\mu} - 1 = 10.1938\%$ (simple compounding)
- Arithmetic mean annual return = mean annual return = average annual return = "expected" annual return $= e^{\alpha} - 1 = 12.3035\%$ (simple compounding)

These returns are all related to each other and can be computed from each other using our equations in section 5.

In our S&P 500 model, prices follow the random walk (Ito process) given by the following two equivalent stochastic differential equations:

$$
\begin{aligned}
\frac{ds}{s} &= e^{0.097070 \times dt + 0.194756 \times dX} - 1 \quad \text{where } dX \text{ is } N[0, dt] \\
\frac{ds}{s} &= 0.116035 \times dt + 0.194756 \times dx \quad \text{where } dX \text{ is } N[0, dt]
\end{aligned}
$$

6.7 The Fallacy of Time Diversification

Financial analysts, columnists, and other pundits are fond of noticing that as time increases, the standard deviation of the annualized return decreases. This is quite true, but it's misleading. They cite this as evidence that the uncertainty of investing in volatile assets like stocks decreases as the time horizon increases. We have seen quite clearly, however, that the uncertainty of the ending value increases rapidly with time. For an investor concerned about the value of his investment at some time in the future, it is the ending value and its uncertainty that is important, not the annualized return. The volatility of the annualized return may decrease with time, but because of compounding, the volatility of the dollar ending value, which is what is important, increases with time. This misconception, which is nearly ubiquitous, is called "the fallacy of time diversification."[5]

The mathematics of this fallacy is easy to see given what we know about probability. Suppose the continuously compounded yearly rate of return on an investment has standard deviation σ. Consider the investment over a period of t years. Let r_i be the rate of return during year i. Assume the r_i are independent and each has standard deviation σ. The average return is $\frac{1}{t}\sum r_i$. The total return is $\sum r_i$.

The standard deviation of the average return is:

$$\begin{aligned}\sqrt{\mathrm{Var}\left(\frac{1}{t}\sum r_i\right)} &= \sqrt{\frac{1}{t^2}\mathrm{Var}\left(\sum r_i\right)}\\ &= \sqrt{\frac{1}{t^2}t\sigma^2}\\ &= \sigma/\sqrt{t}\end{aligned}$$

The standard deviation of the total return is:

$$\begin{aligned}\sqrt{\mathrm{Var}\left(\sum r_i\right)} &= \sqrt{t\sigma^2}\\ &= \sqrt{t}\sigma\end{aligned}$$

It's the total return that matters in terms of the bottom line. This determines how much money we have at the end of t years. The standard deviation of the average return decreases as t increases, but we don't care about that. We care about the standard deviation of the total return, which increases as t increases. This is the confusion that lies at the heart of the fallacy.

[5] For a more extended discussion of this fallacy, see Chapter 8, Appendix C of reference [2].

In portfolio theory, when we diversify a portfolio by mixing together various kinds of assets and asset classes, we learn that if the portfolio components are not perfectly correlated, we can decrease the standard deviation of the expected yearly return.

Novices and people who haven't thought about the problem clearly think that this same phenomenon can be achieved in the time dimension. They argue that as we increase the length of time over which we invest, the standard deviation of the annualized return decreases. Hence, "time diversification." We have seen the flaw in this argument, however. As time increases, the standard deviation of the total return increases, and that's what counts. Hence the "fallacy of time diversification."

6.8 Risk Over the Long Run

The fallacy of time diversification is very widespread. Most people seem to believe that the risk of investing in volatile assets like stocks decreases the longer your time horizon. Indeed, if you read the financial press, popular books on investing, or the literature produced by the mutual fund and other investment management companies, you hear this argument constantly. This problem deserves further exploration.

Zvi Bodie, a Finance professor at Boston University, has shown that with at least one reasonable definition of the notion of "risk," it is possible to prove that, contrary to popular opinion, the risk of investing in volatile assets like stocks increases with time.[6]

Bodie's argument is quite simple. It's reasonable to think of "risk" as "the cost of insuring against earning less than the risk-free rate of interest." This insurance is just a simple put option, and the Black-Scholes equation for option pricing shows that the put option increases in value with time to expiration.

Here's the argument in detail:

Suppose the risk-free rate is r. Let S be the current value of our portfolio. Assume our portfolio's continuously compounded yearly return has a standard deviation of σ. Let E be the future value of S after t years of earning interest at rate r. This is how much we'd have after t years if we just put our money in the bank. Note that $E = Se^{rt}$.

We have two alternatives: We can keep our money in our portfolio, or we can sell our portfolio and put all of our money in the bank. If we keep our money in the portfolio, it is reasonable to think of the risk we run being not doing as well as we would have if we had put the money in the bank instead.

Thus it makes sense to measure the "risk" of our portfolio as the cost of an

[6] See reference [1].

insurance policy against ending up with less than E dollars after t years. This imaginary insurance policy would make up any difference if our portfolio did poorly and failed to earn at least the risk-free rate of return. This insurance policy is nothing more or less than a European put option on our portfolio with strike price E and time to expiration t. So we can use the Black-Scholes equation to get the value of the put option. If an insurance company were to offer us such a policy, this is how much it would cost.

Let P be the value of this put option and let C be the value of the corresponding call option with the same expiration date and strike price. Then by the put-call parity theorem,[7]

$$P = C + Ee^{-rt} - S = C + (Se^{rt})e^{-rt} - S = C + S - S = C$$

Thus the put and the call have the same price, so we can use the Black-Scholes equation for the value of the call to get the value of the put. In the Black-Scholes equation below N is the normal distribution cumulative density function.

$$\begin{aligned} P &= C \\ &= SN\left(\frac{\log(S/E) + (r + \frac{1}{2}\sigma^2)t}{\sigma\sqrt{t}}\right) - Ee^{-rt}N\left(\frac{\log(S/E) + (r - \frac{1}{2}\sigma^2)t}{\sigma\sqrt{t}}\right) \\ &= S\left[N(\frac{1}{2}\sigma\sqrt{t}) - N(-\frac{1}{2}\sigma\sqrt{t})\right] \end{aligned}$$

The expression in square brackets is the area under the normal density function (the bell curve) from $-\frac{1}{2}\sigma\sqrt{t}$ to $+\frac{1}{2}\sigma\sqrt{t}$. This value obviously increases as t increases. So the cost of the insurance, which we've taken to be the risk of the investment, increases with the time horizon, and the proof is complete.

As an example, in Figure 4 we graph risk as a function of time horizon using Bodie's definition of risk for our S&P 500 model, using our estimated volatility $\sigma = 0.194756$ and a portfolio value of $S = \$100$. Note that as $t \to \infty$, the cost of the insurance = the value of the put $P \to S$.

Note also that the risk-free rate r has canceled out of the equation.

For a discussion of this topic of risk and time aimed at a broader non-technical audience, see reference [10].

[7] See reference [9].

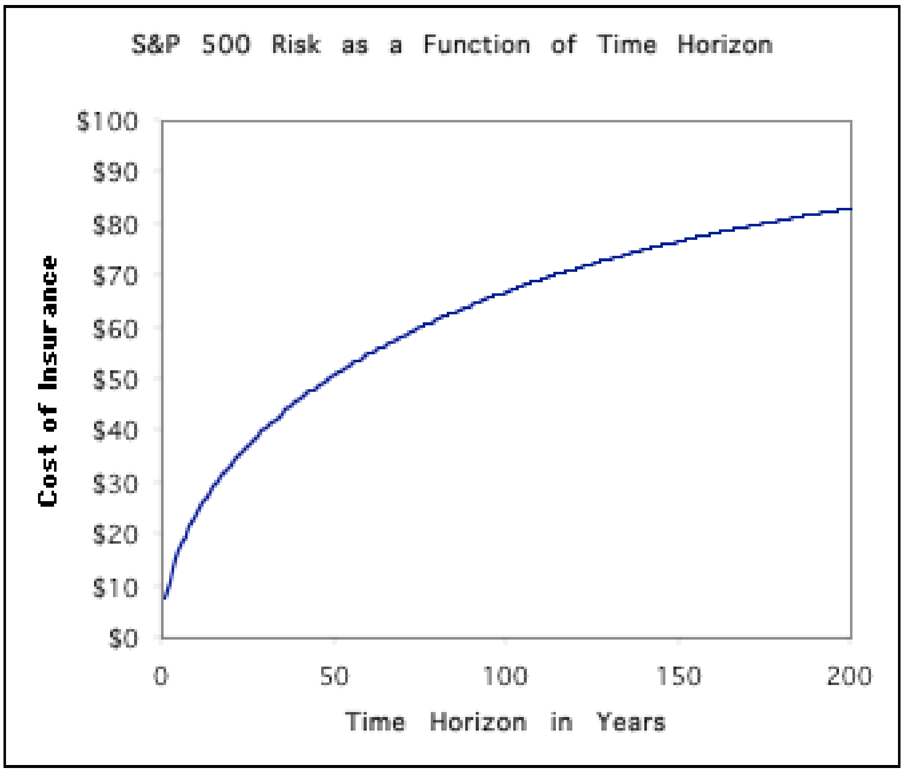

Figure 4: S&P 500 Risk as a Function of Time Horizon

7 Random Walks and Real Markets

How well does the random walk model actually describe how real-life financial markets work?

We can attempt to at least get some insight into this question. After all, our model predicts certain properties of the data. We can compare the predicted values against the actual data and see how well they fit.

As a simple example, our S&P 500 model predicts values for the mean and standard deviation of the yearly simply compounded returns. We can also compute these numbers from the original time series data. Here's the results:

	Model	Data
Mean	12.30%	12.16%
Standard deviation	22.08%	20.35%

The numbers seem to be close. This isn't surprising, since we used the time series data to estimate parameters for the model, so we should expect them to be close.

As another example, our S&P 500 model predicts a cumulative probability density function for the ending values after one year. We can compute the same function using the raw S&P 500 return data and compare them. Figure 5 shows the result, with the smooth curve the model and the "jagged" one the data. Once again, the model and the data appear to be close.

What do these experiments we've performed attempting to "fit" the model to the data have to say, if anything, about how well our simple random walk model describes actual financial markets? The answer is unfortunately "not much." It's clear that the model and the data are not wildly dissimilar, but this is far from being enough to conclude that the model is correct. It's not even clear how one might define "correct!" Answering these questions requires the use of statistical tools which are well beyond the scope of the simple results we have presented in this paper.

Alas, the simple random walk model is in some disrepute these days. Econometricians have in fact rejected the Random Walk Hypothesis.[8]

In section 3 we made three strong assumptions about returns over time. We assumed that they were (1) independent, (2) identically distributed, and (3) have finite variance. All three assumptions are flawed.

The first flawed assumption is number (1), that returns are independent. They do not appear to be fully independent in real markets. We give two examples of this, one for short-term returns and one for long-term returns.

[8]See [5] for a technical statistical treatment of this problem. For a non-technical discussion see [6].

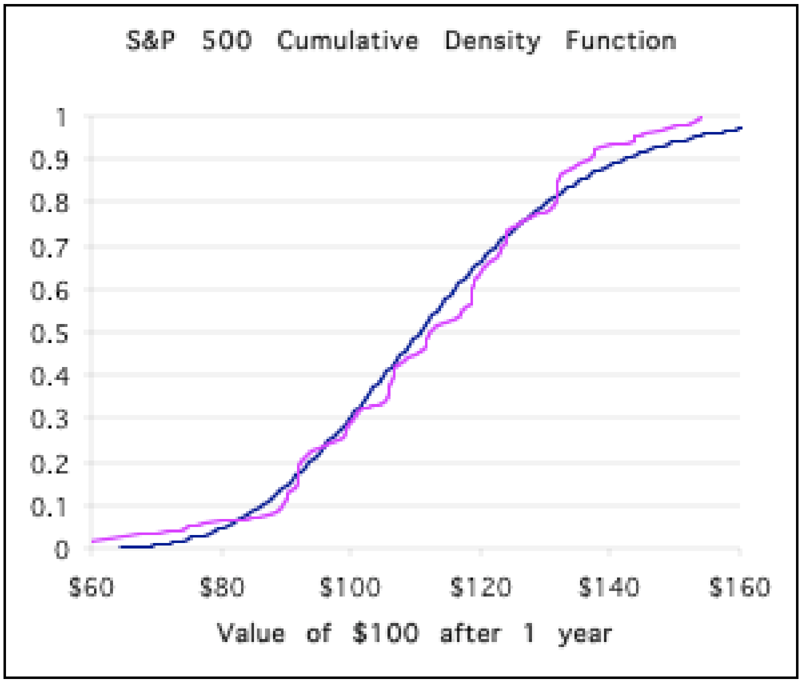

Figure 5: S&P 500 Cumulative Density Function: Theory vs. Reality

There is a small amount of positive serial correlation in short-term returns. This is called "market momentum." For example, if the market goes up one day, it has a slightly greater probability than that predicted by the model of going up the next day too. Before we get excited and try to exploit this with the obvious strategy, we should note that this effect turns out to be so small that any profit made by exploiting it is usually overwhelmed by the costs of trading the stocks we must buy and sell every day to implement the strategy!

Some (but not all) researchers also believe that there may be some amount of negative serial correlation in long-term returns. This phenomenon is called "reversion to mean." This is the opposite of the short-term momentum effect. Over long periods of time, if prices revert to mean, markets which have done better than average in the past tend to do worse than average in the future, and vice-versa. Unfortunately, nobody has figured out a way to tell how long these periods last or how to exploit the phenomenon. (Or perhaps someone has figured this out, in which case she must be off by herself somewhere making piles of money and keeping the secret to herself!) In addition, it is quite difficult to measure this effect because we simply do not have enough very long-term data to analyze reliably.

Thus over the short term actual markets are somewhat more volatile than predicted by the model, and over the long term they may be somewhat less volatile.

The second flawed assumption is number (2), that returns are identically distributed. In particular, there is no good reason to believe that expected returns are constant over long periods of time. Expected returns for risky assets like stocks are a premium that investors demand as compensation for the risk of the investments. The risk of the world in general and of the financial markets in particular is not constant over time. During some periods the world is perceived to be safer, while during other periods it is perceived to be riskier. It is a natural conjecture that expected returns should vary over time – they should be higher during riskier periods and lower during safer periods. It is also possible that this phenomenon of time-varying risk premia is related to the reversion to mean phenomenon discussed above.

Standard deviations are also not constant over time. Markets appear to alternate betweens periods of relative calm with low volatility and periods of relative instability with high volatility. This phenomenon is called *heteroskedasticity.*

Even assumption (3) is flawed – the assumption that the variance of returns is finite. Benoit Mandelbrot and other researchers believe that returns most likely have a "fractal" or "stable Paretan" probability distribution. One of the properties of these kinds of distributions is that they have infinite variance! These kinds of distributions also have the "fat tails" that we notice in the historical data – very large losses and very large gains are more likely than in the lognormal model.[9] Investing in volatile assets like stocks is even more risky in these models than in the lognormal model.[10]

William Sharpe is a Nobel prize-winning economist at Stanford University. On his web site he has a retirement planning worksheet which uses a random walk model similar to ours.[11] He says in the instructions for his worksheet that "it uses a very simple model of investment returns and hence should be considered only illustrative of more sophisticated methods that can be employed for this task." Unfortunately, Professor Sharpe neglects to mention what these more sophisticated methods might be.

For all of these reasons, we should consider the random walk model and the computational techniques presented in this paper to be at best a rough approximation to how financial markets actually behave. We should use this model with caution to do real-life financial planning.

[9] The discerning reader may notice a faint suggestion of the possibility of a "fat tail" for large losses in figure 5 above.

[10] An interesting discussion of these kinds of distributions can be found in Eugene Fama's doctoral dissertation [4].

[11] See reference [12].

References

[1] Zvi Bodie. On the risk of stocks in the long run. *Financial Analysts Journal*, May-June 1995.

[2] Zvi Bodie, Alex Kane, and Alan J. Marcus. *Investments*. Irwin McGraw-Hill, fourth edition, 1999.

[3] Zvi Bodie and Robert C. Merton. *Finance*. Prentice-Hall, preliminary edition, 1998.

[4] Eugene F. Fama. Random walks in stock market prices. *Financial Analysts Journal*, Sep-Oct 1965.

[5] Andrew W. Lo and Craig MacKinlay. *A Non-Random Walk Down Wall Street*. Princeton University Press, 1999.

[6] Burton G. Malkiel. *A Random Walk Down Wall Street*. W. W. Norton, sixth edition, 1996.

[7] Robert C. Merton. *Continuous-Time Finance*. Blackwell, 1990.

[8] John Norstad. The normal and lognormal distributions. http://www.norstad.org/finance, Feb 1999.

[9] John Norstad. The put-call parity theorem. http://www.norstad.org/finance, Mar 1999.

[10] John Norstad. Risk and time. http://www.norstad.org/finance, Oct 2003.

[11] John Norstad. Random walker. http://www.norstad.org/finance, Feb 2005.

[12] William F. Sharpe. Retirement worksheet. http://www.stanford.edu/~wfsharpe/ws/ws_ret.htm.