What Practitioners Need to Know . . .

. . . About Lognormality

Mark Kritzman
Windham Capital Management

When reading the financial literature we often see statements to the effect that a particular result depends on the assumption that returns are *lognormally* distributed. What exactly is a lognormal distribution, and why is it relevant to financial analysis? In order to address this question, let us start with a review of logarithms.

Logarithms

A logarithm is simply the power to which a base must be raised to yield a particular value. For example, the exponent 2 is the logarithm of 16 to the base 4, because 4 squared equals 16. The logarithm of 8 to the base 4 equals 1.5, because 4 raised to the power 1.5 equals 8.

The choice of a base depends on the context in which we use logarithms. For simple mathematical procedures, it is common to use the base 10, which explains why logarithms to the base 10 are called common logs. The base 10 is popular because the logarithms of 10, 100, 1000 and so on equal 1, 2, 3, . . . , respectively.

Why should we care about logarithms? In the days prior to pocket calculators (long before my time), logarithms were useful for performing complicated computations. Financial analysts would multiply large numbers by summing their logarithms, and they would divide them by subtracting their logarithms. For example, given a base of 4, we can multiply 16 times 8 by raising the number 4 to the 3.5 power, which is the sum of the logarithms 2 and 1.5. Of course, you might argue that it would have been easier to multiply large numbers directly than to raise a base to a fractional power. In the olden days, however, an analyst would use a slide rule, which is a ruler with a sliding central strip marked with logarithmic scales.

e

In most financial applications, instead of logarithms to the base 10, we use logarithms to the base 2.71828, which is denoted by the letter e in honor of the famous Swiss mathematician Euler. These logarithms, which are called natural logs and are abbreviated as ln, have a special property. Suppose we invest $100 at the beginning of the year at an annual interest rate of 100%. At the end of the year we will receive $200—our original principal of $100 and another $100 of interest. Now suppose our interest is compounded semiannually. Our year-end payment will equal $225. By the middle of the year we will have earned $50 of interest, which is then reinvested to generate the additional $25.

In general, we can use the following formula to compute the year-end value of our investment for any interest rate and for any frequency of compounding:

$$E = B \cdot (1 + r/n)^n$$

where

E = ending value,
B = beginning investment,
r = annual interest rate and
n = frequency of compounding.

If the 100% rate of interest is compounded quarterly, our $100 investment will grow to $244.14 by the end of the year. If it is compounded daily, we will receive $271.46. And if it is compounded hourly, we will receive $271.81 by year-end.

It seems as though the more frequently our interest is compounded, the more money we will end up with at the end of the year. No matter how frequently it is compounded, though, we will never receive more than $271.83. The limit of the function $(1 + r/n)^n$, when r equals 100%, as n approaches infinity is 2.71828, the base of the natural log.

We can use this result to convert periodic rates of return into continuous rates of return. A periodic rate of return is computed as the percentage change of our investment from the beginning of the period to the end of the period, assuming there are no contributions or disbursements. A continuous rate of return assumes that the income and growth are compounded instantaneously.

From our previous example we know that e, the base of the natural log, raised to the power 1 (the continuous rate of return in our example) yields 1 plus 171.83% (the periodic rate of return in our example). The natural log of the quantity 1 plus the periodic rate of return must therefore equal the corresponding continuous rate of return. For example, the natural log of the quantity 1 plus 10% equals 9.53%. This means that if we invest $100 at a continuously compounded rate of return of 9.53%, our investment will grow to $110. The value 1.10, which we compute by raising e to the power 0.0953, is called the exponential. These relationships are shown below:

$$\ln(1.10) = 0.0953,$$

$$1.10 = 2.71828^{0.0953}.$$

Figure A Continuously Compounded Return as a Function of Time

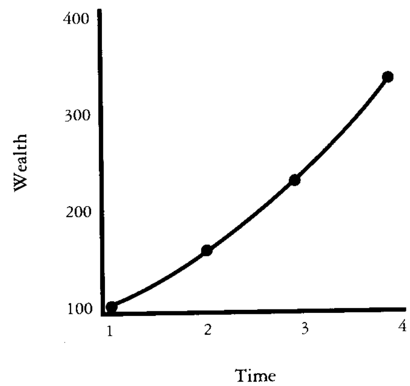

We can also compute the continuous rate of return within a period by subtracting the natural log of the beginning value from the natural log of the ending value. Suppose we start with $100, which is invested so that it grows to $150 after one year, $225 after two years and $337.50 after three years. The logarithms of these values equal 4.6052, 5.0106, 5.4161 and 5.8216, respectively. The difference between each logarithm and the next one equals the continuously compounded return each year, which is 40.55%. This continuous return corresponds to a yearly periodic return of 50%.

If we plot these values as a function of time, we produce a convex curve, as shown in Figure A. The logarithms of these values, however, form a straight line when plotted as a function of time, indicating a constant periodic rate of growth. It is this relation that gives rise to the logarithmic scale, in which equal percentage changes correspond to equal vertical distances. This scale is shown on the left axis in Figure B, with the logarithms shown along the right axis.

Figure B Logarithms of Returns as a Function of Time

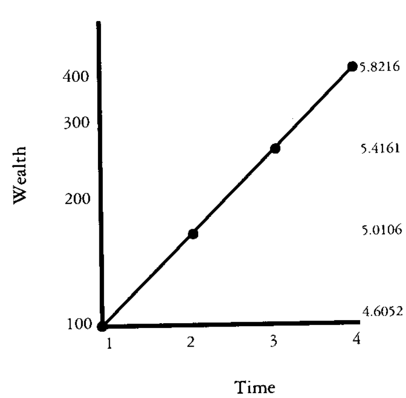

Why Are Returns Lognormally Distributed?

Suppose we invest $100 for one year and suppose that the quarterly returns during this period are 10%, −5%, 15% and −10%. Although the sum of these quarterly returns equals 10%, our $100 investment grows, not to $110, but to $108.16 ($100 · 1.10 · 0.95 · 1.15 · 0.90), yielding a one-year cumulative return of 8.16%. The sum of the quarterly returns does not measure the actual cumulative return we realize. It is the compounding of these quarterly returns that yields the cumulative return.

We can sum the natural logs of the quantities 1 plus the quarterly returns to determine the cumulative return. These natural logs equal 0.0953, −0.0513, 0.1398 and −0.1054, respectively, and they sum to 0.0784. If we raise e to the power 0.0784, we get the exponential 1.0816, which equals 1 plus the one-year cumulative rate of return.

What has all this to do with lognormality? We have seen that the sum of the logarithms of the quantities 1 plus the periodic returns equals the logarithm of the quantity 1 plus the cumulative return. The important distinction is that we sum the natural logs to derive the cumulative return, whereas we multiply 1 plus the periodic returns to derive the cumulative return.

The Central Limit Theorem, which is one of the most important notions of statistical inference, deals with the summation of random variables. A normal distribution, which is symmetric, can be described by its first two central moments—its mean and its variance.[1] Higher moments, such as skewness, have a value of 0 in a normal distribution. Skewness is computed by raising a value to an odd-number exponent; unlike variance, which is computed by squaring a value, skewness can take on either a positive or a negative value. When a large number of random variables are summed, the central moments that are computed with odd-number exponents tend to cancel each other out, leaving mean and variance as the only non-zero central moments.

Because we sum logarithms, the natural logs of the quantities 1 plus the periodic returns are normally distributed. Because these natural logs are normally distributed, and because the exponential of the normal distribution gives the lognormal distribution, the quantities 1 plus the periodic returns, which are the exponentials of the natural logs, are lognormally distributed.

Normal Distribution and Lognormal Distribution

Having said all that, why should we care? Table I shows the estimated probabilities of achieving

Table I Probability of Achieving Target Returns

	One Year		*Five Years*		*10 Years*	
Target Returns	*N*	*L*	*N*	*L*	*N*	*L*
−5%	80%	79%	99%	97%	100%	100%
0%	73	71	96	90	100	96
8%	58	57	74	66	93	72
20%	34	37	5	22	0	14

various annualized holding-period returns for an investment with a geometric mean return of 12% and a standard deviation of 20%.[2]

Given a one-year horizon, it does not make very much difference whether we assume a normal distribution or a lognormal distribution to estimate probabilities.[3] As we extend our horizon, however, the normal distribution overestimates the probability of achieving target returns that are below the expected return. For example, given a five-year horizon, the likelihood of achieving at least an 8% annualized return equals 74%, assuming a normal distribution, versus 66% for a lognormal distribution. Over 10 years, the difference is even greater. Based on a normal distribution, there is a 93% chance of achieving at least an 8% return, while under the assumption of a lognormal distribution, the probability falls to 71%. The normal distribution also *under*estimates the probability of achieving target returns above the expected return. Given a five-year horizon and assuming a normal distribution, the chance of achieving a return of 20% or greater equals 5%, versus 22% for a lognormal distribution. Over a 10-year horizon, the probabilities equal 0% and 14% for the normal and lognormal distributions, respectively.

Based on this evidence, we are well advised to note whether a normal or lognormal distribution was assumed in generating the results of interest to us, especially if they pertain to a multi-year horizon.

Footnotes

1. *For a review of the normal distribution, see M. Kritzman, "What Practitioners Need to Know About Uncertainty,"* Financial Analysts Journal, *March/April 1991.*
2. *The geometric mean is computed as:*

$$\left[\prod_{i=1}^{n} (1 + r_i)\right]^{1/n} - 1$$

where

r = periodic return in year i and
n = number of years.

It is a better description of the true average return than is the arithmetic mean, in the following sense. If we earned the geometric return each year, we would have achieved the same terminal wealth that was generated by the actual yearly returns.

3. *The normal deviate for the normal distribution is computed as:*

$$[(1 + T)^n - (1 + r)^n]/(\sigma * \sqrt{n})$$

where

T = target return,
r = geometric return,
σ = standard deviation and
n = number of years.

The normal deviate for the lognormal distribution is computed as:

$$[\ln([1 + T]^n) - \ln(1 + r) \cdot n]/\sigma\sqrt{n}$$