A Tale of Two Density Functions

by Dick Joss

It has been common in modern academic finance to assume that investment returns may be described by the lognormal probability density function. Using this tool it is possible not only to provide an expected rate of investment return, but a complete distribution of such returns. In short, using the tool one could say that the expected return on stock investments might be 12 percent, but that there is a 30-percent chance that your equity investments could exceed a return of 25 percent for the year. On the down side, it is also possible to say that there is a 30-percentchance that your equity investments could lose money for the year.

To select the lognormal probability density function parameters, finance textbooks provide detailed instructions using the arithmetic mean and sample standard deviation from a set of historical returns. What is often missing, however, is a comparison of the actual historical results, and the expected results provided by the lognormal probability density function. This comparison is not as good as one might expect given the widespread use of this particular model. To illustrate this point, the *2008 Ibbotson and Associates SBBI Yearbook* provides of history of 984 months of stock return data. The chart below compares the distribution of the actual data with the expected distribution provided by the "best estimate" lognormal density function.

As an example of the difference between the two distributions, the actual distribution shows that for 118 of the 984 months (12 percent of the total) stock returns were 5.8 percent or more for the month. Whereas the "best estimate" lognormal density function assumes that 189 out of 984 months (19 percent of the total) will have a return that is 5.8 percent or more. This is a substantial difference. It calls into question the use of the basic lognormal probability density function to describe the historical data.

Actual Data Distribution Compared with Lognormal Assumption

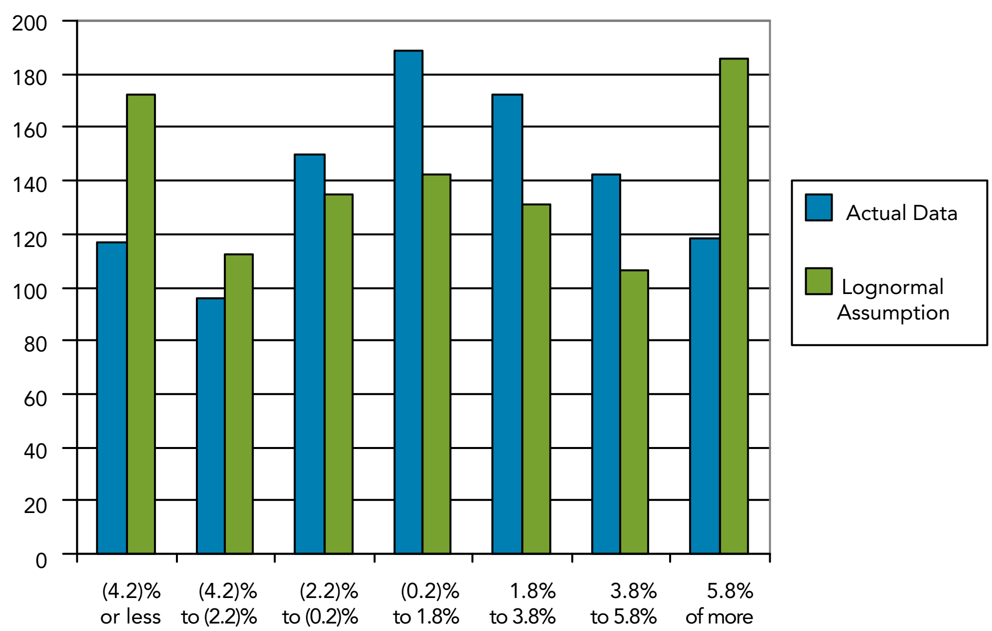

A Tale of Two Density Functions *by Dick Joss*

Actual Data Distribution Compared with Lognormal Assumption

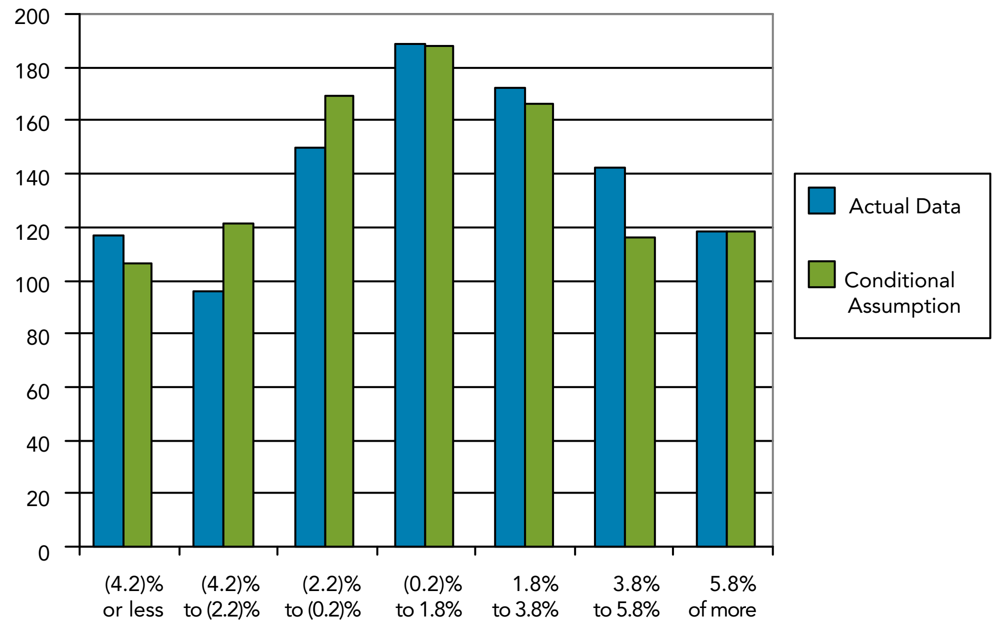

As one contemplates the source of historical investment return data, it is clear that they are periodic observations of a single, long-term historical asset growth. As such, the mathematical theory of probability and statistics would place this single observation at the mean of long-term results, with each of the periodic returns being described by a conditional lognormal probability density function. When this one change is made, the comparison between the actual historical results and those described by the probability density function improves dramatically, as is shown in the following chart.

Not only is the comparison significantly improved, but this one change helps explain the disastrous 401(k) plan results that have been seen. This change in density function causes the best estimate rate of return to change from an arithmetic mean of historical returns to the lower geometric mean of historical returns. Given that employee participants have been led to believe that they would receive the higher arithmetic mean returns, it is not surprising that they are disappointed with the actual geometric mean results.

In addition, the change in probability density functions provides a new explanation for the spectacular collapse of Long-Term Capital Management and the more recent collapses of Bear Stearns and Lehman Brothers Holdings. Until this issue is addressed fully, and corrected, the possibility of similar problems will always be on the horizon, and another “financial crisis” may be just around the corner.

Dick Joss, FSA, PhD is retired and living in Bellingham, Wash. He can be reached at *rrjoss@comcast.net.*