RiskMetrics Group

BenchNote: Practice notes highlighting Wealthbench features and functionality

A New Monte Carlo Simulation Methodology

In response to the unique market events of 2008, RiskMetrics has taken a careful look at the methodology and assumptions behind the Monte Carlo simulation component of WealthBench. We questioned whether we could improve upon the generally accepted approach we have been using. In particular, we were interested in answering the following questions:

- Did the “worst case” simulations in any given year reflect the magnitude of the actual market decline in 2008?
- Given the complexity of real-world markets, could we include additional factors beyond expected risk, return and correlation to make simulations more realistic?

Current Approach

Similar to many Monte Carlo engines, WealthBench assumes that the log normal returns for each asset are normally distributed. For each time period, we convert the arithmetic capital market assumptions into log returns, simulate using a normal distribution, and apply the results to wealth values.

Log returns are widely used for the simulation of investment returns, primarily because they are easy to work with computationally. One characteristic of the lognormal distribution is that it has a positive skew—that is, the upside is bigger than the downside. This skew gets more pronounced as the standard deviation of the distribution gets larger and/or as the time horizon being simulated gets longer. When investment simulation is run across a very short time horizon, (e.g., less than 5 days), this skew is not a factor. When simulation is run across years, not days, the expanding upside becomes a significant issue. As investors have become more concerned with the downside of investment performance, the inherent upside bias of the lognormal distribution seemed more problematic. Therefore, we have explored alternative distribution assumptions that can produce more intuitive results.

Historical Perspective

We began by analyzing historical returns for large cap domestic equities for the period 1871 to 2008 to see how well the lognormal distribution actually fit the historical data and to compare alternative statistical distributions. For these 138 years, the average annual return was 7.87% with a standard deviation of 17.86%.

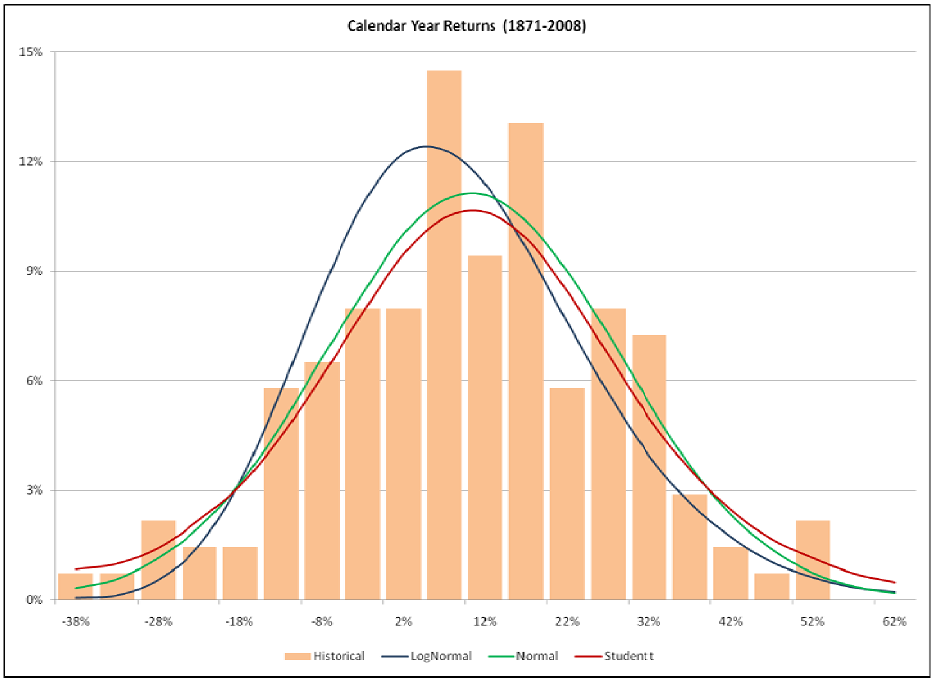

Please refer to the chart, and notice the left part of the data. The lognormal distribution (blue) does not include the lowest historical returns (including 2008). The normal distribution (green) does a slightly better job, but the student-t distribution (red) has the fattest tail of the three—it includes the worst historical returns as well as the best historical returns.

What is a Student-t distribution?

Without going into extensive statistical jargon, a student-t distribution is one of many continuous distributions that are symmetrical but have fatter tails—that is, a slightly larger proportion of random samples drawn from the distribution will be from the lower and upper extremes. The Student-t is widely used in other financial applications for Monte Carlo simulation (RiskMetrics already uses the Student-t as part of its RM2006 risk simulation methodology).

After extensive testing, we have decided to supplement our current log normal distribution with both Normal and Student-t distribution assumptions. While the overall results of wealth simulations will not significantly change, the Student-t distribution addresses the two primary problems with the lognormal methodology. With Student-t distributions, wealth simulations will be less skewed to the upside, and the downside will include returns that are worse, on a relative basis, than those of 2008.

One reason why the lognormal distribution is so widely used for wealth simulation is its characteristic that the one year return can never be less than -100%--meaning that you can't lose more than 100% of your wealth. While it is statistically possible for the normal or student-t methodologies to generate an annual loss of greater than 100%, it will only happen in cases with relatively high standard deviation (With normal distribution, only 1% of returns fall beyond the truncation level at 43% standard deviation; with student-t, only 1% of returns fall beyond the truncation level at 34% standard deviation). In order to make the simulations conform to reality, we have placed an upper and lower limit of the results generated by the simulations so that investment loss cannot be greater than 100% in any one year and your investment gain cannot be greater than 100% plus 2 times the assumed return (if the expected return is 10%, maximum one year gain would be 120%). In the vast majority of cases, these limits will not come into play, but we have included them in order to prevent results that are statistically possible, but practically impossible.[1]

---

[1] With a normal distribution and a standard deviation of 42%, less than 1% of results would be truncated. With a student-t distribution, 1% of results are truncated at 30% standard deviation. Since the tails of the student-t distribution are naturally fatter, we would expect the same truncation to occur at a lower standard deviation. In neither case should the truncation materially impact the simulation results.

Here are three wealth simulations generated with the current lognormal, truncated normal and truncated student-t distributions-- each assumes an annual return of 7% and a standard deviation of 15%.

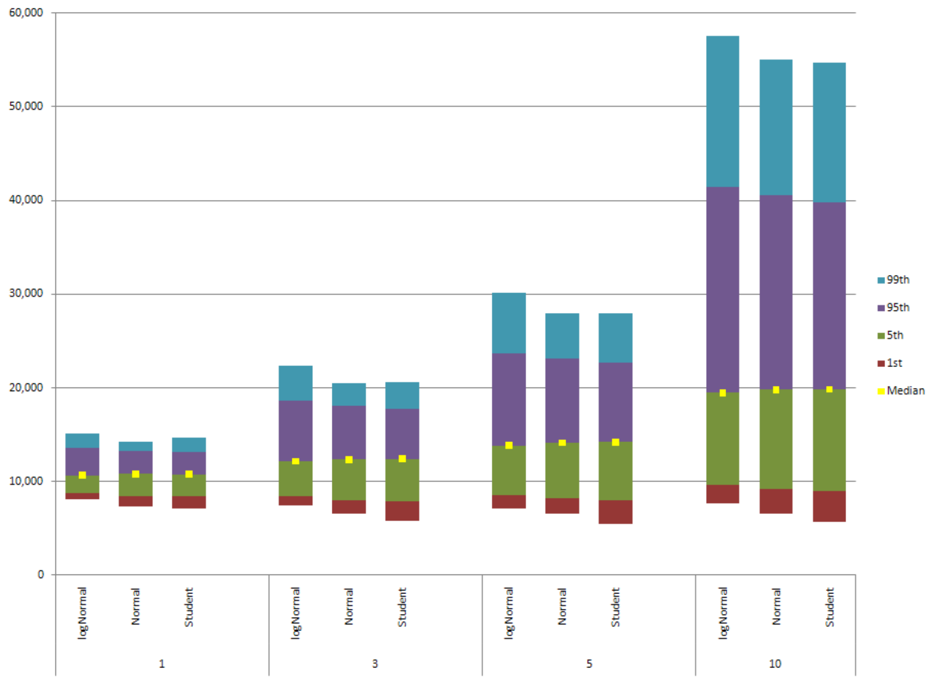

Notice how the medians remain relatively consistent across distribution models; and how, over time, some of the upside bias inherent in the lognormal distribution has been reduced.

New approach to volatility

Our second fundamental change to the Monte Carlo methodology addresses the fact that volatility is not constant over time—that volatility itself has volatility—and that highly volatile periods tend to cluster together. In other words, when the market has a particularly volatile period--like 2008—it does not immediately revert back to normal. Since we will be including more extreme events in the simulations, we felt it was appropriate to address the after-effects of a fat-tail event.

There are various methodologies to model this characteristic—technically referred to as heteroskedasticty. The technique used by RiskMetrics is known as GARCH (Generalized AutoRegressive Conditional

Hesteroskedasticity). While this sounds impossibly technical, it basically means that volatility for a given period is, in part, dependent on past volatility, that it has a statistical “memory”. As a result, volatility will change over time.

Combining Fat Tails and GARCH

GARCH can be added to both the normal and student-t simulation to provide fatter tails. The chart below provides a sense of how much fatter the one-year tails can be when using any combination of new distributions. With a normal distribution, a +/- 2 standard deviation event should occur 2.5% of the time. With the current lognormal simulation, -2 events are only occurring 1.16% of the time. +2 events occur over 3% of the time and no -3 events are captured at all. The new distributions capture events as outsized as +/- 5 standard deviations (albeit once in over 10,000 simulations)

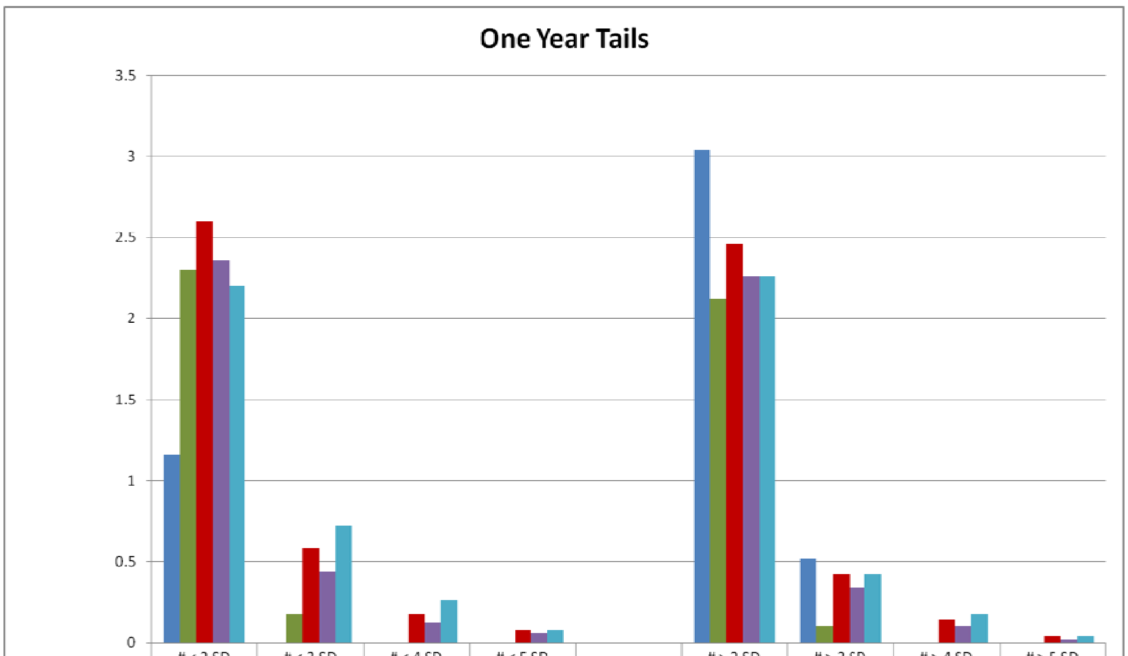

	# < 2 SD	# < 3 SD	# < 4 SD	# < 5 SD		# > 2 SD	# > 3 SD	# > 4 SD	# > 5 SD
LogNorm	1.16	0	0	0		3.04	0.52	0	0
Normal	2.3	0.18	0	0		2.12	0.1	0	0
Student	2.6	0.58	0.18	0.08		2.46	0.42	0.14	0.04
Normal Garch	2.36	0.44	0.12	0.06		2.26	0.34	0.1	0.02
Student Garch	2.2	0.72	0.26	0.08		2.26	0.42	0.18	0.04

An illustration of a wealth simulation is shown below that includes all of the available distributions. Notice that on a one-year basis, all of the new distributions generate more "worst-case" results than the current lognormal approach. Over time, these fat tails diminish, but are still larger than at present

By design, the median end wealth values will not change by a significant amount using the new methodology. The "best" and "worst" cases will be more noticeably different, particularly if displaying 1st and 99th percentile results. Because the multi-period optimization is based on median end wealth values, the allocations and associated simulated wealth values are not expected to significantly change when running a Custom proposal under the new methodology.

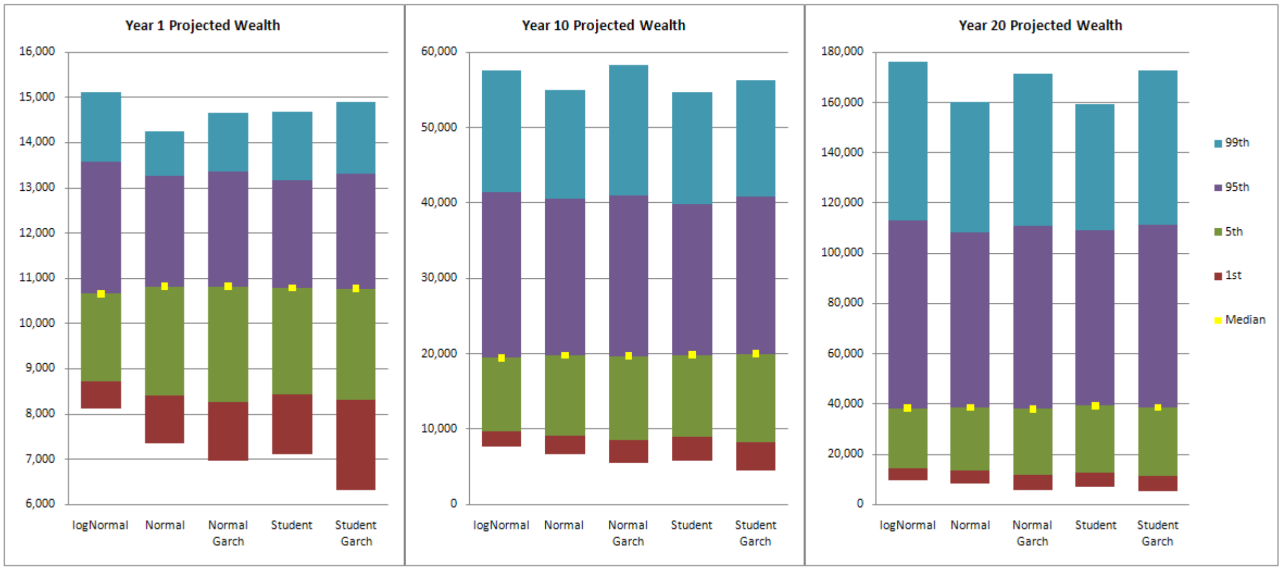

Conclusion

- The current lognormal methodology will still be available.
- Clients will be able to replace the lognormal with normal or student-t distribution, with or without the GARCH component
- The purpose of the new methodology is to include more extreme returns, particularly negative returns.
- There is historical and quantitative justification for replacing the lognormal assumption with a normal assumption.
- Existing simulation results or multi-period optimization results are not expected to materially change if clients use the new distributions/ or if the new distributions are used—median, 25th and 75th percentile results should barely change—only values at or greater that the 95th percentile or at or less than the 5th should be different.
- We will be able to generate comparative results under each methodology in a specific tenant to assist clients in evaluating if and how to implement the new process.