SOCIETY OF ACTUARIES

Investment Section

INVESTMENT FALLACIES 2014

Simulation of Long-Term Stock Returns: Fat-Tails and Mean Reversion

By Rowland Davis

Following the 2008 financial crisis, discussion of fat-tailed return distributions seemed rampant. Although I have no hard data, my impression is that many actuarial and investment consulting firms made sure their simulation engines incorporated some feature that creates fat-tails (e.g. regime switching models, stochastic variance models). Mean reversion, on the other hand, seems to get very little attention. While mean reversion itself may not be treated as a "myth", most model builders avoid any mean reversion feature in their simulations of stock returns. At least implicitly, when these model builders choose not to add a mean reversion feature to their model, they are also making a case that the resulting distributions of long-term returns will be as good without mean reversion as they would be if mean reversion were incorporated. I guess this illustrates what I will call the myth of unimportance. This essay offers some thoughts about why this may be, and then shows how important mean reversion is when simulating long-term stock returns.

Why the limited use of mean reversion? I suspect the main reason is that academic financial economists have not made a strong enough case for mean reversion. In a survey by Ivo Welch (UCLA and Yale) in 2000, only 36 of 102 surveyed financial economists said that they believed in long-term mean reversion for stock returns (17 had no opinion and 49 did not believe). Without stronger support from academics, model-builders might feel they are "out on a limb" if they incorporate mean reversion. But, arguably, the academics who were asked this question might have answered "do not believe" because no one has been able to statistically prove the existence of mean reversion at the usual level of 95% confidence, primarily because there is simply not enough history available (we would need data from a smoothly functioning market from about 1000 AD to meet the needs of the academics).

However, evidence supporting the existence of long-term mean reversion is very strong, even it does not rise to the level of the 95% statistical proof standard. Important early work was done by Poterba and Summers (1988) and Fama and French (1988). This work also established that mean reversion can exist as part of an efficient market. Spierdijk et al. (2010) found strong evidence of mean reversion in the markets for 17 developed economies over the period from 1900 to 2008. They also found that mean reversion was much more pronounced following periods of extreme uncertainty (i.e. when markets had large and sudden price movements). Most investment professionals seem to accept intuitively that mean reversion is probable – and a significant fraction go even further with their belief that profitable trading strategies can be based on mean reversion. Even the actuarial and investment consulting firms that provide stochastic modeling seem to believe in mean reversion, since almost all of them regularly adjust their assumption for future stock returns based on some measure of stock valuation levels (e.g. dividend yield, or P/E ratios). Finally, many plausible explanations have been offered for the underlying causes mean reversion, including ideas based on recent behavioral finance research.

My contention is that the information contained in historical returns needs to be reflected in any good simulation model. Academics in financial economics often approach their work as if it is a branch of mathematics. As a practicing model-builder, I prefer the definition of economics provided by Nobel laureate Thomas Sargent: "Economics is organized common sense." Future stock prices will not unfold as the result of some hidden mathematical process. They will be the result of economic processes and decisions that are entirely based on human endeavors. It may be far from perfect, but past experience is all we really have to provide insight to the future.

So here is how I have proceeded. I start with monthly returns over the period 1926 to 2013. To more closely match the practice of institutional investors, the returns were for a mix of 75% US stocks (broad market) and 25% non-US stocks. From this I use a moving block bootstrap method to remove the potential bias from overlapping time periods. I take random blocks of history for 120 consecutive monthly returns, splice 6 of these together to get a simulated 60 year history with no overlapping time periods, and repeat this 1,000 times. This provides a data set that includes 60,000 annual returns and 3,000 separate non-overlapping 20 year periods. From this data set I extract the shape of return distributions over periods of 1 year and 20 years. For each 60 year set, the shape of each distribution is defined relative to a normal distribution that has the same geometric average return and annual standard deviation as the resampled 60 year data set (i.e. any point on the actual distribution is defined by the number of normal standard deviates from the mean of the normal base distribution). This process gives me targets for the return distributions from any model.

With targets in place, I test various distributions, starting with two that seem to encompass the current practice reasonably well: (1) the traditional log-normal model, and (2) a double log-normal model (i.e. a log-normal model with stochastic variance, which creates some degree of fat-tail risk relative to the traditional log-normal model). Here are two cumulative distribution function ("cdf") charts that show how well these two models match the historically-based target distributions.

They both do a reasonable job of matching the annual return distribution, with the double log-normal looking best. At the 20-year horizon, however, there is significant misfit for both distributions. Here the double log-normal model underperforms the traditional lognormal. I now start to test various mean reversion features to see if the fit can be improved. In this short essay I am not able to discuss all the variations that I tested, so I will jump to a model that seems to work very well. (This iterative testing of parameters reflects my belief that model building is a blend of science and of grind-it-out craftsmanship.) The model is based on the intuition that mean reversion is much more likely following large market moves (e.g. 2008 through 2013), as supported by the findings from Spierdijk et al. (2012). In this model, mean reversion is triggered whenever there has been an unusually large market move over the prior two year period. Downward mean reversion is triggered if the average standard deviate over the prior two years is greater than 1.0 (about 4-5% incidence), intuitively representing the deflation of bubbles. Upward mean reversion is triggered if the standard deviate over the prior two years is less than -0.7 (about 7-8% incidence), intuitively representing the recovery from market

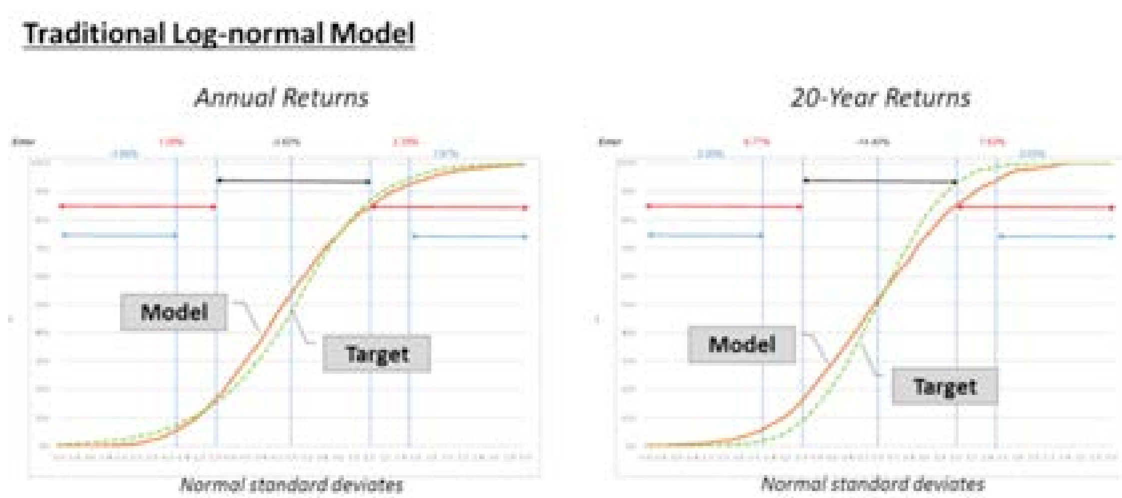

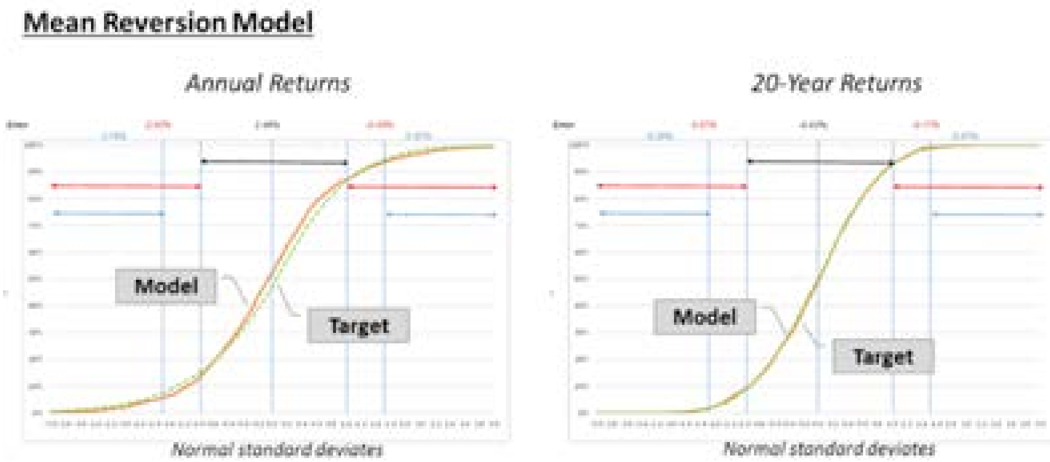

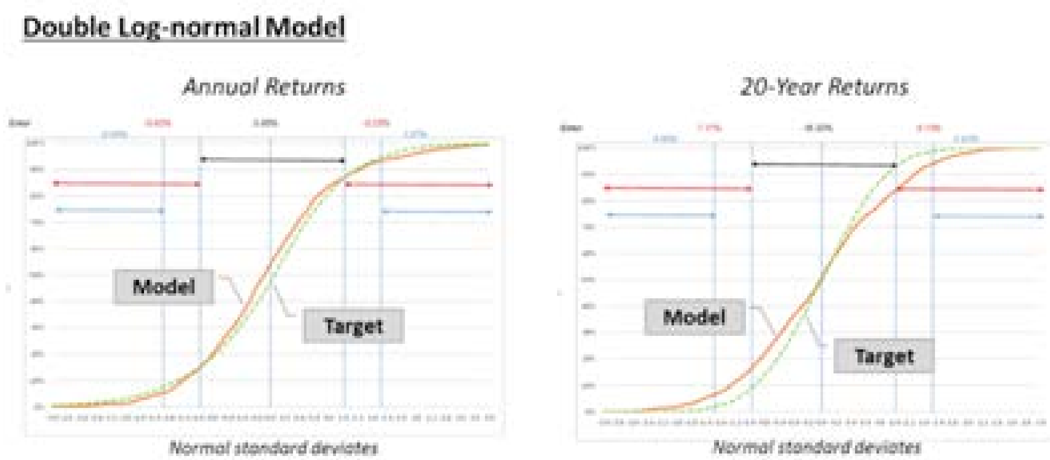

over-reaction to a crisis of some sort. The mean reversion impact is then factored into the returns over the following six years. (Note that my final model also includes a simple regime-switching feature to boost the fat-tail exposure a bit – but the mean reversion feature is, by far, more significant in determining the shape of the distributions.) Here are the cumulative distribution function charts for this model, which illustrate a much better fit for the long-term returns over 20 years. (The distribution of annual returns is roughly similar to the other two models – a little worse than the double log-normal model, but arguably a little better than the log-normal.) I have not tried everything possible, but achieving anything like this fit without mean reversion seems to be impossible, in my humble opinion.

To emphasize the importance of the mean reversion, we can look at the return distributions for two sample portfolios: one with a 50% allocation to equities and a 50% allocation to bonds; the second with a 75% allocation to equities and a 25% allocation to bonds. (For these results I assume expected nominal geometric average returns of 4.6% for bonds and 8.1% for equities – based on an expected 3.5% equity risk premium.)

In all cases, the three models have been set up so that they each offer the same long-term (100-year) geometric average return, and the same standard deviation of annual returns.

Note that the 20-year returns with the mean reversion model have a significantly lower standard deviation, which shows up in the tighter distribution of returns. With a typical investor's focus on downside risk, we can further illustrate the importance of mean reversion by looking at some 20-year shortfall risk probabilities in the following chart

If a simulation model is being used to help an organization set long-term policies (investment, or design) in a way that maximizes return within some specified downside risk tolerance, then the implications of this chart are very significant. Consider a plan sponsor who, for some reason, has established a minimum long-term (20-year) return objective of 2%, and desires a 99% confidence in meeting that target minimum. Then with the traditional log-normal model this sponsor would likely be looking at an investment policy with a 50% equity allocation, and a long-term expected return of 6.75%. Using a double log-normal model they might be at an even lower equity allocation. But the mean reversion model indicates that investing up to 75% in equities would meet their constraint. The expected long-term return would increase from 6.75% to 7.53%. I think this example puts the importance of mean reversion into a framework that all actuaries can immediately understand.

50% Equity / 50% Bond Portfolio (100-year g.a. return = 6.75%)

	Annual returns with			20-year returns with		
	Log	*DoubleLog*	*Mean Reversion*	*Log*	*DoubleLog*	*Mean Reversion*
Mean	7.45%	7.57%	7.29%	6.75%	6.74%	6.79%
S.D.	10.18%	10.18%	10.18%	2.17%	2.18%	1.56%
Percentiles:						
99%	33.53%	36.88%	41.25%	12.28%	12.32%	10.47%
95%	25.15%	26.23%	24.38%	10.46%	10.35%	9.20%
90%	20.17%	19.87%	19.84%	9.56%	9.61%	8.82%
75%	13.59%	12.63%	13.09%	8.02%	8.14%	7.89%
50%	6.64%	6.74%	6.78%	6.66%	6.64%	6.82%
25%	0.64%	1.22%	0.83%	5.33%	5.23%	5.73%
10%	-4.39%	-4.19%	-5.35%	3.98%	4.21%	4.72%
5%	-7.31%	-6.81%	-9.22%	3.36%	3.50%	4.20%
1%	-11.59%	-11.87%	-15.71%	2.15%	1.60%	3.22%

75% Equity / 25% Bond Portfolio (100-year g.a. return = 7.53%)

	Annual returns with			20-year returns with		
	Log	*DoubleLog*	*Mean Reversion*	*Log*	*DoubleLog*	*Mean Reversion*
Mean	8.65%	8.82%	8.58%	7.56%	7.55%	7.61%
S.D.	14.74%	14.74%	14.74%	3.23%	3.24%	2.29%
Percentiles:						
99%	48.52%	51.81%	59.35%	15.72%	15.86%	12.69%
95%	34.69%	36.90%	33.37%	13.11%	12.95%	11.22%
90%	27.91%	27.14%	26.27%	11.72%	11.65%	10.52%
75%	17.79%	16.36%	16.65%	9.51%	9.56%	9.13%
50%	7.49%	7.71%	7.53%	7.40%	7.41%	7.65%
25%	-1.69%	-0.54%	-0.25%	5.46%	5.27%	6.05%
10%	-6.97%	-8.83%	-8.97%	3.44%	3.65%	4.56%
5%	-13.17%	-12.76%	-14.64%	2.49%	2.77%	3.80%
1%	-20.25%	-20.67%	-24.70%	0.78%	0.08%	2.42%

20-year Shortfall Probabilities

	—— 50% / 50% Portfolio: ——			—— 75% / 25% Portfolio: ——		
	Log	*DoubleLog*	*Mean Reversion*	*Log*	*DoubleLog*	*Mean Reversion*
Prob. return less than:						
4%	10.20%	8.30%	3.70%	13.60%	12.40%	6.50%
2%	0.90%	1.50%	0.10%	3.70%	2.80%	0.50%
0%	0.00%	0.10%	0.00%	0.50%	1.00%	0.00%

Citations:

Fama, E. & French, K. (1988), 'Permanent and temporary components of stock prices', Journal of Political Economy 96, 246–273.

Poterba, J. & Summers, L. (1988), 'Mean reversion in stock prices: Evidence and implications', Journal of Financial Economics 22, 27–59.

Laura Spierdijk, Jacob Bikker and Pieter van den Hoek, (2010) 'Mean Reversion in International Stock Markets: An Empirical Analysis of the 20th Century', DNB Working Paper No. 247.

Welch, Ivo, 'Views of Financial Economists On The Equity Premium And Other Issues', Journal of Business 73-4, October 2000, 501-537.

Rowland Davis, FSA is the President at RMD Pension Consulting. He can be reached at *rowland-davis@mindspring.com*.

The thoughts and insights shared herein are not necessarily those of the Society of Actuaries, the Investment section of the Society of Actuaries, or corresponding employers of the authors.