SOCIETY OF ACTUARIES

Investment Section

INVESTMENT FALLACIES 2014

The Fallacy of Independence

By Richard Joss

Between March, 2008, and March, 2009, investment returns on stocks were in the neighborhood of a negative 50%. Investors would have seen half their wealth just disappear due to the large stock market drop. That's the bad news. The good news is that between March, 2009, and March, 2010, investment returns on the very same stocks were in the neighborhood of a positive 100%! Investors would have seen their wealth double. Of course, for investors who stayed in the market for the entire two-year period, the large second return only just made up for the losses incurred in the first year. The investor would have finished the two-year period with an average return of 0%. But even seeing a 0% return, many investors were happy just to have recovered all of the lost wealth.

When I talk with groups of investors about this two-year period, the vast majority feel that in some fashion these two events are related. The 100% return was certainly outstanding, but most feel that it might not have been this high had the market not suffered the 50% swoon the year before. If the drop had been smaller, the gain might have not been as large. In short, the 50% loss played some sort of role in the size of the 100% gain.

In spite of this widely held sentiment expressed by average investors, financial academics have tended to treat these types of investment returns as independent events, like coin flips or dice tosses. In short, they are saying that the huge loss seen in the first return had absolutely no bearing on the magnitude of the gain seen in the second return. This is just like the outcome of one coin flip not having any bearing on the outcome of a second coin flip, or on the coin flip after that, etc.

This assumption of independence underlies much of academic finance and serves as a basis for the comments that Wall Street is the biggest casino of them all, and that investing may be treated like a random walk down Wall Street.

The consequences of the independence assumption can be rather dramatic. Using the independence assumption and the two returns shown above, the academic community concludes that all of the four possibilities shown below are equally likely for any given two-year period:

Return for Year 1	Return for Year 2	Ending Wealth for $1,000
-50%	-50%	$ 250
-50%	100%	$1,000
100%	-50%	$1,000
100%	100%	$4,000
Average Ending Wealth		$1,563

Hence, using the academic model, the "unbiased estimator" for an investment return forecast becomes 25%, as this is the annual rate of return that would produce the average ending wealth of $1,563 shown above. The academics have taken an actual history where real investors had an average return of 0% and translated it into a forecast of a positive 25%. This is rather amazing when you think about it. This assumption of independence turns out to be a pretty big deal.

But what if the assumption of independence is not accurate? Might investors be receiving information that overstates possible returns? Instead of just accepting the assumption of independence, perhaps actuaries should do further digging to see if this assumption is really the best one available. Consider some of the following facts about actual investments in comparison with the independence assumption:

1) If returns were truly independent, then the probability of seeing a negative daily, weekly, or monthly return would be the same no matter how many consecutive negative days, weeks, or months preceded the one in question. But this situation is clearly shown not to be true. Regression analy-

sis always shows that the probability of seeing a negative return decreases as the number of negative returns mount up. The market (not some computer model) is saying that after some period of negative activity, stocks seem to be a better value and the increase in buying will have an upward impact on prices.

2) If returns were truly independent, then the Dow Jones Industrial Average (DJIA) would be expected to have nine negative days in a row about every four years, due just purely to random chance. But the last time that the DJIA was negative for nine days in a row was February 22, 1978 – more than 36 years ago. Either a very, very rare event has occurred, or perhaps returns are not as independent as the academics would like to believe.

3) If returns were truly independent, the widely available history-based data, such as P/E ratios, 52-week highs and lows, etc. would be completely worthless. And the firms supplying this data would be wasting the millions of dollars being spent to create it. Can you imagine a casino spending large sums of money to provide the data on winning roulette numbers or winning poker hands?

4) If returns were truly independent, then the commonly used "Monte Carlo" forecast methods would produce expected distributions of returns which actually match real distributions of historical returns. But when these comparisons are made, the Monte Carlo forecasts tend to produce a much wider distribution of returns that what is actually seen in real markets. A comparison of this difference was shown in *Risks & Rewards* (February, 2012).

It is pretty clear that investment return data is not independent – that actual investors take the history of a given investment into account in making a decision whether or not to buy or sell. And it is this actual behavior of investors (not a computer model) that determines actual returns.

But, if investment return data is not independent data, and the assumption of independence creates problems, what assumption would work any better? When one thinks about actual investment returns, clearly they are periodic observations of a given wealth growth. They are returns given the condition that the wealth grew from A to B. It is a fairly easy mathematical change (but a much more difficult political one) to move from the independence assumption to the conditional one. The Risk & Rewards article mentioned above compared how the distribution of Ibbotson large company stock return data compared with the independence assumption favored by academics, and the conditional assumption offered above. The comparison clearly favors the conditional approach. Not only does this change produce more accurate distributions when compared with actual results, it sets the "unbiased estimator" at the rate of return actually earned on the investment – not some higher arithmetic mean.

Actuaries have a large social responsibility. They help provide for the adequacy of employee benefit trust funds and for the solvency of insurance companies. They need to make sure that their work is as complete and accurate as possible. Practicing actuaries doing their jobs need to rely on the work of academic actuaries to make sure that the theories used are as complete and accurate as possible. It is my hope that the academic actuaries will address this independence vs. conditional issue fully and completely.

I look forward to seeing the results of their studies.

The Fallacy of Independence *By Richard Joss*

Richard R. Joss, Ph.D. is a Resource Actuary (Retired), Towers Watson. He can be reached at *rrjoss@comcast.net.*

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.