SOCIETY OF ACTUARIES

Investment Section

INVESTMENT FALLACIES 2014

Introduction

"Investment Fallacies"

The Investment Section of the Society of Actuaries (SOA) is pleased to release our essay e-book, **"Investment Fallacies."**

This e-book contains 16 topical essays that express the opinions and thoughts of a number of authors on the subject. The origin of word "essay" is from the French verb "essayer", which means to try or attempt. In this way, a written essay frequently tries out an idea or a line of argument. In this sense, an essay is not a narrative account, or a work of fiction, much less a proof at its heart is a conjecture or the expression of an opinion.

As such, and perhaps not surprisingly, essays tend to evoke strong reactions from some readers. This is in fact a sign of their success. In this way, an effective essay can be a spur to both critical thinking and criticism. An essay that does neither could arguably be described as not worth reading at all. It was in this spirit of enterprise that the Investment Section Council solicited the essays debunking investment fallacies, which are compiled in this e-book.

Not everyone will agree with everything set out in these brief works. Our hopes are simply that a) section members find them diverting ... and b) these de-bunking essays may prompt some readers to put pen to paper to de-bunk them in turn ... whether in the future pages of *Risk & Rewards* or in a response to another call for essays someday.

As such, it should be understood that 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.

Introduction

Over the past decade, the investment world has been buffeted by unprecedented events. Many long-standing beliefs or assumptions held by investment professionals may no longer apply to the new realities. At the same time, many common myths and misconceptions that have been previously debunked continue to influence investors today.

Essays have been submitted with the goal of identifying and exposing these fallacies. We plan to share these essays as a resource for investment actuaries and other interested investors later this year.

Types of essays received

We received a wide array of papers containing many dimensions of theories. Some papers received were closer to the rules, clearly starting a well-known practice and patiently arguing why this is a fallacy. For example, the Fallacy of independence, the Bitcoin fallacies or the Best model doesn't win. Some papers took more liberty with the exercise and the fallacy to be debunked was not clearly stated.

As no clear grouping of the papers emerged, the papers presented are in no specific order.

Prizes

The committee decided to grant the following prizes:

1st Place Prize – $500 for Future Equity Returns, by Eric Janecek
2nd Place Prize – $300 for The Fallacy of Independence, by Dick Joss
3rd Place Prize – $200 for The Fable of the Storyteller and the Market, by Steve Scoles

Reading experience

The ultimate goal for this initiative was to generate a pleasant read for our Investment section members and beyond. Reading the papers and discussing their relative strength and weakness have been a living experience for the review committee. This was an opportunity for interesting discussions regarding statements being a fallacy or not, controversial statements to be debated and views to be exchanged. Now this experience is yours.

Feel free to pass along any questions to David Schraub, FSA, CERA, MAAA, AQ, at *DSchraub@soa.org*.

Table of Contents

Table of Contents

Future Equity Returns

By Eric Janecek

Assumptions regarding future equity returns are critical for funding pension plans, designing and pricing variable insurance products and associated guarantees, optimizing investment strategies, and personal financial planning. Most of us assume that equities return an average of 8% to 12% per year. Many actuaries and investment professionals use similar expected returns for their long term projections. What is the basis for future equity return assumptions? What are the sources of equity returns? How reasonable is it to use these returns in long term financial projections? This essay explores these questions and reaches some troubling conclusions.

We will focus on expected average returns on Standard & Poor's 500 Index. The concepts and logic can be applied to other equity indices as well. The S&P 500 just happens to have a lot of historical data.

Historical Equity Returns

Returns for historical periods starting between 1871 and 1995 and ending in 2013 generally average between 9% and 12%. Over the entire 142 year period, nominal returns have averaged 9%. Real, or inflation adjusted returns have averaged 6.8%. This seems to justify the average returns typically referenced in the media and used in many financial projections. Historical returns are often used as justification for assuming future returns of the same amount. We will return to the problems associated with this assumption later.

Sources of Historical Returns

Let's consider the sources of historical returns and whether they are likely to continue in the long term. In doing so, it will be useful to break down equity returns into five components: dividend yield, new shares issued or bought back, changes in the price to earnings ratio (P/E), real earnings growth, and inflation. Share repurchases have roughly equaled new shares issued in recent history and will be ignored as they are unlikely to have a material impact on long term average returns.

Using S&P data from 1871 to 2013, the table below shows the components of the average equity return.

Source	Annual Return Contribution
Dividend Yield	4.4%
Changes in P/E	0.3%
Real Earnings Growth	1.8%
Total Real Return	6.8%
Inflation	2.1%
Total Nominal Return*	9.0%

* *Total return is computed by compounding the individual components.*

Future Equity Returns

Can the historical 9% equity returns be sustained into the future? To answer this question, we can get some insight from four different perspectives: *Statistics, Contributors to Historical Returns, Changes in the Components of Equity Returns,* and finally, a *"Reality Check."*

A Statistical Perspective: From statistics, we know that when calculating a sample average, we would prefer to have 30 to

[1] To be more precise, in order to have statistical support for setting future returns based on historical data we require:

1. Returns over some set of "time periods" such as years, days, or minutes are independent and identically distributed
2. The return distribution of potential future returns will be the same as the past
3. An adequate number of time periods

Clearly the third requirement could be satisfied by using the 390 1-minute returns from the prior trading day. However, this would not satisfy the first two requirements. Even annual time periods are not independent and identically distributed. Just consider business cycles, the Great Depression, World War II, and the 1990's.

100 independent observations. For a fifty year projection, we would prefer to have 30 to 100 fifty year periods, or several thousand years of data. We don't have that. Instead we will have almost three sample periods. We could heroically assume that the last 142 years will be representative of the next 50 years and use a 9% average return assumption. Or, we could assume potential equity returns are independent and identically distributed.[1] We can do better than this by examining the contributors to, and components of equity returns.

Contributors to Historical Returns: Over the last 142 years equity returns have benefited from an unprecedented amount of innovation. Let's consider some of the advances contributing to historical returns.

1. Cheap, abundant raw materials
2. Development of inexpensive, portable energy with oil, gas, and electricity
3. Advances in transportation (automobiles, trains, and planes)
4. Modernization of the financial system
5. Items 1-4 have allowed the combination of materials, labor and capital in central locations contributing to
   a. widespread use of mass production techniques
   b. specialization of labor
   c. increased efficiency
6. Advances in technology, education, health, and longevity leading to increased productivity
7. Advances in farming freeing up labor for other activities
8. High population growth rates
9. Quick communication via telephones, video, and computers
10. Growth in trade allowing specialization and increased efficiency

Will there be comparable advances in the future? In my opinion, it is not likely. Without them, we would expect lower returns.

Changes in the Components of Equity Returns: We know the current dividend yield is 1.9% versus the historical average of 4.4%. This is a decline of 2.5%. Given that the P/E is a little above average, and there is not long term trend in P/E, we might assume future changes in P/E will have little impact on future average returns. Company earnings as a percentage of GDP are near historically high levels. If they revert to normal, we would expect earnings to growth less than GDP.[2] It is therefore unlikely real earnings will grow faster than the historical value of 1.8%. We are left with an anticipated ***future real return of 3.7%.*** Adding the Federal Reserve's 2% inflation target gives a nominal return of 5.8%. Many actuaries and investment professionals currently assume materially higher returns than this. Are we being overly optimistic? This bottom up approach suggests we are.

A Reality Check: Why else so current long term return assumptions seem unreasonable? Consider this: assume we started with $100 and let it grow at what might seem like a low 4.0% real return for 2000 years. What would it grow to? $100,000? $1 million? $10 million? The actual answer is $1.2 trillion, trillion, trillion. This is in today's dollars. How about using the more common assumption of 7% real returns? If you set up a trust of $100, then in 2000 years it would grow to the substantial value of $6,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000. Think about this. Is it even possible? To put it in perspective, if you gathered all of the matter in the observable universe – stars, planets, galaxies, people, plants,

[2] Company earnings have varied between 3% and 11% of GDP. They are currently near their all time high. There is a natural limit to how much of the economy is allocated to equity owners. Earnings are unlikely to go substantially higher as a percentage of GDP, limiting future equity returns.

etc. and somehow converted it all into gold, you still would not have that much wealth. Clearly recent historical returns cannot continue indefinitely.

Future equity returns are important for pension plans, profitability of insurance products, company solvency, asset allocation, and personal financial planning. Now it is certainly possible to have average equity returns of 9% or more for the next several decades. But I would not bet on it for the reasons discussed above. I would argue that a best estimate long term, real return assumptions is around 4% for equities. A 4% real return and 6% nominal return is materially lower than what many people are using. Let's hope their optimism is well placed. Otherwise, we may be in for some unpleasant surprises down the road.

References:

Cornell, Bradford. 2010. "Economic Growth and Equity Investing," Financial Analysts Journal Volume 66, Number 1.

Schiller, Robert, "Stock Market Data," *http://www.econ.yale.edu/~shiller/data.htm,* February, 2014.

Organization for Economic Co-operation and Development, "Labor productivity growth in the total economy," *http://stats.oecd.org/Index.aspx?DataSetCode=PDYGTH,* March 24, 2014.

Bureau of Labor Statistics, "Productivity change in the nonfarm business sector, 1947-2013," *http://www.bls.gov/lpc/prodybar.htm,* March 24, 2014.

Bureau of Economic Analysis, "Current-Dollar and 'Real' Gross Domestic Product," *http://www.bea.gov/national/xls/gdplev.xls,* February 28, 2014.

Board of Governors of the Federal Reserve System, "Why does the Federal Reserve aim for 2 percent inflation over time?" *http://www.federalreserve.gov/faqs/economy_14400.htm,* March 28, 2014.

Eric Janecek FSA, MAAA is an AVP & Associate Actuary at Lincoln Financial Group. He can be reached at *Eric.Janecek@lfg.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.

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.

The Fable of the Storyteller and the Market

by Steven Scoles

Every day the financial media distills the market movements into a simple explanation or two. Lately, the daily market explanation is usually along the lines of: "Joe Confident, chief economist at Big Financial Corp., points to concerns about the Federal Reserve's quantitative easing program as to why the market declined 12 points today. His firm oversees about $432.4 billion."

One of the great investment fallacies that seems impossible to defeat is the "market story" – an explanation of financial markets moves. Furthermore, the market story is almost always along the lines of some external event (e.g. Federal Reserve utterances, Middle East tensions, etc.) being the cause for the market move.

Often the consequences of using these market stories are not particularly significant, but sometime these stories lead to a sort of logic and extrapolation of events that leads people very astray. For example, in early 2009, it was widely viewed that subprime mortgages were the cause of the financial meltdown. Additionally, there were still vast numbers of mortgages in the US that still had their 2-year year rate resets coming up that year. Thus, as the common logic went, if the stock market was going down because of mortgage defaults, more rate resets (with their higher rates) would lead to more mortgage defaults and thus the market would continue going down. However, instead, the stock market turned upwards, leaving behind those who followed the subprime story.

Stories seem fundamental to human communication. For most of our entire existence stories were our only way of communication. Stories usually involve some sort of step-by-step cause and effect relationship. Stories can take a long series of events and reduce it to something you can remember and repeat to others. For most of human life, stories were a great way of understanding things. However, in markets stories are rarely a reasonable approach.

The Sand Pile

Financial markets involve millions of participants (human and algorithmic) with a vast array of objectives, strategies, time frames, and amounts of leverage. The current market price is simply a single output of a vast and complex interactive process. A useful analogy to markets is a sand pile. As more single grains are poured onto the pile, the sand pile builds up with an occasional avalanche. Some avalanches are tiny, some are very large. The avalanches are a result of both an incremental grain of sand and the current state of the sand pile (for example, its steepness). While you could say a particular grain of sand caused an avalanche, the far more important aspect was the internal state of the sand pile. When the sand pile is in a fragile state, any grain of sand would cause an avalanche. There is nothing special about that last grain of sand – the sand pile itself was at a tipping point ready to be tipped. With all of these interacting grains of sands, a story of the sand pile's movements gets real messy, real quick.

With financial market's millions of participants and their various underlying approaches, markets are similar to sand piles – the internal state of the market likely has a far, far greater impact on subsequent market moves than the latest external news event. What's happening with subprime mortgages of whatever the Federal Reserve is doing might be best viewed as a few grains of sand in a much larger sand pile.

A simple and humorous example of the market story was exposed in the oil market a few years ago. A $1.50 spike in the price of oil during the night of June 30, 2009 was eventually found by the Financial Services Authority to be partly due to a futures broker going long 7 million barrels of oil in a "drunken blackout". The position was closed out the next day. When I saw the results of the this investigation, I went to my Bloomberg terminal to see what the explanation of this price spike at the time was: "the rally was prompted by concerns that an attack on a Royal Dutch Shell oil field in Nigeria will impact on the global oil supply." And then in another Bloomberg article, the decline in oil price the next day was due to an "unexpected decline in U.S consumer confidence"!

A more serious current example of the market story is the idea that the US stock market is being driven by quantitative easing – which leads to a sort of logic along the lines of: If the market is going up because of quantitative easing, then the stock market can always be controlled (up or down) by the management of the Federal Reserve's QE program.

Let's say enough market participants believe the Federal Reserve has the back of the stock market. Enough people believing that and taking significant risk (perhaps via leverage) could ultimately lead to a very fragile situation. Interestingly, the growth of such a belief may initially lead to a very strong market not unlike the friction of the grains of sand can initially lead to a very large and stable sand pile.

Overcoming the Market Story Fallacy

The market story fallacy is not easy to overcome. First of all, we are built to jump to quick theories and conclusions of why something is happening (we do it continuously every day in real life) so we need an explanation. Second, humans generally dislike uncertainty and a simple story is one shortcut we take in dealing with the uncertainty of markets. As well, the business of providing market explanations is very lucrative as a way of looking like an expert. Beating the market is a vastly difficult task, but providing confident "explanations" of market moves at least gives the impression of market expertise.

One way to overcome the market story fallacy as far as the daily explanations are concerned is to take two or three weeks of market "explanations" and compare them to the variety of factors given. You quickly find inconsistent explanations (a particular event causing the market to move up one day and other days causing the market to move down). Or you find "evidence" given one day is not referred to on other days (Middle East tensions are often used for the price of oil going up, but are conveniently forgotten when the price of oil goes down).

Overcoming the market story fallacy over longer periods usually requires either a really good memory or an analysis of market explanations from several months or years ago. For example, at one time the Federal Reserve's low interest rate policy was said to be causing inflation (as shown through rising oil prices in 2007 and 2008). However, interest rates were lower several months later as the price of oil dropped dramatically.

And perhaps the best strategy for dealing with the market story fallacy is accepting the uncertainty of markets. There are simply too many interacting forces at play within the sand pile of markets to be able to summarize them into a simple story.

To be sure, the latest news events can have some sort of impact on financial markets, it's just the process of creating a nice story for market moves is often useless at best and very risky at worst.

Steve Scoles, FSA, FCIA is the author of the forthcoming book Fooled by the Market. He can be reached at *steve@fooledbythemarket.com.*

Extended periods of robust macroeconomic growth are healthy for investing

by Paul Conlin

The 2010-2013 economic recovery from the 2007-2009 recession has been weak and even feeble by many historical measures. GDP growth is bouncing around between 1% and 3%, depending on the quarter. Short term interest rates remain at zero, 5+ years after the fall of Lehman Brothers. Long term unemployment is a severe problem, with benefits being extended out to 99 weeks and possibly beyond. And while the unemployment rate itself has steadily fallen, this has primarily due to a sharp drop in labor force participation rates (the lowest since 1978), rather than a return of robust job growth.

The subpar recovery is deemed by many to be a bad thing, and for those unlucky individuals in the vortex of it, it no doubt is. But is it bad for investors, I would say not. Or at least that the alternative is even worse. The investment landscape has changed, in ways I believe most investors are unable or unwilling to face.

The last three periods of robust macroeconomic expansion have all ended badly, all due, I would argue, due to the investment fallout. The 1983 to 1990 expansion resulted in overinvestment in commercial real estate, which caused a macroeconomic shock. The 1993 to 2000 expansion resulted in overinvestment in technology/media/telecommunications, which caused a macroeconomic shock. The 2002 to 2007 expansion resulted in overinvestment in residential real estate, which caused a macroeconomic shock. In all three cases, investors in all asset classes paid dearly when the music stopped playing, and investors in the particular asset class where the overinvestment occurred paid a catastrophic price.

I would argue that the parallels in the three example above are not coincidences, and that there are at least two macro/societal/political events which have caused permanent changes and rendered the old "business cycle" view of the world an anachronism. Investors ignore these at their peril.

1) The revolution in real-time availability of and transparency of financial information causes self-reinforcing vicious circles of new money investment to build and build in favored asset classes until they reach uncorrectable levels.

In the 1983-1990 commercial real estate bubble, the worst investments were made at the tail; new money plowed in around 1988 or later is what ended up sustaining the heaviest losses. At the time, retail participation was little, but life insurance company and commercial bank participation accelerated the most at exactly the peak of the cycle, but in such large amounts that the ultimate crash wiped out all the previous years of gains. The late 1990s internet bubble was of course a retail investor phenomenon, with everyone chasing a, in hindsight, small universe of internet stocks, with 1999 presenting the ridiculous outcome of the entire gains in the Wilshire 5000 being attributable to stocks who reported negative GAAP earnings. And the residential real estate boom was a super-bonanza of retail investor participation thru their primary residence, and institutions piling on via derivatives in an asset class they had historically sat out.

The next asset class to outperform on a consistent basis will draw mind-boggling waves of cash, a situation from which there is no graceful exit. (Perhaps we are on the way there in the $17 trillion U.S. Treasury bond market, although I don't get that sense.)

2) The Federal Reserve no longer has the political stomach to engineer a slowdown before the bad stuff starts to hit.
The story of Alan Greenspan standing by (at least non-verbally), intentionally, during both the internet bubble and the residential real estate bubble is now well known. The legacy of the man once thought the greatest Fed Chairman ever has been permanently tarnished. I believe these same instincts were at work in the late 1980s, although he largely has sidestepped blame for this one, if for no other reason

than it was just so long ago. But I would argue that more was at work than the laissez-faire framework of a single Fed Chairman. Ben Bernanke's term, while admittedly in such an odd time as to make broad conclusions difficult, expressed a strong bias towards expansion, with very low regard for the dangerous consequences of putting the foot full speed on the economic accelerator. As recently as January 2014, he has brushed off concerns of damage to emerging markets by stating that "The Fed is not the central bank of the world."

It pains me to say it, but the 19th century Marxists, at least on this topic, were right. Capitalism cannot control itself; it wants too much of a good thing, and won't stop eating until it to too full and has a belly-ache. Slow economic growth is not what ails us. In fact, we should be grateful it is on hiatus.

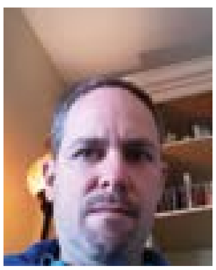

Paul Conlin, FSA, MAAA, is an Actuary at Atena. He can be reached at *ConlinP@aetna.com*.

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.

Mathematical modelling of social phenomena

By Nicholas John Macleod

Introduction

The application of mathematical models to social situations is seen by some as a natural extension of their use in physical science. But there are crucial differences. For example, the lack of controlled and repeatable experimentation stands in sharp contrast to the situation in physics, where theories make precise predictions that can be tested and falsified. A physical theory whose predictions cannot be verified by experiment is normally discarded, but this process of selection is largely absent from social science.

In part that's because the objects of interest in social science are not inanimate particles whose behavior can be expressed in terms of simple laws. They are people - with individual thoughts and the freedom to adopt individual actions.

The normal approach to individuality is to fall back on statistical aggregation. In the classic example where country fair-goers are asked to guess the weight of a pig, the average guess comes out strikingly close to the animal's actual weight. That is a triumph of statistics, but the situation is atypical in that the fair-goers are given no information with respect to earlier guesses. In most real social situations the information flow between the participants is a critical determinant of system behavior.

As might be expected, modelling information flow takes us to a new level of complexity, and comes with its own problems. As models become more complex, it becomes increasingly difficult to distinguish between genuine properties of the system, and properties that arise from the particular assumptions of the model.

Does this mean that mathematical modelling in a social context is a waste of time? I don't think it does, but I do believe that more attention needs to be given to the fundamental differences between social and scientific modelling. In particular, the failure to reject social models whose outcomes differ significantly from their intentions or predictions undermines any claim to scientific method and greatly impedes progress. While we cannot set such strict criteria for acceptance as we do in proper science, that doesn't mean that we should set none.

This is too large a subject to be dealt with in a short essay, so I will try just to highlight some general issues by means of an example drawn from finance. My approach is to set up a straw man – something I've called the Old Model – and to demonstrate that a simple extension of its underlying assumptions leads to distinctly different prescriptions. I don't insist on the details of the Old Model; some people may see it as a caricature, although most will recognize at least some of the elements of conventional investment theory. My point has to do with the need to test the robustness of a model's prescriptions. If a modest change in what are necessarily uncertain or approximate assumptions leads to a radical change in indicated action, the assumptions must be carefully reviewed for dependability.

Old Model

The conventional approach to asset allocation is based on the following ideas[1]:

- There is essentially one state of the world;
- Return variation is fluctuation that reflects new information that by definition cannot be predicted;
- Volatility (the amplitude of fluctuation) measures risk;
- Reduction of volatility depends on low correlation

[1] Here, and in what follows, I'm tacitly assuming that the assets are supporting a funding program.

so finding uncorrelated assets is the key to effective diversification;
- There is a long run positive relationship between expected return and risk;
- You can't time the market.

This model supports the general practice of holding fairly fixed allocations across as wide a range of asset classes as possible. Since return variation is intrinsically unpredictable, there is no point trying to time the market. Instead you should combine the positive risk-return relationship with the volatility-reducing properties of low correlation to maximize the likelihood of achieving your required return. The process can be thought of as starting with the asset with the highest expected return, and progressively shrinking portfolio volatility by adding imperfectly-correlated assets. The volatility reduction comes at a cost in expected return, and relatively few additions account for most of the reduction, so there is a natural point at which the portfolio has the maximum chance of achieving the required return. That's your optimal portfolio and it will remain pretty stable throughout the life of your funding program.

Variation around the long term average return of the portfolio also declines with time, in exactly the same way that it declines with the addition of uncorrelated assets, except that here there's no associated give-up of expected return.

As you expand the range of assets and extend the length of the investment period, return converges around a favorable long run average and the likelihood of achieving the target return over a normal funding period increases to near-certainty.

The idea of trading expected return for reduced risk in order to increase the likelihood of achieving the target return seems to make sense, and the statistical elements of the approach appear to be sound, so why hasn't it worked in practice?

It's helpful here to contrast the statistical modelling used for liability estimation with the use of statistical models on the asset side of the equation.

- Mortality is a natural process that conforms to regular biological and statistical laws, so making allowance for increases in longevity, the statistics of past mortality tend to be reliable indicators of the incidence of future mortality. There are other variables that require estimation (for example, for a final earnings pension scheme, the level of benefits will depend upon the recipients' final salaries, which are not known today) but here again, past experience generally provides a fairly dependable basis for the estimation of averages.

- Finance, on the other hand, is a social activity with occasional regularities, but no fundamental laws. Where there are no well-established laws, models must be justified by their consistency with real world experience. But the asset allocation methodology described above is based on theoretical assumptions about the way markets *should* work, rather than on experience of how they *do* work. Since the prescriptions of the Old Model have not led to the anticipated outcomes, it is falsified by application.

<u>What do I mean by "falsified"? All models are simplistic, and therefore false.</u>

What does it mean to say that one model is better than another when the underlying reality is infinitely more complex than either of them? I think the answer has to do with the models' qualitative prescriptions. The Old Model essentially recommends fixed allocations to a "diversified" (= non-correlated) set of assets. But what if we extend the model to a two-state system? Let's say there are bear markets and bull markets, and that each of those states is sticky, or persistent to some degree. In other words, suppose that market behavior is more like the weather, in that things can change, and when they do, they tend to remain in the new state for some time[2].

This model is only one level more complex than the Old Model, but its implications are entirely different. For example:

- Shifts from bull to bear markets pose a much greater risk than local fluctuation.
- Assets that are statistically uncorrelated may well respond in the same way to environmental shift. As a result, they do not diversify each other.
- It makes no sense to stick with fixed allocations in the face of market change any more than it does to stick with the same set of clothes in all weather conditions. Persistence means that when the world changes, you have to do something about it; risk management should be dynamic, not passive.

Once we agree that market change is real and persistent, the pillars of the Old Model collapse. Risk is not simply volatility, low correlation does not guarantee effective diversification, and fixed allocations are not the best way to fund liabilities. Going from a single state model to a two-state model changes everything. It also explains some of the mysteries of the Old Model; why do outsized losses among assets that were not previously correlated often occur at the same time, for example?[3]

When we generalize the two-state model to cover multiple states, there is no *qualitative* difference in its implications and recommendations; all of the points above still apply. So while a two-state model is obviously far too simple to describe the real behavior of markets in any detail, its qualitative prescriptions are not the by-products of over-simplification.

It's also easy to see how a two-state model can be extended to a multi-state model without introducing any new concepts: a multi-state model is just a string of inter-connected two-state models, so, while it's more complex, it's not fundamentally different. That isn't the case in going from a single-state model to a two-state model, where we have to bring in new mechanisms like **transition** and **persistence** that don't appear in the one-state model. And it's those mechanisms that explain the correlation dynamics and other things that the Old Model can't.

Statistical analysis of recent market activity[4] suggests very strongly that market conditions are persistent. Who can doubt that 2000–2003 was a very different environment from 2003–2007, and that 2008 was different again?

---

[2] This is not to argue that the markets are as simple as the weather. We know what causes the seasons and we understand fluid dynamics, but weather is still difficult to predict. We have a much more limited understanding of financial markets.

[3] This is the "volatility spiking and correlations going to 1" phenomenon, that's unexplained in the Old Model, but perfectly natural in a multi-state model.

[4] And perhaps more important, professional experience and common sense.

The evidence indicates that, far from being as simple as the Old Model suggests, reality is better represented by a highly dynamic, but persistent, multi-state model in which the states are not fixed, there is potentially an unlimited number of them, the degree of persistence – i.e., the stability of market conditions - varies irregularly, and so on.

But even the two-state model—the simplest possible multi-state model—is significantly more complex than the Old Model, and qualitatively quite different. It explains phenomena[5] (outsized losses, coincident losses among uncorrelated assets, etc.) that are not just theoretical mysteries under the Old Model, but real-world events that can damage or destroy a funding program. Following the prescriptions of a model that doesn't even recognize their existence is not what's normally thought of as prudent.

Nicholas John Macleod, ASA, C.Math., FIMA, is a Principal, Bespoke Financial Modelling, Jersey, Channel Islands. He can be reached at *nicholasjohnmacleod@gmail.com*.

[5] It might be more accurate to say that these phenomena arise from dynamics that are built into the structure of a multi-state model.

A real-world approach to Value at Risk

By Nicholas John Macleod

Introduction

A well-known legal anecdote has it that the barrister Sir Edward Marshall Hall, when asked by a judge, "Is your client familiar with the doctrine of *Res Ipsa Loquitur?*", replied, "My Lord, in the remote hills of County Donegal from where my client hails, they speak of little else."

In many ways that captures the problem faced by investment professionals when they come to consider the calculation of Value at Risk (VaR). Far from being able to provide reasonable estimates based on their day-to-day investment experience, they find themselves confronted with controversy concerning the applicability or otherwise of the Normal distribution, and with proposed remedies that encompass Extreme Value Theory, GARCH models, Copulas, etc., none of which falls within their areas of professional expertise.

The difficulties with such mathematical approaches, however, go beyond their inaccessibility to the practitioner. VaR is concerned with the possibility of unusually large losses, so almost by definition, there is little empirical basis for assessing the nature of potential non-Normal goings-on in the extremes of return distributions. As a result, there are no real grounds for attributing any particular structure to that part of the return distribution, whether by GARCH models, generalized Pareto distributions or anything else; whichever approach we choose, our chosen model will tell us exactly what we told it to tell us.

In this essay I'm going to suggest that estimating VaR is not fundamentally a problem requiring advanced mathematics, and that its solution is not to be found by attempting to divine the nature of hitherto unobserved behavior in the tails of asset return distributions. I'll go further and argue that a common sense approach not only provides more realistic estimates of potential loss, it forces us to identify shortcomings in portfolio structure, and provides practical guidance with respect to what to do about them.

Common-sense Value at Risk

The definition of VaR is straightforward: if I can lose $\$X$ or more with probability p , is my VaR at level p. In terms of return, if a portfolio has a probability p of generating a return less than or equal to R, the VaR at level p is R.

It's clear that VaR is related to the distribution of return. (In fact specifying for every value of is the same thing as specifying the return distribution.) But since it's intended to measure risk, it is generally concerned with low-probability, large losses, events that are typically represented as happenings in the left-hand tail of the return distribution.

For real-world examples, we only have to look back a few years. As the events of 2008 unfolded, equity investors began to see a shift towards more frequent and larger negative returns. Then, in September 2008, the MSCI Europe index lost almost 12%. This was an outlier – a return that did not fall within the range of returns observed during the previous five years. And in October 2008, it happened again, and this time the loss was more than 13%. You wait five years for an extreme event, then two come at once!

Or perhaps not. An investor whose expectations were shaped by the previous five years' performance would certainly have experienced the September and October returns as extreme events. On the other hand, a better-informed investor who realized that the world had changed would have seen them as more or less normal returns within a new and much more hostile environment.

The fact that the outlier of September was followed by an even larger outlier in October supports the latter interpretation. Rather than invoking the statistics of extremes, a more modest view might be that risk arises from our own misjudgments. From that perspective, the returns were not extreme events within a normal environment: they were normal events within a new environment we had failed to reckon with.

<u>How does this help us to estimate Value-at-Risk?</u>

The question of how much a portfolio could lose in hostile conditions boils down to

- What sort of market conditions could harm the portfolio, and what would its returns in those circumstances look like?

- How long would it take to recognize that conditions have become hostile, and how quickly thereafter could action take effect?

Simply contemplating these questions brings VaR back into the real world. How we might make a reasonable estimate of potential loss in practice can be illustrated with an example drawn from real life.

<u>Example</u>

Suppose we expect a bond-substitute portfolio that is invested mainly in low-volatility relative-value hedge funds with limited liquidity to return 4% per year with 4% annual volatility.

The pink line in the chart below represents the anticipated mean return path, and the red line marks the 2-standard-deviation lower boundary under these preliminary expectations. An observed cumulative return path that strays below the boundary would suggest that initial expectations are not being met, so we refer to the red line as the *risk signal boundary.*

The portfolio manager's knowledge of the underlying hedge fund strategies enables him to envisage plausible market conditions in which the portfolio might return -8% annually on 4% volatility. In other words, while it may not be all that likely, there is a set of market conditions that constitutes a realistic "risk" scenario associated with the portfolio. If that scenario materializes, the mean return path would follow the blue line shown in the chart.

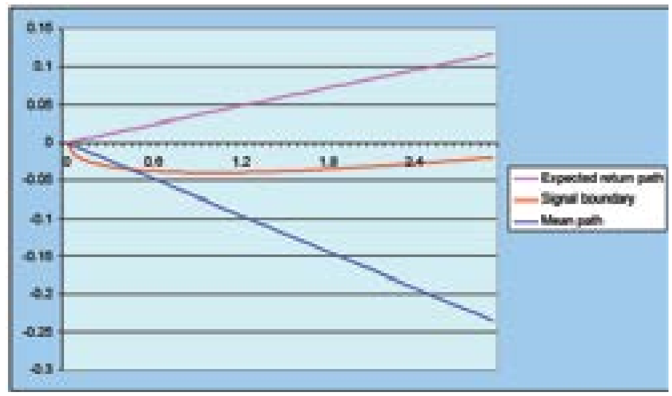

The risk-scenario's mean return path crosses the 2σ signal boundary at about 5 months. At that point we initiate action, but owing to the poor liquidity of the hedge fund investments, it takes effect only after a further year (next chart).

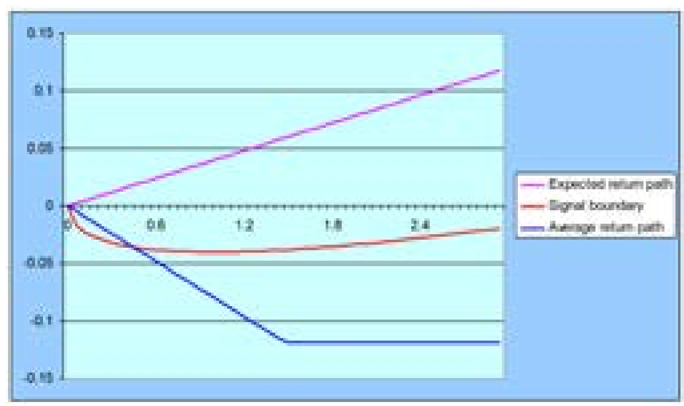

If returns within the -8%; 4% regime follow the average return path, we could expect to lose a total of about 11.5%, made up of 3.5% from inception to signal, plus a further 8% before corrective action takes effect.

But the VaR question has more to do with "How much could I lose *if things go against me?*" than with average loss, so to be safe, the calculation should be based on something worse than the mean path in an unfavorable environment.

The dark blue line on the next chart shows what happens when:

- The actual return path crosses the signal boundary at its lowest point.
- Experience following the signal involves two standard deviations' bad luck[1].

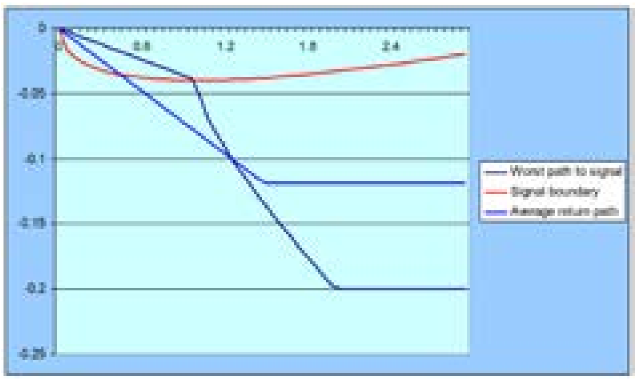

In this case, it takes a year before cumulative return hits the signal boundary — losing 4% in the process — and between Signal and Action we lose a further 16%, for a total loss of 20%, which is clearly unacceptable for a portfolio designed to achieve 4% on 4% volatility[2].

The decomposition of loss into pre-Signal and post-Signal components allows us to identify weaknesses, both in the portfolio, and in the decision-making framework. In terms of the example:

- It doesn't really matter whether the return path crosses the signal boundary at its lower point or somewhere else: the obvious problem is that the time between Signal and effective Action is an entire year, during which we might accrue losses of as much as 16%. The direct solution to that problem is to increase the liquidity of the portfolio.

- We might also find that incurring a loss of around 4% before we even think of taking action is too severe for an ostensibly low-risk portfolio. One solution is to tighten the signal boundary, but that comes at the cost of raising the likelihood of false signals.

- More generally, the idea that a portfolio designed to achieve 4% return on 4% volatility could, under plausible circumstances, annualize at -8% might itself seem inappropriate, in which case the solution is to restructure the portfolio.

Conclusion

The calculation, as described, is obviously a simplification of reality, in that there is typically no single instant at which we decide to take action, and no precise moment at which it takes effect. But the formulation includes the main elements of real risk-management decision procedures, covering questions such as:

- *What market conditions could threaten the portfolio?*
- *How much loss can we tolerate before deciding that observed returns conflict with prior expectations?*
- *Does the liquidity of the portfolio - or our own decision protocol - permit us to act decisively in time to limit losses to an acceptable level?*

These are not easy questions. They require investment expertise and judgment. But they fall squarely within the

---

1 This is not to suggest that return actually follows the path of the dark blue line; it just says that the signal boundary is crossed at its lowest point at time t, and that, at time $t + T$, when action takes effect, things in the meanwhile have gone against us to the tune of two standard deviations.

2 In more conventional terms, a 20% loss over two years on a (4%;4%) portfolio represents a five-standard-deviation event. Even the average loss (with no bad luck) of about 12% is a four-standard deviation event. Each of these numbers is derived here from well-understood, plausible assumptions, but it's hard to imagine that we would have obtained estimates of similar magnitude from a contemplation of the tail risk associated with a (4%;4%) return distribution.

capabilities and experience of investment professionals, and considering them explicitly within the framework of a calculation based on real-world elements helps to restore the responsibility for an investment program to its proper place: in the hands of the program manager.

Nicholas John Macleod, ASA, C.Math., FIMA, is a Principal, Bespoke Financial Modelling, Jersey, Channel Islands. He can be reached at *nicholasjohnmacleod@gmail.com*.

Bitcoin Fallacies

by Larry Zhao

As an investment professional, we have learned many fallacies such as "Sell in May and go away." The question is how we do with these fallacies we know about, in particular, how we apply the knowledge to new concepts, new situations, and new investments.

Introduced by the pseudonymous developer Satoish Nakamoto a few years ago, Bitcoin has brought with it excitement, drama, myths, and fallacies. Some people take it as a déjà vu of the Tulipmania that struck the Dutch in 1636. Some view it as another Ponzi scheme. A prominent economist and Nobel Prize laureate simply called it evil.

What is Bitcoin?

Bitcoin is the first decentralized digital currency designed to facilitate transactions between two parties over a peer-to-peer global network. First and foremost, it is a new technology – a black swan technology that grows out of decades of research in computer security and cryptography by tens of thousands of researchers and scientists globally. Bitcoin is potentially disruptive – just as Google has disrupted the traditional retrieval and catalog system, Bitcoin might seriously challenge the status quo of the current credit/debit and banking system and international remittance markets, because Bitcoin has removed the middle-men facilitating transactions.

How Bitcoin works?

As a payment system running on open-source software, Bitcoin secures transactions using asymmetric key cryptography, which requires a public key and a private key. The public key is distributed widely (i.e., QR codes) while the private key should be kept secret all the time. A secure transaction is done this way: 1) a sender encrypts a payment using the recipient's public key and broadcasts to the Bitcoin network; 2) the payment cannot be decrypted by anyone without the recipient's private key; 3) before the payment is sent, a signature is also created with the sender's private key and the signature can be verified by the recipient using the sender's public key – any tempering of the payment during transition will result in mismatch of the signature.

What are the benefits?

With no third-party intermediary, the cost is zero or near-zero. With no authority to approve, the transaction can be done in minutes between peers at any corners of the globe. Because it is computationally impossible to recover the private key from the encryption, the transaction is secure. Because neither the sender's nor the recipient's information is transmitted, the transaction is free of identity-fraud.

Bitcoin fallacies

Is *Bitcoin* another *Tulipmania* bubble waiting to burst? This is a legitimate question even if the original narrative of Tulipmania was historically inaccurate. After all, unlike gold, Bitcoin has no intrinsic value; and unlike fiat money, there is no backing up central authority. The fact that Bitcoin has limited supply (21 million) does not guarantee the value of Bitcoin in the future because the demand could drop sharply if everyone suddenly loses confidence and stops using it due to unpredictable events. Theoretically its value could drop to zero, like any currency of dissolved states in history. However, Bitcoin derives its value not from the role of digital currency but rather from its *utility* – the usefulness of the Bitcoin system to provide fast, low-or-no-fee, secure, peer-to-peer transactions, and the speculation on future use of the system. As more and more consumers and merchants are using and accepting bitcoins daily and globally, and as more and more Bitcoin tools and technologies are created and improved, it will create a positive feedback-loop and an expanding ecological system, just like elevators, telephones, or Internet. You might see flash crashes occasionally, but Bitcoin more likely than not can survive on the basis of being used entirely as an e-payment system.

Is Bitcoin the dream tool for drug dealers, money launderers, and terrorists to transfer money anonymously without impunity?

No, this is a fallacy. Bitcoin is pseudonymous, not anonymous. Every transaction in the Bitcoin network is tracked and logged permanently, available for any one to see (http://blockchain.info). Bitcoin is significantly easier for law enforcement to trace than cash or gold. Criminals and thugs will continue to use the best tools and technology available no matter what. For migrant workers who go to work in hard jobs in foreign countries, Bitcoin offers a far better alternative than paying 10% or higher fees in order to send money back to their families.

Is Bitcoin a Ponzi scheme?

A Ponzi scheme is a zero sum game – early investors can only profit at the expense of late investors. Bitcoin, however, could have win-win outcomes where early investors profit from the rise in value and late investors benefit from an e-cash system that is inexpensive, fast, flexible and globally accepted. The fact that early investors benefit more does not necessarily make Bitcoin a Ponzi scheme. All good investments in successful companies have this quality. Secondly, Bitcoin does not promise a higher return to maintain a continual stream of investment. Investment continually flows in as people gradually realize its value and potential as they use bitcoins in transactions during everyday life.

Is Bitcoin evil?

Paul Krugman dislikes Bitcoin because of its inherent Libertarian political agenda – to undermine the ability of governments to collect taxes and monitor financial transactions among their citizens. The wonderful features of Bitcoin as a payment system, in his opinion, are simply positive economics – how things work; but on the normative economics level – how things should be – Bitcoin fails, and on this very basis he thinks Bitcoin is evil. As always, Krugman is insightful, because he looks beyond Bitcoin as a payment system but as an idea – a dangerous idea, because no one government can shut it down. What governments and regulators can do is impede its progress and innovation. Bitcoin is a discovery, similar to the discovery of fission, based on which nuclear reactors are built and electricity is generated. Most people focus on the pros and cons of nuclear reactors, or the price and use of electricity, while missing the point that fission in itself changes physics, changes energy, changes worldviews, changes everything. Government can regulate Bitcoin, but cannot make the discovery disappear.

Should you invest in Bitcoin?

Make no mistake. Debunking the fallacies is to offer a balanced view, but not to encourage anyone to invest in Bitcoin unless you truly understand the risks and rewards.

Price volatility is a major impediment to Bitcoin's wide adoption as a viable payment system. Bloomberg data shows that, since inception, the annualized volatility of bitcoin returns in US dollars is about 150%, about 10 times of the S&P 500 index, while the annualized return is about 350%, about 40 times of the S&P 500 index. Its momentum to maintain high future annualized returns is questionable.

But over long term, it behaves like a special asset class with a binary pay-off: it can worth quite a lot or almost nothing.

One investment strategy is to use it as a "tail risk hedge", where a small position invested early on can still offer a substantial return over many years, or lose very little should Bitcoin die or be usurped.

I invest precisely what I can lose 100% of the investment.

The operational risk, however, cannot be underestimated. Misplacing the private key can result in the entire loss of your bitcoins, which cannot be replaced by any organization. Internet hacks, security breaches and counterparty risk at Bitcoin exchanges such as MtGox and BitStamp are perennial headaches to investors.

Larry Zhao, FSA, CERA, CFA, FRM, PhD, is an associate vice president at Nationwide Financial. He can be reached at zhaol1@nationwide.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.

Gambler's Fallacy: Probability of Reversion

by Kailan Shang

An odd event does not necessarily imply an immediate reversion.

Gambler's Fallacy

After a gambler loses a game many times, he/she may mistakenly believe that the chance of losing it again is smaller than normal. However, if the outcomes of the game are independent, it is a false belief. This is called the gambler's fallacy. For example, when tossing a fair coin, the probability of getting five heads in five tosses in a row is small. When making the sixth toss after getting five heads, the player may think that the probability of getting six heads in six tosses is very small. Therefore it is more likely that the sixth toss will get a tail than a head. Here, the player considers the conditional probability as the unconditional probability. The probability of getting six heads in the first six tosses is small. But the probability of getting six heads in the first six tosses given five heads in the first five tosses is not small.

Event	**Probability**
Five heads in the first five tosses	(½).5 = 1/32
Six heads in the first six tosses	(½).6 = 1/64
At least one tail in the first six tosses	**1-(½).6=63/64**
A head in the sixth toss given five heads in the first five tosses	½
A tail in the sixth toss given five heads in the first five tosses	½

Probability of Reversion

The gambler's fallacy can be seen in investment activities as well. Some investors speculate that a reversion of a trend will occur simply because the chance of the long-term trend happening is small. For example, investors may expect a price decrease in the next day after several days' stock price increases. They may be optimistic about a rising interest rate environment after a long period of low interest rates. There might be some in-depth analysis to support their conclusions. But in many cases, it is the gambler's fallacy that leads to them. Sometimes the probability of reversion is not affected by an unlikely trend in the past.

Figure 1 shows the daily close price of Apple Inc.'s stock in the past ten years. It increased a lot during that period but with many fluctuations. Table 2 lists the historical experience of continuous increases or decreases of daily closed stock price of Apple Inc. If the stock price had increased for five days, there are 55 cases that the stock price increased in the next day. And there are 59 cases that the stock price decreased in the next day. Comparing the number of cases that price increased in the next day to that the price decreased, in many situations, it is not obvious that the probability of reversion is higher than the probability of continuum of the trend. Similar to the gambler's fallacy, it does not imply that the chance of having a price decrease in the sixth day is high after having the stock price increasing for five days.

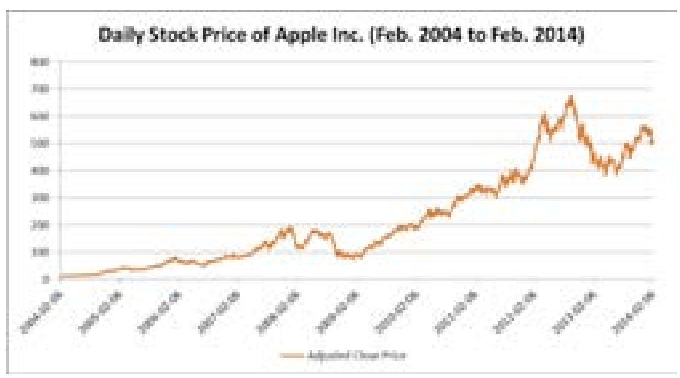

Source: Adjusted close share prices of Apple Inc. from Yahoo! Finance. Close prices are adjusted for dividends and splits.

Psychological Explanation

The gambler's fallacy is not rational, but it is appealing to human cognition. Tversky and Kahneman (1974)[1] explained

Table 2: Daily Price Movement Summary of Apple Inc. (Sept. 7, 1984 ~ Feb. 6, 2014)Source: Adjusted close share prices of Apple Inc. from Yahoo! Finance. Close prices are adjusted for dividends and splits.

# of days with continuous price **increase**	Next Day Price Movement		# of days with continuous price **decrease**	Next Day Price Movement	
	Down	Up		Down	Up
1	942	903	1	856	990
2	451	452	2	494	496
3	237	215	3	274	222
4	101	114	4	106	116
5	59	55	5	56	60
6	30	25	6	38	22
7	8	17	7	13	9
8	7	10	8	4	5
9	5	5	9	3	2
10	4	1	10	1	1
11	0	1	11	1	0
12	1	0	12	0	0

Source: Adjusted close share prices of Apple Inc. from Yahoo! Finance. Close prices are adjusted for dividends and splits.

that the gambler's fallacy is a cognitive bias caused by representativeness heuristic. When people are asked to assess the probability of an event, they compare it to their experiences and knowledge to find out the similarity. And they may expect that short run outcomes should be representative of long run outcomes. But in reality it is not. Using the coin-tossing example, let's compare the mean and volatility of 10 tosses' outcome and those of 100 tosses' outcome. The *Volatility/Mean* of the 10 tosses' outcome is relatively high compared to that of the outcome of 100 tosses. Furthermore, the probability of having heads more than 70% of the tosses for 10 tosses is much higher than that for 100 tosses. A long run outcome is not likely to happen in the short run. Assessing the probability of the short run outcome based on the understanding of the long run outcome will certainly lead to a biased conclusion.

Event	10 Tosses	100 Tosses
# of heads: Mean *(np)*	5	50
# of heads: Volatility ($\sqrt{np(1-p)}$)	1.58	5
Volatility/Mean	.32	.1
Probability that the # of heads is greater than 70% of the number of tosses	5.469%	0.002%

Solutions

It is unlikely to completely remove the gambler's fallacy as it is part of human nature. However, several approaches can be taken to mitigate its effects.

1. Educating the investors about the existence, the causes, and the potential impact of the gambler's fallacy. With an increasing awareness of this cognitive bias, people may adjust their forecast to offset the impact of the

[1] Tversky, Amos and Daniel Kahneman, "Judgment under Uncertainty: Heuristics and Biases." (1974) Science, V185, No. 4157: 1124-1131.

bias intentionally.

2. Conducting analysis to understand the underlying drivers of market change. Fundamental analysis can be used to predict stock price movements. Macroeconomic analysis can be used to forecast the movement of interest rates. Although exploring cause-and-effect relationships is very difficult, it is very useful for prediction. On the other hand, considering the market movement as a random process will make the prediction subject to the gambler's fallacy. In addition, when comparing the predictions using several approaches, it is easier for people to realize and correct their biases.

3. Due to the lack of knowledge, random processes may be used to analyze some issues. In those situations, historical experience other than current status does not need to be provided when asking for people's prediction. This can reduce the impact of the gambler's fallacy.

Changing a cognitive bias can be very difficult and take a long time. But with sufficient training and appropriate tools in place, its impact can be immaterial.

Kailan Shang, FSA, CFA, PRM, SCJP, is Managing Director, Head of Research at Swin Solutions Inc. He can be contacted at *Kailan.Shang@swinsolutions.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.

2 $\sum_{i=x}^{n} \binom{n}{i} p^{i}(1-p)^{n-i}$

n: # of tosses. ***p:*** the chance of getting a head. It is assumed 0.5 in the example. ***x***=70%×n

Invest Fallacy: Active management overall performs different than passive management

by Evan Inglis

We hear it in the media and from asset managers all the time, the idea that certain time periods or certain conditions are better suited to active management.

- "It's a stock pickers' market"
- "Active managers are underperforming the market"
- "Increased volatility is providing opportunity for active managers"
- "Active management outperforms in bear markets"
- "60% of active managers beat their benchmark"
- "Markets that aren't efficient are well-suited to active managers"

Some of these statements can be true, but most investors misunderstand the meaning. This essay will explore the illogic behind these statements and their common interpretation.

Active management does not outperform

There has been a lot of research done identifying that in general, actively managed strategies underperform passive index approaches. Studies also demonstrate that actively managed funds that outperform the market in one period are not likely to repeat that outperformance in the subsequent period. In fact, a prior period of outperformance may be more highly correlated with future underperformance relative to a benchmark[1]. The reasons that actively managed funds underperform simple market-weighted indices include:

- Highly efficient markets
- Higher fund management fees charged to investors
- Greater transaction costs
- Lack of investment discipline by managers
- Winning strategies are copied and lose their ability to generate alpha

When one takes these realities into account along with the realization that the stock market is a zero-sum game, the challenge facing an active manager becomes apparent. The zero-sum game concept means that the market return achieved by a benchmark is made up of all the returns achieved by active managers in the market. Overall gross returns for active managers must be equal to the gross return on an index. Then, because active management costs more, net returns on index funds will be higher.

Implications of the zero-sum game

The zero sum game has been described often enough (by John Bogle and William Sharpe, among others) but the investment world is full of experts and non-experts alike who seem ignorant about it or who choose to ignore it. This essay attempts to highlight the logical fallacy of making pronouncements that presume that there is potential for actively managed strategies, in general, to perform better or worse in certain periods relative to market benchmarks.

This essay is NOT about the simple question, described above, of whether active management outperforms passively managed index strategies over time. This essay tackles the fallacy that active managers, as a whole, have the potential to outperform (or underperform for that matter) a benchmark by choosing certain securities that will outperform the market. This is the fallacy that leaves the media, managers and investors with the impression active management will perform better during certain times, under certain conditions. Note that we will set aside active strategies where management consists primarily of moving in and out of markets, style categories, asset classes, geographies, etc.

Here are some illustrative questions, relevant to this issue:

[1] The Case for index-fund investing, Vanguard, April 2013

- "Can small cap managers outperform a small cap benchmark?
- "Can growth managers outperform a growth stock benchmark when the market is more volatile?"
- "Can emerging market managers outperform an emerging markets benchmark because these markets are inefficient?"

This fallacy is as much about what is unsaid as what is said. It is logically possible for active investment managers, as a whole, to outperform a benchmark, for a number of reasons, including:

- One universe of active managers could outperform at the expense of another universe of active managers that underperforms, as explored below.
- Active managers who select securities that deviate from the relevant benchmark, e.g. when a large cap fund includes some mid-cap stocks without a corresponding adjustment in benchmark.
- Investment managers may find ways to capitalize on illiquidity premiums or other risks for investors who can bear those risks, e.g. managers who overweight credit bonds against the Barclay's Aggregate index (this is very similar to previous bullet point).
- Different managers in the same market may use different benchmarks – e.g. due to differences in the Russell 2000 and the S&P Small Cap 600.

However, these explanations are not provided or even implied by the media or managers. Consumers of the media where these statements are made are left with the impression that certain market conditions allow any of those who search for mispriced securities to do this better or more easily.

Universes of investors

To understand this issue more thoroughly, let's split the universe of investors into three categories:

- Individual investors directly implementing their own strategies
- Institutional investors (pension funds, endowments, foundations, insurance companies, etc.) directly implementing their own strategies
- Fund managers (mutual funds, hedge funds, pooled trusts, separate accounts) investing on behalf of other investors

Together these three groups of investors make up the entire market. Or one can say that they make up the investor universe for any market we want to focus on – small cap, value, Canadian equities, etc. By definition, their combined returns will equal the relevant benchmark.

Better explanations of what happens in the market

Before clarifying how the zero-sum game makes the idea of better market conditions for active management irrelevant, we should acknowledge a few issues:

- Many strategies cannot be easily assigned a specific benchmark. Individual and institutional investors may consciously or unconsciously expose their portfolios to industries, risk factors, countries, etc. Fund managers may blend various approaches and strategies in a single fund such that identifying a true benchmark becomes difficult.
- Taking on exposure to risk factors, different from the benchmark, that generate higher returns can allow a fund to outperform. This can make comparisons to a benchmark almost irrelevant. This is often identified as the reason for the outperformance of so-called "smart beta" or "fundamental indexing" strategies.
- Fees for actively managed strategies may change from one period to another and this can cause differences in relative performance for actively managed funds from one period to another.
- Most actively managed strategies hold some low-yield cash investments.

All of this makes the exercise of adding up all the investment returns in any one "market" somewhat hypothetical. However, the zero-sum game remains conceptually valid. If one identifies any benchmark, it is logically possible to identify all the various investments that make up that benchmark. By definition the sum of the returns for the actively managed money for which the benchmark is relevant will equal the benchmark return.

It is logically impossible for the sum of all active management to perform better than an index fund in any period. It is also impossible for active managers in one period to perform better than active managers in another period, relative to their benchmark (ignoring changes in fees and transaction costs).

Once we realize that the total return on a benchmark is made up of all the investors in that benchmark, we can identify some changes in active management results that actually will arise from period to period:

- The percentage of funds (or investors) that outperform the benchmark may change, but there will necessarily be a corresponding change in the average level of outperformance. If 50% of active managers outperform in one period and only 25% outperform in the next period, then the outperformance in the later period by those who outperform must be by twice as much.
- It is possible for one of the three universes of investors to perform better relative to the other two universes in different periods. For example, fund managers may perform better in one period, at the expense of individual investors who perform poorly in that period.

However, these underlying reasons for the statements commonly made to describe market conditions are not provided, and are almost certainly not understood, by those making the statements, let alone by those reading or hearing the statements.

Markets with more dispersion of returns (note that cross-sectional dispersion of returns should be distinguished from return volatility over time), could provide opportunities for some managers to outperform a benchmark by more than in markets with less dispersion. However, the logical complement to this is that the underperforming managers will also underperform by more. There are bigger potential rewards, but also more risk for active managers in such a market. This hardly seems like the definition of a better market for active management.

It seems unlikely that the media will refrain from continuing to make misleading statements, since this topic is common fodder for many in papers, online and on TV. However, experts should think carefully about this topic and be sure that their own public statements and their research is consistent with the basic logic of the zero-sum game. Otherwise, the investment world is left less able to make sound judgments about allocating their own investments to active and passive strategies.

Evan Inglis, FSA, CFA is a Principal at The Terry Group. He can be reached at *evan.inglis@terrygroup.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.

The Best Model Doesn't Win

By Max J. Rudolph

Actuaries love models. Without models it would be much harder for science to substitute facts for appearances as suggested by John Ruskin. But it is easy to fall prey to ever more complex models, borrowing truths from physics to approximate the underpinnings of the financial world and assuming that past results provide the best information to predict the future.

This essay will argue that models have ceased to be the primary differentiators when discounting contingent events, the lifeblood of the actuary's work. Models provide the elixir that there is one correct answer, where today value is very much in the eye of the beholder. A good modeler can adjust assumptions to generate nearly any result. This raises the importance of an independent peer review.

Efficient market theory (EMT) proponents claim it provides the answer, but in reality it is only a preliminary figure later manipulated by herding, cognitive bias and challenges to the assumption that financial returns are normally distributed. While providing important information, EMT does not define the journey.

Qualitative Analysis and Goals

Models are only the first step. Qualitative analysis is the new driver of value determination. Assumption nuances, competitive advantages, margins of safety and emerging risks are all important in an environment that is complex, uncertain and ambiguous. Seeking out favorable imbalances can reduce risk while increasing returns. Experience matters.

Financial analysis starts by defining risk. Volatility is thought by many to be the primary risk metric. Easy to calculate and comparable between opportunities, mathematicians and traders prefer this type of metric. Portfolio managers working for third parties find it hard to explain. Investment professionals prefer to start with goals, objectives and constraints. While also quantitative, this approach quickly leads to telling a story about what is hoped to be achieved and what restrictions are necessary.

Goals are set based on specific time horizons. So are constraints. Someone might want to maximize savings after 30 years while not going insolvent in the meantime. When using EMT, practitioners tend to forget about the constraints because they are inconvenient. Focusing on constraints can lead you to adopt a slow and steady approach, avoiding leverage and saving more. Is this really a bad thing? Aiming for higher than a modeled "number" provides built-in conservatism in case the markets decide not to cooperate with your goals. This requires an early start and consistent funding of objectives. Defined benefit pension plans would perform better in the long run with this approach, front ending the funding rather than trying to make it up later. An investment strategy should consider the underlying lifecycle of the entity that funds it. An individual must invest during their working lifetime unless they are lucky enough to win the lottery or receive an inheritance. The same is true of a defined benefit pension plan. The plan's lifetime can go many years beyond when the firm funding the plan is active, making it surprising that DB plans are not required to conservatively pre-fund. Those with long time horizons can utilize time arbitrage to profit from those with shorter constraints.

This is not to say that models are not important, just that a recommendation should be arrived at from numerous approaches. There is so much information available in today's environment that it is hard to argue that a model, by itself, provides a comparative advantage. Today we have high government debt and investor leverage, unbalanced trade and sustainability issues unlike any seen previously. How can assumptions developed in stable times be thought to be predictive? Stress tests are much more useful in this type of regime.

A trade implies both a buyer and a seller at a single price. Understanding how your model differs from someone you might trade with/against can be important as you learn and understand your mental biases. Geographic location or access to information and rumors may hurt results rather than aid them. Many have found that turning off the Bloomberg terminal and moving to Omaha leads to improved results!

Models have not proven effective in accurately representing interactions between risks and events. Generally a model does well when the relationship is direct, or linear. It may even do a pretty good job when incorporating second order effects like convexity. When the phase, or regime, changes to reflect higher tail correlations, models unfortunately have a very poor track record. You may be able to show high correlations within a range of outcomes, but predictive timing is no better than rolling dice.

Efficient market theory has driven actuarial models for many years, assuming independent results and a bell-shaped, or normal, distribution. Studies in extreme financial events, and pragmatic experience, tells us that neither is always true. Value at Risk calculations are used extensively by banks, but fall prey to both of these assumption fallacies. For example, a 95% VaR ignores the worst 5% of results and effectively avoids all tail events. Risks are considered as silos or with independence assumed between them. This single metric was more valuable prior to it being used by regulators. Now companies design their product mix around it, leaving the firm susceptible to slightly different risks that are invisible to the metric. When using a single metric, this allows models to be designed backwards, knowing the end result and solving for the assumptions to get the answer you want. While Tail VaR, or CTE, performs better, regulators who seek transparency should ask for the underlying data and run it through multiple metrics to learn where the risks really lie. Another tool that is underused currently is to simply graph the underlying data. This makes it easy to pick out the scenarios that need to be reviewed.

A Better Option

A better option is to develop a range of potential outcomes driven by common sense, using both quantitative and qualitative analysis. Stress scenarios can be built starting with a severe result and reverse engineering the parameters. What outcomes are unacceptable? Can I design hedges to offset those results that meet my other goals and objectives? Can this hedge be designed into a product or do I need to find an external source to reduce this risk? Exposures should be reviewed on a gross basis, prior to any reinsurance or hedging, as a tail event might weaken or eliminate these tools.

For making decisions, cultural theory provides some useful approaches to creating teams. Not everyone thinks about risk in the same way. Being aware of this can help teams work together more effectively. These differences in thought process can be due to background or experience, and leads to unique solutions. This encourages groups, such as boards, to enlist a variety of approaches to decision making. How to implement this is a bit unclear, as multiple viewpoints could lead to paralysis or poor timing decisions as leadership rotation could lag results, but a natural tension between parties can produce better decisions.

While efficient market theory would say that timing does not matter, in reality the decision of when to buy or sell makes a big difference. Goals are not set relative to benchmarks. When I want to retire I don’t care if I outperformed an index fund. I care if I have enough money to stop working. This leads to the obvious conclusion that, without saving, nothing else matters. Most investors, individual and institutional alike, would be better off investing to their benchmarks and minimizing expenses rather than trying to add alpha, beta or any other kind of excess returns. You get most of the way to your goals this way without most of the downside. An opportunity cost approach can lower risk and avoid many cognitive biases.

Historical accounting practices can also get in the way of good decision making. At life insurance companies, for example, assets are treated separately from liabilities on financial statements. This is not how these portfolios are managed, and a better method would be to define a portfolio as being asset only, liability only or a combination of assets and liabilities. This would allow asset-liability management to be reflected in financial statements for these product lines. ALM is a preferred term to liability driven investing, which implies that liabilities are fixed and only investors have levers to improve results.

Conclusion

Decision making is improved when models are transparent and peer reviewed by experts. It is easy in today's uncertain environment for management to hide behind complex models. Regulators accept these models rather than hiring their own experts to peer review them. A checklist review does not add value, whether attempting to identify minimal or best practices.

As Albert Einstein said, "Make everything as simple as possible, but not simpler." Modelers need to concentrate on telling a story rather than building additional layers of complexity. It will ultimately lead to better resiliency and decision making, which was our goal all along. People, and their experience, matter.

Max J. Rudolph, FSA, CFA, CERA, MAAA, is the owner at Rudolph Financial Consulting, LLC. He can be reached at *max.rudolph@rudolph-financial.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.

The Myth of Time Diversification

By Rowland Davis

In 1963 Paul Samuelson published a paper entitled "Risk and Uncertainty: A Fallacy of Large Numbers." Thus was born the phrase: "the myth of time diversification."

The purpose of this essay is not to challenge the accuracy of Paul Samuelson's work, but to challenge the expansive misuse of his findings – an abuse that has substantial implications for actuaries. As an example, a Google search of the phrase yields this quote:

> "It sounds nice in principle, but it's actually an example of the 'time diversification' fallacy. Investments do not become safer the longer they are held. Time reduces the variance in the average annual return, but it actually increases the variance in the cumulative return. In other words, smoothing won't bring more certainty to retirement savings. For any given portfolio, collective DC plans face the same risk-return tradeoff as ordinary 401(k) plans." Jason Richwine in the *National Review* blog

To understand the abuse occurring here, we must return to Paul Samuelson's work. The specific application to investment risk was first developed in his 1969 paper "Lifetime Portfolio Selection by Dynamic Stochastic Programming." It was, in fact, a mathematical proof – of the general nature "**if this**, then **that**", where **that** is essentially the statement that time horizon should not affect an investor's risk tolerance. (The corollary to this is more frequently used – that the risk of stock investing does not decrease with longer time frames.)

Unfortunately, the **if this** conditions are almost universally ignored, and the proof only holds with those conditions in place. There are two important conditions that Samuelson uses to frame the whole analysis: 1) that the investor's utility function is isoelastic (i.e. a single continuous utility function covers the entire spectrum of outcomes, without conditional sensitivity to any particular values of the outcome), and 2) that the only issue at stake is an individual investor's terminal wealth based on the investments alone. In this case, and only in this case, is it wrong to assume that stringing together a sequence of risky bets is superior to a single risky bet (i.e. time does not diversify risk).

Actuarial work involves collective systems, so can the same logic be applied? Is it wrong for a group of investors saving for retirement to collectively take more risk over a longer time frame than they would over a shorter time frame? This essay shows that it is ***not*** wrong to do so in the real world (i.e. free of the narrow constraints on the Samuelson proof).

Since I am not an academically trained economist, I will construct an actual example to make the point. Although the words used are somewhat opaque to a non-economist, Samuelson acknowledges that real world investors might indeed have more risk tolerance in the early stage of their career: "Note: if the elasticity of marginal utility…rises empirically with wealth, and if the capital market is imperfect as far as lending and borrowing against future earnings is concerned, then it seems to me likely that a doctor of age 35-50 might rationally have his highest consumption then, and certainly show his greatest risk tolerance then – in other words be open to a 'businessman's risk.' But not in the frictionless isoelastic model!" (The reference here to a "businessman's risk" is explained elsewhere in the paper as the ability to take more investment risk.) Because the "frictionless isoelastic model" is not very relevant in the real world, the door is immediately open to investment policies that do, in fact, depend on time frame. Target date funds are one simple example, based on the concept of including the value of human capital as part of the investor's wealth.

My example will assume two assets: a safe asset with an expected real return of 2%, and a standard deviation of 5%; and a risky asset with an expected real return of 4.5%, and a standard deviation of 20%. For the Samuelson base case, I use a standard risk averse utility function that meets his

if then conditions: $U(w) = \frac{(w^{\lambda}-1)}{\lambda}$, with λ = -2. With this function, utility is maximized with a risk asset allocation of around 25%. And as Samuelson proved with his equations, a stochastic simulation verifies that this same allocation is the utility-maximizing allocation with both a 10-year horizon and a 30-year horizon.

Now we move into the real world. First we develop a new utility function that reflects an investor (or a group of stakeholders in a collective plan) with a 3% real return target. For this investor real returns in excess of 3% have a decreased marginal value, and real returns less than 1% become painful very quickly. Here is a graph of the utility function I use for this case.

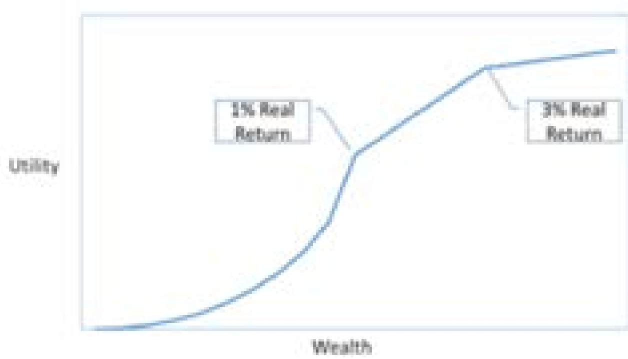

This kind of utility function has been shown by behavioral finance research to represent the way that humans make decisions in the real world (i.e. prospect theory, developed by Kahneman and Tversky).

With this utility function, a 10-year investor will maximize utility with a risk asset allocation of about 20% — very similar to the Samuelson base case. But a 30-year investor will maximize utility with a risk asset allocation of about 60%. For this investor, the time frame does matter, with more risk becoming appropriate over longer time frames. (For a similar example see "The fallacy of large numbers revisited" by De Brouwer and Van den Spiegel, *Journal of Asset Management,* 2001).

Now let us proceed to the issue of human capital. Assume that this investor, seeking a 3% real return, adopts a strategy of dynamic adjustment for his saving plan. After 10 years, if savings fall below 90% of his real return target, he will make additional contributions over the next 5 years with a total value equal to the shortfall relative to the 90% threshold. If savings after 10 years exceed 120% of the real return target, then part of the surplus will be withdrawn. The amount withdrawn is sensitive to the asset allocation, but will always be set so that the expected value of the adjustment process is zero (i.e. expected withdrawals will equal expected additional contributions). The investor is comfortable with this adjustment strategy because his human capital is sufficient to absorb any required additional contributions.

With this dynamic adjustment process, the 30-year investor will now find maximum utility with a 75% risk asset allocation, instead of 60%. Interestingly, even with the standard utility function this adjustment process will move the optimal risk asset allocation for the 30 year investor up to 35%, from the 25% level that applies to the 10-year investor with no adjustment process. Once again, real world details matter when thinking about the relationship between risk and time frame.

Collective systems involve spreading risks among stakeholders and across age cohorts in ways that allow for efficient risk-taking. Human capital is not only recognized, it is pooled – within a single closed cohort, human capital diminishes in value over time, but the aggregate human capital across the full range of cohorts remains constant. Unlike the fund for an individual investor, which builds from a level of zero to ever larger dollar totals, a mature collective fund is expected to remain relatively constant in real terms. A dynamic self-adjustment process (through variable contribution inflows and/or variable benefit outflows) can create a sustainable fund where the risky bet can be repeated time after time with controllable risk. There will always be risk over any specific time frame, but a properly designed system can manage these risks through time in a sustainable way. Risk is no longer measured simply by some value of terminal wealth (as in the Samuelson paper), but by more complicated metrics of ongoing financial risk exposure to various cohorts of stakeholders. Paul Samuelson never said anything different.

The bottom line on this is that critics have the right to say that risks do exist, and need to be carefully measured and managed. And critics also have the right to express their honest opposition to collective systems (i.e. those involving inter-generational risk sharing) on political grounds.

But they do not have the right to invoke Paul Samuelson's proof within any blanket statement asserting that collective systems can't work because they are based on a fallacy. Implicit in any argument of this type is an assumption that a collective system can be simply decomposed into segments consisting of "classical" individual investors – but then they are no longer talking about a collective system, which is far more complicated in its risk dynamics.

Technical Endnote: Samuelson himself acknowledged in a 1989 paper ("The √N Law and Repeated Risktaking" in: Probability, Statistics, and Mathematics, Papers in Honor of Samuel Carlin) three separate cases, using different assumptions, where time frame would change a rational investor's risk tolerance. One of these is the simple one of including human capital in wealth. A second one recognizes that the original argument does not hold if markets are mean reverting (and there is substantial evidence that they are). The third involves an assumption set using a utility function that incorporates some minimum required threshold for terminal wealth, similar in concept to the one used in this essay. Samuelson was well aware of his own if then criteria.

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.

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

Traditional Log-normal Model

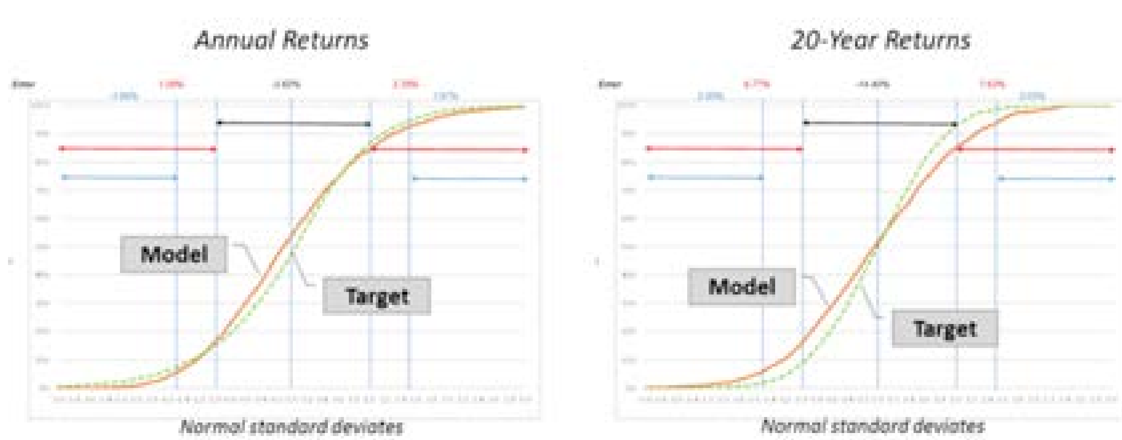

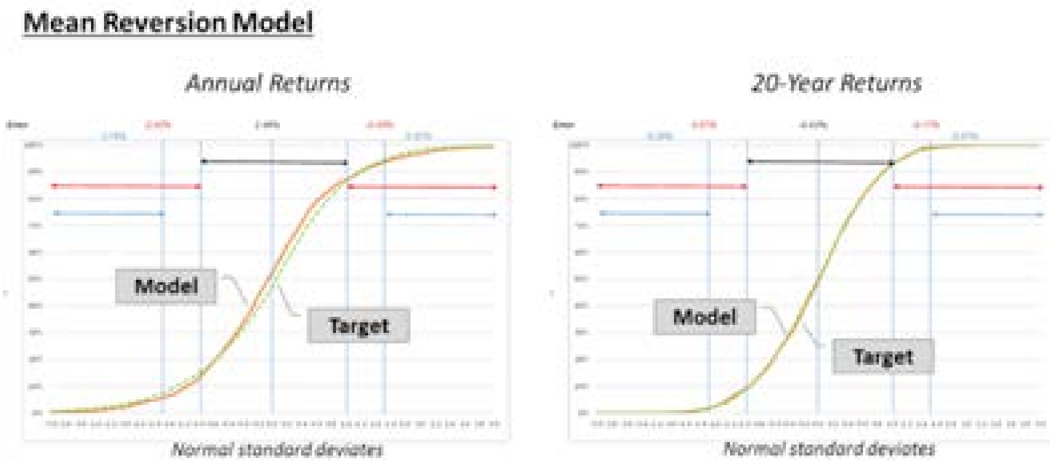

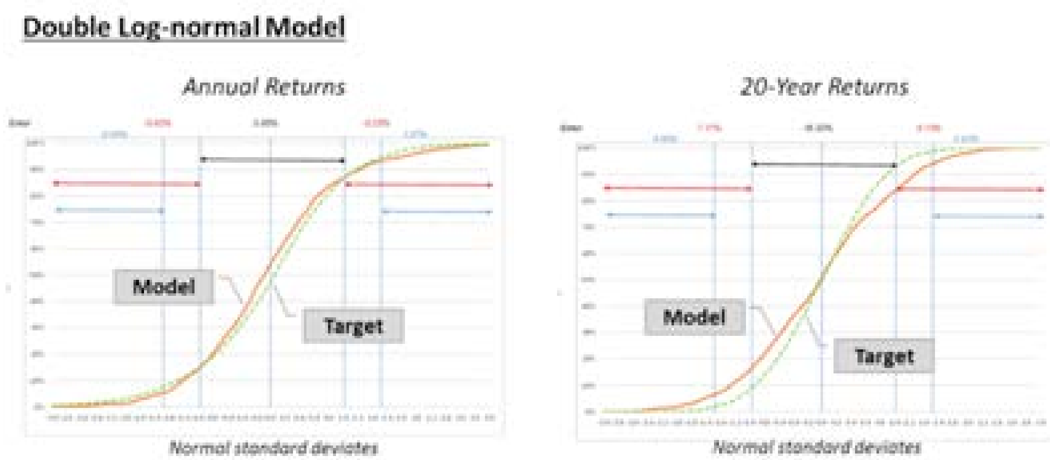

over-reaction to a crisis of some sort. The mean reversion impact is then factored into the returns over the following six years. (Note that my final model also includes a simple regime-switching feature to boost the fat-tail exposure a bit – but the mean reversion feature is, by far, more significant 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%	7.58%	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:		
Prob. return less than:	*Log*	*DoubleLog*	*Mean Reversion*	*Log*	*DoubleLog*	*Mean Reversion*
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.

The "Dollar Bill" Fallacy

by Gary Thomas

Taxes can be applied in many different ways. Some taxes are applied as flat fees (on fishing licenses for example) while others such as highway tolls are loosely tied to the distance driven. Sales taxes and customs duties are generally applied to the value of the product while taxes on property and personal assets are also some function of the value of the property. In all these cases, we instinctively understand how such taxes are applied and arrange our lives accordingly. Furthermore, the formula used to levy these taxes tends to be uniformly applied and doesn't often change.

In contrast, taxes on income tend to be controversial, partly because we all have different types and amounts of income and have strong views about who is more deserving. Also, since these taxes are not uniformly applied, higher taxes for your neighbor might mean lower taxes for you! Policy arguments in this high stakes game are often couched in terms of "fairness".

Because income taxes tend to be applied directly to each marginal dollar of income and one dollar looks very much like any other dollar, it is easy to fall into the fallacy that it is the dollar itself that is being taxed, and not the activity. This "dollar bill" fallacy shows up in several widely repeated arguments related to income and taxes.

The first application of the dollar bill fallacy arises when people take a dollar and follow it through two serial events. We see this in the assertion that estate taxes count as "double taxation" on the grounds that the dollar had been previously taxed when earned as income. We also see this so-called double taxation with corporate dividends when investors are taxed on dividends that have already been subject to taxation as corporate income.

In order to analyze these situations more clearly, it is helpful to take our eyes away from the dollar bill and to look instead at the separate decision points along the way. Every time a dollar passes through a different decision point, the tax at that point has to be evaluated separately from the taxes at other decision points.

Let's look first at estate taxes. After a dollar of income has been taxed, there are a lot of different things that can be done with the remainder – it can be spent or perhaps saved or invested. In all these cases further taxes may be applied. The decision to earn the money is entirely separate from the decision on how to dispose of it. Likewise, estate taxes are not inevitable, but the result of a conscious choice not to spend everything prior to death (that's why they call it estate planning!). Clearly, the thought of paying tax for simply dying must sting; however, in a free country each person chooses their own path in full knowledge of the tax consequences for each step along the way.

A similar perspective can be applied to double taxation of corporate dividends. Each nation can set the taxation of investment earnings of its citizens and it can also set the taxation of its corporations. Citizens of one country are at liberty to invest in shares of corporations from another country. The decision of a corporation to domicile in a country is entirely separate from the investment decisions of its shareholders living in different countries. Indeed, the recent public struggles of governments to figure out how to properly tax Amazon, Apple and other non-traditional companies provide evidence that there is nothing inevitable about this double taxation.

A second, and perhaps more important, application of the dollar bill fallacy arises when considering the appropriate tax rates to apply to different types of income. This issue received widespread attention during the last US presidential campaign when Mitt Romney's 15% tax rate on capital gains was compared to the 35% tax rate supposedly paid by Warren Buffett's secretary. In the eyes of many commentators this was immoral and inequitable and justified the 3.8% increase in the capital gains tax rate that had been used to help fund Obamacare.

But let's look a little more closely at these types of income; as before, the dollars may look the same, but the decisions, effort and risks involved in earning them clearly can be quite different for an investor and a worker, and the application of the two tax rates isn't quite apples to apples.

Firstly the risks associated with different forms of income are not at all comparable. An investor is often taking significant risk that they may not only fail to make any investment income, but may even lose some or all of the investment itself. In contrast, most wage earners know exactly how much they will earn when they walk into work each day; and while some salespeople might not make any commissions at all, they certainly do not run the risk of actually losing any of their own money.

Secondly, a significant component of investment income is comprised of inflation and this is also taxed. According to The Tax Foundation, this hidden tax had the effect of driving up the effective capital gains tax rate on investments since 1950 from an average of 26.4% to 42.5%. Indeed, to take an extreme case, a purchase against the S&P500 index in July 2000 would yield a nominal taxable gain of 18% by July 2013 but a loss in real terms, so that the tax rate was effectively infinite.

Investors must also consider the costs involved in holding investments. A 1.5% annual fee in an actively managed mutual fund would eat away 24% of the value over a 20 year period. Not only does this reduce the potential gains, but it also increases the possibility of a loss.

Against all this, a worker's income requires hard work and a substantial time commitment.

It can be seen from the above that "fairness" is an elusive concept when trying to determine the appropriate level of taxes on investment income relative to taxes on salaried income. Does it matter whether the average person understands these differences and appreciates that the effective inflation-adjusted capital gains tax rate is actually much higher than advertised? Perhaps not, although it is clear that the investor class understands. So how would we expect them to react?

This brings us to a big difference between investors and wage earners – their behavior in response to taxes and other stimuli. The average wage earner will generally keep doing the same job in order to pay down their house, support their family and save for retirement no matter how much tax rates may move. But many investors have significant discretion as to whether to invest and how.

When making an investment decision, an investor will evaluate their potential return net of expenses, inflation and taxes and decide whether the return warrants the risk. If it doesn't, then they may choose to favor current spending over investments and future spending. They could buy more goods and services, take more vacations or perhaps put their money in things that hold their value and can be enjoyed now like art, collectibles and property.

Surely all this current spending would be plowed back into the economy? Yes it might, although there is a big difference between money that is invested in new companies, new factories, and new technology and money that is spent on yoga instructors, gardeners, and maids. One builds the economic infrastructure of a country and the other one doesn't. In the extreme case, potential investors might even choose to take themselves and their wealth out of a country and employ it elsewhere. This has recently happened most notably in France (in response to punitive tax rates) and even in the US where the number of citizenship renouncements

The "Dollar Bill" Fallacy *by Gary Thomas*

have hit an all-time high. In the case of both countries, the numbers may be small but the message is unmistakable.

Clearly, decisions to invest can have a significant and positive impact on society since they can lead to job creation and economic vitality. The mobility of capital in the modern world makes it quite responsive to incentives, and governments everywhere are sharpening their pencils as they figure out exactly what they want to tax and how to make their societies both fair and efficient. It is worth noting that a number of countries encourage investment by allowing workers a limited amount of investment earnings free of income taxes. So there is little doubt about the importance of investing.

So how do we move forward? Economic activity requires application of both labor and capital. The dollar bill fallacy is distorting thinking on several aspects of investing; continued use of faulty logic will lead to sub-par outcomes. If we want to start thinking properly about investment incentives and taxation, then the first step is to move past the dollar bill fallacy and start thinking about how people actually make decisions and respond to incentives. If we don't do it, others will.

Gary Thomas FSA, MAAA is an Assistant Director, Long Term at Bermuda Monetary Authority. He can be reached at *GThomas@bma.bm*.

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.

The Fallacy of the Fed Model

by David R. Cantor, Adam Butler and Kunal Rajani

Managers responsible for asset allocation decisions rely on a variety of models to forecast future equity market returns. These forecasts inform policy portfolios and tactical shifts, and are used for budgeting purposes.

Most equity market valuation techniques rely on comparisons between current equity market values and equity market values observed over many decades in the past. For example, the trailing price to earnings (P/E) ratio is often compared with long-term average (P/E) ratios. James Tobin proposed an adjusted balance sheet measure called the Q Ratio (combined market value of all companies should be about equal to their replacement costs), while Warren Buffett claims to watch the level of aggregate corporate earnings to Gross Domestic Product.

In contrast, the so-called Fed Model is distinguished from other common models by its reliance on a comparison between equities and bonds. Specifically, the Fed Model compares the earnings yield (E/P) on the stock market with current nominal yields observed on 10-year Treasury bonds (Y), so that the value of a Fed Model valuation is calculated as (E/P) – Y.

Proponents of the Fed Model argue that stocks and bonds are competing assets so investors should prefer stocks when stock yields are high relative to bonds, and bonds when bond yields are high relative to stocks. Many augment these assertions by noting that equity prices should reflect the discounted present value of future cash flows; as the discount rate (Treasury yields) declines, so should equity valuations increase. Indeed, strategists might be forgiven for entertaining the above notions given that equity market valuations tracked interest rates quite reliably for over four decades from 1960 through 2007.

Unfortunately, the Fed Model does not hold up under more rigorous theoretical and empirical scrutiny. In fact, as we will endeavor to demonstrate in this article, the Fed Model has very little theoretical support; leads to poor allocation decisions, and; is not significantly predictive of future stock market returns.

The Fed Model is Based on a Faulty Theoretical Premise

While it might appear to the casual investor that the Fed Model deserves attention on the basis of sound intuition, the financial literature is consistent in its condemnation.

Let's take for example the suggestion that, because stocks and bonds are competing assets, investors will compare the yield on stocks, as measured by the E/P, to the nominal yield to maturity on 10-year Treasuries, and favor the asset with the highest yield. Presumably, capital would then flow from bonds into stocks, thus lowering stocks' E/P until equilibrium is achieved.

However, it is not obvious that (E/P) is the appropriate measurement of yield for stocks. Earnings yield as applied in the Fed Model is not comparable to the equivalent bond yield as only a portion of the earnings is actually distributed to shareholders. Rather, the dividend yield or total shareholder yield including share buybacks and share retirement might represent a more comparable proxy.

In addition, Asness (2003) illustrated how yield equivalency would rarely result in equivalent total returns because of the impact of inflation and growth in corporate earnings. Assume nominal bond yields are 8%, the equity market P/E is 12.5 (1/.08), inflation is 6% and expected real earnings growth is 2%. Under the standard Dividend Discount Model, it can be shown (holding payout ratios constant at 50%) that stocks are expected to deliver 12% nominal returns, implying 4% excess returns relative to Treasuries[1].

However, in the event inflation falls to 1% while nominal bond yields fall to 3% (preserving their 2% real yield) real growth rates remain constant at 2%. As a result, nominal earnings growth falls to 3%. Recall the Fed Model assumes that the earnings yield will drop to 3% in-line with contemporaneous Treasury yields, which translates to a P/E ratio of 1/0.03 equal to 33.33. If we feed these new assumptions into our Dividend Discount Model, we observe that expected stock returns have now fallen to 4.5%, just 1.5% more than bonds.

Under the Fed Model, stocks and bonds compete for capital, yet Asness' analysis illustrates how simple shifts in inflation expectations would result in a logical inconsistency, which invalidates the basic premise of the Fed Model. Why should a shift in inflation cause expected returns to stocks to drop by more than bonds if the two should be valued exclusively on the basis of relative yields?

Moreover, why should investors expect stock earning yields to adhere to Treasuries' gravitational pull? Isn't it just as likely that Treasury yields are mispriced, and will correct to the level of earnings yields? This is an especially acute point in the current environment, where central banks have explicitly stated to artificially lower rates across the curve.

Another argument often used to support the Fed Model is that low interest rates suggest a high present value of discounted cash flows and therefore a high P/E. The problem is that all else is not equal when interest rates are low. When interest rates are low, prospective cash flows to investors are also likely to be low. The decline in prospective cash flows offsets the decline in the discount rate. Therefore, it is not necessarily true that low interest rates justify a higher P/E (i.e lower the E/P) .

The Final Arbiter: Fed Model as a Forecasting Tool

Setting aside for a moment the weak theoretical foundation of the Fed Model, we must acknowledge that proponents of the technique appear to have a meaningful empirical argument given the strong relationship between E/P and Treasury yields over the period 1960 – 2007. However, it is worthwhile exploring whether this relationship was unique to the dominant interest rate regime over this period.

In fact, reasonably good data exists for both U.S. equity market E/P and 10-year Treasury constant maturity yields dating back to 1871, and even further with some databases. When this longer period is used, the Fed Model relationship does not hold (Exhibit 1). While the r-squared coefficient for a regression of monthly E/P on 10-Year Treasury yields between 1960 and the present is 0.49, we observe much lower explanatory power in the historical record back to 1871, with an r-squared value of just 0.03. This observation is consistent internationally: analogous data, sourced by Estrada (2005), for several other large countries and demonstrated that the insignificant statistical link between E/P and government bonds is universally persistent .

---

1 Under the Dividend Discount Model, the expected return on the market equals the current dividend yield plus the long term nominal growth rate of dividends. The dividend yield can be expressed as the payout ratio multiplied by earnings. If we assume a constant percent of earnings then growth rate of dividends equals the growth rate of earnings. We can then express the return on the market to equal: payout ratio multiplied by the earnings yield plus the growth in nominal earnings

2 This also ignores changes in the risk premium associated with stocks. The risk premium can also be time-varying and affect the pricing of stocks.

3 In fact, if the P/E ratio in the numerical example given above remains at 12.5, not 33.33 as implied by the Fed Model, the 4% expected return of stocks over bonds would actually be preserved.

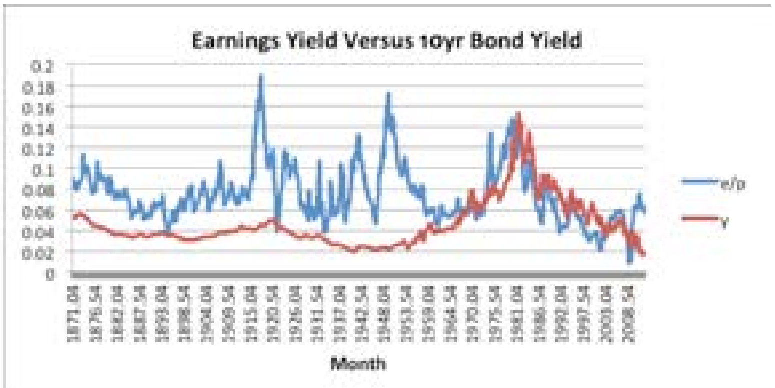

While regression analysis implies a spurious and non-stationary relationship between earnings yields and Treasury yields, the true arbiter of validity must be how well the Fed Model forecasts stock market returns. To test, we regressed forward total nominal and real returns to stocks over a variety of forecast horizons against contemporaneous Fed Model values. For comparison, we also regressed forward returns against simple trailing E/P ratios with no adjustment for the level of interest rates (Exhibit 2).

Exhibit 2

	Regression of Nominal S&P Returns Against Fed Model and Earnings Yield							
	1871-Present				1960 - Present			
	E/P		E/P-Y		E/P		E/P-Y	
S&P return	coeff	R^2	coeff	R^2	coeff	R^2	coeff	R^2
1yr	0.95	1.9%	0.32	0.3%	1.03	3.0%	0.00	0.7%
5yr	4.64	5.7%	0.94*	0.4%	4.90	7.2%	-5.04	3.7%
10yr	13.50	18.2%	3.33*	1.6%	15.54	32.7%	-3.76*	0.6%

*The coefficients obtained also have a p-value higher than 0.1 and are thus not significant

The regression analysis shows the Fed Model has minimal predictive ability over time horizons of 5 and 10 years. In fact, univariate regression using just the E/P does a much better job in forecasting future stock market returns. If anything, adjusting for the level of interest rates destroys any predictive ability achieved by using the simple earnings yield alone.

One other way to demonstrate the fallacy of the Fed Model in making useful investment decisions is to perform a decile analysis. We sorted Fed Model readings into deciles and calculated the average forward returns to stocks over 1, 5 and 10 year horizons for each decile. If the Fed Model exhibited forecasting ability, we would expect to see a somewhat linear relationship between starting Fed Model valuation level and average forward returns. In fact, there is no obvious linear relationship whatsoever (Exhibit 3).

Exhibit 3

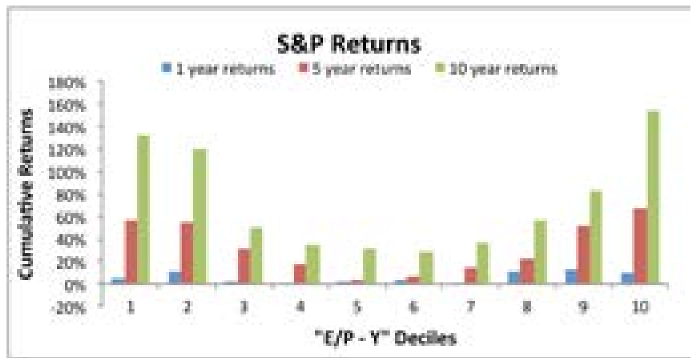

From Exhibit 3 we see that nominal stock market returns are high when the Fed Model indicator signals extreme levels of equity market *under-(decile 1)* or over-valuation (decile 10). There may in fact be a meaningful signal there, but clearly it is inconsistent with the theoretical foundations of the model.

Perhaps the Fed Model's most profoundly misguided signal came in 1982. The Fed Model suggested the market was fairly priced precisely when more reliable indicators suggested markets were cheapest on record. Of course, subsequent returns over horizons from 1 through 20 years were well above average.

Conclusion

The Fed Model implies that high stock market multiples are not a cause for concern for investors because these multiples are justified by low interest rates. Unfortunately, investors relying on such logic to invest in the stock market are likely to be very disappointed in the coming years. While low interest rates may *explain* why investors assign such high stock market multiples, low rates do not *justify* such high multiples.

Investors would be better served by heeding the many more reliable valuation metrics currently signaling caution. Moreover, those responsible for institutional portfolios should prepare for a lower return future for equity markets from current levels.

References:

Asness, C. *Fight the Fed Model*. Journal of Portfolio Management. Fall 2003

Buttonwood's notebook. *A Misleading Model*. Economist.com. August 3, 2013

Buttonwood's notebook. *Burying the Fed Model.* Economist.com. November 29, 2012

Estrada.J. *The Fed Model: A Note*. IESE Business School. November 2005

Hussman, J. *Inflation, Correlation, and Market Valuation*. May 29, 2007 Weekly Market Commentary

Hussman, J. *Long-Term Evidence on the Fed Model* and *Forward Operating P/E Ratios*. August 20, 2007 Weekly Market Commentary

Hussman, J. *Explaining is Not Justifying*. July 11, 2005. Weekly Market Commentary

Keefe, T. *Breaking Down the Fed Model*. Investopedia.com. February 22, 2013

Leuthold Group. *The Fed's Stock Valuation Model*. June 1999

Ritter, J. *The Biggest Mistakes We Teach*. Journal of Financial Research. Summer 2002

David R. Cantor CFA, FRM, ASA is Director at PwC in New York, NY. He can be reached at *david.r.cantor@us.pwc.com*

Adam Butler, CFA, CAIA is Senior Vice President and Portfolio Manager at DundeeGoodman Private Wealth, Toronto, Canada. He can be reached at *adbutler@dundeegoodman.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.

The Misconceptions of Retirement Risks

by Dimitry Mindlin

Financial theory suggests that a *cash flow should be discounted at a rate that is consistent with the cash flow's inherent riskiness*. This seemingly uncontroversial statement, when applied to pension plans, has raised a great deal of controversy in recent years.

There are good reasons to believe that the financial commitments of most pension plans are "low-risk." However, the discount rates selected by most pension plans are higher than today's "low-risk" rates. This disparity is especially pronounced for public pension plans. Most economists reportedly consider such selections as a violation of the principles of finance. The highest echelons of the economic hierarchy appear to wholeheartedly support this perspective:

> *"While economists are famous for disagreeing with each other on virtually every other conceivable issue, when it comes to this one there is no professional disagreement: The only appropriate way to calculate the present value of a very-low-risk liability is to use a very-low-risk discount rate.*[1]*"* Donald L. Kohn, Vice Chairman, Federal Reserve.

This perspective has profound consequences for pension management in general and pension investing in particular. Mindful of these consequences, many pension practitioners do not support the perspective despite the alleged consensus among economists. Overall, this perspective is in conflict with the current prevailing practices in the pension industry.

The purpose of this article is to demonstrate that this perspective is theoretically suspect and illuminate certain built-in misconceptions. These misconceptions include, but are not limited to, the five misconceptions outlined in the next section.

The key messages of this article are the following. The utilization of "low-risk" discount rates to discount "low-risk" financial commitments is a choice, not a requirement supported by a sound economic theory. The claim that there exists "the only appropriate way to calculate the present value" is little more than an attempt to apply certain principles of finance beyond the scope of their applicability.

Five Misconceptions

Misconception #1: Plan Sponsor Risk vs. Plan Risk

The primary risks embedded in bond pricing and actuarial valuations are fundamentally different. The financial health of the *bond issuer* is at the heart of bond pricing. In contrast, the financial health of *the pension plan* – not necessarily the plan sponsor – is at the heart of a typical actuarial valuation.

The primary risk embedded in bond pricing is that *the bond issuer* may not make all required payments. The primary risk embedded in actuarial valuations is that a *particular value of pension assets* may be insufficient to make all required payments.[2] These risks are clearly different and may require different discounting procedures. Those who believe that there is "the only appropriate way to calculate the present value" explicitly disregard a multitude of challenges and risks that may require different discounting procedures.

The perceived connection between "low-risk" pension commitments and "low-risk" discount rates is an example of a valid concept that becomes inapplicable when taken out of its proper context.

---

1 The Economic Outlook, Vice Chairman Donald L. Kohn, Federal Reserve, the National Conference on Public Employee Retirement Systems Annual Conference, New Orleans, Louisiana, May 20, 2008, *http://www.federalreserve.gov/newsevents/speech/kohn20080520a.htm*

2 The value of pension assets includes the existing assets and present value of future contributions.

Misconception #2: Pricing vs. Funding

The objective of a pension plan is to fund its benefits. The funding objective is much more expansive than the objective to price the "matching assets" for the plan's accrued benefits. Most pension plans are sophisticated investors that utilize a broad variety of financial assets. "Matching assets" represent a small segment of the assets available for pension investing.

The price of an asset is relevant only if the asset is under consideration for investing. Pricing pension benefits is essentially an asset allocation decision, not a theoretical requirement. Forcing all pension plans to evaluate specific assets regardless of the plans' policy portfolios makes little sense.

The uncritical application of the principles of bond pricing to pension funding is another example of a valid concept that becomes inapplicable when taken out of its proper context.

Misconception #3: Economic Foundation

The suggestion that pension benefits must be priced similar to tradable bonds is based primarily on the Law of One Price. This law essentially states that identical assets should have identical prices. There are stringent conditions, however, that must be satisfied for this law to be valid. These conditions include the tradability of the assets – long and short. It is well-known that this law is not necessarily valid when these conditions are not satisfied.

Pension benefits are non-tradable and non-transferable, thus the Law of One Price is inapplicable. Moreover, even if we considered pension benefits as "assets held short," then these illiquid "assets" and their liquid counterparts generally should have different valuations.

The economic foundation for the requirement that pension benefits must be priced similar to tradable assets is shaky at best and probably non-existent.

Misconception #4: $100 of Stocks Is the Same as $100 of Bonds

The suggestion to discount pension benefits by "low-risk" rates is in part based on the notion that economic value of pension benefits does not depend on the manner these benefits are funded. This notion is often expressed as "$100 of stocks is the same as $100 of bonds."

While an auditor may support the notion that $100 of any asset is still $100, the capabilities of $100 of stocks and $100 of bonds to fund a future cash flow are quite different. Moreover, the prices of otherwise identical derivative-based instruments for stocks and bonds are generally different as well. Therefore, even if we consider today's asset prices only, *$100 of stocks is not the same as $100 of bonds.*

The notion of "*$100 of stocks is the same as $100 of bonds"* is yet another example of a valid concept that becomes inapplicable when taken out of its proper context.

Misconception #5: "The Only Appropriate Way"

It is hard to find an object that has one measurement that is clearly superior to all other measurements. Virtually all objects have multiple attributes that require multiple measurements. Yet, the proponents of "the only appropriate way to calculate the present value" essentially have designated pension benefits as an object that is uniquely different.

Once again, this perspective is unsubstantiated. The claim that pension benefits allow only one "appropriate" measurement defies common sense.

Conclusion

In recent years, many pension practitioners have been sharply criticized for not embracing the so-called "financial economics" perspective on pension plan management. The acceptance of this perspective would have dramatically changed the prevailing practices of pension management. There are reasons to believe that these changes would have negatively impacted the retirement security of numerous plan participants and put the DB system under additional pressure.

This criticism continues to this day. This author hopes that this article would be helpful to pension practitioners and bring some clarity in the ongoing debate regarding better pension management.

Finally, the following quote should serve as a reminder to those who believe that the alleged consensus among scientists is a valid scientific argument.

> *"Science is not a democratic institution. Scientists do not resolve their disagreements by plebiscite, acclamation, voice vote, or any other democratic means. To a courteous scientific debate, scientists contribute books and scholarly articles, which gain recognition via the quality of their contents. In the presence of quality academic publications, any "consensus" declaration is needless. In the absence of quality academic publications, any "consensus" declaration is useless. Either way, the claim "every economist knows this" is an inconsequential line of reasoning as well as a clear sign of weakness of one's arguments."*[3]

It remains to be seen if the abovementioned consensus can withstand a close scrutiny. This author is exceedingly skeptical, thank you very much.

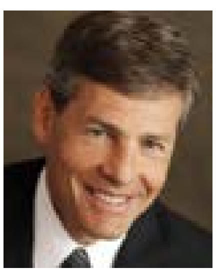

Dimitry Mindlin, ASA, MAAA, PhD, is president of CDI Advisors LLC. He can be contacted at *dmindlin@cdiadvisors.com*

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.

3 Mindlin, D. 2010. The "Financial Economics" Debate, CDI Advisors Research, April 2, 2010, *http://www.cdiadvisors.com/papers/CDIFinancialEconomicsDebate.pdf.*

On the Validity of Common Portfolio Return Assumptions

by Dimitry Mindlin

The Society of Actuaries has commissioned an independent Blue Ribbon panel to issue "recommendations for strengthening public plan funding." The report issued by the commission contains a number of practical recommendations that should improve public plan funding and management. The Actuarial Standard Board is likely to consider these recommendations in the development of actuarial practice standards.

One of the panel's major findings is the recommendation to focus on "median expected future investment conditions" and "median expected outcomes." In particular, the panel recommends to use "the median expected return" as the discount rate. These recommendations significantly affect the calculations of actuarial present values, contributions, and funded status.

Yet, while the logic of the report is reasonable, its language is occasionally imprecise and open to (mis)interpretations. The wording of some statements may imply certain relationships between the concepts utilized in the report that are actually not true. Given the significance of these concepts, this paper highlights these important issues.

Median Return vs. Geometric Return

The report's key statement regarding the selection of the assumed rate of return is the following:

> *"The Panel believes the assumed rate of return should be set at the median expected return, which should be based on the geometric mean return."*

Taken at face value, the part of this statement that claims a relationship between the median and the geometric returns is problematic. Generally, the median and the geometric returns are not the same. Normal distributions would represent one example of this observation.

But let us give this statement the benefit of the doubt and view it in the context of the current practices in the pension industry. Most pension plans use forward-looking capital market assumptions (CMA) that specify the expected return, volatility, and correlations between the major asset classes under consideration. These CMA are used to calculate the expected return and volatility of portfolio returns. Furthermore, there are robust estimates of the geometric expected return based on the expected return and volatility of return.[1]

The calculations of median returns, however, require additional assumptions that deal with the shape of return distribution. One of the most prevalent assumptions of this kind is the assumption of lognormal *portfolio returns*. Under this assumption, the median and geometric returns are the same.[2]

The assumption of lognormal portfolio returns, however, has a glaring mathematical problem. While the assumption that asset class returns are lognormal creates no mathematical problems, the distributions of *portfolio* returns – linear combinations of asset class returns – are not necessarily lognormal. Generally, a linear combination of lognormals is not lognormal.

Technically speaking, the assumption of lognormal portfolio returns represents a lognormal *approximation* of linear combinations of lognormals. This approximation is based on matching the first two moments of the underlying distribution. The key question is, how good is this approximation?

To answer this question, let us assume that all individual

---

[1] See Mindlin [2010] for more details.

[2] See Mindlin [2011] for more details.

asset classes have lognormal returns and examine the impact of the *lognormal* portfolio return assumption on the key measurements of portfolio returns – the arithmetic expected return, the geometric expected return, and the median return. The choice of these measurements was driven primarily by their role in the selection of discount rates.

To estimate median returns, this paper utilizes the following approach. Given a portfolio, conventional CMA, and the assumption of lognormal *asset class* returns, we calculate the first *three* moments of the portfolio return. Then we design a known distribution that matches these three moments (this methodology is called CDI3 in this paper).[3] The median return for this distribution is compared to the median return for the lognormal distribution that matches the first two moments of the portfolio return.

Let us consider three asset classes (A1, A2, and A3). The conventional CMA for these asset classes are presented in the Appendix. We consider six portfolios – from aggressive to conservative. Exhibit 1 presents the results for these portfolios.

Exhibit 1 demonstrates that the lognormal assumption for all portfolios can significantly overestimate the median values. Exhibit 1 also shows that, taken at face value, the statement regarding the connection between median and geometric returns is still problematic. Pension practitioners that wish to use the median portfolio return as the discount rate should avoid computational shortcuts and utilize more comprehensive approaches.

Exhibit 1

	Portfolio 1	Portfolio 2	Portfolio 3	Portfolio 4	Portfolio 5	Portfolio 6
A1	60%	55%	50%	45%	40%	35%
A2	10%	20%	30%	40%	50%	60%
A3	30%	25%	20%	15%	10%	5%
Total	100%	100%	100%	100%	100%	100%
Arithmetic Return	8.80%	8.20%	7.60%	7.00%	6.40%	5.80%
Geometric Return	7.06%	6.84%	6.58%	6.26%	5.90%	5.48%
Volatility	19.70%	17.34%	15.01%	12.74%	10.54%	5.50%
Lognormal Median						
	7.06%	6.84%	6.58%	6.26%	5.90%	5.48%
CDI3 Median						
	6.65%	6.36%	6.07%	5.79%	5.51%	5.22%
The Difference = Lognormal Median – CDI3 Median						
	0.42%	**0.49%**	**0.51%**	**0.48%**	**0.39%**	**0.25%**

[3] The design of this distribution and the moment- matching technique involve certain technicalities that are outside of the scope of this paper.

Investment Conditions vs. Outcomes

Let us continue giving the abovementioned statement the benefit of the doubt. Let us assume that the "median expected return" means the long-term median return, not the portfolio median return. The following statement support this conjecture:

> *"Plans should be using rates of return that they believe can be achieved over the next 20- to 30-year period with a 50 percent probability."*

Technically speaking, this statement is based on the observation that the long-term accumulated asset value of today's $1 is approximately lognormallly distributed. Therefore, the "long-term" median return is close to the geometric return.[4]

There are several problems with this logic. First, this observation is generally invalid for multiple payment cash flows (in addition to today's $1). Even if the accumulated asset value of every payment is lognormal, the sum of lognormals is generally not lognormal. Second, if the portfolio return distribution is not lognormal, then the short- and mid-term accumulated values may not be close to lognormal. Yet, they may be responsible for a substantial portion of the present value. In both cases, the relationship between the "long-term" median return and the geometric return is unclear.

But the biggest problem is reflected in the following statement.

> *"In practice, this means that funding should at a minimum provide for benefits if the median expected future **investment conditions** occur. By focusing on the median **expected outcomes**, the adequacy concept considers both return volatility and those scenarios in which **investment return** assumptions are not realized."* [5]

Dear reader, did you notice a quick journey from investing to outcomes and back to investing? This statement implies that median returns generate median outcomes. Even if it is true for each payment, the sum of medians is not necessarily equal to the median of the sum.

To illustrate this issue, let us consider the following numerical example: ten end-of-year contributions of $1 and their accumulated value after ten years. We assume that Portfolio 3 (50% of A1, 30% of A2, 20% of A3) is utilized in all years. For simplicity, let us assume that portfolio returns are lognormal.[6]

We calculate the deterministic accumulated value ($13.54) of these contributions using the median return 6.58%. Then we calculate the first three moments of the stochastic accumulated value and design a known distribution that matches these three moments (the CDI3 methodology). The median value for this

Exhibit 2

Median Return	Accumulated Value	CDI3 Median	Implied Return	The Difference
6.58%	$13.54	$13.65	6.75%	0.17%

---

1 See Mindlin [2011] for more details.

2 Emphasis added.

3 The reader may notice, that the lognormal assumption used in this section is inconsistent with the message of the previous section. In this section, we use the lognormal assumption for simplicity. The technical details required to evaluate stochastic accumulated values for non-lognormal portfolio returns are outside of the scope of this paper.

distribution is $13.65. This accumulated value implies 6.75% return, which is higher than the geometric return 6.58%. The results are summarized in Exhibit 2.

Thus, median investment conditions and median outcomes are not necessarily closely connected.

Conclusion

Some popular approximations may have convenient features and, at the same time, generate considerable errors. These approximations should be properly identified and disclosed. Pension practitioners should exercise caution with these approximations and utilize more comprehensive approaches to the calculations of measurements of portfolio returns in particular and outcomes of retirement programs in general.

APPENDIX

Asset Class	Arithmetic Return	Volatility	Correlations		
			A1	A2	A3
A1	8.00%	16.00%	1		
A2	4.00%	5.00%	0.2	1	
A3	12.00%	35.00%	0.9	0.1	1

References:

Mindlin, D., [2010]. On the Relationship between Arithmetic and Geometric Returns, CDI Advisors Research, CDI Advisors LLC, 2010, *http://www.cdiadvisors.com/papers/CDIArithmeticVsGeometric.pdf.*

Mindlin, D. [2011]. Present Values, Investment Returns and Discount Rates, *CDI Advisors Research*, CDI Advisors LLC, 2011, *http://www.cdiadvisors.com/papers/CDIDiscountRate.pdf*

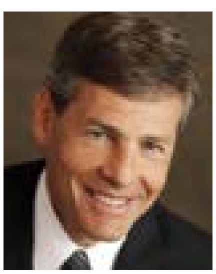

Dimitry Mindlin, ASA, MAAA, PhD, is president of CDI Advisors LLC. He can be contacted at *dmindlin@cdiadvisors.com*

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.

SOCIETY OF ACTUARIES

475 N. Martingale Road, Suite 600
Schaumburg, Illinois 60173
SOA.org