Cuffelinks
Connecting investors with ideas

Edition 70, 11 July 2014

This Week's Top Articles

Making sense of performance statistics

David Bell

Whether finance or sport, people love talking statistics. Where cricketers talk batting averages and run rates, fund managers talk information ratios, hedge fund managers talk Sharpe's and the academics talk t-statistics. What do all these statistics mean, can we reconcile them, are they useful, and which statistics should we place most credence in?

The starting point for performance comparisons is often alpha, which for this article I define simply as performance relative to the market. Alpha can be positive or negative. Comparing the alpha of two fund managers may not be a fair comparison as they make take different levels of active risk and so they are not evenly matched.

Manager skill is better assessed using risk-adjusted return statistics, but different parts of the industry use different statistics to describe their performance. This is where our information, Sharpe and t-stats need some defining.

Information ratio

This ratio compares annualised realised alpha against the volatility of that alpha (known as tracking error or active risk). The statistic is most relevant for traditional fund managers who manage against an index. It addresses the question: what is the manager's ratio for converting active risk into active return in a benchmarked fund?

$$Information\ Ratio = \frac{Active\ Return\ (ann.)}{Active\ Risk\ (ann.)}$$

Sharpe ratio

This ratio, designed by Nobel laureate William Sharpe, focusses on the total excess return (return in excess of a risk-free return, typically cash returns) against the total volatility of returns. It is most commonly used by hedge fund managers. Hedge fund managers are not benchmarked and have greater flexibility and so are assumed to 'own' all the risk they take. It is considered appropriate to compare their total excess return versus the total volatility of fund returns.

$$Sharpe\ Ratio = \frac{(Total\ Return\ (ann.) - Risk\ Free\ Rate\ (ann.))}{Volatility\ (ann.)}$$

t-statistics

t-statistics are statistical scores used for testing the significance of a result found in an empirical (data-based) analysis. t-statistics are similar to the better known z-scores associated with normal distributions. The t distribution is considered more appropriate than the normal distribution when it comes to sample data. If a result is significant, meaning that it is unlikely to have occurred by chance, then the t-statistic would be high.

Bringing statistics together

Unfortunately most statistics have shortcomings. I always recommend the use of multiple statistics and a qualitative overlay by an experienced professional is a superior approach when analysing funds. On one hand, the information ratio and the Sharpe ratio do not distinguish between skill and luck. The t-statistic indicates the chance of the return-for-risk observation occurring randomly but doesn't focus on the scale of the outcome. Academics may find a statistically significant result but it may be one that generates only a small level of active returns which does not even offset fees – hardly a significant finding in a real world environment.

In a broad way, we can bring this all together. First note that the information ratio (IR) and Sharpe ratio (SR) are based on annualised data and we can think of them both as excess return-for-risk statistics. If they are used in their appropriate contexts then we can compare the two against each other. From here I will show how we can use the t-statistic to assess the probability of these results, allowing us to assess the potential of whether a result was simply random luck experienced by a manager with no skill.

The key relationship is that the t-statistic is equal to the return-for-risk statistic multiplied by the square root of time. We then use this statistic to determine the probability of this return-for-risk outcome over this time period. I do not detail this calculation to keep things simple; it is the findings which I think are more important.

The chart below demonstrates the probability of different return-for-risk outcomes over varying time periods to be generated by a manager with no skill.

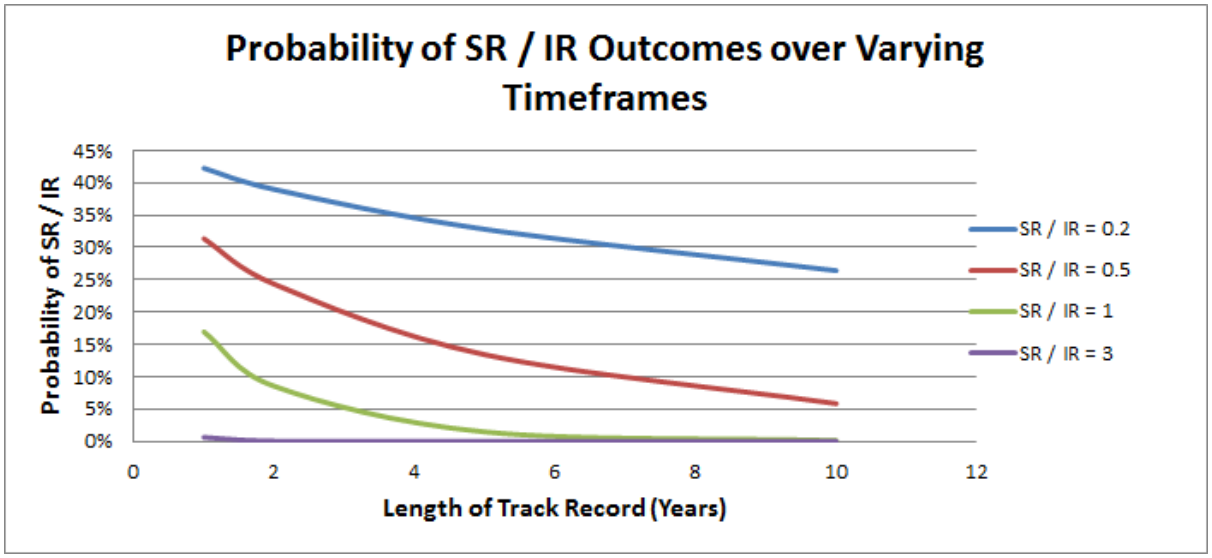

It is easier to understand the above chart using an example. Consider a fund manager with an information ratio of 1. If their track record is only one year in length then there is a 17% chance that this is just a random outcome generated by a manager with no skill. If the track record length is two years then the probability is around 9%, the five year probability number is less than 2%. If the track record is 10 years, the chance of generating a long term information ratio of 1 is very close to zero.

We can see the importance of the length of track record: for instance it is more likely that a manager with a 0.5 information ratio over five years has skill than a manager with a higher information ratio of 1 generated over one year.

We can also see how information ratios and Sharpe ratios of 3 are highly unlikely. In such cases you should dig deeper to understand why the return-for-risk ratio is so high.

Focus on length of track record

One important take-out is that comparing the information ratio or Sharpe ratio of fund managers with different track record lengths is a flawed approach unless you make a time adjustment. My rule of thumb has always been that a long term (greater than 10 years) information or Sharpe ratio of 0.5 is worthy of respect.

While interesting and useful this should be just one part of your analysis toolkit. I encourage caution when it comes to the use of performance statistics. For instance, these statistics assume that excess returns are (broadly) normally distributed, that they are not taking any binary bets in their portfolios (such as a hugely concentrated single stock or sector bet), and that the same teams, styles and processes have remained in place during the fund's existence.

After all, batting averages generated in different eras against different bowling attacks on different pitches are not perfectly comparable.

During a 12 year amateur cricket career David Bell average substantially less than 10 with the bat. In mid-July 2014, David will become the Chief Investment Officer at AUSCOAL Super. He is also working towards a PhD at University of NSW.

Diversification's focus moves to matching future needs

Harry Chemay

This series in Cuffelinks on investment diversification has focussed on the mathematically precise analysis of expected returns, risk and correlations. But modern portfolio theory assumes a world without fees and taxes where all investors have the same time horizon and access to the same information. Investors are able to interpret and act in the same way and have identical unbiased expectations regarding the future. In the real world of investing, these assumptions are too simplistic.

Consider investment forecasting. Correlations, rather than being static, change over time, and risk (however defined) is no more stable. Volatility is itself volatile, as the below chart illustrates:

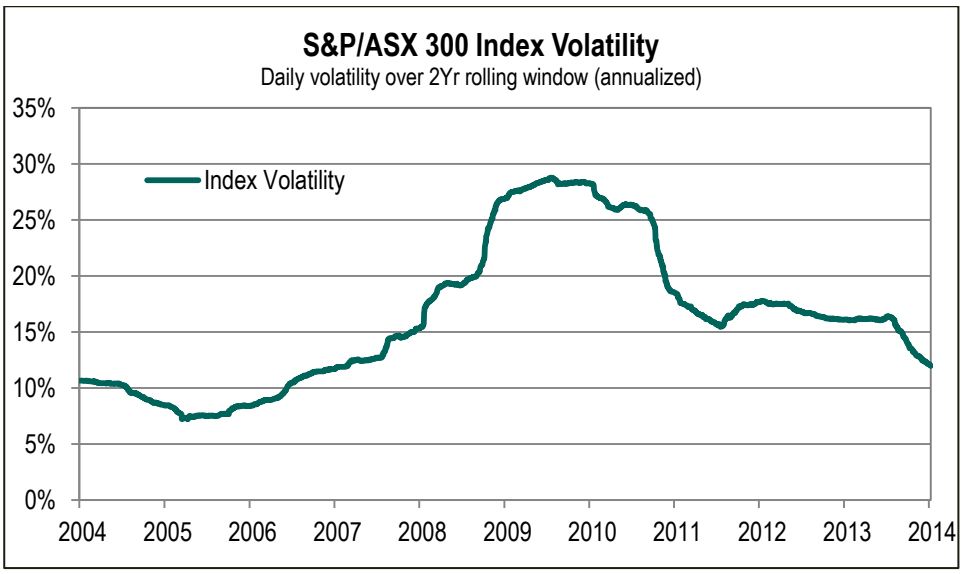

Data supplied and analysed, and chart provided, by S&P Dow Jones Indices

The previous article used the terms risk and volatility interchangeably. Why? Because modern portfolio theory holds that volatility of returns is the most appropriate measure of risk. But is this the way people actually view risk?

In over a decade of advising individual clients, nobody asked me about their portfolio's standard deviation, or how it sat relative to some theoretical efficient frontier. Clients had a keener interest in the change in portfolio value between review meetings, and paid far more attention when these were significantly negative than equivalently positive. In behavioural finance, this asymmetric concern is known as loss aversion. Therefore, let's put to one side the neat world of modern portfolio theory and consider instead how diversification can be applied to real-world retirement planning.

Framing retirement objectives appropriately

Why bother saving for retirement at all? We do so to smooth our lifetime consumption. If we did not (with charity and social security offering inadequate safeguards), we would swing from exuberant spending in our working years to relative poverty in retirement. Modern portfolio theory forces a single-period *risk/return frame* onto individuals, when focus might be better directed toward a multi-period *consumption frame.* A schism exists in the understanding of risk between superannuation trustees and members. Trustees view risk via a text book definition of volatility (standard deviation). Members see risk as a failure to generate sufficient purchasing power in retirement to allow for a preferred level of consumption. Whose view of risk is more relevant? Whose risk is being managed?

Funding retirement consumption

It is possible to estimate the value today of the future cost of retirement. It is the present value sum of each year's expected cost of living for the number of expected retirement years.

Consider an example of a recently retired 65-year-old male. Using the current Association of Superannuation Funds of Australia (ASFA) retirement standards for a single person ($23,283 p.a. for a modest lifestyle and $42,254 p.a. for a comfortable lifestyle), the total retirement cost today sums to $336,000 for the modest lifestyle and $610,000 for its comfortable equivalent. A 65-year-old female would require $374,000 and $679,000 respectively, due to higher life expectancy.

Whilst these numbers are sensitive to inflation and discount rate assumptions, and subject to some variability due to heterogeneous later-life health care costs and longevity risk, they provide a valuable insight into retirement expenditure on average.

Armed with a measure of retirement cost, we now have a basis for comparing these prospective liabilities against retirement assets. To do so we need, however, to consider the totality of assets capable of funding retirement.

It is unlikely that the average retiring 65-year-old male will have $610,000 in superannuation. APRA data currently suggests $151,000 as a more likely balance. Such a large superannuation balance may not, however, be necessary for two reasons:

1. The government age pension

The age pension is effectively a government-backed lifetime indexed annuity. One recent study estimated the value to life expectancy of the full age pension is $377,000 for a 65-year-old male. As some 80% of retirees will receive at least part age pension, it will continue to remain an important 'shadow retirement asset' (despite the changes foreshadowed in the government's 2014/15 Budget).

2. Other non-superannuation assets

Non-super assets such as shares and property play an important role in real-world retirement funding. A recently released Melbourne Institute/Towers Watson working paper calculated median wealth (excluding the family home) for those aged 65 – 69 years at around $389,000. Critically, non-super assets account for over 67% of total financial wealth.

Diversification in an asset-liability framework

Putting all the pieces together, it is possible to consider retirement planning as an asset-liability matching exercise comprised of various layers as depicted below:

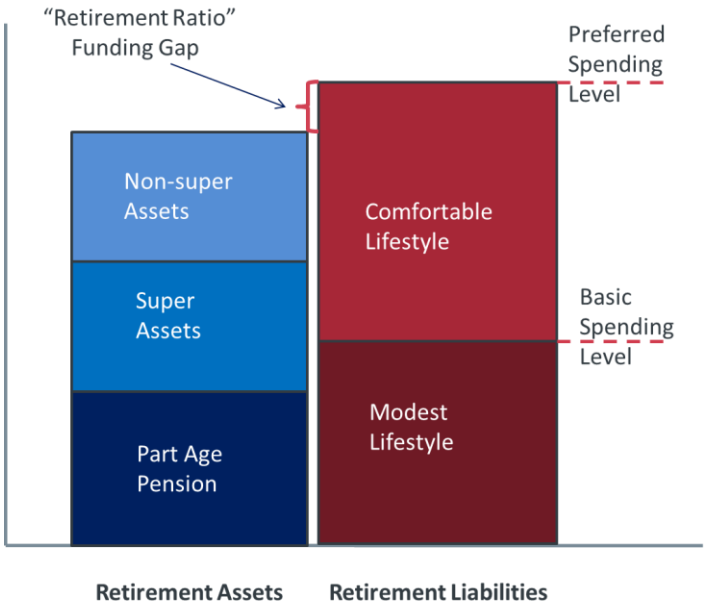

The aim of retirement planning becomes the attainment of a 'retirement ratio' of at least 100% by the preferred retirement age. Any combination of four levers can be manipulated to achieve (or maintain) fully-funded status; savings rate, retirement age, target retirement income and investment risk.

Diversification's role changes in an asset-liability paradigm. The investment objective moves from risk/return optimisation to matching the nature, duration and variability of retirement liabilities (or needs). For couples this would ideally incorporate differing life expectancies and age pension entitlement.

There is an obvious link here to Defined Benefit (DB) retirement plans, where the provider assumes the risk of meeting a comfortable retirement lifestyle. The challenge is that these plans have been replaced by Defined Contribution plans, and this recent article made the case for retaining some DB features. In the Netherlands, where DB funds still dominate, the average pension fund has an allocation to growth assets of 24%, whilst in Australia it is around 68%. The Dutch objective is to fund long-term retirement cash flows. Australia's focus remains primarily on accumulating lump sums payments and shorter-term investment returns.

The challenge for the Australian superannuation sector is to move from a 'to retirement' mindset to a 'through retirement' mindset within a member-centric consumption frame. As Nobel laureate and pioneer of the lifetime consumption approach, Professor Robert Merton, has opined: *"sustainable income flow, not the stock of wealth, is the objective that counts for retirement planning"*.

Harry Chemay is a Certified Investment Management Analyst who consults across both retail and institutional superannuation, focusing on post-retirement outcomes. He has previously practised as a specialist SMSF advisor, and as an investment consultant to APRA-regulated superannuation funds.

Australia's default: the winners and losers from bonds

Ashley Owen

Part 1 of this story contained a brief outline of government debt – both domestic and foreign - and various ways in which governments can avoid repaying their debts in full, through default, restructure and/or inflation. We also looked at the level of Australian government debt compared to the current government debt crisis, and how it compares to Australia's debt level today.

In Part 2, we looked at how Australia's big default and debt restructure occurred, which bond holders were rescued, and which were forced to take a 'haircut' on their interest and principal repayments.

Here, in Part 3, here we look what it meant to bond investors, including the returns achieved by bond investors before, during and after the debt default and restructure. We will focus on returns to local investors on domestic Australian debt since they were the ones directly affected.

We see that money was made for bond holders, and even those who hung on and were hit with the 'haircut' restructure.

Returns were good for patient investors

The first chart shows monthly returns (top section) and cumulative returns (lower section) on domestic long term Commonwealth bonds from the start of 1929 to December 1932.

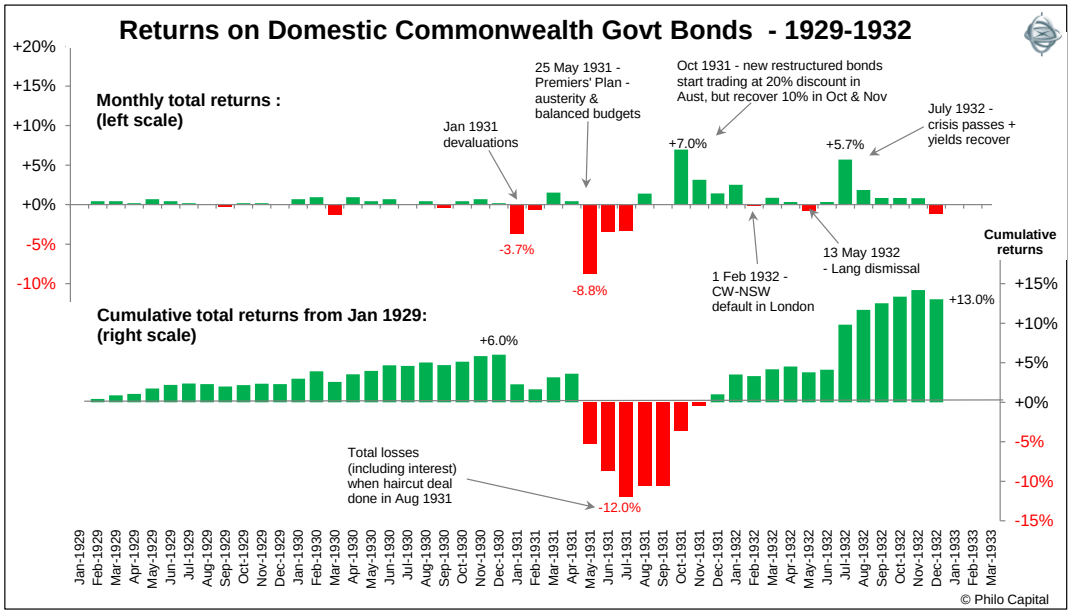

Bonds returned 13% in total over the four years, made up of 19.3% from total interest payments minus a 6.3% capital loss.

The next chart shows prices for long term bonds in both markets: domestic bonds in the Australian market (green line) and foreign bonds in the London market (red line). The lower section shows cumulative returns from interest on domestic bonds (blue bars) and cumulative capital value (red bars).

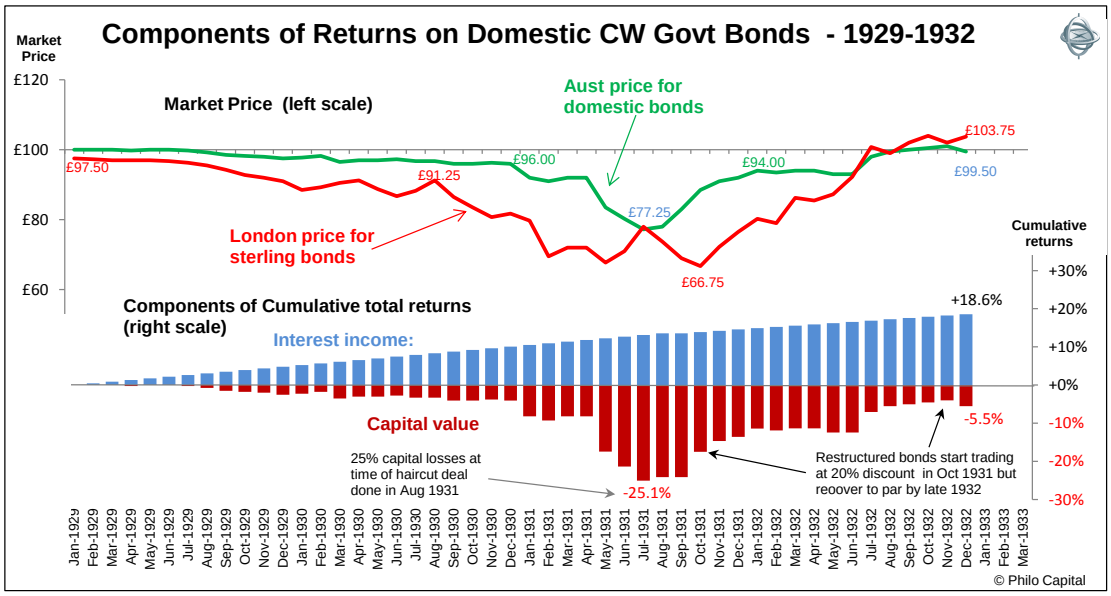

Patient investors who held onto their bonds, took the haircut deal, bought new replacement bonds, and then held them until the end of 1932, would have achieved total returns (i.e. interest plus capital gains/losses) of 13%.

A 13% total return over a 4 year holding period doesn't sound very good, but after inflation (or *deflation* actually) it amounted to a healthy 30% real return because consumer prices had *fallen* by 17% over the period, and wages had fallen by around 10%. In terms of the real spending power of investors' money after inflation, the 30% return over 4 years equated to 6.8% per year compound return after inflation. Not a bad return for patient investors, even after the default and restructure in 1931-1932.

A major contributor to the good returns was the fact that the new replacement bonds did not sell for par at auctions. They were sold at significant discounts to par because investors were still nervous about the government's ability to service even the new lower interest payments. These discount purchase prices contributed to good returns as market prices quickly rose back up near par during 1932 as confidence returned and yields declined from crisis levels.

Winners and losers

On the other hand, investors who bought domestic 10 year bonds in December 1930 when yields were an attractive 6.4%, and then sold out in the panic in July 1931 when and the government was talking about default and bond prices had crashed 20% from £96 to £77, with yields soaring to nearly 13%, would have lost 18% on their investment.

Likening this to the current European debt crisis, investors who lend money to Greece at 100 cents in the dollar before the crisis (mainly European banks) and then panic sold in early 2012 at less than 20 cents in the dollar before the restructure was announced, would have lost 80% of their money. These were permanent losses. The big winners in Greek debt were the numerous hedge funds that bought up Greek bonds from the panic sellers at prices as low as 15 cents in the dollar in early 2012, and then doubled their money when the old bonds were exchanged for new bonds an average of around 33 cents.

There are winners and losers in all markets. Even when governments default on their debts there is money to be made by investors who resist the temptation to panic sell in a crisis (and also resist the temptation to panic buy in a boom for that matter).

In Part 4 we will look at the returns from the broad stock market versus the government bond market during the crisis. We see how the impact of the Greece-like default and restructure of government bonds affected bond returns, compared to the impact of the 1929 crash had on share returns.

Ashley Owen is Joint CEO of Philo Capital Advisers and a director and adviser to the Third Link Growth Fund.

Dive or stay: the biases of goalkeepers and portfolio managers

Mariana Araujo

Millions of fans around the world have tuned into to this year's World Cup in Brazil. Whether you love it or hate it, the penalty kick is one of the most exciting plays in football. As the striker approaches the ball, often with the outcome of the game hanging in the balance, the goalkeeper has a split second to decide what to do. It's not unlike the plight of portfolio managers in today's fast-paced market.

Dive left? Dive right? Stay standing in the centre between the goal posts? The odds aren't good: fewer than one in five penalty kicks are not converted at this level of play.

In 2006, the World Cup Final between Italy and France came down to a penalty kick shootout. On the first kick, the French goalie chose to dive to his right. However, the shot from the Italian striker went straight down the middle. Had the French goalie stayed at home, the outcome of that shot, and perhaps the game, may have been different.

According to one study, goalkeepers choose to dive nearly 94% of the time.[1] In response to the relatively even distribution of kicks between the goalposts, however, and the greater chance of saving those in the middle, goalkeepers who stand and defend the centre may experience a better outcome. Simply put, not taking action may be the best course of action.

Due to what behavioural researchers call *action bias*, a goalkeeper is expected to act. In the case of a penalty kick, the norm is to dive. A scored goal is perceived to be less disappointing when it follows action. Innate self-confidence, years of training and the crowd's expectations further contribute to this suboptimal decision. If the goalie dives, he feels that he did his best to stop the ball, and so does almost everyone else.

Investment managers often fall into the same trap of action bias, trading frequently, with confidence that this action adds value. And whether the trades ultimately prove to be right or wrong, the manager who trades frequently *looks* like he's doing something to generate results. This is one of many behavioural traits contributing to widespread short-termism in the markets.

In recent years, the average holding period for a stock has dropped to about seven quarters (and many studies claim it is much shorter). All too often the concept of buy and hold investing has been subsumed by short-term trading strategies. Many of these trading strategies, which rely on top-down macro-economic calls, are often no better at predicting the future direction of the markets than the goalie who tries to guess which way the shot is going.

Such a short-term bias creates an enormous time-horizon arbitrage opportunity for individual and institutional investors who are willing to take a long-term view. Over very short time periods - say, one week - the average difference between the best- and worst-performing stocks usually comes down to a few percentage points. Move out to one-year and you will begin to see stocks that significantly outperform in any given year. However, as they say, there is no free lunch and many of these high-flying stocks will often see market sentiment turn against their lofty valuations and find themselves at the bottom of the league tables the following year.

By contrast, if you look at the performance dispersion between the best and worst stocks over a five-year period, the numbers becomes quite meaningful. Simply put, over the long-term, the cream rises to the top, with the top 10% of stocks outperforming the bottom 10% by over 160 percentage points. And a common thread among managers who consistently generate long-term results is a strong buy-and-hold mentality. Managers who look to invest in companies that are well-positioned to generate growth over multi-year time periods have the courage to do nothing when short-term trends and negative headlines have the traders running for the exits.

Portfolio managers can lengthen the investment horizon by avoiding the temptation to trade frequently, choosing instead to hold securities for longer periods. Though portfolio managers and goalkeepers are prone to act, an awareness of this action bias may help them recognise that inaction can be an optimal strategy. And deciding to hold the position has the potential to result in a better outcome for their clients — and fans.

Mariana Araujo is a Sao Paulo-based equity research analyst for MFS Investment Management.

1 Bar-Eli, Michael, Ofer H. Azar, Ilana Ritov, Yael Keidar-Levin, Galit Schein, "Action Bias among Elite Soccer Goalkeepers: The Case of Penalty Kicks," *Journal of Economic Psychology*, January 2007. According to this study, kickers distribute shots much more equally across the left, centre and right thirds of the goal. The authors suggest that if goalkeepers altered their strategy to stay in the centre, then kickers might adjust and send more shots to the left or right.

Just what is a re-contribution strategy?

Monica Rule

You may have heard about a 're-contribution strategy', but do you really know what it is and how it works? Often when an adviser or accountant provides an idea that will reduce a tax liability, you accept it without really knowing how it is achieved. That's what you pay them for – right!

Let me explain what a re-contribution strategy is.

A re-contribution strategy is where you withdraw money from your SMSF and re-contribute the money back into your SMSF. Before you can do this, you need to be able to access your money by satisfying one of the following conditions of release:

- reached your preservation age (55 to 60 years) and retired from your employment
- reached age 65 (working or not)
- ceased work temporarily as a result of physical and/or mental illness
- ceased work as a result of permanent incapacity
- experienced a terminal medical condition
- accessed money under severe financial hardship grounds
- accessed money under compassionate grounds via the Department of Human Services
- accessed money under the transition to retirement arrangements.

Of course it is unlikely that a person accessing their money under financial hardship or compassionate grounds would be considering a re-contribution strategy.

If you access money in your SMSF without satisfying at least one of the conditions of release, then you will be in trouble with the Australian Taxation Office which regulates SMSFs.

The two main reasons why advisers may suggest a re-contribution strategy is to:

1. Reduce the tax payable on your superannuation pension, especially if you are under the age of 60
2. Lower the tax payable on benefits paid to your beneficiaries in the event of your death.

Money in your SMSF is comprised of two components. One component is the tax-free component which is made up of non-concessional contributions received by your SMSF. The other component is the taxable component which is made up of concessional contributions received by your SMSF and earnings from SMSF investments. Under the superannuation and income tax laws, superannuation benefits (pension and lump sum) paid to you are subject to a proportion rule which requires your benefit to be paid in the same proportion as the tax-free and taxable components of your superannuation interest in your SMSF.

For example, if your SMSF is comprised of a 60% taxable component and a 40% tax-free component, then your superannuation benefit, when paid out, must retain the 60% taxable component and the 40% tax-free component.

A withdrawal and re-contribution strategy involves withdrawing or accessing your superannuation entitlements that consist of the taxable and tax-fee components and re-contributing some or all of the money back into your SMSF as a non-concessional contribution (i.e. all tax-free). This increases the amount of tax-free money in your superannuation account which provides tax savings if you are accessing a pension while under the age of 60. It may also mean large tax savings when you pass on your superannuation savings to your non-dependant beneficiaries after your death.

This is because the taxable component of a pension benefit received by a person under the age of 60 is taxed at the person's marginal tax rate less a 15% tax offset. Converting the taxable component to a tax-free component increases your tax-free pension income.

When you pass away, your beneficiaries who are over the age of 18 or are non-dependant will also receive a greater portion of your death benefit without having to pay tax.

Watch that you don't exceed your non-concessional contributions cap. In addition, if you are aged 65 to 74 you will need to be at least working 40 hours in a period of 30 consecutive days to be able to make non-concessional contributions into your superannuation fund.

It pays to understand how things work so you can better discuss with your adviser.

Monica Rule worked for the Australian Taxation Office for 28 years and is the author of 'The Self Managed Super Handbook – Superannuation Law for Self Managed Superannuation Funds in plain English'. Monica is presenting a series of SMSF seminars and for details see www.monicarule.com.au.