BETWEEN THE IVORY TOWER AND THE COAL FACE - PRACTICAL ASSET ALLOCATION

Introduction

This paper covers the main steps in the asset allocation process, it discusses how those with a purely academic mindset approach the task, how most practicing planners do it and a middle ground which combines the best of both approaches while avoiding the manifest problems inherent in each approach.

The need for a rational, integrated process

Few dispute the need for process when making any investment decisions. A typical asset allocation process consists of five steps

1. Return forecasting - estimating the expected returns of each asset class for the period in question
2. Risk estimation - estimating the risk of both individual assets as well as the way that assets behave compared to others, or how correlated assets are
3. Optimization - determining what portfolio gives the best return for any particular level of risk
4. Profiling – determining what is the suitable level of risk for any particular individual, institution or circumstance
5. Implementation – actually putting it into practice

Different parts of this process can be done by one person, or it can be divided up. However, regardless of who actually does each step, without a process the results achieved are likely to be haphazard at best. Unfortunately, as we shall see, merely having an apparently sensible process is no guarantee of good results either. If the different parts of the process do not fit together in a well integrated way, or if the assumptions underlying the process are irrational, then poor results will almost certainly follow.

Asset allocation is a core competency of financial planners

While asset allocation does not provide the 90% of investment returns that is often claimed, no one doubts that it is important and, done well, provides the key linkage between an investor's objectives and the portfolio design. It is a core competency for planners.

For the remainder of this paper we will look at how the different steps of this process may be handled; for better or worse.

Return forecasting

The art of financial markets forecasting has a well earned reputation for unreliability. Nonetheless it is not possible to carry out asset allocation without some form of forecast, be it formal or implied. Even the statements, 'equities outperform in the long term' or 'equities are more risky than bonds' have forecasts embedded in them. Given this situation it is better to formally approach the task of forecasting returns, but to do so in full knowledge of the reliability or otherwise of the techniques used.

Academics and the efficient market

The state of the art with academics goes along these following lines.

> Markets are efficient which means that assets with higher levels of risk will have higher expected returns. The amount of the return is estimated by adding the risk free rate (normally either the bond rate or the cash rate) to a premium to reflect the risk of the asset class. Because that risk premium cannot be directly measured ahead of time, the best estimate is the historical difference in returns between, say, equities and cash for the equity risk premium, or property and cash for the property risk premium.

So, if over the past 20 years equities have returned 10% more than bonds we assume that over the next, say, 10 years equities will again return bonds plus 10%.

Unfortunately, this approach has some major drawbacks.

- Even using quite long-term historical data, say 10 to 30 years, this approach has proved to be stunningly unsuccessful. Over these types of timeframes this sort of methodology tends to produce the highest forecast returns at times when the returns are about to be the lowest and vice versa. This is clearly illustrated in the charts below.
- Assets without 10 years of data aren't even considered as part of the mix, because of the inability to forecast returns.

Fig 1 : Australian Equities: Using the past to forecast the future

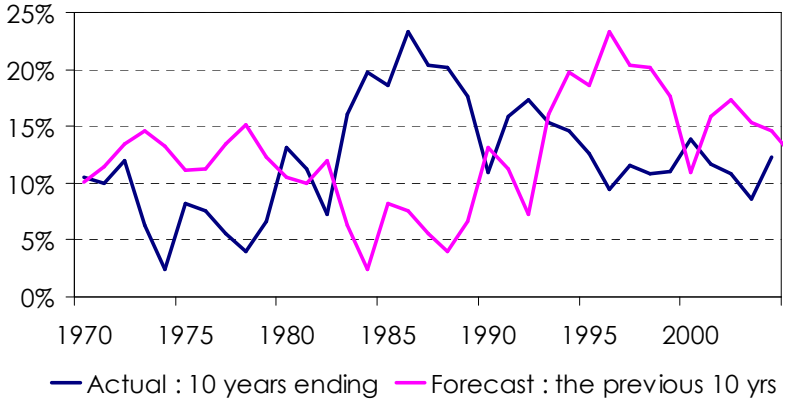

Fig 2 : US Equities:
Using the past to forecast the future

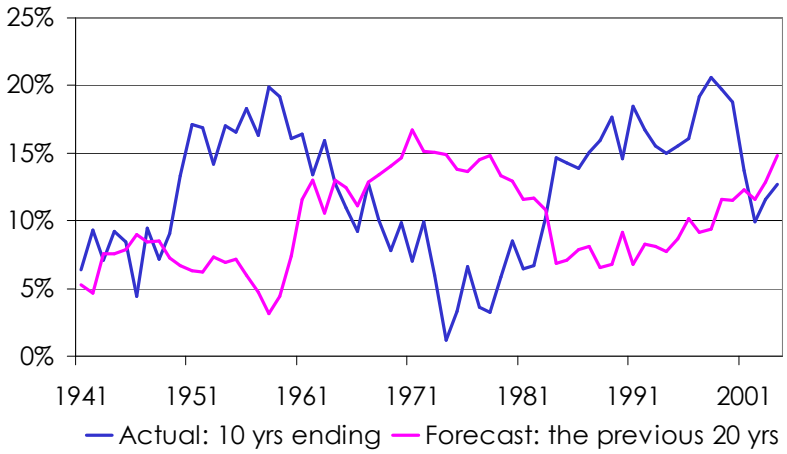

Chasing past returns

In the so-called real world few practitioners, that is investors, fund managers and financial planners, believe in efficient markets but still behave as if they do. It's called chasing past returns, which is really just a less sophisticated version of the academic approach. The idea goes something like this; short term market moves are notoriously hard to predict, even the academics say that you can't pick markets, so we might as well invest in with the latest top performing assets, and stay away from the poorly performing assets.

If you are in the business of managing other people's money this is quite a safe strategy. You will rarely be fired for being invested in assets that have been hot and turn sour (who could have known?) whereas if you invest in assets that have performed poorly in the past and continue to perform poorly in the future (how could you be so stupid???) client retention becomes difficult to say the least.

However, while chasing past returns is a strategy that may be safe from a business perspective, it is any thing but safe from an investment perspective. As many planners and investors have found, it is a reliable recipe for producing terrible returns.

A forward looking approach

A much better approach is to seek to firstly understand what drives returns and to then to sensible forecast those drivers. One such approach is the Occam's Razor methodology developed by John Bogle[1] , named after Sir William of Occam who declared the simplest explanation is generally the best. And this is indeed a simple approach, and all the more powerful because of it.

[1] John Bogle. Investing in the 1990s: Remembrance of Things Past and Things Yet to Come. Journal of Portfolio Management, Spring 1991, pp. 5-14

This approach to forecasting decomposes market returns into three drivers of returns

- Income
- Growth in income
- Effect of changing valuation ratios

The three elements can then be simply added together to produce remarkably reliable long term return forecasts.

Returns = Income + Growth in income +Effect of changing valuation ratios

Or for those that prefer equations

R= Y + G + V

Where

Y is the current investment yield, a known quantity; hence no forecasting is required for this input.

G is the annualized growth in income or earnings for the asset, for

- Property it is growth in rents
- Equities it is growth in Earnings Per Share
- Fixed interest growth is zero, by definition!

V is the Valuation effect; it is the compound effect of an increase or decline in PE ratios or yields on the value of the asset.

So for equities, over a one year period

V= (PE at end of period / PE now) –1

For example if PEs rose from 10 to 12 then V = 12/10 -1 = 0.2 or a 20% increase

For longer time periods we use the compounded growth rate;

$$\mathbf{V= (PE\ at\ end\ of\ period\ /\ PE\ now)^{1/10} –1}$$

Using the same example over 10 years

$V= (12/10)^{.1}-1 = 1.0183-1= +1.83\%pa$

Over long periods this method has proved to be a remarkably useful forecasting tool because over longer time periods, say 10 years, EPS growth rates become quite predictable and PE and other valuation metrics tend to revert to the mean.

A simple mechanical back test of this methodology can be done using retained earnings as a forecast of real EPS growth and assuming PEs stay at their starting level over rolling 10 year periods. While this is a more simplistic approach than

one would use in real time, when back testing it is good practice to keep subjective measures out of any calculations as far as possible.

Nonetheless the track record of this simple device is quite impressive as is shown for the US and Australian equity markets in the charts below. While not perfect by any means, the general pattern is more than useful and a vast improvement on simply using historical returns as a basis for forecasts.

Fig 3 : Australian Equities
Actual v forecast, 10 Yrs ending

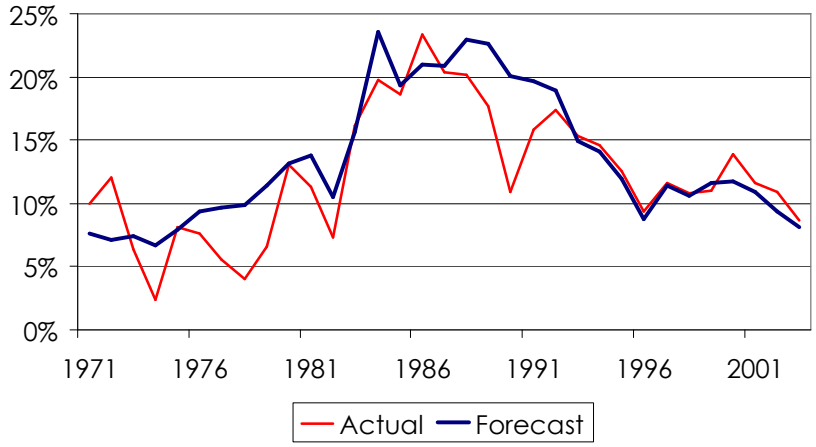

Fig 4 : US equity returns
Actual v Forecast, 10 yrs ending

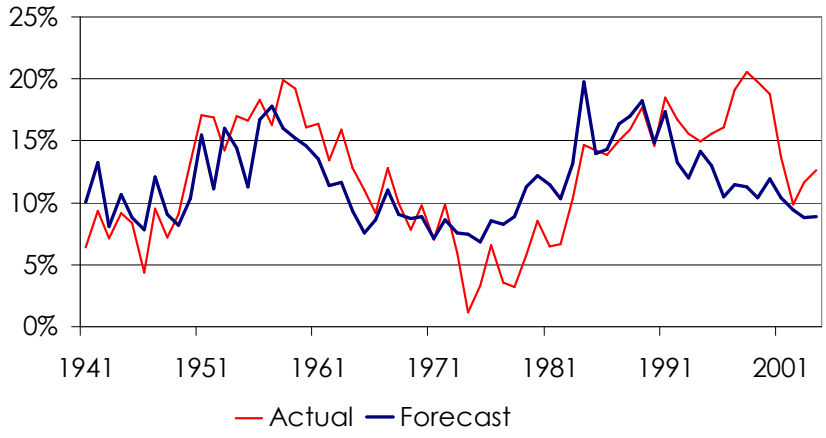

The original Bogle paper is strongly recommended reading for any who are interested in learning more about this method. Once grasped, this basic methodology is simplicity itself, can be used on a large number of assets and is something that every competent planner should fully understand.

Risk estimation

Estimation of risk has long been problematic, not least because there is little agreement on how it should be defined let alone measured or estimated. In this section we look at some alternative ideas about risk.

Academics generally define risk as standard deviation or tracking error

There are a couple of ideas running through this.

- If markets are efficient, then returns in one period don't have any impact on returns in the next period, (what academics describe as returns being serially uncorrelated.) if this is the case then the measured short-term volatility of assets is directly related to the long term uncertainty of those asset's returns. Furthermore, if we measure the short term correlation coefficients of those assets (i.e. their tendency to move together or to move independently) we can mathematically estimate the risk of an entire portfolio of assets
- The second idea is more arcane but important to be aware of. Based on a number of assumptions including efficient markets, no transaction costs and perfect information flows, it is possible to mathematically 'prove' that the so called market portfolio, i.e. the index, is the highest returning portfolio for a given level of risk. As a consequence any departure from the index involves additional risk and therefore tracking error is a measure of risk we should all worry about.

Institutional money managers love the second concept in particular because it provides a definition of investor risk that lines up with their own business risk

Unfortunately there are a number of practical problems with these concepts of risk, particularly in a financial planning context.

1. Long term risk and short term volatility are often unrelated, and therefore short term volatility can be a very poor measure of risk for many common asset types
    - 10 year government bonds have substantial short term volatility but no long term uncertainty; a ten year bond bought at a yield of 6% will return 6%pa over 10 years, regardless of short term ups and downs
    - Returns from investments like mortgage trusts or debentures show no volatility; until they default. In other words short term volatility ignores credit risk
    - Returns from unlisted investments (private equity, property syndicates and so on) show little or no short term volatility; but are certainly not low risk
2. Correlations vary over time and, in particular, dramatically increase during downturns. That is, when trouble strikes, everything seems to go down together. Diversification benefits seem to magically disappear when most needed.
3. Very few people believe all the assumptions required to make tracking error a real measure of investor risk. And even fewer believe that it is a genuine risk faced by clients; in the end it's just a business risk management device and should be recognized as that and only that.
4. Clients don't understand these risk definitions, which makes them unhelpful in explaining risk.

Planners often use the income/growth split as a proxy for risk.

In an attempt to make the whole exercise simpler many planners use a risk measure which effectively says all income assets are low risk, all growth assets are high risk, and the risk of a portfolio is governed by it's split between income and growth assets. While this approach certainly does achieve the goal of simplicity and it is a concept that clients can grasp, it does have some real problems if taken too far.

- It assumes that all income assets are of equal risk and all growth assets are of equal risk. That is, junk bonds and cash are of the same level of risk, and Listed Property Trusts and emerging market equities are of the same level of risk.
- What are growth assets and what are income assets? Why is it that LPTs which give about 7% of income and 2% of growth are deemed to be growth assets?
- The measure gives no real indication of risk; how far can the portfolio fall in the short term? How bad can long-term returns be?

Risk of not meeting objectives

This is probably the most complete definition of risk. Unfortunately it does not lend itself to easy measurement and estimation, particularly given that it will vary from person to person and from situation to situation.

Nonetheless the risk of not meeting objectives can be quantified and in a way that is ideally suited to a financial planning context. The steps involved are

- Understand the clients objectives and work out what returns are required to achieve them
- Work out what can get in the way; typically it will be insufficient, long-term, real, after tax returns, but it also may an unpalatable level of short term volatility.
- Determine the probability of meeting the client's targets. A good proxy for this measure of risk is the long term return uncertainty. This can be estimated again using the Occam's Razor approach. Essentially the Occam's Razor approach to estimating risk involves thinking about different economic scenarios and using the Occam's Razor approach to estimate returns in those scenarios. This will give a range of possible long term outcomes and by then attaching probabilities to these scenarios one arrives at a return distribution. This may be done in a quite simplistic way or very thoroughly. In either event this technique provides a much better window into possible futures than most other risk estimation techniques.

Building Optimal Portfolios

Once forecasts for return and risk have been determined, the next step is to work out an optimal combination of assets. Financial planners typically outsource this step of the process to their dealer group or external research house. The following describes how researchers typically go about the task. While some of this may again seem arcane, planners should at least be aware of how their research suppliers arrive at their recommendations. Furthermore, an understanding of some of the issues involved will help planners develop sensible optimization approaches of their own.

Markowitz style, mean variance optimization

The classical academic approach is to use a mathematical optimization process known as Mean Variance Optimization (MVO). This involves putting ones estimates for expected (or mean) returns, variance (or volatility) and correlation coefficients into an optimization program to work out which portfolio has the maximum return for any level of variance. The full set of such portfolios, for each level of risk, is described as the Efficient Frontier. This is shown in a stylized way in the chart below, where each point represents the expected risk and return outcome for a different portfolio and the red line represents the efficient frontier.

Figure 5 :The Efficient Frontier

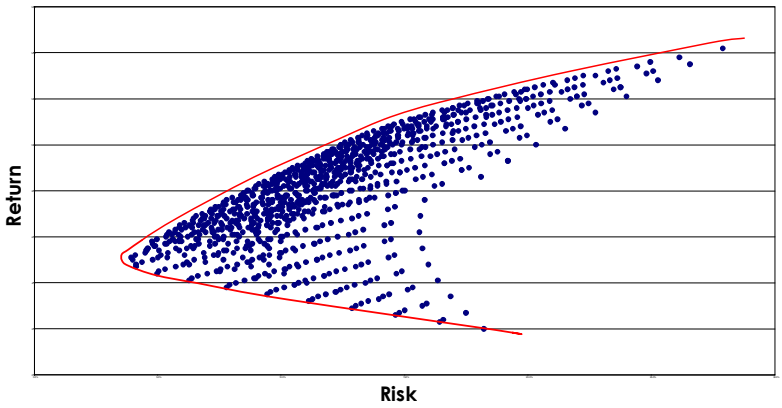

This is the basic framework for portfolio construction that has been almost universally adopted by academics for the past 50 years since its development by Nobel Laureate Harry Markowitz.

Unfortunately, even assuming one was happy with their return and risk forecasts this methodology has some severe practical difficulties,

The first practical difficulty is that the MVO process behaves as if there is total confidence in the forecasts, and has no problem taking extreme positions. John Reckenthaler, global head of research at Morningstar describes MVOs as "investment terrorists, pursuing fringe ideas with absolute conviction in their own rightness." Thus a portfolio which could have been chosen from five asset classes

may end up being recommended as 10% Australian equities, 50%, in private equity, 40% in cash, and nothing in the other two asset classes, say fixed interest and international equities.

The second practical difficulty stems from the extreme sensitivity of the output to small changes in inputs. In the example above, a small increase in the forecast for, say, international equities may change the optimal exposure to 20% Australian Equities, 60% international equities, 20% in cash and forget everything else. Not so much garbage in, garbage out, as molehills in and mountains out!

These problems have made the MVO all but unusable in practice, even if it remains the optimizer of choice for academics.

Optimization as practiced by fund managers, researchers and asset consultants

Most fund managers and researchers use MVOs, but with some crucial adjustments.

- Limits on exposures. Because the MVO tends to take extreme positions, those who have to make real decisions tend to put in limits such as maximum weights to certain asset classes, fixed weights to income and growth assets, or a fixed allocation to international assets.
- The forecast inputs are varied until they give a 'respectable' looking asset allocation result. (Seriously!)
- Portfolios are optimized subject to a limit on tracking error, in this case defined as the difference between the managers asset allocation and that of the average manager

The problem with these all these approaches is that they tend to drive asset allocation towards that of the average fund, towards allocations that look conventional. Thus when asset classes get severely mispriced the optimizers don't really get much chance to do their stuff; they are too severely constrained.

For example, for much of the time between 1995 and 2005 many managers unconstrained optimizers were recommending high weights to Listed Property Trusts, up to 30 to 50% of a balanced funds. However, most managers put in limits which capped exposures to LPTs at around 10%. Why? Largely, it seems, for competitive reasons. Having 20 to 30% more in property than ones competitors runs the risk of making the fund manager or consultant appear very uncompetitive in the event of a long equities bull market. Understandable when looked from the manager or asset consultant's perspective, but hardly a useful model upon which planners should base decisions.

Optimization as practiced by planners.

In practice most planners rely on the recommendations of fund managers and research houses as starting point, and then overlay their insights to make tweaks up and down. Fine as long as the starting point is reasonable, but we have seen that there are many reasons why this might not be the case.

Robust portfolio optimization techniques

These are a new breed of optimizer. They essentially take the key weakness inherent in MVOs and turn it into a strength.

The main weakness of the MVO approach is its sensitivity to small changes in assumptions. This arises because there are a great number of very different portfolios with similar risk and return characteristics near the efficient frontier. A small change in assumptions and the portfolio that was previously *near* the efficient frontier moves *onto* the efficient frontier. The expected return and risk of the old and new efficient portfolios has hardly changed at all, because the change in the assumption was small. Just the composition of the mathematically best portfolio is very different.

The key point in all of this is that there are a large number of very different portfolios with similar risk and return characteristics. Given the limits of accuracy of all of the forecasts that serve as the inputs to the MVO, all these portfolios can be considered to be equally efficient. This is an extremely important observation.

The chart below indicates the idea; all the portfolios between the red lines can be considered of equal investment merit, and there are lots of them in the band.

Figure 6 : The Efficient Band

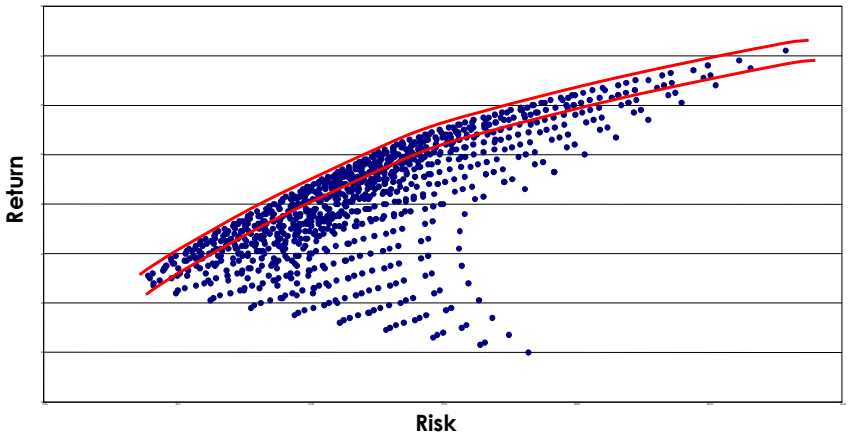

What the new robust optimizers do is search among the many sound portfolios within this band to identify those with other desirable characteristics such as

- those which are most likely to do reasonably well over a wide range of economic scenarios,
- those which involve the lowest amount of transactions to implement
- those which are least sensitive to changes in assumptions

These types of optimizers have many advantages over MVOs.

- They tend to be much more diversified
- The capture the value of investment insights because they don't need limits and constraints in order to produce sensible looking portfolios (even though the results are often very different from conventional portfolios)
- They tend to produce stable recommendations, small changes in forecasts produce small changes in recommended portfolios
- They assume the forecasts are not perfect and take that into account

Profiling- getting the right portfolio with the right level of risk

Assuming we have determined a set of model asset allocations of different risk levels, the next task is to work out which one to use for any particular client or application. This is no trivial exercise and has created much debate within the industry.

Academics and Utility functions

The academics say that in order to find the right portfolio for a particular situation you use a utility function; in other words solve for a formula like;

Maximise; $E\{U(R)\} = E\{R\} - \lambda\sigma^2(R)$

Where $E\{U(R)\}$ is expected utility, $E\{R\}$ is expected return, λ is risk tolerance and $\sigma^2(R)$ is forecast variance.

The challenge is determining the value of λ, the investors risk tolerance. The academics can simply assume that if you have a client with a risk tolerance of say 2.5, then the portfolio will look so. In practice this is not a trivial issue, determining λ is not intuitively obvious and there is little guidance on this subject.

Nonetheless this could be a promising approach if the other flaws in the approach could be resolved, i.e. forecasting returns, defining and forecasting risk; overconfident optimization….

Industry practice

Needless to say the utility approach has little resonance with the practicing investment community. Nonetheless risk profiling has become a contentious issue, with practices ranging from a focus on financial risk tolerance, to a focus on emotional risk tolerance, all to the way through to some approaches that are a completely misguided combinations of the two.

Emotional risk tolerance can best be described as the client's ability to cope with the stress caused by volatility and uncertainty during the investment journey. Financial risk tolerance is their ability to cope with poor long-term financial outcomes - the investment destination.

By way of illustration, imagine a stockbroker who retires after 40 years in the industry, he has seen it all and fully understands and is totally comfortably with sharemarket risk. Should he invest 100% in the sharemarket?

Those who focus on emotional risk tend to say that very high sharemarket exposures are just fine for this type of investor. The journey will be fine, as he does understand and is comfortable with the risks involved with the sharemarket.

But what if the ex-broker said that his absolute bare minimum lifestyle requirements were such that required a real return from his investments of 2%pa, and his desired lifestyle required a real return of 4%pa? He is now describing his financial risk tolerance. What he is saying is that he would not be *financially* able to tolerate the consequences of long-term real returns below 2%pa, which, as any Japanese investor can tell you, is a real possibility for an investor with 100% exposure to the sharemarket.

In practice the right level of risk to be adopted by the client is the *lesser* of these two risks. It's no point of being comfortable with the journey if there is an unacceptably high probability of not arriving at the destination. Similarly there is no point in heading towards the right destination if the client is not on board at the finish.

Both risks have to be addressed separately. Which brings us to the misguided approaches, ones that attempt to measure both tolerances at the same time and end up with some sort of crude average.

For example, the client fills out a questionnaire which asks questions such as;

- What is your age; 30-40, 40 -50, 50-60, 60-70, 70+?
- When will you retire 10+ years, 5-10 years, 0-5 years, already retired?
- Are you generally conservative, aggressive etc?

The first two questions have very little to do with the clients emotional risk tolerance, but have a lot to do with their financial risk tolerance. However as financial risk questions they are extraordinary blunt instruments. As all planners know, deferring retirement for two years can often dramatically alter an individual's financial situation and therefore their financial risk tolerance. This type of survey doesn't even distinguish between 1 and 5 years to retirement.

It then mixes up the data collected on financial risk tolerance with the data on emotional risk tolerance, to get a mixed score that lines up with a model portfolio of some kind, without reference to the client's objectives. While this approach is probably better than no attempt at profiling, it probably misses the mark more often that it hits it.

Current best practice

Best practice seeks to understand both financial and emotional risk tolerance and ensures that the portfolio is within acceptable bounds for each.

There are excellent commercial systems available which assess both sets of risks; and understanding the client's relative tolerances creates further opportunities for portfolio design. For example, if emotional risk tolerance is the limiting factor then using direct fixed interest rather than fixed interest funds, or property syndicates rather than LPTs can reduce volatility without impacting long term returns. If financial risk tolerance rather than emotional risk tolerance is the limiting factor for a particular client then it will be necessary to work on the underlying structure of the investments or spending plan in order to get a satisfactory outcome.

Implementation

Let's say that we have worked through the first four steps of the process and arrived with a suggested asset allocation for a particular client. We now have to consider how this actually gets implemented. And now a whole new batch of problems arise, (except for the academics who don't have to worry about implementation.)

Practitioners have to worry about a whole raft of real world issues such as their client's biases and preferences, their own biases and preferences, the current portfolios of clients and the capital gains tax and cost consequences of any transactions. They should also worry about whether the way they choose to implement investment in any asset class will reflect the broad underlying asset class. Is a large holding in a single share from a client's employer representative of the risk inherent in a broad exposure to the Australian equities market? How much value will an active manager add? These types of issues are all part of the equation.

In practice most planners start with a model portfolio and try to implement in bands around the model portfolio, i.e. if the model portfolio has a 30% exposure to International Equities they try to keep the International equities exposure between 25 and 35%, while accommodating client preferences and existing holdings. And, if one is comfortable with the model portfolio and all the underlying assumptions that it took to get there(!), then that is not a bad approach. But it can be improved upon substantially.

A Disciplined and flexible approach to implementation.

If the planner focuses on expected portfolio outcomes and not on portfolio structures there can be a surprising amount of flexibility to design portfolios. By outcomes we mean the expected return and the expected risk of the portfolio, after costs.

Consider the two portfolios in the table below; is it worth the client moving from the current portfolio to the model portfolio?

Figure 7 : Impact of transaction costs on optimal allocations

Asset class	Current Portfolio	Model Portfolio	Model Portfolio at band edges
Australian equities	20%	40%	35%
International equities	10%	30%	25%
Property	40%	10%	15%
Fixed Interest	20%	10%	15%
Cash	10%	10%	10%
Forecast return before transaction costs	7.7%pa	8.0%pa	7.9%pa
Forecast return after transaction costs	7.7%pa	7.5%pa	7.6%pa

If there were no costs, maybe. But there are costs, and after taking them into account any expected return pick up more than disappears. Even moving to the edges of the model portfolio bands doesn't save the situation. There are many such situations where a theoretically optimal portfolio is sub optimal in the real world.

To understand why this occurs refer back to Figure 6. In the Robust Band there are a myriad of portfolios which have very different asset allocations, but have very similar return outlooks. The job of the asset allocator is to choose the portfolio that best suits the client's particular circumstances including their current portfolio and their beliefs and preferences. In the example illustrated in Figure 7, what if the client was an experienced property investor and was totally comfortable with their property exposure but felt far less sure of the sharemarket? Why make the change?

Furthermore, even if the model portfolio showed a modest, say 0.1%pa, pickup in expected returns after transaction costs, is it worth the cost to chase that small pickup? What is the confidence level of the return forecasts, are you absolutely certain that you will get the extra 0.1%? Because you can be absolutely certain about paying the costs.

This approach of staying in the robust region provides both flexibility and discipline. But it does rely on

- Having forecasts that you can rely on
- Using measures of risk that relate to the risks faced by clients
- Having portfolios that are optimized or designed to meet client objectives
- Taking into account transaction costs and taxes

In short a carefully constructed integrated approach to asset allocation is required to achieve reliably achieve client goals.

ACTION STEPS

1. Those practicing asset allocation should thoroughly understand the how the asset allocations they use as the basis for their recommendations are arrived at. They should have at least a passing knowledge of the underlying theory, the key assumptions behind that theory and its limitations.
2. All forecasts of risk and return should be forward looking. Chasing historical returns is worse than useless. Develop a thorough understanding of the Occam's razor forecasting process or an equivalent. The original john Bogle paper can be found at www.farrelly.com.au in the Articles page.
3. Risk management should be focused at real risks faced by investors. A discussion on this issue can be found in Towards A better Measure of Risk. Tim Farrelly, Portfolio Construction Journal, Spring 2004, Volume 1 Issue 2. Also available on www.farrelly.com.au
4. Portfolios should be designed to meet cash flow needs. The major purpose of most portfolios is to fund some future liability such as retirement, keep that uppermost in mind when building portfolios.
5. Transaction costs and taxes are important and should be factored into all decisions. How do the expected gains compare with the costs involved to achieve them?
6. Concentrate on investment outcomes rather than portfolio structures.

ABOUT THE AUTHOR

Tim Farrelly is Principal of **farrelly's**, Australia's only specialist asset allocation researcher specializing in asset allocation. Tim brings to the task of asset allocation a unique combination of analytics, understanding of financial markets, knowledge of capital market history and insight into the practical requirements of financial planners.

He is on the Editorial Board of the Journal of Portfolio Construction and chairs the Inquisitor Program at the Portfolio Construction Forum.. Tim is a sought after speaker and a frequent presenter at FPA annual and state conferences on a range of topics including capital market history, risk management, and portfolio construction.

Prior to founding **farrelly's** in 2004, Tim was an Executive Director of Macquarie Bank Ltd, and Director of Macquarie Investment Management Ltd.(MIML) At various times during his 14 years at Macquarie he sat on the MIML Asset Allocation and Risk Committees, and was responsible for distribution of the Bank's products through third party financial planners and stockbrokers. While at Macquarie, Tim was responsible for the booklet 'Understanding Risk to Meet Your Financial Goals', which is jointly published by Macquarie and the FPA and has become an industry standard, and the Long Term Forecasting program in 2000 which foreshadowed the bearmarket in US equities.

Between 1981 and 1986 Tim was head of research for the Monitor Money Corporation, where he was responsible for asset allocation and manager selection.

Tim can be contacted at tim.farrelly@farrelly.com.au