MEAN VARIANCE TACTICAL ASSET ALLOCATION PROCESS

Peter Mancell and Corey Edwards[1]

July, 2009

Abstract

A comparison of performance was undertaken for two investment portfolios. One portfolio adopted a tactical asset allocation approach which was determined by excessively good or poor market returns when compared to the mean of Australian and International Equity market returns. The other portfolio adopted a passive approach to asset allocation. Both portfolios were rebalanced to their target asset allocations monthly.

The monthly returns of the ASX 500 Index and the MSCI World Index ex-Australia (Unhedged) from January 1980 to December 2008 were recorded and plotted, which led to the calculation of a series of 5 year rolling returns for each asset class.

These rolling returns were then compared to both their respective means across the period measured and an estimate of the long term annual average return of each asset class.

By looking at the variance of the rolling returns from the means, a series of asset allocation rules were formulated. These rules determined the tactical asset allocation at any given point in time.

The tactical asset allocation rules were based upon 'bands' determined by over or under performance for each of Australian Equities and International Equities.

The 'band' adjusted mean variance asset allocation portfolio (referred to as the MVP hereafter) is then compared to a passive rebalancing portfolio (PP), attempting to find improved portfolio performance.

[1] Disclaimer: This paper contains opinions of the author's but not necessarily of the Asset Class Investors Association, and does not constitute or represent a recommendation of a security, strategy or investment product. The opinions of the authors are subject to change without notice. This paper is prepared for educational purposes only and should not be considered investment advice. This paper may not be reproduced without the express consent of the authors. The information presented in this paper was only current at creation and should not be relied upon without further data analysis in future.

Mean Variance Asset Allocation Process Hypothesis

That a tactically re-balanced portfolio using a set of band-determined asset allocations (established by the mean variance of Australian and International Equities) will outperform a passively managed portfolio[2].

Background and Assumptions

For the purposes of this exercise, the initial portfolio asset allocation will be 40% Defensive Assets and 60% Growth Assets (Shares).

The cash will be held in A$ and deliver a return equal to the UBS Bank Bills index.

Of the 60% growth assets initial split will be 35% Australian and measured by the ASX500 Index. The International will be 25% and measured by the MSCI ex-Australia and will remain un-hedged.

The trigger point bands are set either side of the mean return for each Equity asset class during the data sample period available.

The initial trigger point for the overweight Equities position will be when the 5 year rolling return from Australian or International Equities reaches 50% below the long term mean and then the position will be increased at 75% below the mean.

The initial trigger point for the underweight Equities position will be when the 5 year rolling return from Australian or International Equities reaches 50% above the long term mean and then the position will be increased at 75% above the mean.

Naturally, Australian and International Equities may not move in a synchronised manner, and as such, a matrix of overweight and underweight positions was created to facilitate the potential series of outcomes. It is possible therefore to have as an example an overweight Australian Equity position and an underweight International Equity position, if the relevant rules are applied within the MVP.

Transaction costs have been applied to both portfolios, with a buy/sell spread of 0.50% and 0.60% assumed for Australian Equity and International Equity respectively.

[2] It should be noted that both the PP and the MVP will be rebalanced monthly, however the MVP will be balanced to conform with the asset allocation rules provided by the 'band' it is currently placed within.

Investigation Process

Step 1.

The longest data series available for both Australian Equities and International Equities (Unhedged) were compiled (Jan 1980 to Dec 2008)[3]. This data was collected in the form of monthly returns and used to compile a MR5AA (monthly returns 5 year annualised average). For 29 years of data there were 348 data points.

After data was compiled, the mean and standard deviation of all rolling returns was computed. All data was then exported to Excel.

Step 2.

All rolling returns were then compared to the mean of the entire data series. Upon collection of all the variances from the long term mean, a distribution was calculated. This distribution placed each of the results in a band around the long term mean.

Step 3.

From the data all bands of 'trigger points' were identified.

Step 4.

An asset allocation hypothesis was then constructed, where by a difference of amount x% between the mean and the rolling return resulted in an adjustment of y% in the given asset class allocation, to or from defensive assets.

These asset allocations were then compiled into a tactically rebalanced mean variance portfolio (MVP), which was then compared to a passive rebalanced portfolio (PP). Performance of each portfolio was calculated, and transaction costs removed.

The difference determined if the rule would have been of benefit to the investor.

Step 5.

An in-series analysis of three 15 year time blocks was then conducted, to attempt to remove from the data the effects of any one significant event.

Step 6.

The process was then repeated using an assumed long term mean for each asset class, rather than the mean of the data series. This provided a more realistic practical model.

Step 7.

Finally, various other scenarios were modelled.

[3] Monthly returns data provided from the DFA Returns Program

Data

The Australian Equity rolling returns period data presented a distribution with a positive skew[4], while no pattern of distribution appears evident in the International Equity rolling returns data. Note the large difference in standard deviation sizes, reflecting a much greater dispersion of returns.

Histogram 1.

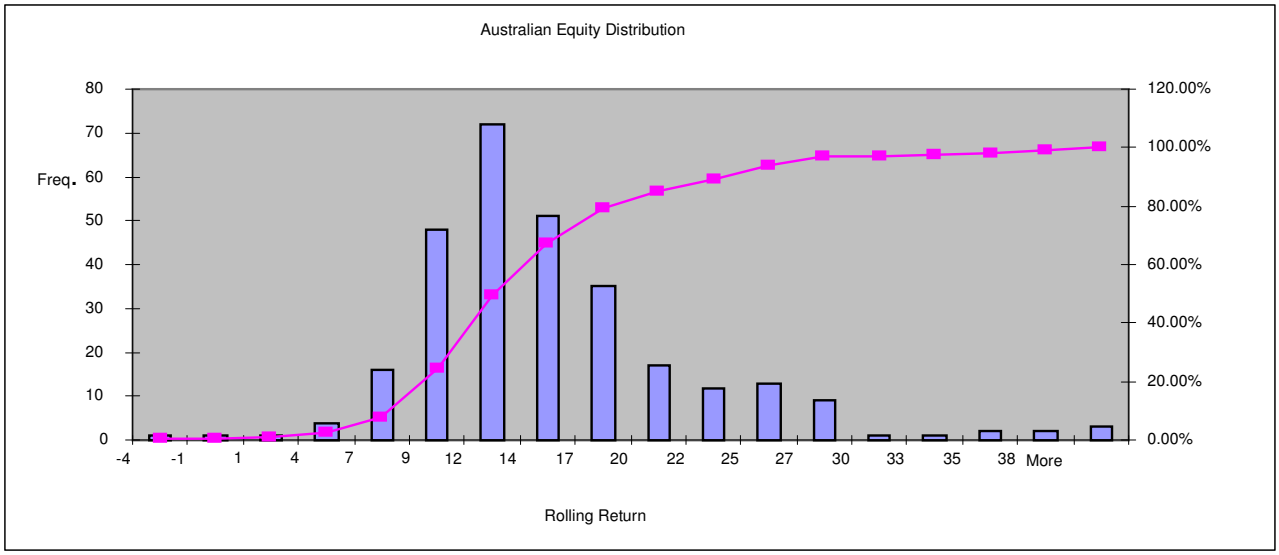

Mean	13.41
Standard Deviation	6.62

Histogram 2.

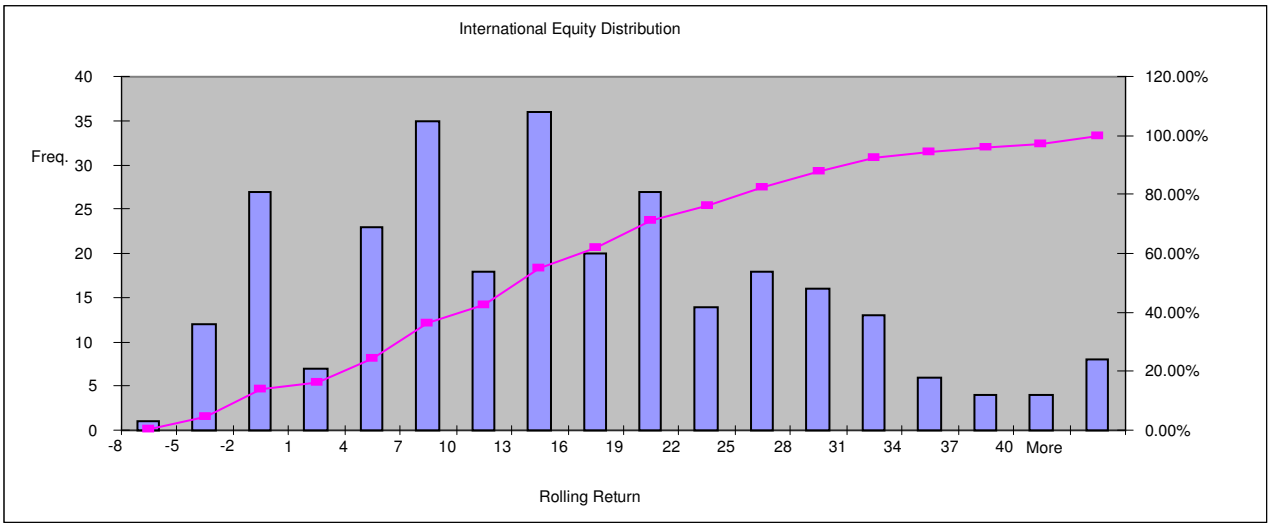

Mean	12.93
Standard Deviation	12.15

[4] No formal comparison to a normal distribution was made. The data for AE does appear to be close to approximating a normal distribution, usually defined as approximately normal if approximately 68% of the values are within 1 standard deviation of the mean (mathematically, $\mu \pm \sigma$, where μ is the arithmetic mean), approximately 95% of the values are within two standard deviations ($\mu \pm 2\sigma$), and approximately 99.7% lie within 3 standard deviations ($\mu \pm 3\sigma$). The unusual distribution of IE could be expected to be caused by currency fluctuations, however this has not been investigated.

5 Year Rolling Returns Data

Data was available for both MSCI World Index ex-Australia (Unhedged) and ASX 500 Index from January 1980 to December 2008. 5 Year average annualised returns were calculated on a rolling monthly basis, and plotted across the available time period.

The 5 year rolling return mean was calculated, as well as 'bands' around the mean at the 25%, 50%, 150% and 175% of mean levels. Also plotted were actual monthly returns (See Chart 1 & 2).[5]

Chart 1.

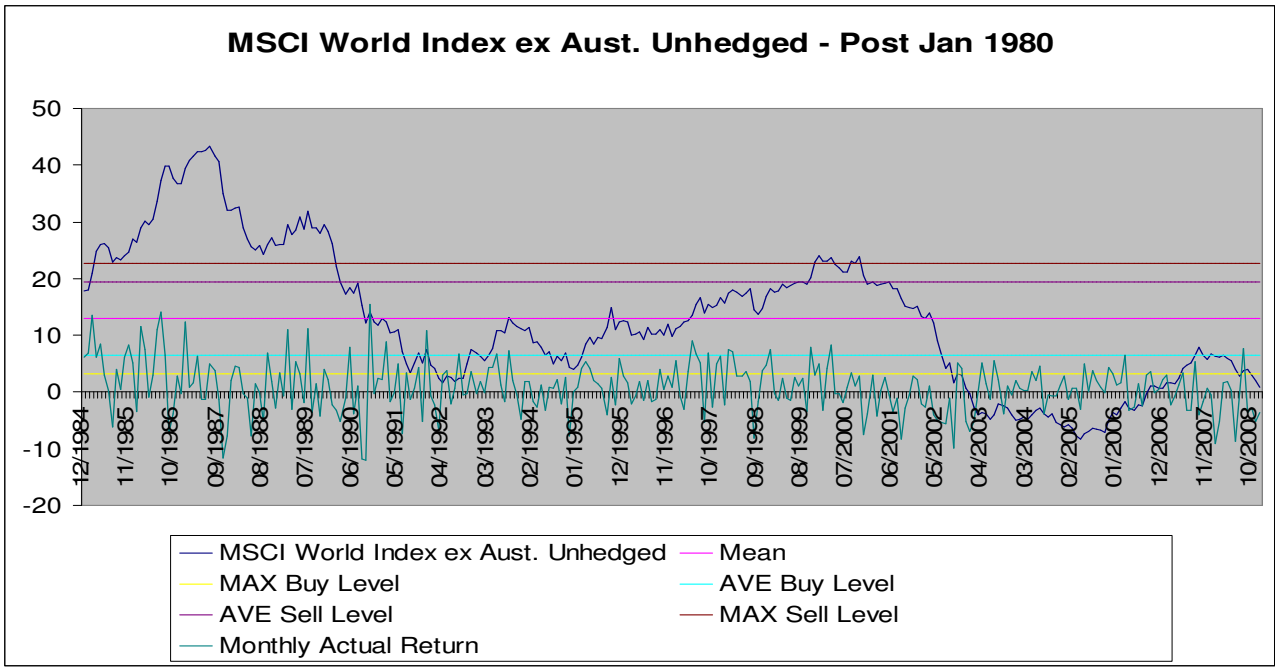

Chart 2.

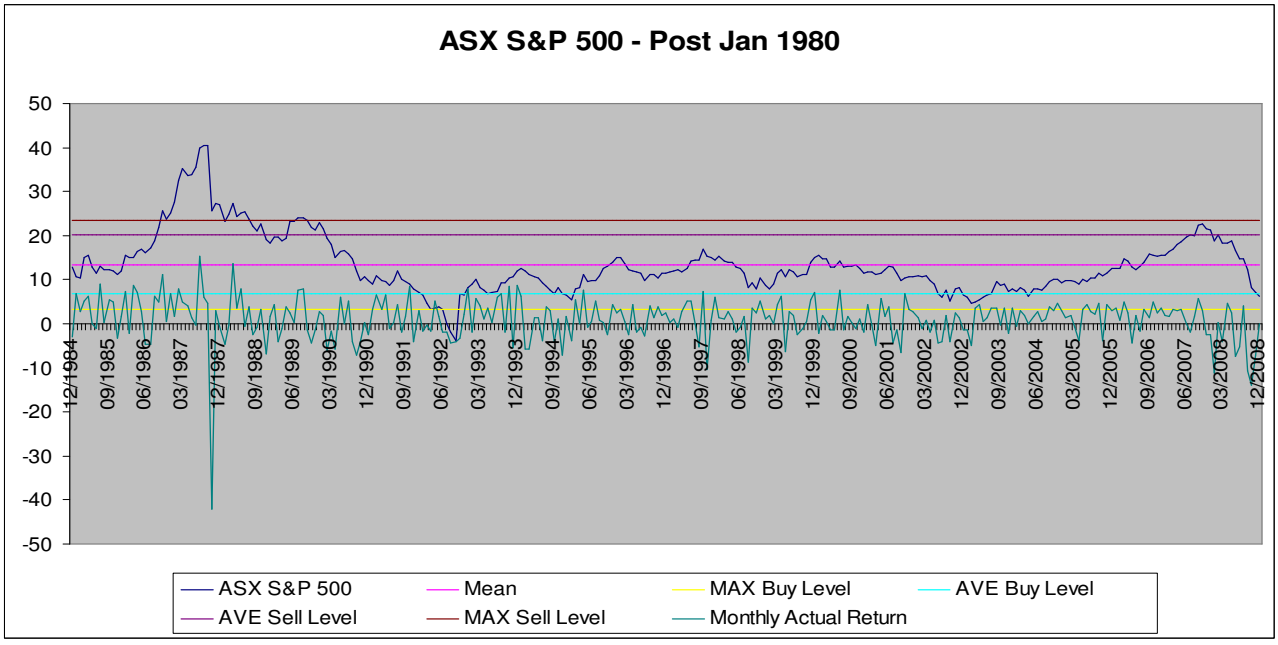

[5] Recall the mean of the MSCI World Index ex-Australia across the time period was 12.93124%, while the mean of the ASX 500 Index was 13.41296%

Trigger Point Bands

The Trigger Point for adjustment to the asset allocation is the movement of the 5 year rolling return compared to the mean. When the 5 year rolling return passes a trigger point into a new 'band', a new asset allocation is to apply, until the 5 year rolling return moves into a different band.

The number of asset allocation adjustments is determined via the amount of occasions the average passes into a new 'band'. Each of these occasions is highlighted on the charts below by a red mark. These indicate a transaction being made to amend the asset allocation to the level specified for that particular band.

Chart 3.[6]

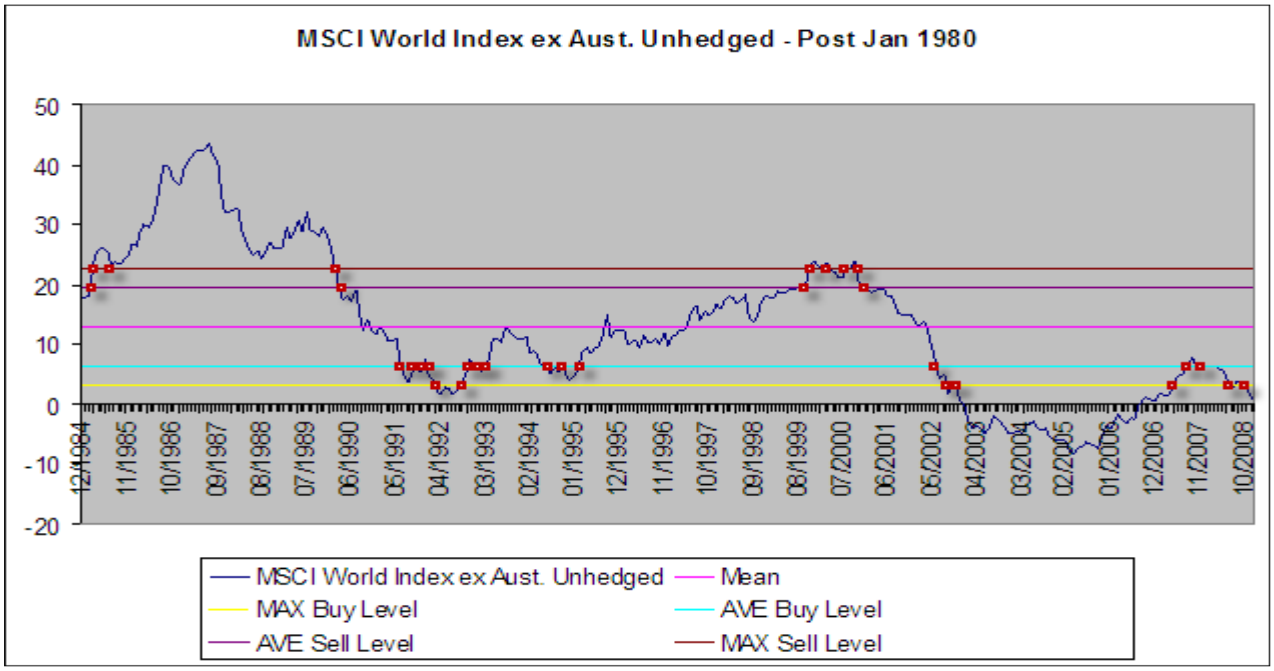

Chart 4.

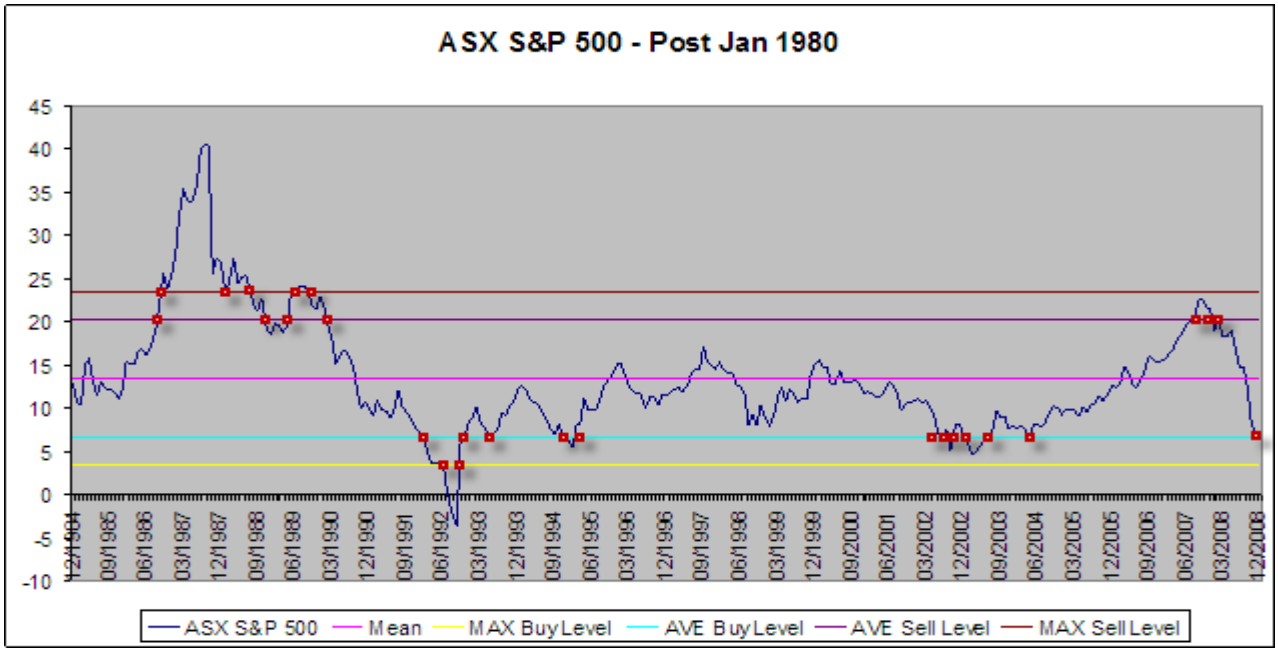

[6] Lower quality graphic in Chart 3 and Chart 4 due to image capture process utilised to insert trigger points.

Asset Allocation Matrix

The Trigger Point Bands, Allocation Impact and the Initial Asset Allocations[7] have all been set as per the tables below. These have been set at this level to attempt to find a balance between minimising the amount of transactions and still capturing significant movements.

The Asset Allocation Impact shows the adjustment in terms of the overall portfolio movement to or from an asset class.

For example, should the ASX 500 Index 5 year rolling average fall to below 25% of the mean, the 35% portfolio allocation to Australian Equity will rise to 55%.

Trigger Point Bands @ Percent of Mean	
MAX Buy Limit	25%
AVE Buy Limit	50%
AVE Sell Limit	150%
MAX Sell Limit	175%

AE Allocation Impact	IE Allocation Impact
20%	20%
10%	10%
-10%	-10%
-20%	-20%

Initial Asset Allocation	
Cash	40%
Australian Equities	35%
International Equities	25%
Total	100%

The matrices below highlight all possible permutations of our portfolio given the listed constraints above.[8]

It should be noted that the UBS Warburg 90-Day Bank Bill is used as the cash proxy.

Asset Alloc. AE/IE	MAX Buy IE		AVE Buy IE		Neutral		AVE Sell IE		MAX Sell IE	
MAX Buy AE	55%	45%	55%	35%	55%	25%	55%	15%	55%	5%
Cash Remaining	**0%**		10%		20%		30%		**40%**	
AVE Buy AE	45%	45%	45%	35%	45%	25%	45%	15%	45%	5%
Cash Remaining	10%		**20%**		30%		**40%**		50%	
Neutral	35%	45%	35%	35%	35%	25%	35%	15%	35%	5%
Cash Remaining	20%		30%		**40%**		50%		60%	
AVE Sell AE	25%	45%	25%	35%	25%	25%	25%	15%	25%	5%
Cash Remaining	30%		**40%**		50%		**60%**		70%	
MAX Sell AE	15%	45%	15%	35%	15%	25%	15%	15%	15%	5%
Cash Remaining	**40%**		50%		60%		70%		**80%**	

Growth Asset Alloc.	MAX Buy IE	AVE Buy IE	Neutral	AVE Sell IE	MAX Sell IE
MAX Buy AE	**100.00%**	90.00%	80.00%	70.00%	**60.00%**
AVE Buy AE	90.00%	**80.00%**	70.00%	**60.00%**	50.00%
Neutral	80.00%	70.00%	**60.00%**	50.00%	40.00%
AVE Sell AE	70.00%	**60.00%**	50.00%	**40.00%**	30.00%
MAX Sell AE	**60.00%**	50.00%	40.00%	30.00%	**20.00%**

[7] All of these features are adjustable via use of the excel spreadsheet calculator.
[8] The coloured areas on the above matrices represent the distribution of growth assets in any given situation, with green implying overweight compared to initial allocation, blue equal, and red underweight.

Results

Chart 5 illustrates the comparison between the MVP and the PP. Also included here is a passive non-rebalance portfolio (PP–NR). This portfolio was created initially with the same asset allocation as the PP but was allowed to grow without adjustment over the full time period.

It can be seen that the only significant event that was predicted and dealt with by the tactical rebalance process was the 1987 crash. In most other situations the adjusting does not provide substantial outperformance to the portfolio. Note also that the outperformance of the MVP at this point avoided the much of the major rally as well as the subsequent fall.

Upon looking at the returns data, it seems to show that momentum pulls in both directions, causing the tactical asset allocation process to move too early. The boom in the data series tended to rise far higher and longer into the band limits, with the MVP missing much of the peak[9].

The declines also moved much deeper and persisted longer, with the allocation moving the MVP to be more exposed to several of these deeper lows, before reducing the amount of growth assets held when the asset class has begun to rise again.

In this way the MVP tended to be inappropriately invested at crucial points in the particular asset class returns cycle where larger gains or losses accumulated.

This can be seen in the Chart 5, where the difference in the cumulative growth of $1 of the PP ($10.41) over the MVP ($9.98) over the full period was $0.43, or 4.13% lower. The PP had a final value of $9.61.

Chart 5.

MVP vs. PP vs. PP-NR

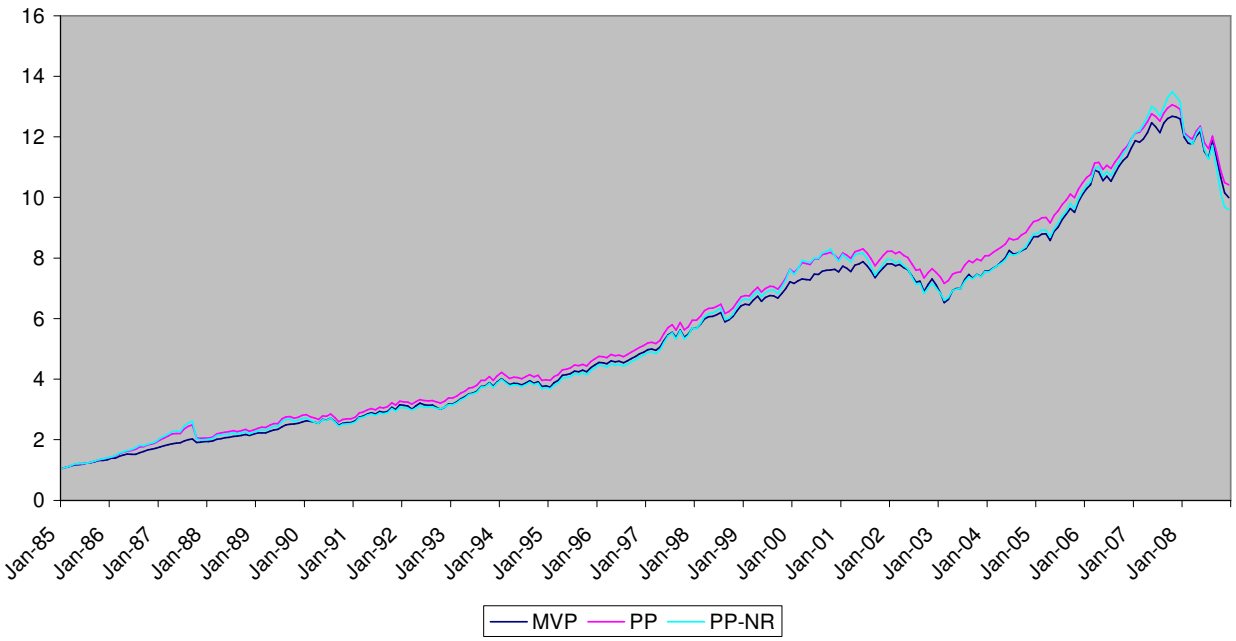

[9] The boom referred to in the data series predominantly revolves around the lead up to 1987, while the decline is the aftermath.

In-Series Analysis

Unsurprisingly the 15 year comparison beginning in 1980 appeared very similar to that of the initial stages of the total time period. The major difference statistically is that the mean being used to make comparisons was the internal mean. i.e. that of the 15 year period being reviewed, rather than that of the entire time period. Note that transaction costs are applied to these scenarios in the same fashion as the full period results.

Chart 6.

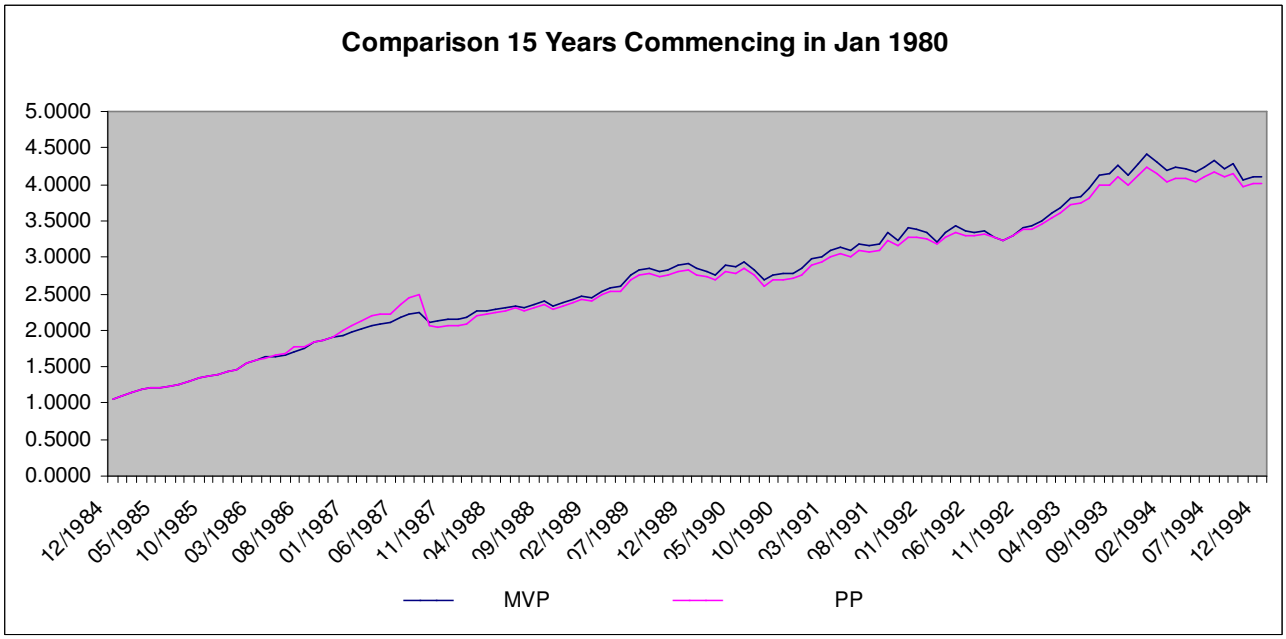

The middle period comparison shown in Chart 7 provides more evidence that the PP was likely to perform close to the MVP.

Chart 7.

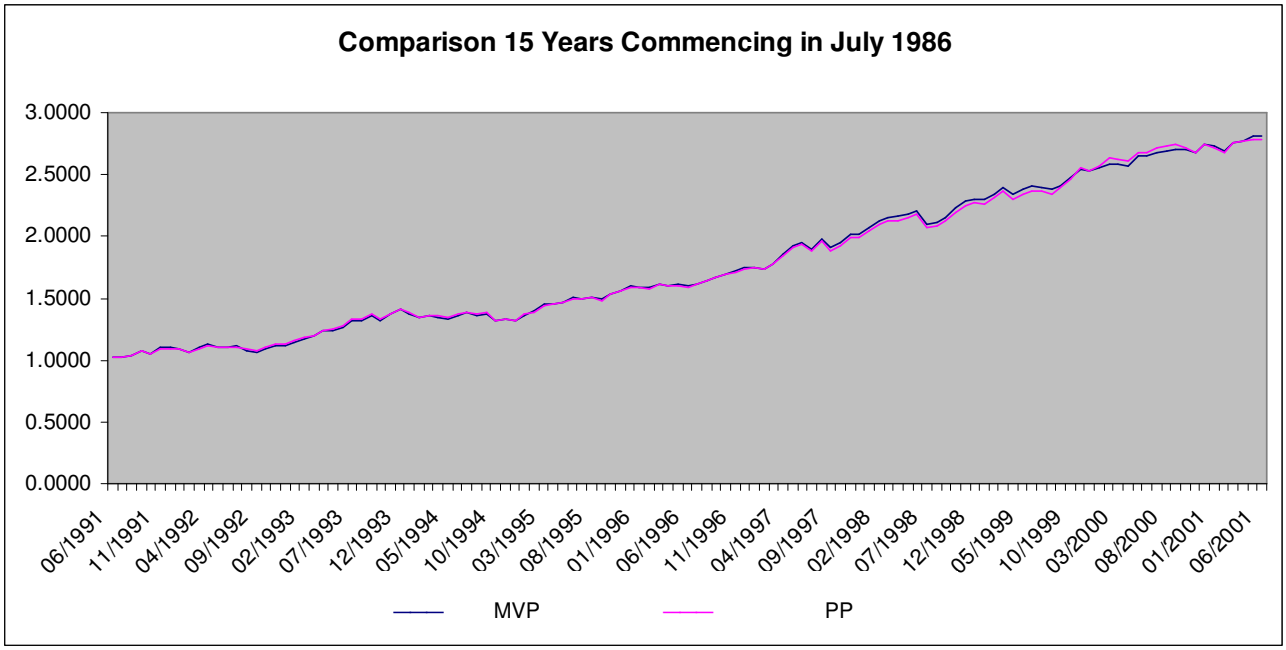

The comparison in Chart 8 highlights that before transaction costs and taxes the MVP actually seemed to provide a small benefit over the final 15 year time period.

Chart 8.

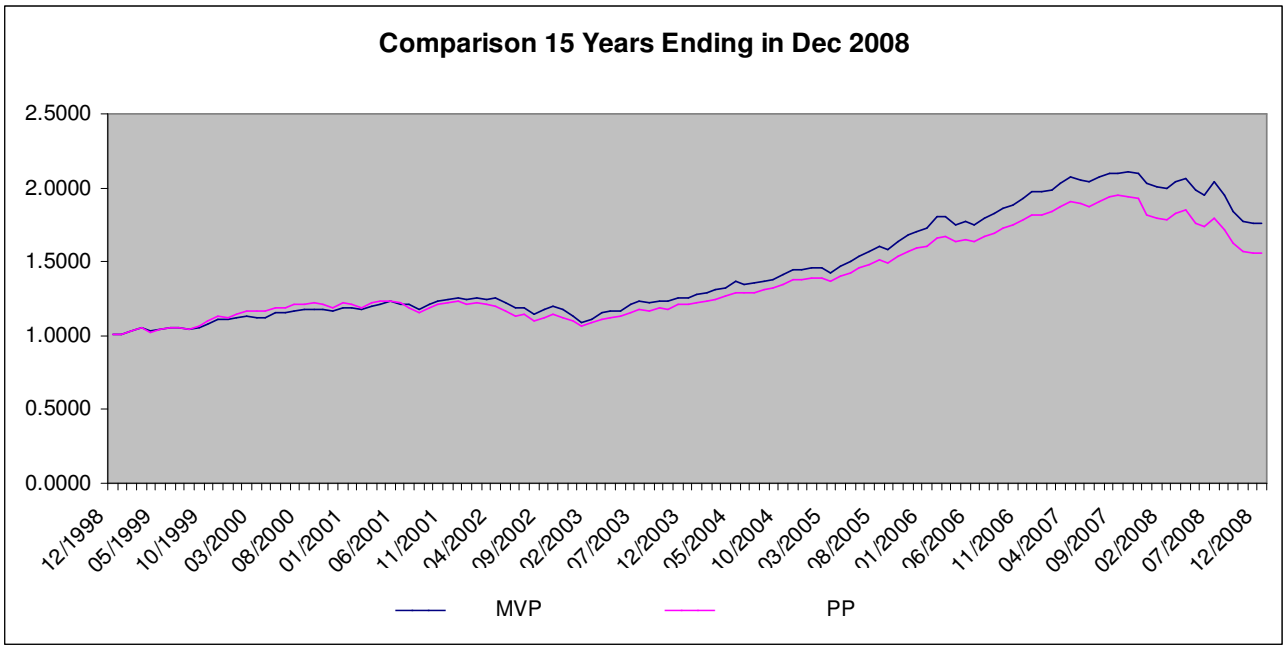

Mean Set @ 13%

A notable flaw with the creation of this process was the use of the mean of the entire 28 year period being applied to make asset allocation decisions that occur at the beginning of the period. It would be impossible to have this information prior to December 2008.

It is also apparent that when the mean of a period is being utilised to ascertain if a period of time is mean reverting, the answer will be yes.

To provide a more realistic investigation of the process, the mean of the period was replaced with a specified mean at 13%. This represents an approximate long term return figure for both Australian and International Equities.

Upon testing using this specified mean, the results seen in Chart 9 were produced.

A slight improvement in performance was seen, with the final cumulative growth of $1 figure now equalling $10.01 for the MVP, an improvement of $0.03. The two passive portfolio results were unchanged from $10.41 and $9.61 for the PP and PP-NR respectively.[10]

Chart 9.

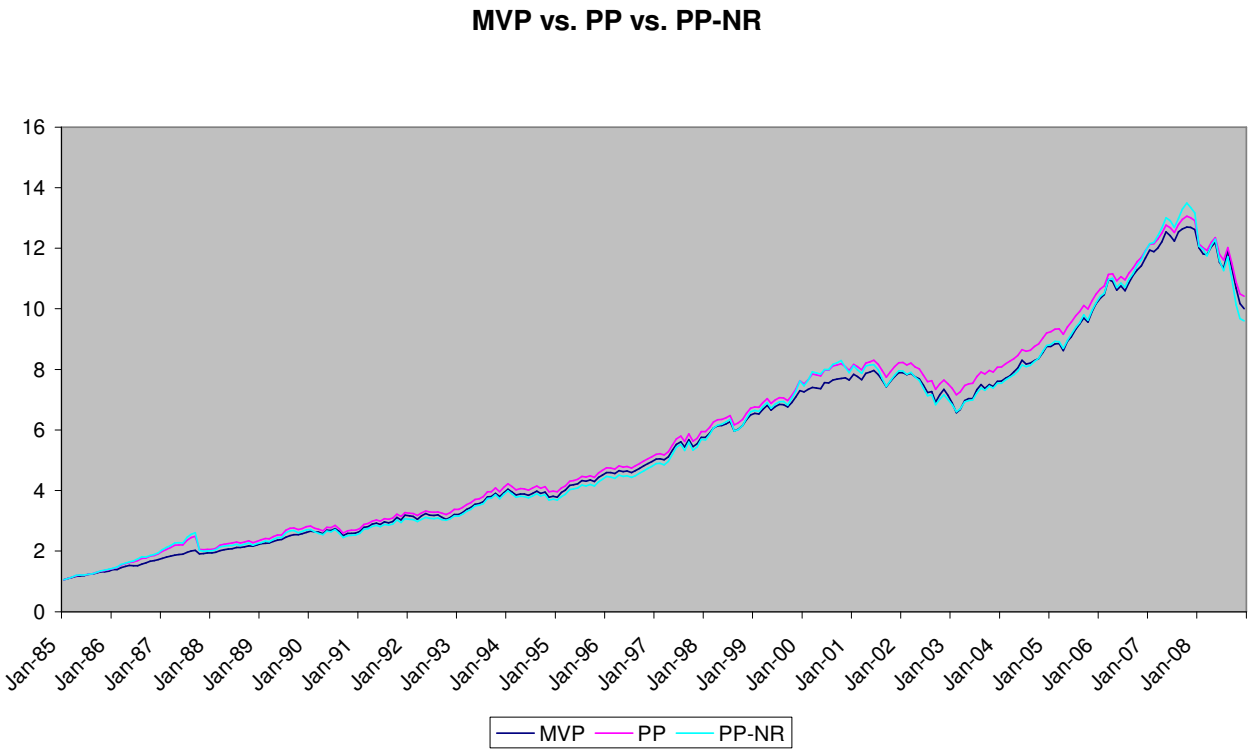

[10] This has the effect of reducing the shortfall from the PP-NR portfolio from 4.13% to 3.84%.

Practical Applications – Tax and Rebalancing

The process described, while workable as a theory, suffers from several major problems when considered for practical application.

Taxes would be a significant cost. This cost could greatly increase the small deficit seen to occur in comparison to the passive rebalance portfolio over the entire data series.

Both portfolios are assumed to rebalance to their particular asset allocations monthly. For the PP this implies that the asset allocation never varies, while for the MVP it conforms to the asset allocation applicable at the given 'band'. While statistically this does not present a problem as both are being treated equally, in reality rebalancing is highly unlikely to take place monthly.

Variables and Manipulation

The ability of the process to be manipulated also casts doubt over the validity of any disproving conclusions. The amount of variables in the process makes it very difficult to test all possible variations.

Data Sample Size

A major concern for this process is the inability to perform out of sample testing of the process due to the data sample size. The 15 year in-sample testing is the only method available to attempt to remove the influence of major events that tend to 'skew' results. In-sample testing is still flawed however. Ideally testing would be able to be conducted over an entirely fresh data series.

Data Analysis

Further steps in analysing the data would likely be focussed on regressing the asset class data series returns against the 5 year annualised average performance of the class, to test the existence and strength of any relationship. This is not possible at this stage due to resource limitations, but if completed would likely involve using a GARCH process due to the nature of the data. This may then provide some more concrete answers as to the ability of the long term predictive capacity of the average.

Final Note on Concerns

At the time of publishing this document, we have received a report from David Surridge and Nigel Walker (with the comments and assistance of Jim Davis) confirming the validity of our conclusions and also extending our analysis to deal with some of the limitations and concerns raised above. This report and its findings are commented on in the Further Analysis section.

Conclusions

Results drawn from the MVP process make it impossible to support the hypothesis, but also impossible to entirely disprove it.

On balance, more of the evidence is contrary to the hypothesis.

It is possible to say that using the current variable choices, the process did not support the hypothesis.

Interestingly, the process did show two 15 year periods of slight improvement and also high levels of variability in performance upon manipulation of the variables.

It will only be viable to conduct future tests outside of the current data sample, as the 'home sample' bias already present in the results is significant.

Another possible extension of the process may be to include other asset classes. It should be noted however that this will significantly complicate the analysis.

At this stage however, it must be stated that the process does not produce a reliable short term indication of future market movement, nor how an asset allocation can be constructed to take advantage of any under or over pricing of asset classes apparent in the market. At best it has managed to meet passive market performance.[11]

[11] Acknowledgement must be provided to Dimensional Fund Advisers (DFA) for utilisation of the Returns Program software.

Further Analysis

A version of this report was provided to David Surridge, Nigel Walker and Jim Davis[12] and they proceeded to test both the robustness of our model and also to extend the analysis.

Their analysis can be found in the paper titled '*Under Mean Variance Assumptions can Tactically Rebalanced Portfolios outperform a Passive Approach?'*[13].

Their findings on the Australian situation that was analysed within our paper were supportive of both our method and conclusions. In response to our concerns and limitations they further extended our work to look at both a wider array of variable change and also conduct testing of our model outside Australia.

Within Australia 100 different MVP's were tested, with variable adjustments occurring to the mean, trigger points and the length of rolling returns used to construct each portfolio. 29 out of 100 portfolio's were found to have a result that was statistically significant[14]. In each of these 'significant' results, the PP was found to have outperformed the MVP.

In an aside, testing was also conducted on the impact of removing the monthly rebalancing. Under these rules asset allocation change to conform to the 'rules' only took place when a trigger point was reached in the MVP, and not at all within the PP. This PP was the same therefore as the PP-NR created within our work. No statistically significant results were found, with the PP slightly outperforming the MVP.

To eliminate the sample bias that may have been present in the Australian findings further testing was conducted by creating a MVP based upon the United States data[15]. Over the entire period the results were found to be a slight improvement in results by the MVP over the PP (7.72% annualised return compared to 7.69%), however because of the very low statistical significance the evidence is insufficient to support the result.

In an attempt to isolate potential periods of dramatic over or under performance, the data series was then broken into two discrete components, Jan 1931 – July 1951 and July 1979 – Sep 2001. In the early period, which included the Great Depression, the MVP portfolio had outperformance that bordered upon being statistically significant with a t-stat of 1.81[16]. In the latter period the results were reversed, with the PP outperforming the MVP but with no real statistical significance.

In conclusion, the mean variance process was found to be as ineffective in outperforming the passive process in the longer US time period as it was in Australia.

[12] Jim Davis has been a professor of finance at Kansas State University and held positions with Arthur Andersen & Co. and Phillips Petroleum. He has also written and reviewed for many financial publications. Jim holds a Ph.D. from the University of Illinois, an MBA from DePaul University, and a degree in finance from Oklahoma State University.
[13] Surridge, D. & Walker, N. (2009). *Under Mean Variance Assumptions can Tactically Rebalanced Portfolios outperform a Passive Approach*. DFA Australia Limited. Sydney, Australia.
[14] Statistical significance here is determined by the value of the related t-statistic being above 2.00 in absolute value.
[15] In the US calculations, only two asset classes were considered. US equities as represented by the S&P 500 Index and Cash as represented by US Treasury Bills. US data used extended from January 1931 through to December 2008.
[16] Actual outperformance of the MVP over the PP was 7.15% to 5.77% annualised return.