January 2002

Abstract

This paper presents a new approach to portfolio optimisation, which we call generalised mean-variance (GMV) analysis. One important case of this approach is based on what we call the stocks m-tile but which is often referred to as its quantile. If $m = n$, where n is the number of stocks, m-tile membership becomes rank. We consider our rank based GMV analysis to be in effect, the rank equivalent of conventional Markowitz Mean Variance analysis. The first stage of this process is to generate rank probability statistics using, historic data, Monte Carlo analysis or direct user input. The second stage is optimisation based on those rank statistics to calculate recommended portfolio weights.

The approach we take to optimisation uses state preference theory to derive an objective function that can be minimised using standard quadratic programming techniques. The paper outlines a number of advantages of this method which include, a more intuitive fully diversified (or minimum risk) position on the efficient frontier with all the portfolio holdings equally weighted. It also results in more stable portfolios due to reduced sensitivity to the perfect substitute problem, as well as the well-known robustness of rank statistics to the presence of outliers in the data. We also introduce a very intuitive set of summary statistics which provide some protection against undetected concentrations of risk such as those that can occur when working with assets which have a very non normal forecast return distribution.

The disadvantage of the approach, is that if we use ranked mean and ranked variance in the search for robustness, it throws away some of the information available in the conventional analysis. Also there is as yet little practical experience of using this approach, hence additional care must be taken with the design of the forecasting and portfolio construction process until sufficient confidence has built up in its use. It is worth noting that our GMV approach could use a mix of a ranked mean and a conventional variance to construct portfolios, or indeed other combinations, so that the above disadvantage can be reduced.

Key Words: Mean Variance Analysis, Diversification, Portfolio Construction, Forecasts.
JEL Classification: G00, C53, C52, C15

Generalised Mean-Variance Analysis and Robust Portfolio Diversification

Stephen M. Wright[1] and S.E. Satchell[2]

1. Introduction

Mean Variance (MV) Analysis is widely accepted as the best way of analysing and explaining the benefits of diversification of holdings across a portfolio of assets at least in principal. In addition the MV framework is tractable and allows us to incorporate constraints, tilts, inequalities, and indeed all the features of linear and quadratic programming. Together these benefits make MV analysis popular both with teachers of financial market theory and with system implementers within the investment technology industry.

At the same time, among practitioners, and specialists in financial market theory, MV analysis and the portfolios that result from MV optimisation continue to attract a steady flow of detailed criticism. To quote R.O. Michaud (1998), "the basic problem is that MV portfolio efficiency has fundamental investment limitations as a practical tool of asset management". Four major problems that occur in practice are discussed next:-

Firstly, it is often difficult for practitioners to produce forecasts in the form required for MV optimisation. They will often prefer to forecast relative return between assets, or wish to restrict their forecasts to those assets currently impacted by "big picture" issues. Also many fund managers prefer to forecast rank rather than linear return. Turning these alternative forms back into the format required by standard mean variance analysis can be an inelegant and error prone compromise.

Secondly, our MV optimal portfolio can be highly sensitive to the exact value of the return forecasts. This problem is magnified by the fact that the forecasting process is known to usually produce results which are highly inaccurate and noisy. This leads to undesirable instability in recommended portfolio holdings .see Merton (1981) among many others. This instability is then further compounded by the fact that, the correlation coefficients used in standard mean variance analysis can themselves be worryingly unstable over time.

[1] UBS Warburg, 1 Finsbury Avenue, London EC2M 2PP.
[2] Faculty of Economics and Politics, University of Cambridge, Cambridge CB3 9DD

Thirdly, many investors feel uneasy about the use of variance as a risk measure. The most counter intuitive feature is its equal penalisation of gains and losses, see Sortino and Foresey (1996). However, it is also possible for returns on a minimum variance portfolio to be dominated by a nominally riskier portfolio which while more volatile, still always produces a higher return. Also, depending on the choice of asset set, if one of the assets has much less volatile returns than the others, then a minimum risk portfolio in a mean variance sense will be heavily concentrated in this one asset rather than diversified across a wide range of holdings.

Finally, it is usual when using MV optimisation to exclude assets with a highly non normal distribution of return (e.g. options, or some of the dynamic trading strategies executed by hedge funds) as the conventional summary statistics do not fully capture the distribution of the return on these assets and hence the portfolio statistics could conceal undesirable concentrations of risk. Similarly less liquid assets such as property or private equity show artificially low volatilities which lead to an over allocation in mean variance analysis.

Given these problems, it is not clear which aspects of MV analysis to retain and which one should jettison. It is desirable to use the linear quadratic framework of the MV world without limiting ourselves to the specific choice of mean return and variance or tracking error of return as the parameters which drive the investor decision function. This approach has already been adopted in the literature. Wang (1999) uses a MV model to solve multiple benchmark problems. Chow and Kritzman (2001) convert MV analysis into value at risk based capital allocations. We plan to take this further by putting all these cases into a general framework. In particular, we show how non-parametric statistics can be incorporated into this framework. The approach we use is based on state preference theory, see Copeland and Weston (1988), we derive an objective function which can be minimised using standard quadratic programming techniques.

We will then show how this generalisation of the mean variance framework allows a more robust portfolio construction process that is less sensitive to noise in the input parameters and we will introduce an innovative set of summary statistics that provide a very intuitive representation of downside risk.

At the heart of the investment management process is a simple question "In what proportion should I hold this range of assets given my expectations about future risk and return on each one". In principal, the process for arriving at this optimal set of portfolio weights has been firmly established since Markowitz introduced mean variance analysis, see Markowitz (1959).

As one might expect, this deceptively simple question hides a vast range of complications. In practice, in order to apply classical mean variance analysis you have to make a number of simplifying assumptions. Many of these assumptions are routinely violated in practice.

The first assumption is that risk and standard deviation of expected return are synonymous. In reality investors appetite for upside and downside risk are very different! If the forecast return on all your assets can be adequately modelled by a normal distribution then upside and downside risk are equal and this is merely a semantic quibble. However, investors are becoming increasingly concerned about the asymmetric behaviour of return on many asset types, hence alternative approaches have been developed to address this problem, see Sortino and Forsey (1996). Usually these approaches require accurate modelling of the higher moments of the return distribution hence they are extremely sensitive to limitations in the next assumption.

The next problem is that estimating the parameters used in mean variance optimisation (expected return, standard deviation and correlation of return on all assets) is a process which is plagued by data problems. In particular, correlation estimates can be quite unstable with the diversification effect calculated in normal times reducing dramatically when most needed as the markets go through a turbulent period! Suggested improvements involve separately modelling normal and abnormal market behaviour and the mixing the results in some way, see Chow et al (1999), or employing less summarised risk statistics see Embrechts (1999), or Gardner (2000). Equally, the return forecasts used will typically be subject to estimation error. This can lead to further instability in the recommended holdings in fact the process of optimisation has been referred to as one of error maximisation! . See Jorion (1992).

The final problem is that providing an excess return forecast for each asset is not a very effective format for capturing the insight of professional investors into likely market movements. Usually investors will feel more confidence in some forecasts than others. Typically, they will feel more confident about relative return forecasts than absolute return forecasts, and they would prefer to only forecast return for that limited range of assets currently being affected by "the big picture" while leaving the remainder to set to a neutral value. This has lead to use of Bayesian approaches to building up a forecast from multiple partial views; see Black and Litterman (1992). While this can be very effective, the mathematics involved can be intimidating, and the detailed implementation decisions made are critical to achieving well-behaved intuitive results.

From an investment practitioner's perspective, our initial simple question has turned into a vast specialist subject where the best approach is highly dependent on the detailed circumstances in which the optimisation is done. In the many investment organisations who can afford to develop (or commission) the correct level of specialist expertise, this is not a problem. For many others, a two-culture situation develops. Practical fund managers distrust and dislike the black box characteristics of the usual diversification approaches. Specialist quantitative analysts dislike and distrust the apparently "ad hoc" nature of many investment decisions. For a final group, the value added by formal risk management is so outweighed by the costs and complexities of implementation that they adopt one of a range of alternative heuristic approaches to diversification.

Most attempts to date to make the assumptions inherent in mean variance analysis more closely reflect day to day realities of the investment world have usually involved ever more sophisticated mathematics. Unfortunately, for many people the complexity of the mathematics is a barrier to acceptance in its own right. An alternative is to use a diversification technique which is based on mathematics which is inherently less sensitive to noise in the data, less dependant on the assumed form of the forecast return distributions, and with built in assumptions which are essentially simple to understand. Robust statistics addresses all of these issues with the potential penalty of not being able to use all available information. In addition this approach builds on and formalises the established practice among a subset of fund managers who actively use ranking approaches in their forecasting and portfolio construction processes,

In Section 2, we present a discussion on what we call generalised mean-variance analysis which attempts to put the above problems and approaches into a general framework. The procedure we advocate to replace expected returns is m-tile membership. The concept of m-tile membership means what m-tile does the stock belong to when ranked over the Universe of stocks. If $m = 10$, for example, we are asking what decile the stock belongs to.

In Section 3 we detail our "mean-variance" approach whilst in Section 4 we present details of more Monte Carlo investigations and empirical implementations. Conclusions follow in Section 5

2. Generalised Mean–Variance Analysis

In modelling decision making for an organisation, we can afford to be a little more hazy than in modelling the decision making of an individual, where the accepted wisdom is to use a variant or generalisation of expected utility theory. The reason for this laxity is the fact that we have very little clear guidelines as to how to aggregate the preferences of individual stakeholders into the decision function of the organisation. To take a simple example of a company with an employee pension plan, the interest of the average shareholder typically conflicts with the interests of the pension plan members. Conflicting interests may be resolved via the use of game-theoretic notions, but such resolutions usually depend upon a set of auxiliary assumptions describing the behaviour of the individuals playing the game.

The preliminary remarks above justify, in our view, presenting a firms decision function in terms of an $(n \times 1)$ vector of positive attributes a associated with the n investible assets in the Universe, together with a positive definite matrix $C(n \times n)$. Then, for a given set of portfolio weights $\underset{\sim}{\omega}(n \times 1)$, the firm maximises

$$\underset{\sim}{\omega}'\underset{\sim}{a} - \lambda \underset{\sim}{\omega} \underset{\sim}{C} \underset{\sim}{\omega} \qquad (1)$$

where λ represents the trade-off between the attribute of the portfolio $(\underset{\sim}{\omega}'\underset{\sim}{a})$ versus the risk of the portfolio $(\underset{\sim}{\omega}'\underset{\sim}{C}\underset{\sim}{\omega})$. For obvious reasons, we call such an analysis generalised mean-variance analysis. Together with some additional constraints, such as $\omega' e = 1 \quad or \quad 0$, where $e = (1,1,1,1,1,1)'$, or $\omega_i \geq 0$ (long-only positions) we have a conventional quadratic programming problem.

We next turn to the choice of attribute and the choice of $n \times n$ measure $\underset{\sim}{C}$. Necessary features for $\underset{\sim}{a}$ would be that more $\underset{\sim}{\omega}'\underset{\sim}{a}$ is desirable for the firm and that $\underset{\sim}{a}$ is approximately linear so that if $r_p = \frac{1}{2}r_1 + \frac{1}{2}r_2$, then $a_p \approx \frac{1}{2}a_1 + \frac{1}{2}a_2$. We say approximately linear because an attribute that is almost linear but reasonably easy to measure and/or forecast should lead to better portfolios than an attribute, such as the expected rate of return, which is exactly linear and very difficult to forecast.

One attribute of considerable interest is m-tile membership where *m* is a divisor of *n*, the number of stocks in the Universe. If *m = 10*, for example, this tells us what decile of the Universe of stocks we expect the stock to lie in. If *m = n*, then the attribute is the expected rank. At first glance, it might seem that the theory of order statistics might help us advance our analysis. Sadly, that theory is based on the assumptions that the *n* returns are a random sample, i.e. independent identically distributed random (iid) variables. Equity, returns are anything but iid. To illustrate the type of mathematical issues set *m=2* and *n = 4*, then for stock 1 we can compute the probability stock 1 is in the top two of the four stocks. Denoting R_i as the return of stock *i*, call this event A, then denoting Probabilities by *P( )*,

$$\begin{aligned} P(A_1) &= P(R_1 > R_2 \text{ and } R_1 > R_3) + P(R_1 > R_3 \text{ and } R_1 > R_4) \\ &\quad + P(R_1 > R_2 \text{ and } R_1 > R_4) \end{aligned}$$

Now analogously, we calculate $P(A_i), \quad i = 1,4,$ and retain the two stocks with the largest $P(A_i)'s$ then, in the population rather than the sample, we would have defined what we mean by top-half stocks. If we took many samples from our Universe then we could construct sample estimators of $P(A_i)$, thus we could identify and estimate top-half membership. Of course with real data the changing time varying nature of return distributions inhibits this.

In the population, decile (or *m-tile*) membership will be partly linear. For weights ω_i, $\Sigma\omega_i = 1$, if assets $(1,\ldots,k)$ belong to *m-tile j*, then, $\max(R_i) = \Sigma\omega_i \max_{i\in 1,\ldots k}(R_i) \geq R_p \geq \Sigma\omega_i \min_{it}(R_i) = \min(R_i)$ so that R_p also lies in *m-tile j*. By the same argument if $(1,\ldots,k)$ lie between *m-tiles j*, and *k* so will $R_p = \Sigma\omega_i R_i$ also lie between *m-tiles j* and *k*. It should also be clear that *m-tile* membership is not fully linear as the following example demonstrates.

Consider stocks 1 to 6 with returns 10, 9, 8, 7.5, 7.4 and 0 and $m = 3$, thus stocks 1 and 2 are in the top ter-tile, stocks 3 and 4 are in the second ter-tile and stocks 5 and 6 are in the third ter-tile. Suppose we construct a portfolio of .8 stock 2 (rank 1) plus .2 stock 6 (rank 3), the portfolio rank is $.8 \times 1 + .2 \times 3 = 1.4$, thus the portfolio is a rank 2 asset according to linearity. However, its return is $.8 \times 9 + .2 \times 0 = 7.2$ which is in the third ter-tile, since assets 5 and 6 have returns 0 and 7.4

3. The state preference theory approach to portfolio construction.

In conventional mean variance analysis, we use correlation matrices as a key intermediate variable when calculating the optimum. When we do this, we are assuming that knowing the mean, standard deviation, and correlation of the return on all the constituents of the portfolio is all that is needed to fully describe the risk characteristics of any portfolio built from these constituents. In practice, the real distribution of return may often be fat tailed, skewed, and/or discontinuous as will the multivariate probability distributions.

Unfortunately it is not obvious that the mean variance results derived for conventional correlation apply equally to rank correlation statistics. Even worse, there are alternative non parametric statistics that we might think of using. (i.e. Spearman's rank order correlation or Kendall's Tau) Hence the need to go back to first principals in order to prove our method.

In the state preference model, uncertainty takes the form of not knowing what the state of nature will be at some future date. To the investor, each security is a set of possible payoffs each one associated with a mutually exclusive state. Once the uncertain state of the world is revealed, the payoff on each security is determined exactly. This is a very flexible way of modelling complex valuation and decision taking processes. This very flexibility is also its main problem. There are usually an infinite number of states (high/low) hence the usual assumption of normal distribution of return in order to make the mathematics tractable.

When considering *m-tiles*, we can consider all cases from $m = 2$ to $m = n$ (rank). We shall investigate the rank case next. If we assume that the states of nature are adequately represented by the rank order of the portfolio then for an "n" asset problem the number of states of nature has reduced from infinite to "n" factorial.! This total falls even further as we reduce the resolution of the calculation from rank to m-tile. If we then further assume that the actual probabilities observed are a sample drawn from a distribution which is continuous and with limited magnitude of first derivative then we do not need to exhaustively evaluate all "n" factorial states in order to optimise our weights. However, by sampling repeatedly, the probability that stock *i* belongs to m-tile *j* can easily be determined with any degree of accuracy. In effect this generates a time series for which we can find a formula for the "mean" and "variance".

For values of *m* that are reasonably large, it is clear that conventional historic data cannot be simply applied to compute rank probabilities as there are not enough degrees of freedom. However, as mentioned above, given assumptions about the data-generating process (DGP), for example, that returns $\underset{\sim}{r}$ are multi-variate normal with mean vector $\underset{\sim}{\mu}$ and covariance matrix $\underset{\sim}{\Sigma}$,

$$\underset{\sim}{r} \sim N(\underset{\sim}{\mu}, \underset{\sim}{\Sigma}) \qquad (2)$$

we can employ Monte-Carlo methods to compute rank probabilities by simulation. The benefit of the above approach is that we can replace (2) by more complex assumptions, i.e. we can model non-normality, extreme returns, conditional volatility etc. without complicating matters unduly; all we need to be able to do is to simulate the DGP.

Once we have evaluated the p_k, we can set up our quadratic optimisation. x_k is the return (i.e. rank) of the portfolio on each of "m" possible states. p_k is the probability) of each of the "m" possible states. ω_i is the weight of each of the "*n*" assets. r_{ik} is the rank of each asset in each state

$$x_k = \sum_{i=1..n} \omega_i r_{ik} \qquad \text{(m-tile value)}$$

$$\bar{x} = \sum_{k=1..m} \left(\sum_{i=1..n} (\omega_i r_{ik} p_k / m) \right) \qquad \text{(average m-tile value)}$$

$$\bar{x}^2 = \left(\sum_{k=1..m} \left(\sum_{i=1..n} (\omega_i r_{ik} p_k / m) \right) \right)^2 \qquad \text{(squared m-tile value)}$$

$$x_k^{\;2} = \sum_{i=1..n} \left(\sum_{j=1..n} (\omega_i r_{ik} \omega_j r_{jk}) \right) \qquad \text{(squared root value)}$$

$$v = \left(\sum_{k=1..m} (x_k^{\;2} p_k) - m\bar{x}^2 \right) / (m-1) \qquad \text{(rank variance)}$$

$$\theta = \bar{x} - \lambda v \qquad \text{(objective function)}$$

$$\theta = \bar{x} - \lambda\big(\sum_{k=1..m}(x_k^{\,2} p_k) - m\bar{x}^2\big)/(m-1) \qquad \text{(substituted for variance)}$$

$$\begin{aligned}\theta = &\sum_{k=1..m}\Big(\sum_{i=1..n}(\omega_i r_{ik} p_k / m)\Big) - \lambda\Big(\sum_{k=1..m}(x_k^{\,2} p_k)\Big)/(m-1) \\ &+ m\lambda\Big(\sum_{k=1..m}\Big(\sum_{i=1..n}(\omega_i r_{ik} p_k / m)\Big)\Big)^2/(m-1)\end{aligned}$$

$$\begin{aligned}\theta = &\sum_{k=1..m}\Big(\sum_{i=1..n}(\omega_i r_{ik} p_k / m)\Big) - \lambda\Big(\sum_{k=1..m}\Big(\sum_{i=1..n}\Big(\sum_{j=1..n}(\omega_i r_{ik}\omega_j r_{jk})\Big)p_k\Big)\Big)/(m-1) \\ &+ m\lambda\Big(\sum_{k=1..m}\Big(\sum_{i=1..n}(\omega_i r_{ik} p_k / m)\Big)\Big)^2/(m-1)\end{aligned}$$

As this is a standard quadratic objective function, it can be minimised using all the conventional quadratic programming algorithms.

In the above we consider *m-tile* rank as our "return" and rank variance as our "risk". Actually, hybrid procedures could be used. For example, if we thought that tracking error/conventional variance was a sensible risk measure, then we could use this in conjunction with a ranked return. In mixing different characteristics, care needs to be taken in the determination of λ. This is not, however, an insurmountable problem as we can take the λ of the market portfolio based on the mix assumed much as is done in conventional mean-variance analysis where you choose the λ which makes the FTALL share optimal. Then by varying λ we can assume that we are more or less risk-tolerant than the market representative agent.

4. Implementation and Simulation

In order to test the properties of the ranking approach to mean variance analysis, a number of simulations were run to demonstrate different aspects of the approach. The first test was to look at the "perfect substitute" problem. This occurs where some assets in a portfolio have such similar forecast risk characteristics that the optimiser sees little extra risk in moving a holding from fully in one of the pair into fully in the other as the relative forecast return for the pair of assets changes sign. Unconstrained, this results in large long short bets for such asset pairs, and sharp changes in recommendation over small change in forecast return.

To demonstrate this, we chose a three asset portfolio containing French equity, and French and German bonds. Historically the latter two assets have been closely correlated. The first plot below shows the changing holdings recommended by conventional mean variance analysis (at a constant tracking error) as the forecast return varies over half a percentage point. The second chart below shows the same test done with rank optimisation. (Tracking error in this case is only approximately constant hence the slight "wobble" in the lines)

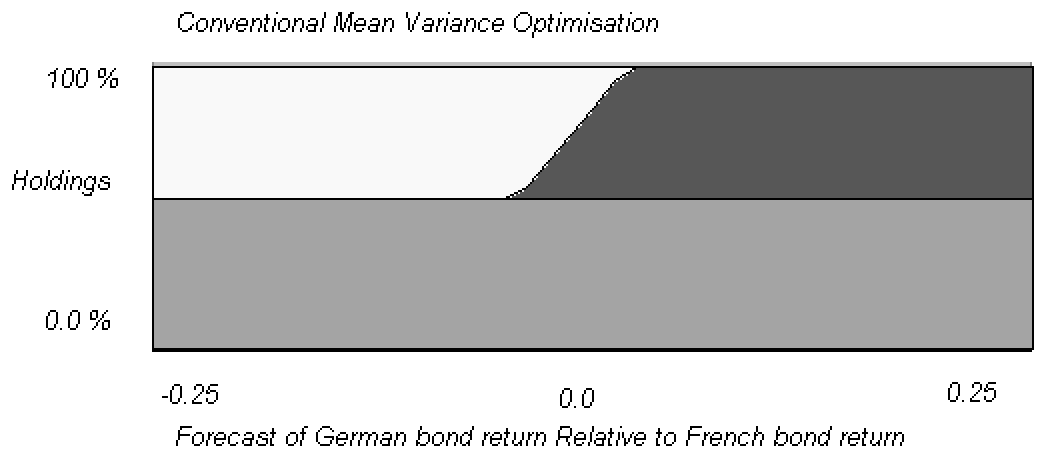

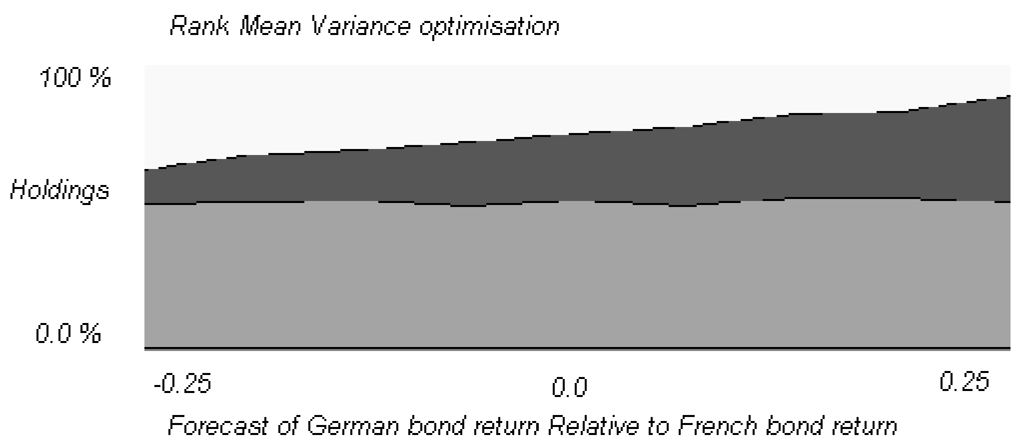

The second issue with mean variance analysis is the problem of defining risk. Taking a typical global asset allocation portfolio of European, Asian and American equities and bonds looked at from a US dollar perspective, the efficient frontier calculated using five years of return history from 1986 to 2001 shows recommended holdings which vary along the efficient frontier as shown in the first chart below. The minimum risk portfolio is in reality not very well diversified, consisting as it does largely of US bonds. This represents a very concentrated exposure to forecasting errors.

The same exercise undertaken using rank optimisation (the second chart below) results in a minimum risk portfolio with equally weighted holdings as by definition placing the same bet on every asset must result in zero volatility of average rank return. This is intuitively a much more diversified position than the standard mean variance analysis above.

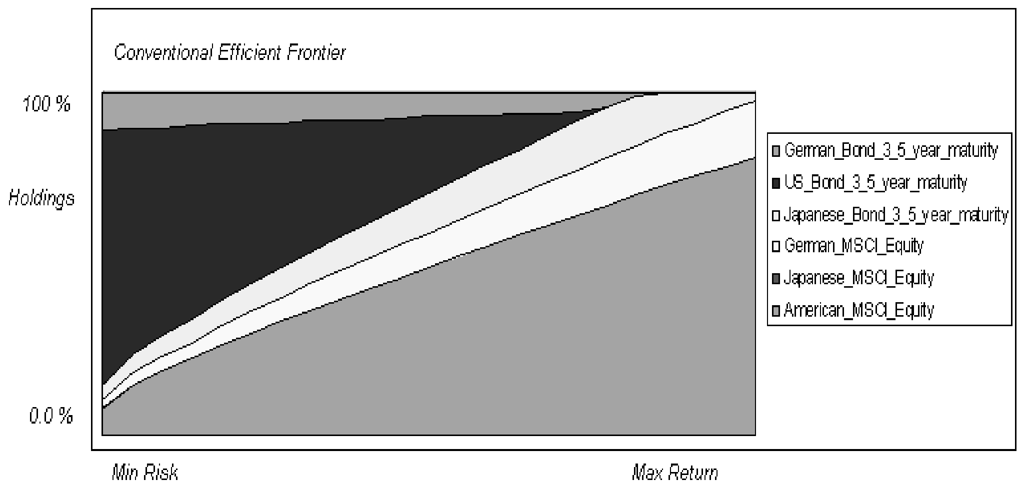

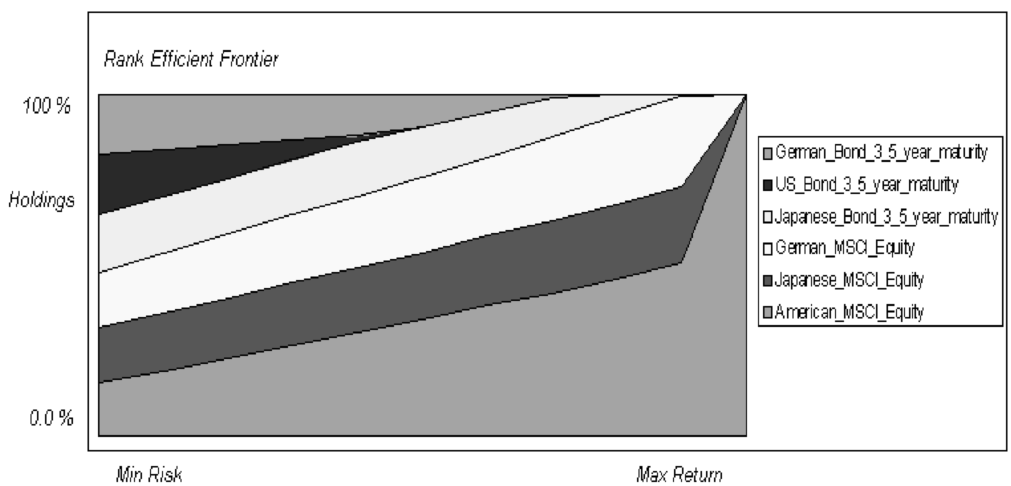

The third issue is volatility of recommended holdings caused by noise in the correlation matrix over a period of time. The chart below shows the same set of assets with constant forecast returns but now at a fixed tracking error position recalculated monthly over the period 1998 to 2001. This was calculated using three years of monthly data, exponentially weighted with a weighting half life of one year) This can be seen to routinely produce substantial short-term movements in recommended holdings that are highly undesirable from a practical investment perspective.

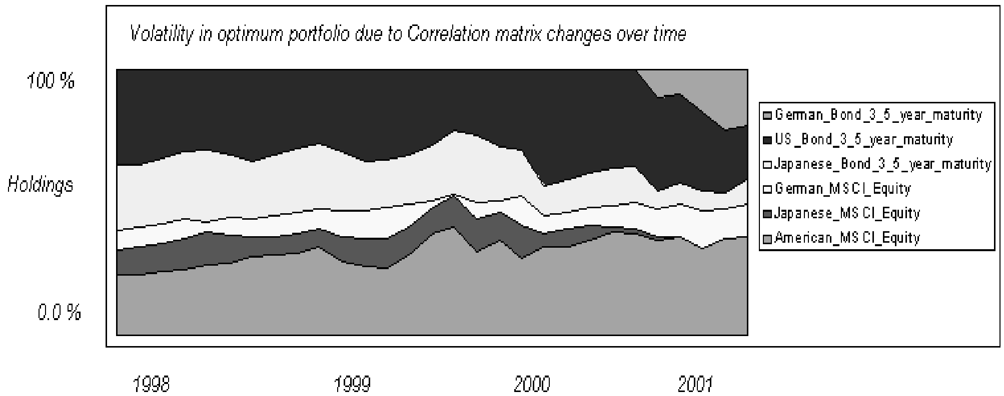

Volatility of coefficients in the calculated correlation matrix is an inevitable consequence of the stochastic nature of the time series being compared. In an idealised world, where there is an underlying "true" value of correlation, the distribution of the estimated values around that "true" value can itself be estimated using a knowledge of the method used to calculate the correlation coefficients. The width of this distribution is then usually taken as a measure of the significance of the estimated coefficient.

In linear correlation, occasional large values of return (fat tails,) and uncertainty about the form of the joint probability distribution create uncertainty as to the formula to use to calculate the significance of the correlation estimate. In rank correlation, there is no uncertainty over the distribution of returns from which any given sample is drawn (it is always drawn from the set of integers 1 to n). Hence the correct formula to calculate the significance of the correlation is known. As the joint probability distribution of the underlying variables approaches a bivariate normal distribution, the significance of the linear estimates becomes as good as those produced from the rank correlation. In fact the significance formula becomes the same. Viz:- $T = r * \text{sqrt}((N-2) / (1- r^2))$. We shall delay discussion of the properties of significance tests for ranked data to a later paper.

Another way of saying the same thing is that for any given level of underlying real correlation, and sample size, as the underlying distribution deviates more and more from bivariate normal, the volatility of the estimated linear correlation coefficient increases relative to the volatility of the rank correlation. Given that these deviations from normality are unknown and potentially large and time varying, the implicit use of rank correlation rather than linear correlation must reduce this source of volatility in recommended holdings.

Given a hypothetical samples size of a hundred data points and historic distribution of return for the above asset set, this could reduce the volatility of correlation coefficients very significantly.

The final problem with traditional mean variance analysis is forecasting return in a way that is intuitive to a typical fund managers. Capturing this insight into likely market movements is an art not a science. In due course, a range of alternative methods are likely to be developed to capture this information. However all of these methods need to convert the managers input, in whatever form it is most conveniently entered, into a standard set of summary statistics used for analysis. In mean variance analysis these summary statistics are the means, and covariance of return. The equivalent summary statistics in our approach is the rank (or m-tile) probability matrix illustrated below.

This m-tile probability matrix is a very simple way of describing intuitively the likely rank of each asset. There is one column per asset and in each column the probability of that asset being in each quantile is indicated in the relevant row. In the rank case the matrix is a matrix of probabilities, hence all the rows and all the columns must sum to 100% Such a matrix is called bi-stochastic. We give an example below.

Rank Probability Matrix

USB	CVB	EAFE	R1000	R2000
16.47	23.96	5.94	17.19	36.44
16.65	23.18	11.44	29.99	18.74
16.23	23.41	26.66	20.14	13.56
19.74	19.64	27.24	19.88	13.50
30.91	9.81	28.72	12.80	17.76

The above matrix, which details an $m = n = s$ case, tells us for example, the probabilities that R2000 is ranked first is 36.44 per cent. The row sum adding to one simply means that the sum of probabilities that different stocks could be first should add to 1. Likewise the 2^{nd} row is the sum of probabilities that different stocks could be second, which again adds to 1. The column numbers, say the second column, CVB, tells us the rank probabilities of CVB, which again add to 1. One can deduce immediately from such a column the expected rank of CVB.

An alternative graphical representation is shown below where the height of each band of grey is proportional to the probability of that asset being in that quantile.

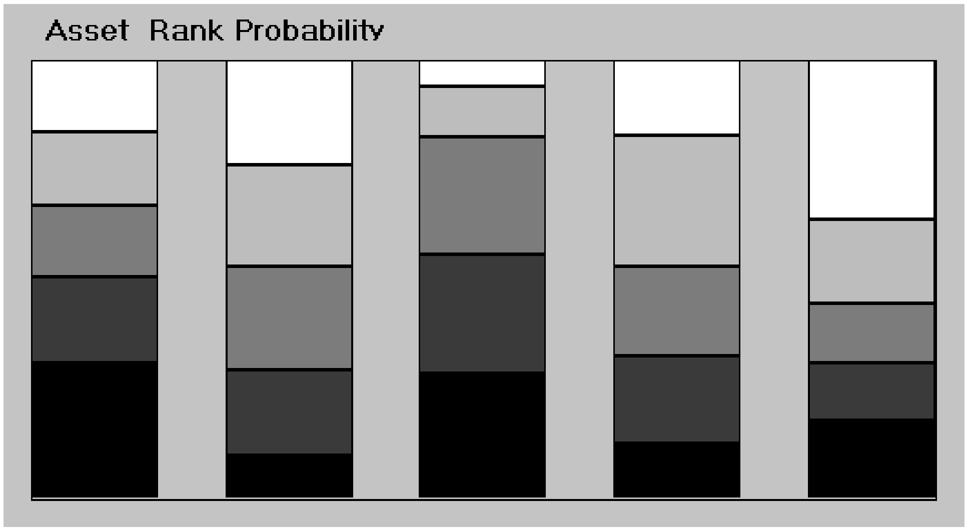

As with covariance matrices, these m-tile probability matrices provide summary statistics that give a valuable insight into likely investment out turn. They implicitly contain joint probability information. They are straightforward to calculate from the state probability values, and are easily interpreted.

Any optimisation process, inherently optimises an average expected value. Given that return is frequently highly non normal (even if we choose to ignore this fact in the interests of tractable computation), this average could hide concentrations of downside risk which are highly undesirable. The m-tile probability matrix allows this downside probability to be observed in a very direct way.

When combined with the fact that the significance of rank correlation is much more robust to non normality in the return distributions, the presence of this safeguard makes it practical to include assets in a portfolio being diversified which have a very non normal distribution of forecast return.. This considerably extends the range of asset types and forecasting methods which can be reliably employed in this type of diversification exercise.

In particular, scenario forecasts are more naturally handled in this environment as unlike standard mean variance analysis which requires that you forecast a single vector of mean returns. This approach allows you to postulate a range of alternative scenarios and allocate probabilities to each. Monte carlo simulation enables you to combine these scenarios to generate a forecast of the probability of different assets outperforming each other.

5. Conclusion and suggested further work

It can be seen from the preceding section that the recommended holdings produced using GMV analysis are more stable than those calculated using standard mean variance analysis.

Choosing an appropriate definition of risk is ultimately a matter of your particular circumstances and preferences. Defining, the minimum risk portfolio such that it is the equally weighted set of holdings has the attraction that it is more consistent with our intuitive ideas of full diversification than are portfolios produced using the minimum volatility criteria.

In addition, this approach is much less sensitive to problems caused by non normal forecast return distributions than classical mean variance analysis. The ability to display the rank (or m-tile) probability matrix provides an easily interpreted safeguard against hidden concentrations of risk. The fact that the significance level of the rank correlation is independent of the return distribution, avoids the problem with classical mean variance which can lead to unreliable recommended portfolios if the significance of the calculated linear correlation coefficient proves to be particularly low.

As we have retained the linear quadratic framework of the MV world, we can incorporate constraints, tilts, inequalities, and indeed all the other operationally convenient features of linear and quadratic programming. However we have also retained the simplifying assumption that the return distribution is adequately represent by its (rank) mean and standard deviation. Hence if users have a marked preferences for the shape of the probability distribution of forecast portfolio rank return as well as for its mean and standard deviation, they should compare this GMV approach to those based on optimising downside deviation, value at risk, or stochastic dominance. Simple extensions of our approach can allow optimisation based on the m-tile equivalent of these criteria.

The main problem using this approach is the other side of the robustness coin. The standard mean variance model has beatified excess return and tracking error as measures of performance and risk. It will take considerable effort to establish alternative non parametric risk and return measures in general use. However, in practice this may not be the major problem that it first appears. While the holdings are optimised on the basis of the m-tile (or rank) statistics, the results can still be reported in terms of tracking error, excess return etc. hence achieving the best of both worlds.

References

Black, F. and Litterman, R. (1990) Global Portfolio Optimisation, *Financial Analysts Journal Robust Diversification* (September/October).

Chow, G., Jacquier, E., Kritzman, M., and Lowry, K. (1999). Optimal Portfolios in Good Times and Bad, *Financial Analysts Journal* (May/June).

Chow, G. and Kritzman, M. (2001). Risk Budgets, *Journal of Portfolio Management*, Winter pp. 56-61.

Copeland and Weston (1988). Financial Theory and Corporate Policy. *State Preference Theory* (3rd Edition, Ch. 5). Addison Wesley Publishing (ISBN 0-201-10

Embrechts P., McNeil A.J. and Straumann D. (1999) *Correlation: Pitfalls and Alternatives*, RISK 12(5), 69-71.

Gardener, D., Bowie, D. Brooks, M. and Cumberworth, M. (2000) *Predicted Tracking Errors Fact or Fantasy?* Faculty and Institute of Actuaries Investment Conference Hatfield Heath 25-27 June (2000).

Jorion, P. (1992) Portfolio Optimisation in Practice, *Financial Analysts Journal* (January/February).

Markowitz, H.M. (1959) *Portfolio Selection: Efficient Diversification of Investments*, John Wiley, New York.

Sortino, F.A. and Foresey, H.A. (1996) On the Use and Misuse of Downside Risk, *Journal of Portfolio Management.*