An overview of ordered mean difference portfolio technology

Roger J. Bowden*

The ordered mean difference construction provides a metric for investor surplus in the form of a simple spreadsheet construction relating the returns on a designated security or portfolio to those on a given benchmark. The metric itself is well grounded in utility theory, and represents alternative attitudes towards risk in the form of risk profiling against a range of possible investor preferences. This paper reviews the basis of the construction in non technical terms and shows how it can be applied to a range of problems in portfolio practice and related investment or risk management situations. These cover performance measurement, mispricing and anomalies, CAPM testing, risk profiling, portfolio enhancement, hedging, portfolio choice algorithms, and stochastic dominance. Discussion is illustrated with examples drawn from recent research papers in the area.

Key words: Enhancement, equivalent margin, hedging, ordered mean difference, portfolio performance, risk profiling, utility generators, investor surplus, return anomalies.

* Victoria University of Wellington and Kiwicap Research Ltd. Email is roger.Bowden@vuw.ac.nz or roger.Bowden@kiwicap.co.nz Phone is * 64 4 472 4984 or 4986, fax 4983. Thanks go to Dawn Lorimer and Joe Cheung for comments and data.

I Introduction

Portfolio practitioners are familiar with mean variance analysis and its uses. As an aid to portfolio decision making, mean variance has some highly useful features, and its trade offs between expected reward and risk, such as the efficient frontier, are relatively simple to understand. Unfortunately, the limitations of mean variance analysis are almost equally well known. It suffers from a number of conceptual or empirical limitations which collectively make it a poor foundation for investor choice, most notably those relating to one or other of the quadratic utility function or the normality of return distributions[1].

Recently an alternative body of techniques have been developed which are not subject to these limitations. Collectively, these may be called ordered mean difference (OMD) technology. Some of the building blocks used in OMD analysis are recent, such as the equivalent margin or investor surplus metric, though these have references to a number of similar measures in welfare economics or project evaluation. Likewise, constructs such as the zero capital approach to enhancement are natural outgrowths. Other elements, such as the utility generators, exploit representations that have been known for many years, but whose consequences or applications have not been fully realized in isolation.

The OMD methods are non parametric in nature, although semi-parametric assumptions of some sort are necessary if one wishes to add confidence limits. However very conservative bands can be obtained using simple linear regression. The computations are quite simple, indeed can be executed off a simple Excel or Lotus spreadsheet.

The range of OMD applications has been growing and a number of recent research papers have appeared that document the manner and extent of these applications. They span such areas as performance measurement, mispricing and anomalies, CAPM testing, risk profiling, portfolio enhancement, hedging, portfolio choice algorithms, and stochastic dominance. The purpose of the present paper is to draw all these together with a non technical exposition of the methods and the concepts underpinning them.

[1] The normality can be slightly relaxed to elliptical distributions but even this does not address the fat tails property that one observes in practice. For a useful review of the historical debate over mean variance, see Bigelow (1993).

Discussion throughout is informal, and the reader is referred to the original papers for more complete references and additional technical details.

II The building blocks

OMD measures are technically concerned with what happens at the margin of the portfolio, so that one speaks of efficiency conditions for the investment of the last dollar. In practice one can say quite a lot more than that by calling on one or more of the building blocks, but the margin is a good place to start in expositing the principles.

2.1 **The *equivalent margin*; or *investor surplus*; or *optimised deprival value***

These are alternative names for the same concept. They all have semantic references[2] to other areas of economics or accounting, which is itself a source of conceptual comfort. Suppose I have a security (of return r) that I am considering adding to my portfolio (of return R). Often we shall refer to these as simply 'security r', meaning the security whose return is r (similarly for R). How could I measure the added welfare? In particular, could I give this a return dimension, rather than say, 'utiles', or any other proposed dimension of welfare measurement?

This could be done by means of the following conceptual experiment. Suppose that instead of return r, the return was $r - \tau$, where τ is a notional tax of some kind. Instead of the original return of 10%, you might get only 7%, so $\tau = 0.03$, or 30% on the original dollar amount, though we stick to the return dimension. Now ask yourself what tax you would be prepared to pay to still consider adding the security to your portfolio. If the answer is precisely zero, the security is priced just right so far as your existing portfolio is concerned, meaning that it is in portfolio equilibrium. If the answer is positive, $\tau > 0$, this means that you should add some more of the security (or add some if you don't already have some). If $\tau < 0$, this is a signal that you should go short in the security. The implied tax rate is given the above names.

[2] The equivalent margin comes from welfare economics; it measures the amount of income you would have to be compensated to give up a designated benefit. Consumer surplus is technically the area under a demand curve, but has close relationship with the compensating idea; it is a measure of the intrinsic 'economic rent' or 'monopolistic rent' that a consumer has to pay for the more highly valued portion of his consumption. Optimised deprival value comes for management accounting or capital budgeting, and is concerned with the valuation of capital assets at current replacement value rather than historic cost.

The formula[3] used to obtain the equivalent margin indicates that it represents a weighted sum of the return differences, where the weights are proportional to the marginal utility evaluated at the benchmark return. This means that the weights are higher if the state of the world turns out to be one where times are bad for the benchmark. Thus if r –R is expected to be high whenever the benchmark return is low or negative, then security r will be the more highly valued.

Suppose that the equivalent margin for security r against benchmark R comes out to be 5% for a given investor. The interpretation would be that a new dollar invested in security r is worth 5 % extra over and above the same dollar invested in an additional holding of the benchmark portfolio R. Equivalently, it would be worth 5% for the investor to take a dollar out of the last unit of security R and reallocate it to security r. This does not mean that the expected return of r is 5% above that of R; the 5% is what it is personally worth to our investor. Notice that the equivalent margin is a relative measure, in this case of the marginal worth of one security relative to another. A dollar of additional capital is to be tied up in one or other of the securities or portfolios, and the issue is which is the best way to go. Later we shall also discuss a variant in which investment requires zero capital (the enhancement version), and in this case the interpretation is a bit different.

2.2 The utility generators

One problem with the equivalent margin, or investor surplus, is that the outcome will depend upon the particular investor utility function for money. Some investors are more risk averse than others. It would be very useful if we could say that there was a canonical class of utility functions, in terms of which any given utility function could be represented. By saying what the tax rate would be for one of these *generator* utility functions we might be able to make statements about what it would be for any arbitrary utility function.

[3] The original formula was developed in Bowden, R.J. (1992); its implications are summarized in Bowden, R.J. (2000), which developed the OMD empirical construction:

The required utility generators have a special simple form[4], portrayed in figure 1 below (solid line). This has been plotted against alternative values of some return R, which is equivalent to plotting it against alternative sums of money represented by $(1+R)W_0$ where W_0 is invested capital. There are just two linear segments, one rising at rate unity, and the other at rate zero, i.e. horizontal. The breakpoint is at the return number P. One might think in terms of a target return: investors like more and more money up to some point (P in return terms), after which they just don’t care. The utility generator is a bone fide risk averse utility function (that it evidently has negative values will not matter). If you increase P to a new node P', you will get a new utility function, denoted in figure 1 with dotted lines. This one is less risk averse than the one centered at node P. Falling P denotes higher risk aversion. Positive P is the normal state of affairs.

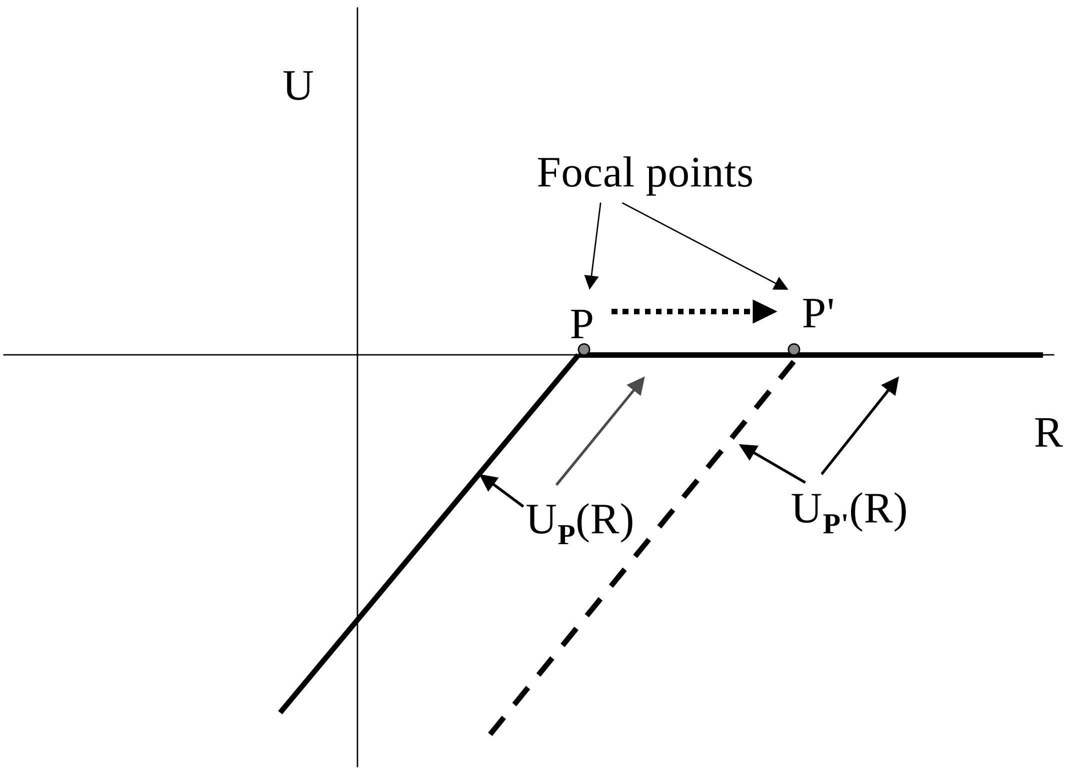

Figure 1: The utility generators

[4] These special functions appear as early as Markowitz (1959),also Fishburn (1997), while implications for stochastic dominance were explored in Russell and Seo (1988) and Bowden (2000).

The spanning property[5] means that every risk averse investor can be thought of as a collection of 'gnomes', each of whom has a utility function of the above form, with a particular value of the node P. The risk preferences of the investor determine the relative weighting given to alternative gnomes in his psychological make up. Thus if I am more risk averse than you, it is because the risk averse gnomes, those with low P values, are more heavily weighted in my risk profile.

2.3 The OMD schedule

Let R be the return on any suitable benchmark, against which we wish to compare or evaluate security of return r in some way. Suppose we were able to work out the equivalent margin τ for the gnomic utility function centered at P. We could call this τ_P or $\tau(P)$, to indicate that it depends upon the particular generator centered at P. If you now varied the nodal point P, and did the same, you would get an entire schedule, or a function of τ against P. This is called the theoretical OMD schedule.

The theoretical OMD schedule is given by the formula:

$$\tau(P) = E[r - R \mid R \leq P] \;,$$

which is the truncated or running mean of the difference in return across values of the benchmark less than or equal to the given number P.

Figure 2 illustrates a theoretical OMD schedule and how to interpret it. It can slope any way, and take any shape – the one shown has a positive slope and cuts the horizontal axis just once. The horizontal axis now has a dual interpretation:

(a) As values of the benchmark return R;

(b) As values of a risk index P. It can be shown that in optimal portfolio choices involving r and R, all risk averse investors can be indexed by a single value of P. The case P = ∞ corresponds to a risk neutral investor and P = 0 to an extremely risk averse investor. Thus the sense of increasing risk aversion is opposed to that of the actual benchmark return.

---

[5] The spanning property, in the sense used here, is implicit in the original formulation of stochastic dominance theory in Rothschild and Stiglitz (1970). The explicit representation is contained in Bowden, R.J. (2003b), though many of the operational consequences were proved in Bowden (2000).

As illustrated, investor A is rather risk averse. This investor will have a negative investor surplus from holding an extra unit of r and will be better off shorting r against the benchmark R, or downloading, if security r already features in portfolio R. Investor B is much less risk averse and will gain from adding security r at the expense of a unit of the benchmark R. The surplus has a return dimension and can be measured on the vertical axis.

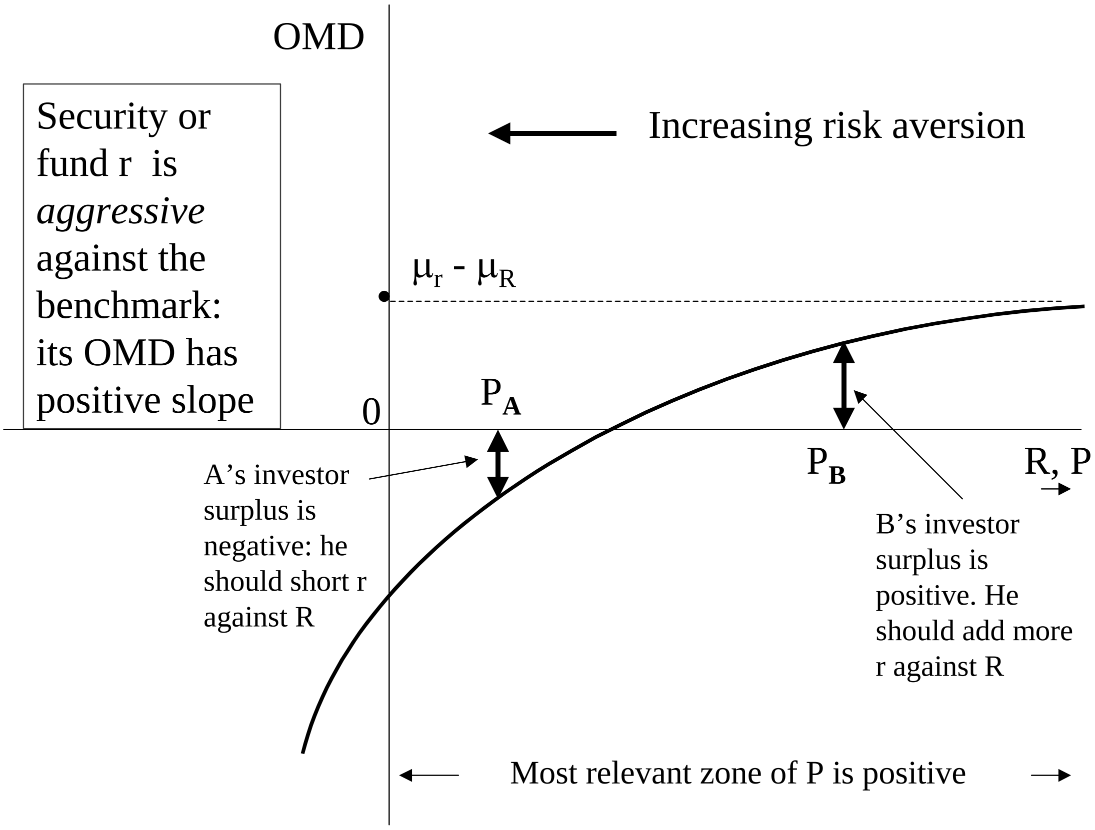

Figure 2: Interpreting the theoretical OMD schedule

In practice, one usually works off the *sample OMD schedule*, in which we start with observations on r and another set of observations on R. The operation itself takes a few moments with an Excel spreadsheet. Table 1 shows the sequence. The original values are tabulated in the first two columns. First reorder the observations via ascending values of the benchmark R. To do this, you can use *data/sort* from the main menu. Next, take the difference r-R . Finally take the moving average of the differences, which is done in column 6. Thus the number 0.014 is the average of the first two return differences; the number 0.012 if the average of the first three, and so forth. (You can accomplish this in Excel by use of the *average( )* function, anchoring the first number of the range with an absolute address).

			Sample OMD calculation				
r values	R values	reordered by R		r-R	running mean difference		
		r	R		OMD		
0.032	0.011	-0.021	-0.045	0.024	0.024		
0.048	0.047	-0.026	-0.030	0.004	0.014		
0.028	0.018	0.004	-0.005	0.009	0.012		
0.004	-0.005	0.017	0.001	0.016	0.013		
0.027	0.008	0.027	0.008	0.019	0.014		
-0.026	-0.030	0.032	0.011	0.021	0.015		
-0.021	-0.045	0.028	0.018	0.010	0.015		
0.017	0.001	0.048	0.047	0.001	0.013		

Table 1: The OMD schedule - computation

The final step is to graph column 6 against the benchmark return R. In doing, so you have graphed the estimated equivalent margin schedule τ (P), with alternative values of P. The latter are taken to be just the same as the alternative values of R, so we might equally well write τ(R) instead of τ(P). However, it is better to preserve the distinction between P as a risk index and R as a value of the benchmark return. The discussion below contains several such graphs. The schedule can slope upwards, downwards or in any way, and it might or might not cross the horizontal (R) axis. How the schedule behaves determines what one can infer from it.

The three items just described constitute a basic OMD toolbox. There are additional elements that one can draw on according to context, and these will be described in due course. Likewise, one can add further detail to the OMD plots, such as statistical confidence bands. These are always wider for very low values of the benchmark, the reason being that there are only a few values of R available to the computation at this point; hence the sampling variability is greater. The diagrams below contain some examples. Adding these bands can be done in Excel by using the regression wizard, so that all operations can be done in the same spreadsheet.

III Some applications

The first applications of the OMD arose more or less directly out of the framework just described. The starting point was the result that any utility function could be expressed in terms of some suitable weighted combination of the utility generators. This meant that the equivalent margins were also expressible as a weighted combination of those for the utility generators, taken over different values of P, with all the weights positive. Suppose that the OMD plot lay everywhere above the horizontal axis. This would mean that no matter what the utility function actually was, provided only that the investor was either risk averse or risk neutral, the investor's surplus must be positive. One would not actually need to know his or her utility function to be able to make such a judgment.

3.1 ***Performance***

Take a benchmark return defined in some way, and consider the return on a designated fund. How would we know whether the fund manager is adding value relative to the given benchmark? One way of deciding this is to ask whether if the fund were traded in the market along with the benchmark, you would want to hold some units of it in addition to the benchmark portfolio. If not, then we could say that the fund adds nothing, so that the manager possesses no special investment insight. The neatest way of measuring this is via the equivalent margin: is it zero, positive or negative?

As it stands, the equivalent margin is not really a measure of performance as such, but rather of 'spanning'. It is a marginal concept rather than the average or total metric we would like if we are to consider investing in just one or the other of r, R. But it certainly does capture the manager's special ability, in a way that Jensen's alpha cannot. It is in fact quite easy to get situations where Jensen's alpha gives perverse negative results where the manager shows excellent market timing ability; this is the substance of the well known Dybvig-Ross critique[6]. The equivalent margin is much harder to 'fool' in this way.

As earlier remarked, the equivalent margin is specific to the particular investor utility function. However, we can do an OMD plot, as above. If the fund's return has an OMD

---

[6] A good market timer will show returns that go up by more than the market when the market rise and go down by less when the market falls. This creates a nonlinear regression of the fund on the market. Jensen's alpha is a linear construct, creating conflicts. For the original discussion, see Dybvig and Ross (1985).

that is always above the horizontal axis, then we could say that the fund would add value for any risk averse investor. In addition, we can do a risk profiling, according to which we can decide whether the fund adds value for investors in different risk categories. Suppose, for instance, that the fund advertises itself as conservative in stance. But if you find that the fund adds value only for very risk neutral investors, then you could at least say that the fund manager is not living up to his or her claims. The idea of risk profiling is explored further below.

3.2 Mispricing, anomalies, CAPM

The OMD plots can have all sorts of shapes or slopes. However, if the slope is uniformly positive, then we would say that security r is *aggressive* against the benchmark; and *defensive* if the slope is negative. This terminology is reminiscent of the CAPM model, and indeed if a CAPM holds, then a security with beta >1 will lead to a positively sloping OMD plot against the market as benchmark, negative sloping if beta < 1. Moreover if a CAPM does hold, all the OMD schedules for the different securities, expressed in the form of excess returns relative to the risk free rate, should cross the horizontal axis at exactly the same point, so you would get a geometric pencil sort of effect, with rays of different slopes all emanating from the same point[7].

Figure 3 below is an OMD plot of a selection of 14 U.S. stocks against the S&P500, using monthly total return data measured as excess returns relative to the 30 day CD rate. Closer inspection reveals that most of the plots do more or less cross in the same region, though the presumed common point is in fact not quite common. Note, however, some exceptions, highlighted in bold: IBM is always above the horizontal axis and MMM is virtually so, just touching at the far right. The implication would be that IBM and MMM have apparently been underpriced by the market, at least over the period under consideration. That does not necessarily mean that these stocks were, or are now, a good buy. But it might warrant further exploration as to why, and whether this state of affairs might persist.

[7] The theory of OMD with CAPM and related testing is developed in Bowden, R.J. (2002).

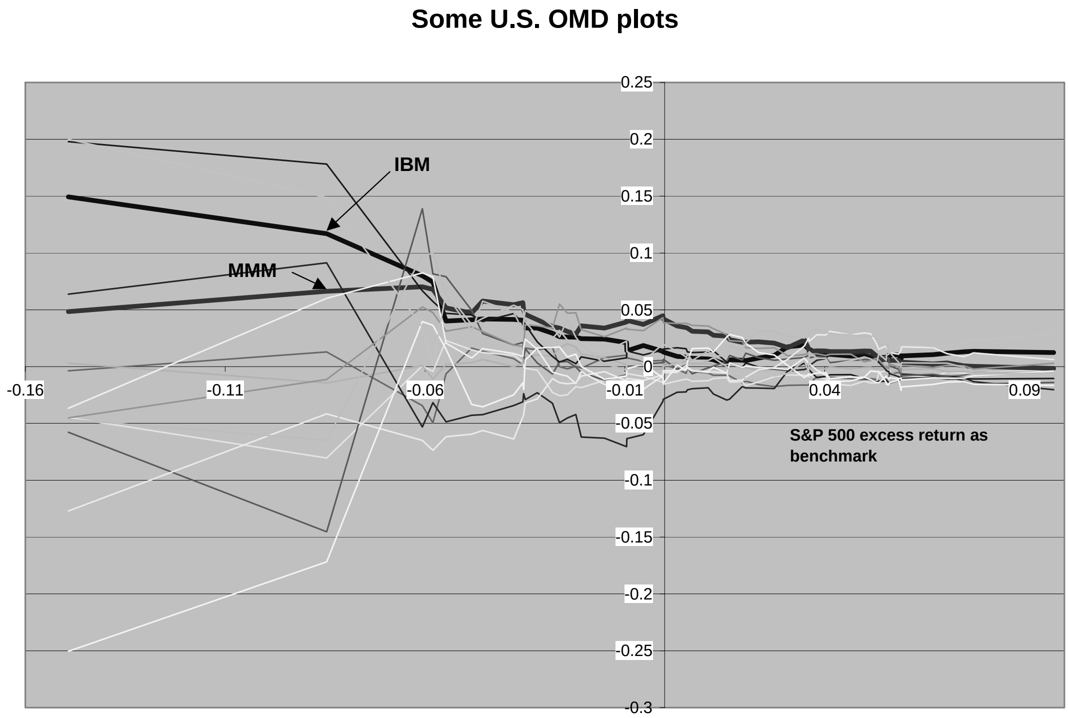

Figure 3 : US monthly OMD schedules, March 1996- Jan. 2001

(Data from Lorimer and Rayhorn (2003))

Risk profiling

As earlier noted, the ordinate of the OMD at each point R represents the investor surplus to a gnome whose P-node is focused at that point (recall that the gnomes are the elementary investors whose utility functions coincide with the utility generators). If the OMD schedule at any point R is below the horizontal axis, this indicates that the gnome whose node P=R will have negative investor surplus, so such a gnome would want to get rid of r relative to R or else go short. This suggests that we could use the entire profile of the OMD schedule to characterize how investors of different risk profiles would react to the security against the given benchmark. Points along the horizontal axis have a dual meaning: one as the return R value for the benchmark portfolio; and another as the P value for the representative gnomes, the sense running in the opposite direction.

To give this additional meaning, we would like to calibrate the risk aversion. One way to do this might be to use the market itself for the purpose; to obtain the representative market gnome, so to speak. If a CAPM holds, this is quite easy. We can identify the market gnome as the investor whose generator node P coincides with the common crossing point of §3.2 above. An easy way to decide what this might be is to plot the running mean of the excess market returns and see where this crosses the horizontal axis. The required running mean is constructed after a similar fashion to figure 2 except that we use just R instead of the difference r – R.

Figure 4 below is the graph corresponding to the data of figure 3. The crossing point is P≈ 5.60%. So any investor who considered himself as more risk averse than the representative market investor would act as though his gnomic node P < 0.056, while less risk averse investors would have P>0.056. This means that we could tell how an investor would react to any given security in terms of the sign of its OMD in different regions.

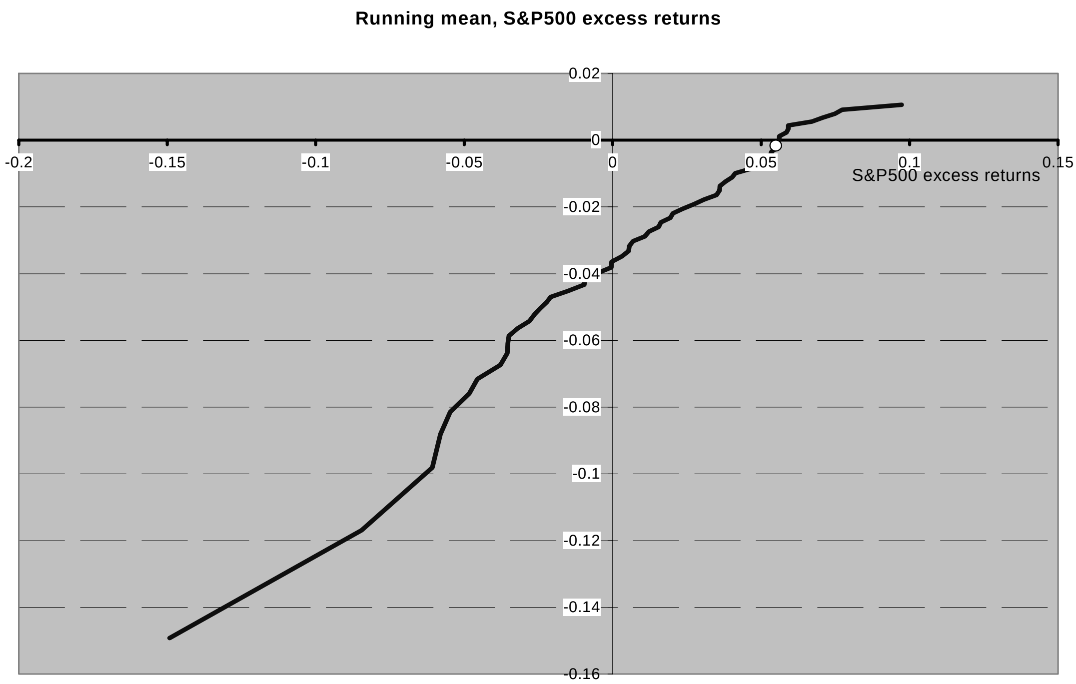

Figure 4: Running mean of the market on itself

We can use risk profiling to explore some issues about attitudes to stock classes in a strategic portfolio approach, using alternative possible holding periods. Figure 5 below does this for small cap versus large cap stocks over a long period using Ibbotson[8]

[8] Data is taken from the Ibbotson *Stocks, Bonds and Bills Yearbook* for 2002.

monthly data from 1922 to 2001. So one is asking whether the smallcap index adds investor surplus to the largecap index[9].

The carat symbol just underneath the horizontal axis marks the point where the S&P500 excess return crosses the horizontal axis, so we could take this as the representative market gnome. The OMD plots indicate that the market gnome would find investor surplus from holding smallcap as well as largecap for any of the holding horizons. Smallcap would add more for less risk averse investors, and do so for shorter holding period horizons. Very risk averse investors should not invest at all in smallcap stocks, no matter what their holding horizons are, at least within the range 1-5 years.

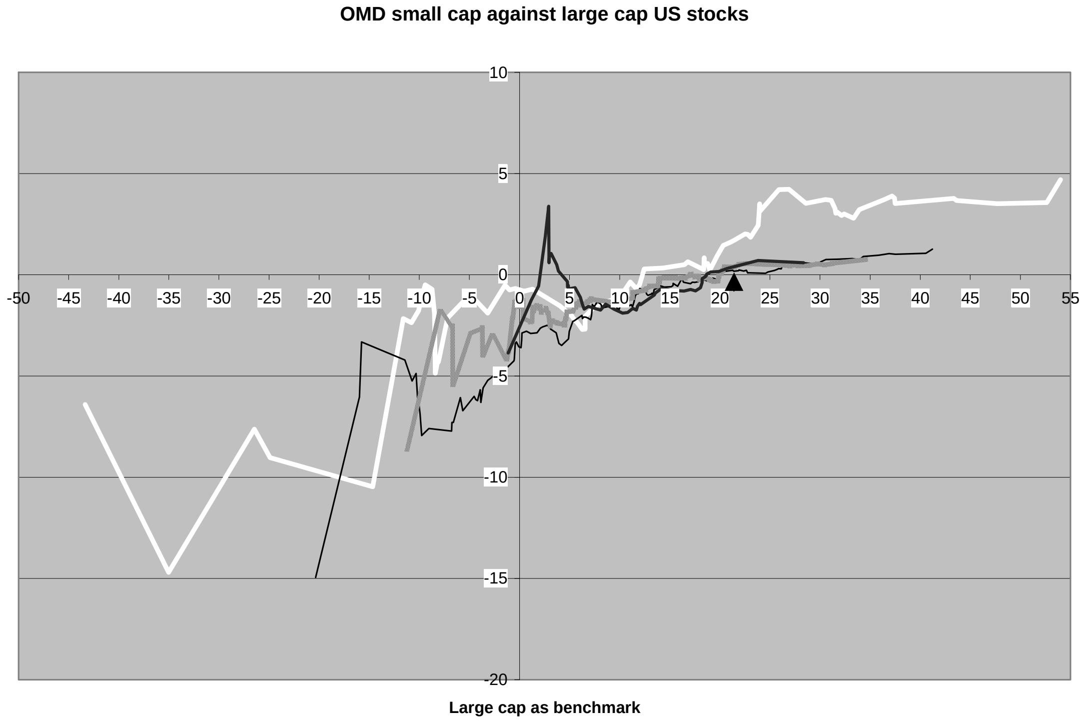

Figure 5: Risk profiling and holding period analysis

An interesting facility is to examine the effect of changing risk attitudes. This is a common justification for portfolio insurance: In the upswing, wealth is increasing and risk premiums are declining, while on the downswing the reverse is true and investors are becoming more risk averse. This means that the effective P value is changing, first rising and then falling. So one can examine OMD plots to see whether the proposed security is

[9] For further detail and discussion, see Bowden (2003d).

changing in its portfolio value over the cycle. In the context of figure 5, for instance, small cap investments would be a better bet on the upswing just for this reason (at least if one believes in the risk premium story).

3.3 Portfolio enhancement

The term portfolio enhancement can cover a variety of functions. Commonly it is used in the context of outsourcing aspects of the fund manager's task. For instance, the manager may decide to outsource foreign exchange exposure management to specialist 'currency overlay' managers. It can also be used where the manager employs derivatives to adjust for the risk or other characteristics of the portfolio, as in portfolio insurance. Or it can refer to the practice of adding elements of differing credit or other characteristics, on either a temporary or permanent basis. Particularly where outsourcing is contemplated, the expected return increment from such an enhancement is sometimes referred to a 'portable alpha', with reference to Jensen's alpha as the implied performance metric.

Outsourcing often involves instructions to the enhancer that the function should be accomplished with no formal capital allocation. So will enhancement based on zero cost instruments such as forwards and swaps. The zero capital[10] aspect is a useful device to isolate the true incremental effect, one that can be employed even where less formal portfolio reallocations away for the benchmark or base portfolio are involved. In the latter context, one would imagine going long in the proposed enhancement portfolio financed by going short in the benchmark portfolio, so the net effect entails zero capital.

The earlier reference to the 'portable alpha' suggests that the same measurement process could be carried out by using the OMD of the proposed enhancement security or portfolio against the base or benchmark portfolio. The only modification necessary is occasioned by the zero capital to be devoted to the enhancement. The enhancement OMD is now computed by taking the running mean of the enhancement return, ordered by the base portfolio. In the illustrative computation of figure 2 of section II, it is the running mean of column 3. The plot of this against the base as benchmark is called the enhancement OMD. The requirement is simply that it be positive over the desired region.

---

[10] The zero capital framework, together with the OMD enhancement, is contained in Bowden (2003a).

Suppose that the result comes out to be 2%. One could say that this constitutes a net or absolute 2% gain to the portfolio. If security or portfolio *r* is an enhancement, it needs no capital, so there is no opportunity cost. This differs from the standard version of the equivalent margin, in which the opportunity cost is the return on a dollar in the benchmark or base portfolio, so a relative return is involved. In the enhancement context, the 2% is a 'serependity' or 'pennies from heaven' return from knowing how to enhance.

Figure 6 illustrates in the context of portfolio enhancement. It portrays the benefits or otherwise of offshore investment and associated currency hedging to two investors, one domiciled in the US and the other in the UK. The US investor has a base position in the US index. The first question is whether a dollar in the UK index will add value on a natural return basis, i.e. one hedged against further value changes due to foreign exchange. The second question is whether this should or should not be hedged against the currency risk. The associated graph indicates the net gain from the currency. The unhedged offshore position is obtained as the vertical sum of the two.

Looking at the upper graph, one can see that to the US investor, the UK index investment does not add much value; the OMD schedule is either negative or very flat. Indeed, the negative zone tells us that very risk averse investors should not add the UK index, or go short. On the other hand, the lower graph tells us that for the UK investor, the US index will add value uniformly over all risk preference zones. In this respect, however, the investment should definitely remain hedged, as the currency component adds negative value over all zones.

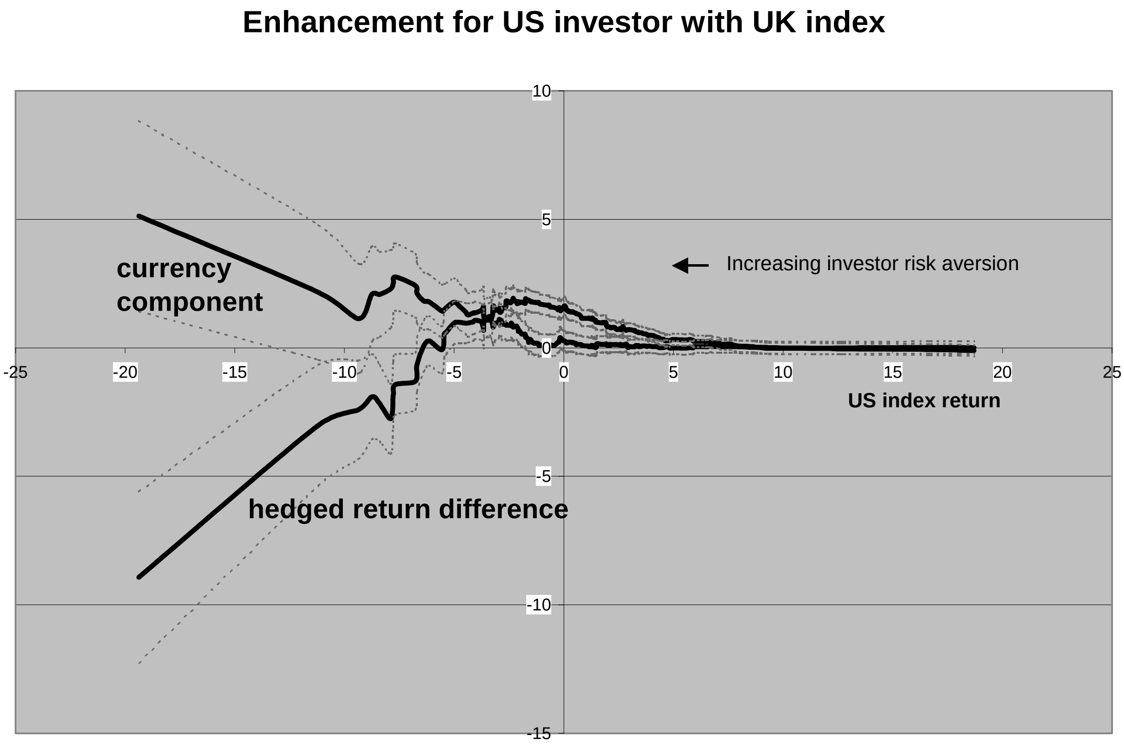

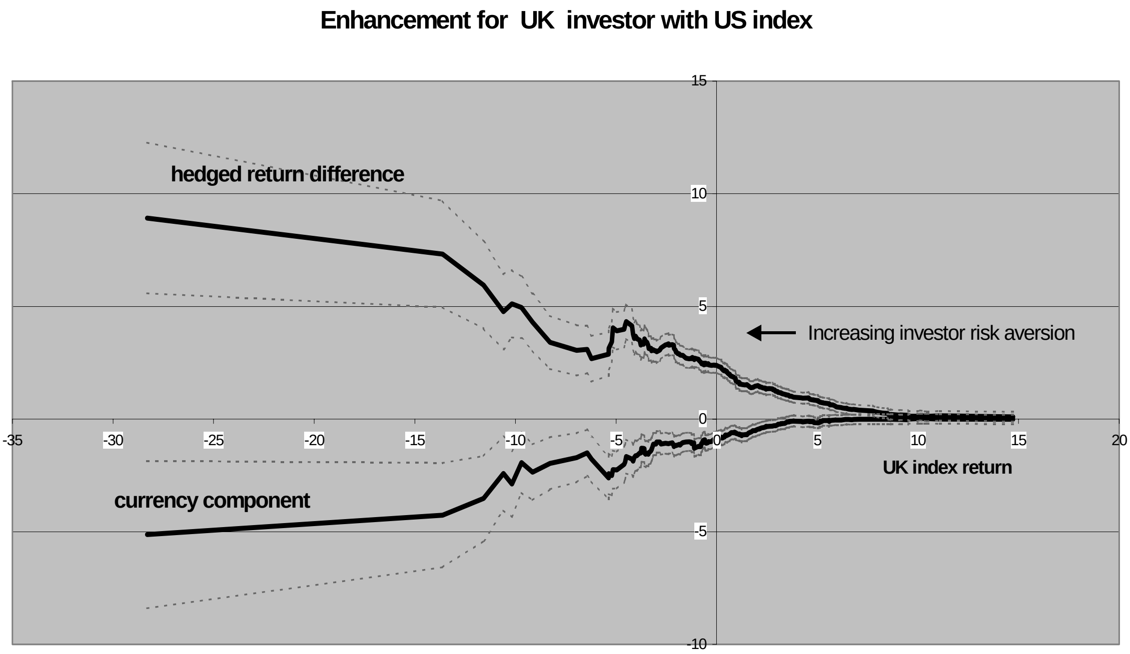

Figure 6: Offshore enhancement OMD's

3.4 Hedging

Hedging can be regarded as enhancement, so the remarks in the foregoing also apply to more or less routine hedging practice. Specific interest attaches to the question of whether the optimal hedge ratio should be independent of the manager's risk preferences[11].

Empirical hedge procedures typically adopt a symmetrical mean square error immunization criterion, wherein positive values of the hedging basis or net hedged return are penalized equally along with negative. The least squares, or econometric hedge, which proceeds by estimating a presumed underlying regression relationship, is naturally adapted to this loss function. If both the hedge instrument (e.g. a commodity future) and the position to be hedged (e.g. the commodity spot price) are normally distributed, then the optimal hedge is indeed theoretically independent of the manager's risk preferences, and is identical with the econometric beta.

In many circumstances, however, a symmetrical loss function might be excessively limiting or statistically inappropriate. An instance might be where a required value at risk type lower boundary is in danger of being violated, so that the manager is hypersensitive to low values. Ordered mean difference techniques can be used to adjust the econometric hedge up or down according to the user's perception of sensitive zones of net return, in a variant on risk profiling. If s is the return on the position to be hedged (e.g. a spot price) and f is the return on the hedge instrument (e.g. a future), the hedging basis is $u = s\text{-}\beta f$, where β is the hedge delta. If the OMD schedule for f against u as benchmark is positive in the sensitive return zone, then the recommendation is to increase the hedge ratio relative to the econometric hedge β.

3.5 Computing optimal portfolios

In mean variance portfolio analysis, the efficient frontier is the locus of portfolios that maximize expected return for a given standard deviation. This locus is one dimensional, meaning that all such portfolios can be obtained just by varying a single parameter. For instance, one can find a point on the efficient frontier by imagining a notional risk free rate and looking for points of tangency. So the notional risk free rate would be the single parameter that one varies to trace out the entire efficient frontier.

---

[11] Formal conditions for hedge invariance, together with further discussion, are given in Bowden (2003c).

A similar yet more powerful result can be obtained by regarding the utility generator P as the notional parameter. Fix a particular value of P and find the portfolio that maximizes expected utility for the utility generator with focus P. Now repeat over all possible values of P. The resulting set of portfolios will be one dimensional, in this case parametered by P.

What sort of frontier would this provide? The answer is the set of all utility maximizing portfolios, regardless of what the utility function might actually be, provided only that it is risk averse. In other words, a global efficient set, regardless of whether or not mean variance applies. The investor's optimal portfolio, whatever it is, will automatically be in that set. All that the manager would need is to decide just which one, effectively what value of P he or she is operating off, or the equivalent gnome for this manager. Revelation experiments can be designed to help with that, though no settled technique yet exists.

The process of actually solving for the maximising portfolio for the utility generator at P is quite easy. It amounts to solving a very simple linear programming problem, one where the objective function has just the two linear segments[12]. Most standard LP programmes can do this.

3.6 Stochastic dominance

The OMD techniques have a close relationship with stochastic dominance, which ultimately stems from the fact that the utility generators also generate the cumulative distribution functions used in stochastic dominance theory. The OMD schedules yield a simple test for second order stochastic dominance (SSD).

To test whether security r is SSD over security R, we simply redo the OMD calculation as in figure 2, but this time we reverse the benchmark; so we plot the running mean of the difference r-R reordered by r rather than by R. If the resulting OMD schedule is everywhere above the horizontal axis, then we know that security r is SSD over security R. Note that the benchmark choice is important in this process. Thus Figure 3 suggests that IBM might be dominant over the market. But if the test is carried

[12] We have to solve it subject to the constraint that the portfolio proportions must add to unity. The technique will not work if further portfolio constraints such as non negativity are required. The algorithm itself is set out in Bowden, R.J. (2003b).

out using IBM as the benchmark, instead of the market, the resulting OMD schedule now crosses the horizontal axis, which does not indicate stochastic dominance.

Stochastic dominant efficient portfolios are those that are not dominated by any other portfolios. The global efficient portfolio set described in section 3.6 has this happy property; it is also the SSD efficient set[13].

III Concluding remarks

Although it has less restrictive implied assumptions than mean variance analysis, the OMD techniques have a few limitations which they share with mean variance. As currently formulated, they require one to assume efficient markets, which means that successive observations are serially independent. It is also assumed that over the chosen window of time, there is some stable law generating the observations, so that data distributions remain the same over the data window. They are also quite demanding in terms of the number of available observations, though this can be partially overcome with the use of bootstrapping techniques. However, limitations understood are limitations allowed for in the interpretation of results from OMD analyses. The latter are often revealing enough to warrant further investigation and applications.

Most experience thus far has been in bilateral applications where the benchmark is well understood, and one wishes to explore issues of performance or enhancement. Applications to full portfolio construction are less developed, though the algorithm itself is a straightforward application of linear programming. An issue here is the optimal design of revelation experiments to decide just where on the efficient frontier is the preferred point, or in OMD terms, to identify the focus for the representative gnome for a given client. A further extension might be to encompass multi period portfolio planning. The utility generator construct can certainly be extended in this way.

---

[13] The relationship of OMD and related constructs with stochastic dominance has been noted by several authors, starting with Shalit and Yitzhaki (1994), whose ‘absolute concentration measure’ can be renormalized to yield the OMD schedule. For direct use of the utility generators to derive SSD efficient portfolios, see Bowden (2003d), while the OMD emerges as a dual in the method of Post (2001).

References

Bigelow, J.P. (1993) `Consistency of mean-variance analysis and expected utility analysis: A complete characterisation', *Economics Letters*, 43, 187-92.

Bowden, R.J. (1992) 'Welfare economics and fund performance: the equivalent margin and associated measures', School of Banking and Finance, University of New South Wales.

Bowden, R.J. (2000) 'The ordered mean difference as a portfolio performance measure', *Journal of Empirical Finance*, 7,195-223.

Bowden, R.J. (2002) 'Ordered mean difference benchmarking, utility generators and capital market equilibrium', School of Economics and Finance, Victoria University of Wellington; forthcoming, *Journal of Business*, March 2005, 444-*.

Bowden, R.J. (2003a), 'The zero capital approach to portfolio enhancement and overlay management', *Quantitative Finance*, 3 (August), 251-61.

Bowden, R.J. (2003b) 'Utility spanning and the ordered mean difference envelope', revised, School of Economics and Finance, Victoria University of Wellington. [www.kiwicap.co.nz]

Bowden, R.J. (2003c) 'Utility-invariant linear hedges', School of Economics and Finance, Victoria University of Wellington. [www.vuw.ac.nz/sef/, research projects]

Bowden, R.J. (2003d) 'Portfolio risk profiling via the bootstrap: holding period analysis for small cap versus large cap stocks', School of Economics and Finance, Victoria University of Wellington. [www.vuw.ac.nz/sef/,research projects]

Dybvig, P.H., and S.A. Ross (1985)'Differential information and performance measurement using a security market line', *Journal of Finance*, XL, 383-99.

Fishburn, P.C. (1977) 'Mean link analysis with risks associated with below target returns', *American Economic Review*, 67,116-126.

Lorimer, D.E. and C.R. Rayhorn (2003) *Financial modeling for managers: with Excel applications*, Corballis, Oregon: Authors Academic Press.

Markowitz, H.M. (1959) *Portfolio selection*, New York: Wiley.

Post, T. (2001) 'Testing for stochastic dominance with diversification possibilities', Research Report #117, Erasmus University, Rotterdam; forthcoming, *Journal of Finance*.

Russell, W.R. and T.K. Seo (1988) ‘Representative sets for stochastic dominance rules’, in T.B. Fomby and T K. Seo (eds) *Studies in the economics of uncertainty: In honor of Josef Kadar,* New York: Springer Verlag.

Rothschild, M. and J.E. Stiglitz (1970) ‘Increasing risk I: A definition’, *Journal of Economic Theory*, 2, 225-43.

Shalit, H. and S. Yitzhaki (1994) ‘Marginal conditional stochastic dominance’, *Management Science*, 69, 670-84.