Revised June 2002

Ordered mean difference benchmarking, utility generators and capital market equilibrium

by Roger J. Bowden*

Abstract

Originally designed as a fund performance measure, the ordered mean difference construction is extended to characterise zero surplus situations such as a capital market equilibrium generated by arbitrary risk preferences. This enables non parametric testing for whether CAPM applies and the detection of pricing inefficiencies or anomalies from historical data, including international capital market segmentation. Any risk averse utility function can be decomposed into a weighted average of elementary put option profile or 'gnomic' utility functions, which collectively generate the OMD areas. The risk profile of the investor can be summarised in terms of a representative gnome, as can the market risk premium.Pricing efficiency turns on whether such a representative gnome exists.

Key words: Benchmarking, CAPM, capital market segmentation, equivalent margin, fund performance, generalised distributions, investor surplus, market efficiency, ordered mean difference, risk premium, running mean operator, utility generators.

* School of Economics and Finance, the Victoria University of Wellington, P.O. Box 600 Wellington New Zealand email roger.bowden@vuw.ac.nz . This paper was delivered at the New Zealand Finance Conference Feb.2-3 2001, and the author is grateful to participants for their comments. Thanks go to Glenn Boyle, Graeme Guthrie and a referee for some helpful comments.

I Introduction

Ordered mean difference (OMD) fund performance technology (Bowden 2000) exploits the areas generated by regression of the fund on a benchmark to create measures of investor surplus. It develops an earlier idea of Merton (1981), who pointed that effective market timing was similar to a put option profile. When the market return declines, the fund return falls by less, and the resulting regression is non linear. Classical measures of statistical dominance ignore the temporal associative aspect. For instance, in second order stochastic dominance (SSD), all timing information is discarded and only the marginal returns distributions are compared. In the regression framework, on the other hand, information is derived from the conditional returns of the subject asset given the return on the benchmark. OMD techniques exploit the conditionality aspect, but are themselves non parametric in nature, and undemanding as to implementation.

In addition to its use in performance measurement, OMD technology has also been utilised to solve for the SSD efficient portfolio set and related aspects in efficient portolio design (Post (2002), Bowden (2002)). The present paper further extends the scope to a market equilibrium context. It is shown that OMD techniques can be adapted to test for capital market efficiency, spanning such topics as testing for CAPM, detection of historical pricing anomalies, and international capital market segmentation.

Application to CAPM testing follows from a key property of the OMD schedules of the security returns against the market. If a CAPM holds, all OMD schedules should cross the horizontal axis, and do so at the same point. Because they are non parametric in nature, OMD methods are not vititiated by a failure to specify what happens under the alternative hypothesis. Thus even if a CAPM is revealed not to hold[1], the OMD schedules remain a source of information about the nature of the pricing inefficiency. The OMD methods are essentially tests of whether the subject asset is or is not already spanned by the benchmark portfolio. The first contribution of the paper is to develop these themes, and in doing so to exhibit a range of empirical applications not canvassed in the original source.

An associated objective is to extend the OMD theory itself It is shown that any risk averse investor can be represented by a spectrum of elementary investors (‘gnomes’) each with a put option payoff profile (the utility generators) with a different strike price or ‘focal point’. Investors differ in the weights assigned to the different utility generators. Thus investor B is more risk averse than investor A if B has a higher proportion of utility generators with lower foci. We show that each investor can effectively be replaced by a representative gnome, and the gnomic focal point used an an index for interpersonal comparisons of risk appetities. Moreover, the financial general equilibrium as a whole can be imagined to be created from just a single representative market gnome, with the market risk premium generated by that single utility generator. Additional theoretical development is concerned with risk premium theory, where a generalisation is discussed of the Rubinstein (1974) version of the risk premium.

The scheme of the paper is as follows. Section II commences with a preview of OMD technology. We introduce some new material formalising the OMD operator and also exhibiting the relationship with aggressive and defensive security profiles relative to a given benchmark. Section III develops the connection with investor utility, risk preferences, and the utility generators. It exploits the the idea of the ‘representative gnome’, and extends the theory to interpersonal utility comparisons. Illustrative material is drawn from national stockmarket performance, where it is suggested that the Australasian stock markets over the past decade were not fully integrated with the U.S. market. Section IV specialises the theory to the case of a CAPM market equilibrium, where the benchmark is the market portfolio We show what should happen to OMD schedules if a CAPM holds, and how the market risk premium is generated in terms of the representative option style utility generator. These properties can be used to detect historical inefficiencies or pricing anomalies. Application to some New Zealand security return data reveals behaviour that is not consistent with a CAPM model of market pricing. Section V concludes, and there are three short appendices elaborating computational or technical aspects.

.

II Benchmarking regressions and value generation

Theoretical OMD based welfare comparisons are based on the regression of the subject fund or security return r upon a suitable benchmark return R. Issues as to the choice of benchmark and the precise meaning of the welfare comparison will be discussed below, but for expositional purposes, r and R will provisionally be taken as arbitrary. For brevity we shall often use the return symbol (r or R) to refer to the security itself (so that 'security r' means 'the security whose return is r ').

The theoretical regression $e(R)$ is defined as the schedule of conditional expectations $E[r \mid R]$ expresssed as a function of R. Thus one can write

$$(1) \qquad r = e(R) + \epsilon \; ; \qquad e(R) = E[r \mid R], \qquad E[\epsilon \mid R] = 0 .$$

There is no necessity for the conditional expectation function $e(R)$ to be linear; indeed in detecting managerial market timing ability, $e(R)$ may be quadratic or have option- like profiles (eg Merton 1981 or Dybvig and Ross 1985). A strong approach to superiority of return r over the benchmark R is cast in terms of the difference $e(R) - R$, so that for any given realisation of the benchmark R , one can expect return r to do better. The welfare assessment utilises the area between the 45^0 line and the theoretical regression, which one can provisionally regard as cumulative investor surplus. As a criterion for superiority, it is too demanding to require $e(R) > R$ uniformly, for there may be temporary blips below the 45^0 line that make little difference to an overall picture of the welfare gain. One might therefore consider the area between the regression line of r on R and R itself, accumulated up to some given point P. If this is positive, over all values of P, then one could conclude that fund r is superior to the benchmark R.

The OMD family of measures utilises the above area, with some minor adaptations. First, the benchmark observation at R has to be weighted with its relative frequency. Let $F(R)$ denote the cumulative frequency of the R observations, ie. $F(R) = prob\ (R \leq r)$ and for simplicity, imagine that this has a data density $f(R) = F'(R)$. Thus the deviations $e(R) - R$ have to be weighted with their relative frequency, $f(R)$. Second, we take the running mean of the area elements up to the chosen point P, effectively dividing the area up to P by the number of oservations with $R \leq P$.

In summary, the theoretical ordered mean difference at any chosen point R =P is given by

$$(2a) \qquad t(P) = \frac{1}{F(P)} \int_{-\infty}^{P} [e(R) - R] f(R) dR \; .$$

It might now be possible for $e(R)$ to dip temporarily below the 45^0 line without necessarily making $t(P)$ negative. However, if $t(P) \geq 0$ for all P, it would be concluded that the fund is superior to the benchmark.

Expression (2a) can be written in the more concise form:

$$(2b) \qquad t(P) = E_P[e(R) - R] = E_P[r - R] \; ,$$

where the running mean expectation operator $E_P[\bullet]$ is the ordinary expectation with respect to the truncated density $f_P(R)$. The latter is defined by

$$f_P(R) = \frac{1}{F(P)} f(R) \quad ; \quad R \leq P$$
$$= 0 \quad ; \quad R > P.$$

The second equality in (2b) is shown by utilising the decomposition (1).

To compute the average area, one does not need to know what the theoretical regression function $e(R)$ actually is, or even to estimate it as such For most practical purposes, the following non parametric measure will accomplish the same thing :

$$(3) \qquad OMD(P) = \hat{t}(P) = \frac{\sum_{i=1}^{i=\#P}[r_i - R_i]}{\#P} .$$

In formula (3), the symbol # P denotes the number of observations R_i for which $R_i < P$. The numerator indicates that the difference between r and R should be summed only over those observations, once the observations have been ordered in ascending order of R values. Expression (3) constitutes a non parametric estimate of the ordinates of the OMD function $t(P)$, motivated by the second equality in (2b). It is equivalent to first estimating the theoretical regression e(R) by using a non parametric kernel conditional expectation estimator and then taking the running mean of the result. Alternatively the function $e(R)$ can first be estimated by parametric means, e.g. by specifying a polynomial form, followed by integration. However, the version (3) has the virtue of simplicity: Appendix A shows how the calculation may be accomplished with an Excel spreadsheet. For further estimation theory see Bowden (2000).

Aggressive and defensive security-benchmark relationships

The running mean operator $E_P[\bullet]$ will be used frequently in what follows. Applied to an arbitrary differentiable function $\phi(R)$, the theoretical running mean of the function values satisfies the error correcting property:

$$(4) \qquad \frac{d}{dP}E_P[\phi(R)] = \frac{f(P)}{F(P)}(\phi(P) - E_P[\phi(R)]).$$

Hence the running mean will consistently be increasing if current values lie above the running mean to that point. An important special case is the running mean of R itself, which will be written for brevity as

$$(5) \qquad \mu(P) = E_P(R).$$

Inspection of expression (4) with $\phi(R) = R$ will show that the function $\mu(P)$ is always increasing with $lim_{P\to-\infty}(\mu(P) - P) = 0$ and $\lim_{P\to\infty}\mu(P) = \mu_R = E[R]$.

Applied to the 'excess function' $\phi(R) = e(R) - R$, property (4) motivates the following definition.

Definition *A security or fund of return r will be said to be aggressive with respect to a benchmark of return R if the OMD schedule is increasing*, *i.e.* $t'(P) \geqslant 0$, *and defensive if the OMD schedule is decreasing*, *i.e.* $t'(P) \leqslant 0$.

[Remark: Here and in what follows, the signs $\leqslant$ or $\geqslant$ applied to pointwise comparisons will be taken to mean that strict inequality holds over some interval or point P of non zero probability measure with respect to F(P)].

The above definition is motivated by conventional CAPM contexts, where return R is that on the market; because the asset with return r is supposed to be in equilibrium with the market, its theoretical regression is linear. In this context, the excess function is $e(R) - R = (\beta - 1)R$, where β is the linear regression coefficient of the security on the market. An aggressive security is conventionally one for which $\beta > 1$, and a defensive security has $\beta < 1$. The slopes of the OMD schedule t(P) will certainly reproduce this behaviour (see section IV).

However, the definition also allows for non equilibrium or other benchmark situations where the theoretical regression may be non linear. A particular application might have the theoretical regression $e(R)$ having a slope less than unity in some small region, but an overall sense of slope greater than unity. One could describe the function $\phi(R) = e(R) - R$ as increasing 'on the average' if its current value was always above its moving average to that point, so that the right hand side of expression (4) would always be semipositive. The characterisation of aggression or defensiveness in terms of the OMD is more forgiving of temporary aberrations. It also turns out to accord with attitudes to risk and how they determine preferred security profiles (see section III).

Illustration

Figure 1 illustrates with OMD schedules for the Australian and New Zealand stock market indices against the U.S. over the decade of the 'nineties. The respective indices are the ASX All Ords gross, the NZSE-40 gross, and the Dow Jones Industrial Index (DJIA). Monthly return data run from January 1990 to December 1999. All returns are imagined to be currency hedged, so this is a pure index returns comparison. The NZ market OMD worm[2] lies wholly below the horizontal axis, suggesting that it would have paid investors to go short the NZSE-40 index against the DJIA over the decade. The Australian market fares better though does not dominate the Dow Jones index, crossing over once. The dotted lines above and below the respective OMD schedules are conservative[3] one standard deviation bands. The Australian market is revealed as defensive against the US, which is a reversal from the seventies (see Bowden 2001). The NZ market lies wholly underneath the horizontal axis, so that investor surplus is negative over all values. The latter property means that every investor, no matter what their preferences, should have shorted the NZ market in favour of the U.S. Such interpretive aspects are considered further in what follows.

Figure 1 OMD worms for Australia and NZ against U.S. Benchmark

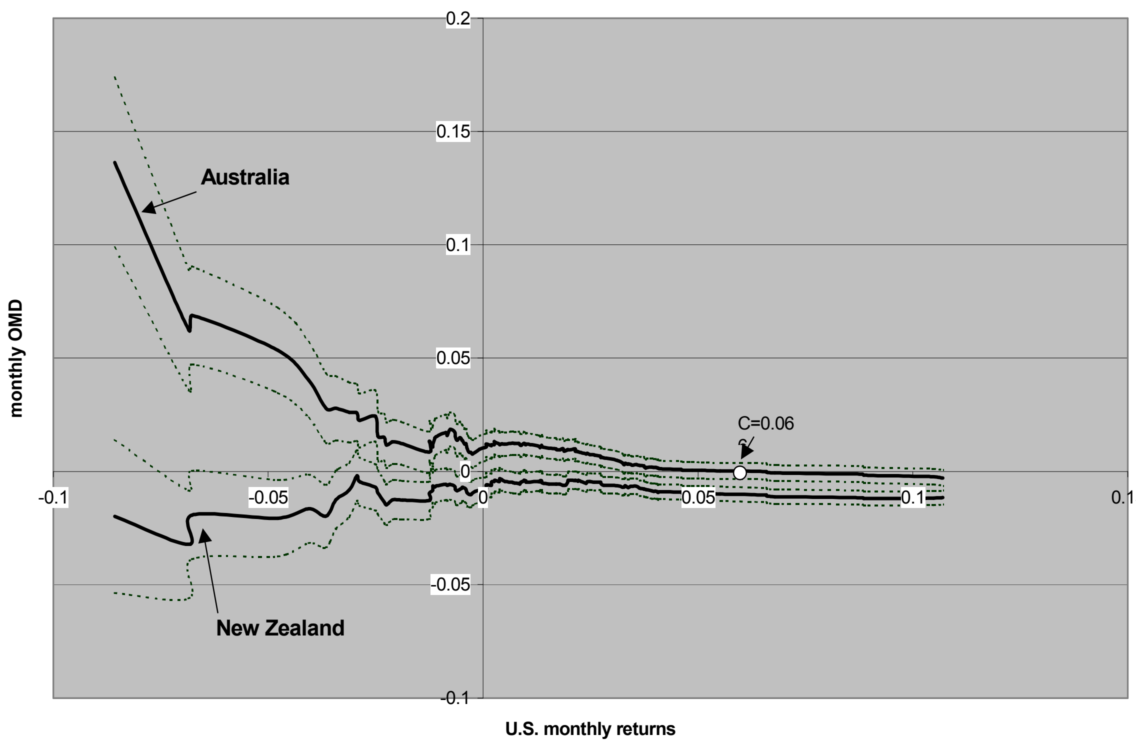

III Investor utility and equivalent margin

The portfolio theoretic basis of OMD technology lies in relative efficiency. Given a benchmark return R and a subject security or fund (of return) r, we will describe the benchmark portfolio R as efficient relative to r if there is no utility value in adding further units of r to the benchmark portfolio. The benchmark might, for instance, already include optimal amounts of the subject security, or the benchmark may already span the payoff of the subject security. If there is utility value in going long or short in return r, the implication is that the existing benchmark portfolio is not efficient for the investor, and the manager could do better by adding or subtracting additional units of the subject security. Performance benchmarking for internal performance and reward systems commonly compares the manager's portfolio return with the return on a designated benchmark portfolio; the latter may be chosen as a passive portfolio incorporating the desired risk profile for the fund. If the manager can add value by combining with the benchmark portfolio R, this implies that the latter is not efficient relative to r. Because the manager always has the alternative of simply investing in the benchmark, a better portfolio can be constructed than the benchmark, and the fund manager has added a source of value.

Thus performance measures of the present kind relate to personal relative efficiency. The same ideas can be utilised to study any presumed capital market equilibrium. In this case, a given security might not be in relative equilibrium with the market portfolio for the representative investor, and one could then say that the presumed market equilibrium R is not efficient as a whole. For the moment, however, the discussion will remain general as to context.

The equivalent margin

Published indicators of fund performance commonly rely on the growth in value of money invested with the subject fund. Over a long period, and where returns were uncorrelated over time, this would amount to maximisation of the expected one period utility function taken as the logarithm of wealth. In what follows we shall remain more general, simply because investors do differ as to their risk appetites, even though the expected log of wealth will remain an important particular case. Thus if R_p denotes the return on some portfolio, then we consider the maximisation of expected utility $E[U(R_p)]$ as the welfare criterion, where U is a conventional (concave, Von Neumann Morgenstern) utility function for a risk averse investor.

To assess the surplus attributable to the fund, we use a further idea drawn from welfare economics. Imagine that the return r on the fund is to be taxed, the return on the benchmark R remaining untaxed. Thus the fund return will now be $r - t$, where t is the proposed tax (or just a certainty equivalent type penalty). Suppose we form a new portfolio from the subject return r and the given benchmark R. What is the tax t that the investor would be willing to bear before he or she will discard the fund from the combined portfolio, i.e. its portfolio proportion $x(t) = 0$? That tax rate is the measure of how much the opportunity to invest in the fund means to the investor. Notice the dimensions of this welfare measure: It is itself a rate of return, so the welfare losses or gains are given an equivalent rate of return dimension, which is readily comprehended by the investor or manager.

The portfolio return is $R_p = x(r - t) + (1 - x)R$. If expected utility is maximised for a given tax rate t, than the optimising portfolio proportions $x = x(t)$ will depend upon the tax rate t. Imposing higher t is always more unpleasant as a matter of certainty, and the optimising portfolio proportions $x(t)$ will plot as in figure 2. In this figure, two funds are depicted, of returns r_A and r_B, against a common benchmark. The natural (i.e. without the tax) portfolio stance in each, in combination with the benchmark, is indicated by $x(0)$. As indicated, the natural position is long (L) in A and short(S) in B. Fund A derives value from going long, and in this case, the tax would be positive (arrow to the right), while fund B derives value from going short, and in this case the tax is negative, meaning that the investor would have to be paid to give up his short position. One could claim that B is inferior to the benchmark, while A may or may not be superior to the benchmark.

Figure 2 The equivalent margin

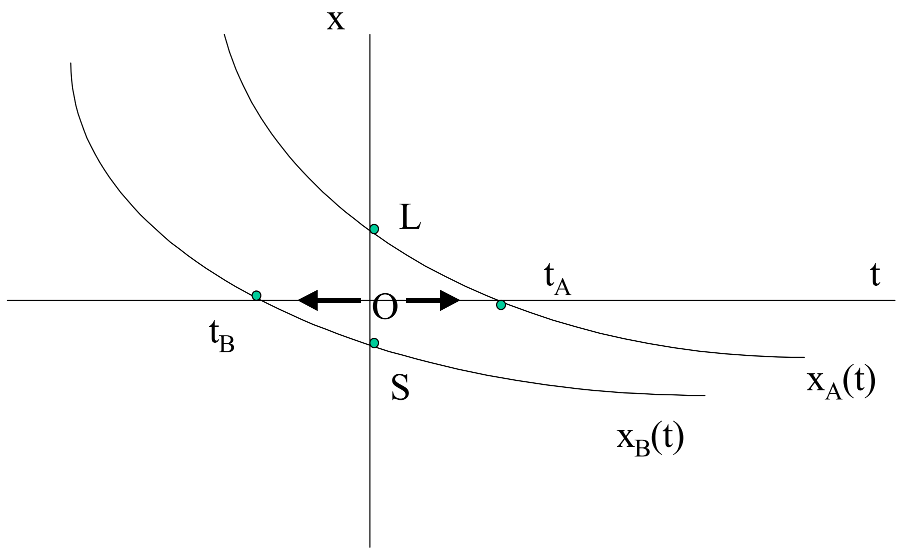

Let $U'(R)$ denote the marginal utility, evaluated at the benchmark R. Using the iterated expectation, one can explicitly solve for the equivalent margin as

$$(6) \qquad t_U = E[\pi(R)(r-R)] \ ,$$

where $\pi(R) = U'(R)/E[U'(R)]$ is a series of positive weights that sum in probability to one. The equivalent margin therefore weights the return differences according to the marginal utility of benchmark realisations. States of the world in which the benchmark performs badly but the fund performs well are rated especially highly, depending on the degree of risk aversion of the investor. In the form (6), the equivalent margin connects to existing portfolio measures[4] and also to the theory of financial general equilibrium (e.g. Duffie 1992). In the latter, if the utility function U refers to the representative agent for the market as a whole, then the $\pi(R)$ are called 'state price deflators', and a value $t_U \neq 0$ indicates that the fund is not spanned by the market. If it were actually traded, an arbitrage would exist in which the investor would go go long or short in security r against the market R.

Returning to the general benchmark context, security or fund r will be in relative equilibrium with the benchmark R if $t_U = 0$. In terms of figure 2, there is no incentive to add or subtract holdings of r additional to any that might already be implicitly represented in the benchmark portfolio. In the case of a CAPM equilibrium, the condition $t_U = 0$ for each individual security r in relation to the market return R provides a simple way of deriving the CAPM equation $\mu_r = \rho + \beta(\mu_R - \rho)$ where ρ is a risk free or zero beta rate.

Another special case of the relative equilibrium principle arises in deriving a risk premium for a representative security, which for future use we shall identify with the benchmark R. Consider a risk free asset that pays a certain return $r = \mu_R - \theta$, for some positive number θ. Setting the number θ such that the risk free rate is in personal equilibrium ($t_U = 0$) with the uncertain return R yields a risk premium:

$$(7) \qquad \theta = -\frac{E[R-\mu_R)U'(R)]}{E[U'(R)]} \quad .$$

The special case where return R is normally distributed yields[5] the Rubinstein (1974) form of the risk premium $\theta = -\frac{1}{2}\sigma_R^2 E[U'(R)]/E[U''(R)]$, where $\sigma_R^2 = Var(R)$. So one could call the more general expression (7) the generalised Rubinstein risk premium. Appendix B is a diagrammatic illustration showing how the comparison looks in mean variance space between two investors, one taken to be more risk averse than the other. If return R represents the market return in a CAPM equilibrium, and the utility function U is that of the representative investor, the corresponding θ is the market risk premium.

Relationship with the OMD

Suppose that U is any investor utility function with marginal utility of wealth diminishing to zero as W→ ∞, and let t_U denotes the equivalent margin welfare measure as above. The connection with the OMD schedule of section II stems from the theoretical relationship

$$(8) \qquad t_U = \int_{-\infty}^{\infty} w(P)t(P)dP \ ; \quad w(P) = -\frac{1}{2}\frac{U(''(P)]}{E[U'(R)]}F(P) \quad ,$$

where the semipositive weight function $w(P)$ sums to 1 over P. The weight function depends upon (i) the risk aversion properties of the given utility function; and (ii) the cumulated probability mass $F(P)$ of the R distribution at the indicated point $R = P$. Expression (8) decomposes the equivalent margin into effects due to the asset returns (the $t(P)$), and effects arising from the particular risk preferences of the agent doing the evaluation (the $w(P)$).

As the OMD schedule $OMD(P)$ is the non parametric estimate of $t(P)$, the equivalent margin welfare measure can be regarded as some weighted average of the values of the ordered mean difference schedule, where the weights depend upon particular investor preferences. Possible sampling error aside, one could therefore say that if the OMD schedule has been computed as in section II and is always positive, then the equivalent margin for any risk averse investor will be positive and the fund will add value relative to a position in the benchmark alone.

Indeed, the theoretical OMD function $t(P)$ is itself the equivalent margin for a special sort of utility function, one in which utility is linear for $R = P$ and zero thereafter, as in figure 3 below: $U_P(R) = R - P$ for $R \leq P$ and 0 for $R > P$. These are the utility generator functions, which can also be written as $(R - P)SF(P - R)$, where $SF(\bullet)$ is the unit step function. The generator at P be regarded as the payoff profile to the writer of a put option with strike price P. Alternatively the generator can be viewed as a 'target return' type of utility function (see Fishburn 1977) where the investor is concerned to obtain a target return P, is indifferent to values of R in excess of this, and is adversely exposed to the extent that the return falls linearly below the target. In what follows, it will be convenient to refer to the point P as the 'focal point' for the utility function $U_P(R)$.

For fixed P, each $U_P(R)$ is a risk averse utility function and the risk aversion diminishes as the focal point P moves to the right (dotted lines in figure 3). Intuitively, if the position of the density $f(R)$ is fixed, then as P moves to the right (e.g. to P') the utility function $U_P(R)$ effectively becomes more and more linear relative to the empirical distribution of R values, and the risk premium declines. Evaluating[5] the generalised Rubinstein risk premium (7) for the utility function $U_P(R)$ gives

$$(9) \qquad \theta(P) = \mu_R - \mu(P) \quad ,$$

where the running mean $\mu(P)$ of R is defined by expression (5). As the running mean is always increasing in P, the risk premium attaching to $U_P(R)$ is monotone decreasing[6].

The interpretation of expression (8) is that the equivalent margin t_U for an arbitrary utility function U is made up of a weighted sum of the equivalent margins for the generators U_P.One could imagine such as investor as a population of type P investors ('gnomes') with utility functions U_P , the population proportions of type P gnomes depending on the risk aversion of the parent investor at each point P.

Figure 3 Generator utility function U_P

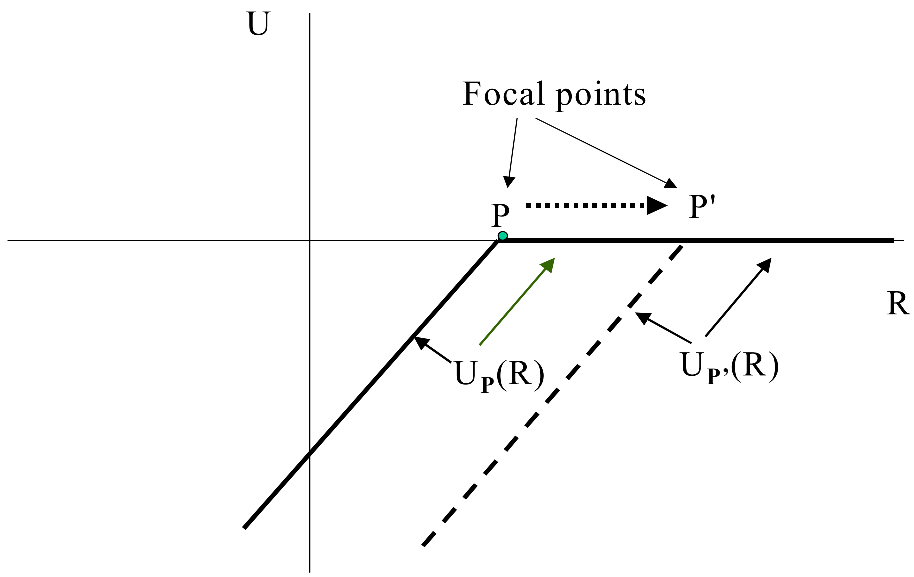

Interpersonal risk profiles

The form of expression (8) suggest that the weighting pattern w(P) can be used to differentiate the risk appetites of different investors. Denote by

$$W(P) = \int_{-\infty}^{P} w(p)dp$$

the cumulated weight distribution. If $W_A(P) \leqslant W_B(P)$ for all P, then one might claim that investor A is less risk averse than investor B. The following lemma supports such a claim.

Lemma *Suppose that* $lim_{P\to-\infty} U'(P) = 0$. *Then*

$$\text{(10)} \qquad \frac{f(P)}{F(P)} W(P) = -\frac{1}{E[U'(P)]} \frac{d}{dP} E_P[U'(R)] \quad .$$

[Proof: Integrate by parts the weight w(P) as defined by expression (8), using equation (4)].

The right hand side of equation (10) represents a slope of the running mean of the normalised marginal utilities at each value of P. The denominator term $E[U'(P)]$ amounts to a normalising factor to allow for the scale indeterminacy of a Von Neumann Morgenstern utility function. Thus the lemma tells us that for two investors facing a common risk $F(P)$, the slope of the running mean of their respective marginal utilities determines the integrated weights at each point P. A more risk averse investor is one for whom the normalised marginal utilities diminish faster.

Additional support comes from the generalised Rubinstein (GR) risk premium. This is defined as the number θ such that $t_U = 0$ with $r = \mu_R - \theta$ (see expression (7) above) . Applying this condition to equation (8) yields

$$\theta = \int_{-\infty}^{\infty} W(P)\mu'(P)dP \quad .$$

The running mean $\mu(P)$ for R is always upward sloping. It follows that if $W_A(P) \leqslant W_B(P)$, all P, then $\theta_A < \theta_B$, bearing in mind our convention on the sign $\leqslant$. In other words, investor B has a higher

GR risk premium.

It should be remarked that the condition $W_A(P) \leqslant W_B(P)$, all P, is only a sufficient condition for investor B to be more risk averse than investor A. However, the idea of interpersonal risk aversion comparisons using the weight pattern allows an interpersonal comparison that is not limited to parametric comparisons based on particular families of utility functions.

The representative gnome, and interpersonal risk aversion comparisons

As the weights $w(P)$ is expression (8) are semipositive, it follows from the second mean value theorem in differential calculus that there must exist a value P_* such that

$$(11) \quad t_U = t(P_*) = \int_{-\infty}^{\infty} w(P)t(P)dP \ .$$

This means that for any given investor, there exists a representative options style utility function $U_{P*}(R)$ that generates exactly the same equivalent margin for a given investment and benchmark. So there must exist a ‘representative gnome’.

Expression (11) also suggests that an investor’s risk aversion can be summarised by the focal point $P*$ as a sufficient statistic for the investor’s risk appetite. A more risk averse investor B should have a representative gnome with $P_{*B} < P_{*A}$ for the less risk averse investor A. The following result shows that this is certainly true if the OMD schedule $t(P)$ is monotonic, either increasing or decreasing.

Proposition 1 *If the OMD schedule $t(P)$ is monotonic, and $W_A \leqslant W_B$, then $P_{*B} < P_{*A}$.*

Proof. Integrating expression (8) by parts for both A and B and subtracting,

$$t(P_{*A}) - t(P_{*B}) = -\int [W_A(P) - W_B(P)]t'(P)dP.$$

Given monotonicity, there are two possible cases.

Case (a): $t'(P) \geqslant 0 \Rightarrow t(P_{*A}) > t(P_{*B}) \Rightarrow P_{*A} > P_{*B}$.

Case (b): $t'(P) \leqslant 0 \Rightarrow t(P_{*A}) < t(P_{*B}) \Rightarrow P_{*A} > P_{*B}$ again.

Hence in both cases, $P_{*B} < P_{*A}$ as required.

□

Figure 4 illustrates. There is imagined to be a common benchmark and two assets̩ one aggressive and one defensive. Investor B is revealed to be more risk averse than investor A, with investor C intermediate betwen the two The two securities are in personal equilibrium with the benchmark in the case of investor C. For investors A and B, however, neither of the two securities are in personal equilibrium with the benchmark.The less risk averse investor A will prefer to go longer in the aggressive security and the more risk averse investor B will prefer to go longer in the defensive security.

Figure 4 Interpersonal risk indexing via representative gnomes

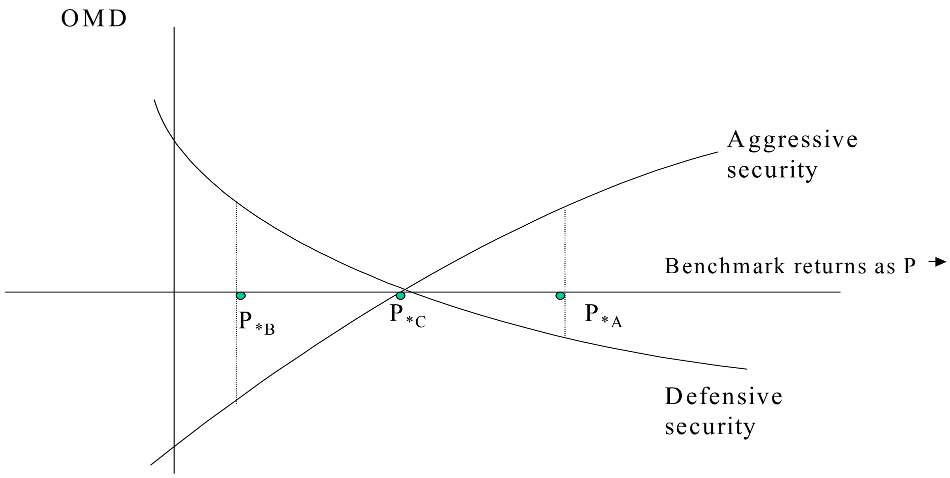

Illustration

Continuing with the national stockmarket comparison of section II, examination of figure 1 suggests that even though Australia does not dominate the U.S. benchmark in OMD terms, the more risk averse risk averse investors might consider Australia to have investor surplus over the U.S. Such an investor would give the high return values very little weight, and would be influenced instead by the superiority of the Australian market over zones of low benchmark return. For such investors, the risk index P_* is less than the crossover point C. Less risk averse investors with $P_* > C$ would add value by going short the Australian market against the U.S. benchmark.

To check such a conclusion in parametric terms, one can compute the equivalent variation for a family of utility functions $U(R) = U(R;\xi)$ allowing increasing risk aversion via a parameter ξ, and see if the t_U values diminish as this happens. A possible family is provided by constant relative risk aversion functions:

$$U(R) = \frac{1}{\xi}(1+R)^{\xi} \;;\xi \leq 0\;.$$

The case $\xi = 0$ corresponds to the log utility function $U(R) = \log(1+R)$. Setting the relative risk aversion coefficients $\xi = 0, -1, -2, ...$gives progressively higher risk aversion than logarithmic.Table 2 gives the outcomes. The equivalent margins are all as annualised percentage returns.

Table 2 Equivalent margin (t_U) values: Australia, with U.S. as benchmark
(as annualised percentage returns)

coeff.of risk aversion (ξ)	equivalent margin $t_{U\xi}$
0 (log)	-3.30%
-1	-2.63%
-2	-1.96%
-3	-1.28%
-4	-0.60%
-5	+0.10%

A risk neutral investor, corresponding to gnomic P= ∞, would have $\xi = 1$, with $t_U = \mu_r - \mu_R = -3.97\%$, the difference between the two means over the period. Investors with risk aversion coefficients in the range usually considered to be normal, namely 0 to -2, would evidently have elected to go short the Australian market against the U.S. as benchmark. However, for progressively higher values of the aversion coefficient, the willingness to go short disappears, with the zero at $\xi = -4.91$, corresponding to a gnomic P = 79%, the annualised point C in figure 1.

The risk aversion needed for investors to go long in the Australian market is evidently very high. A more reasonable hypothesis is that the Australian and U.S. markets cannot be viewed as an integrated whole, perhaps for reasons of habitat preference. Bearing in mind the domination of the NZ market in figure 1, this would indicate that neither of the Australasian markets was fully integrated with the U.S., in both cases showing evidence of international segmentation.

IV Market equilibrium relationships

The important special case where the benchmark is a market return enables tests of whether or not the past returns on any given security are consistent with a presumed capital market equilibrium. If the security is revealed to have non zero investor surplus, this suggests either that the security is not actively traded, or else that the the market as a whole has been revealed as inefficient. One should not hasten to dismiss market efficiency just because one or two minor securities have revealed unutilised possibilities for investor surplus in conjunction with the market portfolio. But if a significant number of such anomalies have been revealed, then there is a case that either the market is inefficient or else that the chosen representation of the market equilibrium is flawed. Thus the focus of the present section shifts from personal to market equilibrium.

The representative market gnome

In financial general equilibrium theory (e.g.Duffie 1992), market equilibria can be described in terms of the the risk premiums generated by a representative agent, the latter obained in terms of an aggregation of the individual agent utility functions. Rubinstein(1973) showed that in a CAPM style equilibrium with R_m the market return normally distributed and ρ a risk free rate, the market risk premium could be obtained as

$$(12) \quad \mu_m - \rho = -\sigma_R^2 \frac{E[U''_*(R_m)]}{E[U'_*(R_m)]}$$

where U_* is the utility function of a representative agent for the economy. The relationship readily generalises to non normal market returns, in which case the Rubinstein risk premium has to be replaced by the generalised version (expression (7) of section III with $R = R_m$).

In efficient capital market equilibrium, each individual asset r should trade so that its equivalent margin evaluated by the representative agent, is precisely zero. Thus $t_{U*} = 0,$ for all assets This means that for any asset r, the representative gnomic P_* must be such that $t(P_*) = 0$; so that the theoretical OMD schedule $t(P)$ crosses the horizontal axis at the point P_* . Thus the point at which the OMD schedule for each asset crosses the horizontal axis assumes special significance. In the next section it is shown that this can be turned into a diagnostic for the existence or otherwise of a CAPM equilibrium. If such an equilibrium does hold, then it can be generated in terms of the representative gnomic utility generator, so that the representative utility function U_* can be replaced by U_{P*} . The argument as it stands does not show that the crossover point P_* is necessarily the same for all assets However, this turns out to be true for a CAPM equilibrium, so that a CAPM market gnome can be said to exist.

All assets in CAPM equilibrium: the once cross over rule

If a given security is already in a CAPM with respect to the market as benchmark, then its return r will on average plot in a straight line against the market return R_m, the familiar 'characteristic line' of CAPM theory. (Henceforth, with the context understood,we replace R_m with just R). The theoretical regressions $e(R)$ are straight lines, the slope of the line depending upon the beta (β) of the fund against the market. If $\beta > 1,$ the areas underneath the regression are at first negative then become positive after $R = \rho,$ where ρ is the risk free rate or the expected zero beta portfolio return. If $\beta < 1,$ the sign pattern is reversed.

Bearing in mind that the theoretical OMD function $t(P)$ is a running mean up to any given point $R = P$, one might expect that the schedules $t(P)$ will change sign, but only after the point $R = \rho$, and then only once. In that case, every security does the same thing and the collection of OMD schedules form a pencil of curves all crossing at the same point P_*, illustrated in figure 5b below. A formal demonstration of the once crossover property is as follows:

Theorem 1 *If all securities trade in a CAPM equilibrium*: (*a*) *Their theoretical OMD schedules* $t(P)$ *plotted against the market return cross the horizontal axis just once*, *and at the same point*. (*b*) *The common crossing value is given by the solution* P_* *to the equation* $\mu(P) = \rho$, *where the function* $\mu(P)$ *is the theoretical running mean schedule of the benchmark return distribution*. (*c*) *The crossing value* P_* *is greater than the risk free rate* ρ . *It is greater than the mean* μ_R *of the market return if* $\mu(\mu_R) < \rho,$ *and less than* μ_R *if* $\mu(\mu_R) > \rho.$

Proof: Suppose that security r is a fully traded part of the market portfolio R, and a CAPM model applies with a risk free rate ρ. From standard CAPM theory, it follows that

$$r - \rho = \beta(R - \rho) + \varepsilon \; ; with \; E[\varepsilon \mid R] = 0,$$

so that $e(R) - R = -(1 - \beta)(R - \rho)$. Recalling expressions (2a,b) of section II, this means we can write

$$(13) \qquad t(P) = -(1 - \beta)[\mu(P) - \rho)$$

where $\mu(P) = \int_{-\infty}^{P} R \frac{f(R)}{F(P)} dR$ is the mean of the truncated R distribution for $R \leq P$. As a function of P, the truncated mean (running mean) has $\mu(P) \leq P,$ also $\lim_{P \to -\infty} \mu(P)/P = 1,$ and $\mu(P) \to \mu_R$ as $P \to \infty$. From the error correction property, it follows that $\mu'(P) = -\frac{f(P)}{F(P)}[\mu(P) - P] > 0$ and also that $\mu''(P) < 0,$ so that the running mean function $\mu(P)$ for the market return is upward sloping and concave.

Considering equation (13), suppose P_* is such that $t(P_*) = 0$, so that $\mu(P_*) = \rho$. Hence at P_* it must be true that $P_* = \rho + \mu'(P_*)/(\frac{f(P_*)}{F(P_*)} > \rho$. As $\mu(P)$ is monotonic so is $t(P)$, and it must cross the horizontal axis just once, to the right of R=ρ.

Turning to the relationship of the crossing point P_* with μ_R, the concavity of $\mu(.)$ implies that

$$(14a) \quad \mu(\mu_R) \le \mu(P_*) + (\mu_R - P*)\mu'(P_*)$$
$$(14b) \quad \mu(P_*) \le \mu(\mu_R) + (P_* - \mu_R)\mu'(\mu_R)$$

Combining (14a,b) with $\mu(P_*) = \rho$, it follows that

$$(\mu_R - P_*)\mu'(\mu_R) \le \mu(\mu_R) - \rho \le (\mu_R - P_*)\mu'(P_*)$$

which in turn gives $P_* > \mu_R$ if $\mu(\mu_R) < \rho$ and $P_* < \mu_R$ if $\mu(\mu_R) > \rho$.

□

Figure 5a illustrates the concepts of the above proof, while figure 5b shows how a whole assemblage of OMD schedules for different securities would plot against the market. If $\beta < 1$ then the theoretical OMD schedule $t(P)$ for any security will cross the horizontal axis from above and it must do so to the right of the risk free rate, while if $\beta > 1$ the schedule will cross from below, also to the right of the risk free rate. In each case, the theoretical OMD schedules for each security r tend to $t(\infty) = \mu_r - \mu_R$ where $\mu_r = E[R] = \rho + \beta(\mu_R - \rho)$.

If the CAPM applies, there is therefore a once cross over rule for the OMD , from above or below depending on whether the beta is greater than or less than unity. In either case, the representative investor's equivalent margin t_{U*} is precisely zero, so that surpluses for one region are compensated exactly by deficits from the other.

Figure 5a Once crossover rule for OMD in CAPM

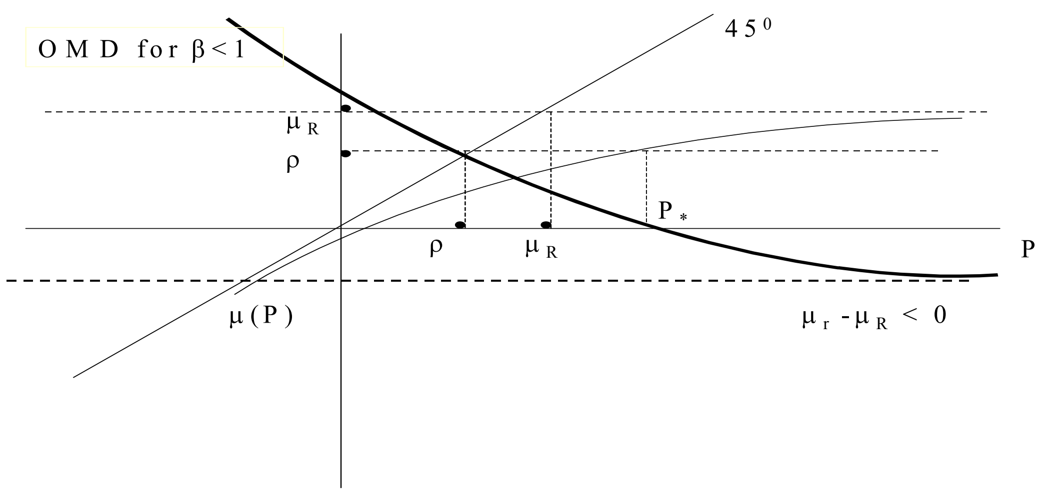

Figure 5b OMD pencil in CAPM equilibrium

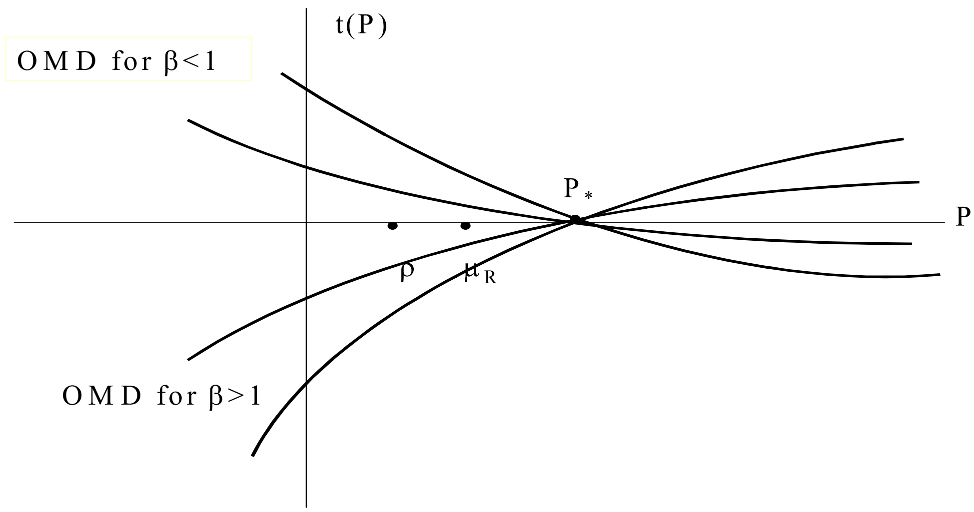

The crossing point and the representative utility generator

Because the crossing points for the different assets are one and the same, the option style utility generator $U_{P_*}(R)$ will suffice to determine the resulting risk premium for the economy. The following result is a special case, namely when the market returns are normally distributed, but it has illustrative value.

Proposition 2 *Suppose that market returns R are normally distributed. Then the market risk premium is given by*

$$(15) \quad \mu_R - \rho = \sigma_R^2 \frac{f(P_*)}{F(P_*)} = -\sigma_R^2 \frac{E[U''_{P_*}]}{E[U'_{P_*}]} \ ,$$

where $f(R)$ and $F(R)$ are the density and distribution function of the market returns.

Proof. The utility generator for arbitrary P can be written

$$U_P(R) = (R-P)SF(P-R)$$

where $SF()$ is the unit step function defined by $SF(x) = 0 \ if \, x \le 0, = 1 \ if \, x > 0$. Hence

$$E[U_P(R)] = \int_{-\infty}^{\infty} (R-P)SF(P-R)f(R)dR \, .$$

It follows that

$$\mu(P) = P + \frac{1}{F(P)} E[U_P(R)] \ .$$

From Theorem 1 , the CAPM equilibrium is therefore characterised by

$$(16) \quad E[U_{P_*}(R)] = -F(P_*)(P_* - \rho) \ .$$

If R has a normal distribution, then[7]

$$(17) \quad E[(R-P)RF(P-R)] = (\mu_R - P)F(P) - \sigma_R^2 f(P) \ .$$

Combining (16) and (17) with $P = P_*$ yields the first equality in the required result (15).

The second equality in (15) needs some more advanced mathematics. The sense in which the derivatives of the formally non differentiable function $U_P(R)$are defined is that of the Heaviside or Temple - Schwartz calculus of generalised functions (e.g. Lighthill 1958) according to which $SF'(x) = \delta(x)$, the Dirac delta function . Appendix C gives the essential conventions and results, and the second inequality is a straightforward consequence.

□

Returning to the general context, the above analysis supplies two crossing points of potential interest in empirical CAPM verification. The first is the crossing point for the presumed CAPM linear conditional regressions. This occurs at the point $R = \rho$, or equivalently at the point of zero excess returns relative to the risk free rate. The second is the OMD crossing point. Traditional CAPM testing methods set up the null hypothesis:

$$H_0 : \ r - \rho = \alpha + \beta(R - \rho) + \varepsilon \ ,$$

and proceed to test whether $\alpha = 0$, jointly or across all securities. This amounts to testing whether Jensen's alpha is zero. This procedure suffers from a lack of model specification under any alternative hypothesis. If H_0 is not correct, because the true theoretical regression is not linear, then the fitted statistics for the above equation are meaningless and reveal little or no information. Indeed, if the security being tested is an equity fund, it is well known that good market timers can produce apparently negative values of Jensen's alpha (Dybvig and Ross 1985).

The OMD pencil diagram avoids this sort of difficulty, as it is essentially non parametric in nature. Even if the true theoretical regression is non linear, the diagram will have meaning and one can proceed to examine the crossing points for consistency with what is predicted by the CAPM model. Thus the non parametric character of the OMD methodology is an advantage, as is the cumulation of regression errors implied by the running mean nature of the OMD calculation.

If a presumed CAPM equilibrium is established, characterised by a crossing point P_* , then $t(P_*) \neq 0$ for a given security could be taken to indicate a historical anomaly or inefficiency, favourable to a long position if $t(P_*) > 0$, to a short position if $t(P_*) < 0$.

Some operational aspects

(a) In practice, the risk free rate or zero beta rate will itself vary over time and it is more convenient to use excess returns, in the form $\tilde{r} = r - \rho$, $\tilde{R} = R - \rho$, $\tilde{P} = P - \rho$.The theoretical OMD schedule in CAPM is now $t(\tilde{P}) = -(1\text{-}\beta)\mu(\tilde{P})$ and the cross over point is defined by $\mu(\tilde{P}_*) = 0$. The cross over point $\tilde{P}_*$ can be greater than or less than zero depending upon whether the running mean of the excess returns, when evaluated at the mean excess return $\mu_{\tilde{R}}$, is less than or greater than zero, respectively.

(b) It is desirable to provide one or two standard deviation bands on either side of the fitted OMD schedule, to indicate the sensitivity of its position to sampling error. In Bowden (2000 p208) it is suggested using the conditional residual error based on $s(\hat{t}_j) = \hat{\sigma}_e/\sqrt{n_j}$,where $\hat{\sigma}_e$ is the fitted equation standard error in the regression $e(R)$ of r on R, and t_j is the jth value of the OMD, i.e. position n_j in the ascending ordering of the R values. An upper bound for σ_e for the true theoretical regression could in practice be obtained by choosing $e(R)$ to be linear. Most interest will usually centre on the crossing point and here n_j is interpreted as the number of R observations less than or equal to the crossing point.

Illustration

The above methods were applied to a time series of 83 monthly returns on a set of 15 larger stocks quoted on the New Zealand stock exchange, chosen from the period Jan. 89 to Nov. 95, a period of relative stability in the definition and composition of the index and the stocks themselves. An earlier study by the author (Bowden 1998) using time series methods had suggested that a reasonably stable CAPM might have existed existed over this period, so a conclusion of this kind can be checked out by using OMD technology.

The running mean function for the market is graphed in figure 6. The mean monthly excess return for the market index (NZSE-40) over the period was 0.001265 (i.e. 0.013%), and at this point the value of the running mean schedule for the market excess return was aproximately -0.044. Proposition 1 adapted for the excess returns data format therefore indicates that for CAPM to hold, all the security OMD schedules should cross the horizontal axis together, and the crossing point should be greater than zero.

In the event, 5 of the 15 stocks appeared to obey this requirement, crossing at a monthly return of close to 0.05. Figure 7 illustrates with the best 4 (the fifth, Brierley Investments crossed at 0.02 and is not illustrated). Figures in brackets asociated with the names of the stocks are the estimated beta values, with the asterisk denoting significance at 5%. Standard deviation bounds for all these stocks evaluated at R=0.05 were very small, of the order of 0.007-0.009. Another group of 4 stocks crossed roughly together but at monthly returns between -0.02 and -0.075. The reminder showed unclear or anomalous behaviour. Indeed, two stocks did not cross the horizontal axis at all, suggesting market inefficiencies, and the crossing point for a third was clearly only a matter of higher sampling variability in the low return area. These three stocks are graphed as figure 8. Nor did the one standard deviation bands for these stocks span the horizontal axis. Overall, these findings did not support the hypothesis that a New Zealand CAPM existed over the period.

Figure 6 Running mean graph for the NZ market index

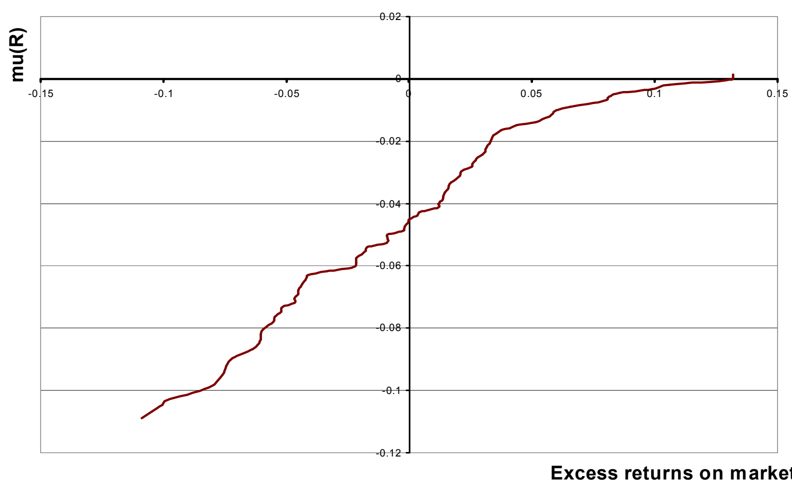

Figure 7 Some CAPM consistent plots

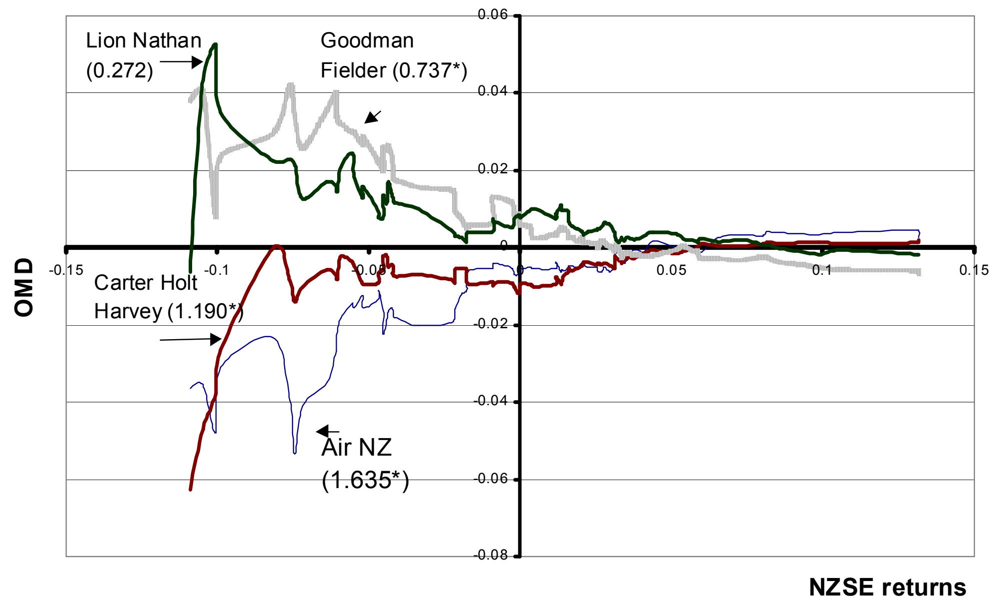

Figure 8 Anomalies in NZ OMD excess returns

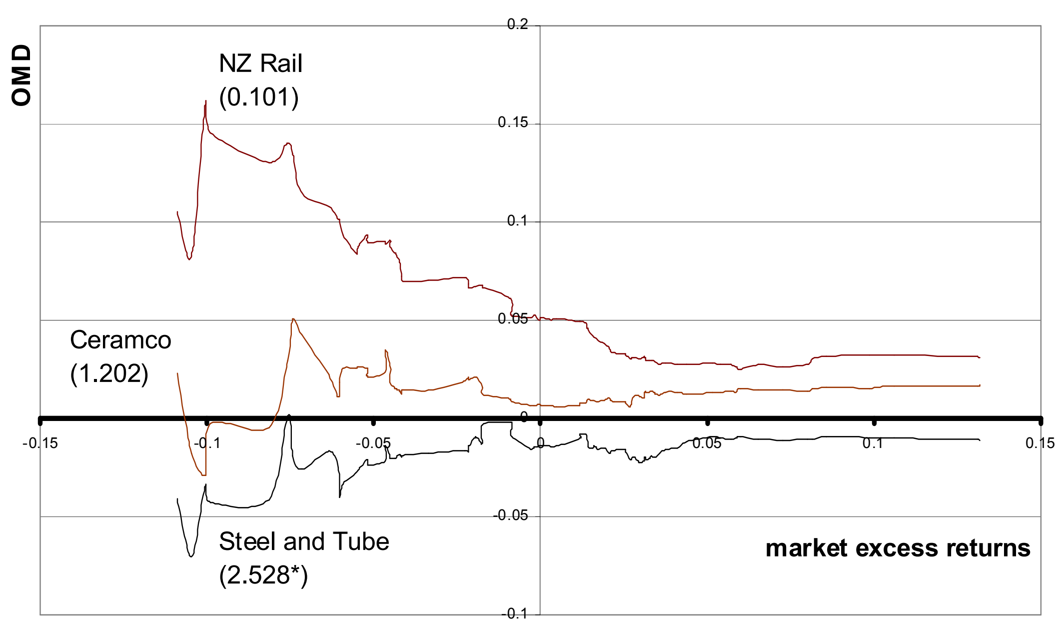

V Concluding remarks

Fund performance comparisons only rarely turn up cases where one fund dominates its benchmark according to classical stochastic dominance measures. One exception was provided by figure 1 of section II, where the NZ OMD graph lay wholly beneath the horizontal axis. In most cases, however, one will not find that the given fund is dominant over the benchmark. Where the OMD graph can then help is in determining whether a dominance region exists for particular types of investor, according to their degree of risk aversion. The results should then be squared up against the professed purpose of the fund. For example, if the fund advertised itself as a growth fund, one would not be too concerned about negative values of the OMD graph for low values of the benchmark, but would be if the negative zone was for high values. Conversely, a good income type fund whould be expected to have a negatively sloping OMD curve.

A second area in which OMD technology can potentially help is in testing for individual departures from presumed CAPM relationships. In this respect, conventional CAPM testing, such as the Fama -French methodology, is inevitably aggregative in nature. It is possible to accept a null hypothesis that a CAPM applies even when some stocks do not obey such a pricing model. In other words, conventional aggregative tests cannot examine whether an individual security is itelf in CAPM relationship with the chosen market benchmark. As the illustration of section IV showed, the OMD technology can help on the individual level. Even if a CAPM relationship held for most of the other stocks, the fact that a particular stock did not adhere might be taken to indicate that a pricing inefficiency held. The stock is either underpriced or underpriced, according to the position of its OMD curve relative to the others.Thus the OMD technology can help to identify pricing anomalies, in a way that exploits the extra generality of non parametric methods, which are relatively unaffected by the maintained hypotheses characteristic of parametric methods.

Perhaps the most interesting theoretical finding is the extent to which portfolio decisions and risk premiums can be 'deconstructed' using the utility generators, even up to the level of an entire capital market equilibrium. Interpersonal comparisons of risk appetites can be indexed by the focal points of the representative gnomes for each investor. It may be that the utility generator calculus extends to wider contexts of decision making under risk in other areas of risk and insurance, incorporating non marginal optimality conditions in place of the marginal ones used in context such as portfolio analysis[8]. One might then be able to generate utility functions of any desired risk appetite by varying the focal weights $w(P)$, and by doing so lessen the dependence on the simple but restrictive families of utility functions in common use for such purposes. This remains a potentially fruitful area for further research.

Footnotes

[1] It should be stressed that the OMD techniques do not live or die by the reported demise of CAPM (as in Fama and French 1992,1993). Indeed the empirical results of section IV show that the OMD technology supports a 'beta is dead downunder' thesis from quite another direction. Apart from CAPM, however, the OMD methods are applicable to any situation where a natural benchmark exists, for example in fixed interest debt portfolio benchmarking.

[2] In one day cricket, the 'worm' is a moving picture of the run rate (progressive average) against a straight line benchmark created by the previous performance of the opposition over their innings. American readers who like musicals but not cricket will find a similar idea in the Rodgers and Hammerstein song 'Anything you can do, I can do better,' from *Annie Get Your Gun*.

[3]The error bands in this example are derived by assuming that $e(R)$ is linear, deriving the

equation standard error σ and replaing this by $\sigma/\sqrt{\#P}$ to capture the effect of the running mean on sampling variation. The linearity assumption will amount to an overestimate of the true equation error STD σunless the regression function happens to be linear, eg as in a CAPM. Hence the error bands are conservative.See Bowden (2000) for discussion of sampling variability issues.

[4]For related measures see Chen and Knez (1996), Grinblatt and Titman (1989).

[5] This can be shown by using Price's lemma for functions of normally distributed random variables (see Bowden 1997).

[6]Alternatively, one could define the Pratt-Arrow type risk premium $\theta > 0$ by the equality $E[U_P(R)] = U_P(\mu_R - \theta)$. For this to be satisfied it must be that $\theta > \mu_R - P$, in which case the equality can be written as $\theta = \mu_R - P - E[(R-P)SF(P-R)]$, where $SF(\bullet)$ is the unit step function: $SF(x) = 0$ $if\, x \leq 0, = 1$ $if\, x > 0$. Differentiating, it follows that $d\theta/dP = -(1 - F(P) < 0$.

[7] This can be shown either from first principles or else by using Price's lemma for functions of normal random variables (Bowden 1997), utilising $SF'(x) = \delta(x)$, the Dirac delta function.

[8] The GR risk premium is an example of a marginal construct, as it refers to the marginal effect of adding a unit of security r to an existing holding of the benchmark R. In mean variance space it involves points of tangency - see for example figure B1 Appendix B below. On the other hand, the Pratt-Arrow style risk premium refers to the effect of replacing return R by a certainty equivalent $\mu_R - \theta$, so one solves the implicit equation $E[U(R)] = U(\mu_R - \theta)$. This does not entail a marginal relationship; in mean variance diagrammatic terms, one reads off θ from intersection of the relevant utility indifference curve with the vertical axis, rather than from a point of tangency.

Appendices

Appendix A OMD spreadsheet

Table A1 shows how the calculation is done using a simple spreadsheet. The far right hand column (#6) tabulates the values of the ordered mean difference function. The OMD schedule is a graph of column 6 against column 2.

	Table 1A Sample OMD calculation					
r values	Rvalues	reordered by R			running mean difference	
		r	R	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	

Appendix B Illustrating the generalised Rubinstein risk premium

In the mean variance diagram below, investor B is more risk averse with a steeper trade off between σ and μ. At portfolio of return R, investor A will be in personal equilibrium with a GR risk premium of θ_A. However with an effective risk free rate of $\mu_R - \theta_A$, investor B will want to combine portfolio R with more of the risk free asset to arrive at portfolio R$'$. Hence at A's portfolio R, we must have $t_{UB} > 0$. Only for a risk free rate of $\mu_R - \theta_B$ will investor B be in personal equilibrium. So $\theta_B > \theta_A$, capturing their relative risk preferences.

Figure B1 Comparative GR risk premiums

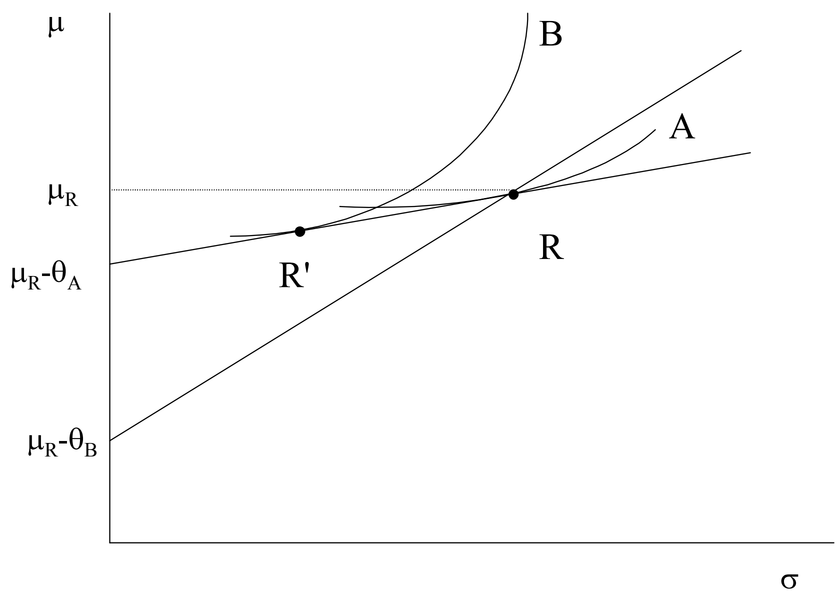

Appendix C Basic results of generalised function theory: differentiation and integration

The Heaviside calculus referred to in section IV involves use of the Dirac delta function and step functions, and associated constructs. The step function is generally taken as

$$H(x-a) \quad = \quad \begin{cases} 1 & if \;\; x > a \\ 1/2 & if \;\; x = a \\ 0 & if \;\; x < a \end{cases} \quad .$$

However, the same results will obtain if the definition used in the text is employed, wherein $SF(x-a) = 0$ for $x \leq a$ and unity otherwise. The step function has a formal derivative

$$\delta(x-a) \quad = \quad H'(x-a) \quad .$$

The sense in which the derivative is taken is that of the Temple-Schwartz theory of distributions or generalised functions (see Lighthill (1958)). The step function may be regarded as the limit of a suitable probability distribution function (e.g. the normal) centred at $x = a$ as the variance becomes small, and the Dirac delta function can be regarded as the corresponding density, ultimately taking the form of a spike of infinite height at x =a . Moreover the delta function can itself be formally differentiated as $\delta'(x), \delta''(x)$, and so forth, again in the limiting sense just described. The most important property is the 'filtering property': for suitably 'smooth' functions f,

$$\int_{-\infty}^{\infty} H(x-a)f(x)dx = \int_{-\infty}^{\infty} f(x)dx;\ \int_{-\infty}^{\infty} \delta(x-a)f(x) = f(a);\ \int_{-\infty}^{\infty} \delta^{(i)}(x-a)f(x)dx = (-1)^i f^{(i)}(a)$$

In the context of Proposition 2, $U_P(R) = (R-P)SF(P-R)$. So $U_P'(R) = SF(P-R) - (R-P)\delta(P-R)$ and $U_P''(R) = -2\delta(P-R) + (R-P)\delta'((P-R)$. Obtaining $E[U_P'(R)]$ and $E[U_P''(R)]$ is then a straightforward application of the above rules for integrations involving the generalised functions H (here replaced by SF) and the Dirac delta.

References

Bowden, R.J. 1995 Production organisation and risk control when market instruments are available. *Management Science*, 41, 1073-1082.

Bowden, R.J. (1998) CAPM and availability bias: when is 'near enough', good enough? Mimeo, School of Economics and Finance, Victoria University of Wellington.

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

Bowden, R.J. (2001) National performance worms over the 90's: Too far downunder? *ASX Perspectives*: *Journal of the Australian Stock Exchange*, 1Q 2001, 23-27.

Bowden, R.J. (2002) Utility spanning and the ordered mean difference envelope, School of Economics and Finance, Victoria University of Wellington.

Chen, Z. and P.J. Knez (1996) Portfolio performance measures: Theory and application. *Review of Financial Studies*, 9, 511-55.

Duffie, D. (1992) *Dynamic Asset Pricing Theory* New Jersey: Princeton University Press.

Dybvig, P.H. and S.A. Ross (1985) Differential information and performance measurement using a security market line. *Journal of Finance*, XL, 1749-78.

Fama, E.F. and K.R. French (1992) The cross section of expected stock returns. *Journal of Finance*, 47, 427-66.

Fama, E.F. and K.R. French (1993) Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 17, 3-56.

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

Grinblatt, M. and S. Titman (1989) Portfolio performance evaluation: Old issues and new insights. *Review of Financial Studies*, 2 , 393-421.

Lighthill, M.J. (1959) *An Introduction to Fourier Analysis and Generalised Functions*. New York: Academic Press.

Merton, R.C. (1981) On market timing and investment performance. I: An equilibrium theory of value for market forecasts. *Journal of Business*, 54, 363-406.

Post, G.T. (2002) Testing for stochastic dominance with diversification possibilities. Erasmus University, Rotterdam; forthcoming, *Journal of Finance*.

Rubinstein, M. (1973) A comparative statics analysis of risk premiums. *Journal of Business*, 39, 119-38.

Rubinstein, M. (1974) An aggregation theorem for securities markets. *Journal of Financial Economics*, 1, 225-244.

Varian, H. R. (1992). *Microeconomic Analysis* 3rd edition. New York: Norton.