February 2003

Portfolio risk profiling via the bootstrap: holding period analysis for small cap versus large cap stocks

by Roger Bowden*

Abstract

Strategic modes of portfolio management commonly focus on longer period holding returns, with a number of associated recommendations, amounting to 'riding the risk premium'. The ensuing data problems in examining such a claim can be handled by bootstrapping the available set of returns, which avoids specific distributional assumptions such as log normality. Using ordered mean difference techniques, it is shown that the holding period should be considered in conjunction with risk preferences. Depending upon the latter, small cap stocks can be either beneficial or detrimental relative to a large cap benchmark.

Keywords: Bootstrapping, holding period returns, ordered mean difference, risk profiling, small cap, strategic portfolios

* The author is Professor of Economics and Finance at the Victoria University of Wellington, PO Box 600, Wellington, New Zealand, ph * 64 4 472 2984, fax 4983, email roger.Bowden@vuw.ac.nz or roger.Bowden@kiwicap.co.nz . He is grateful to Alex Fowler of Strategic Asset Management Ltd and Ibbotson Associates for discussions and data.

I Introduction

Holding period returns refer to the geometric average return over a specific horizon. In portfolio analysis and performance measurement, the typical horizons range from 1 to 10 years. The empirical analysis of longer term holding horizons will always be handicapped by the effective reduction in the number of independent observations. For instance, with the historical Ibbotson data used in the present study, 76 annual observations become reduced to just 7 non overlapped 10 year holding periods. A variety of methods are used to overcome this problem, depending upon the context. One is to use overlapped returns, which amounts to using a moving ten year frame, but destroys the statistical independence of the ten year returns, even given a maintained hypothesis of efficient markets. Another is to assume that the geometric average returns constitute a sampling statistic and use the formulas for the mean and variance for the associated sampling distribution. Drawbacks to this are first, the need to assume or pretest for a specific returns distribution such as log normality; and second, the rather limited parametric analysis that results, being effectively confined to mean variance analysis, with its well known shortcomings. A more complete risk analysis needs more than just two summary portfolio statistics, and the ordered mean difference techniques used in the present analysis are such an instance.

An alternative line of attack is simulation based on the available data. In the case of one dimensional distributions, one can use straightforward montecarlo techniques to generate further independent returns numbers from the computed empirical distribution function. But portfolio analysis is always at least two dimensional. For example, one may be interested in adding another asset class to a given benchmark, and in this case a bivariate distribution is involved, with observations on two variates in the same time period. . It is suggested that bootstrapping is a more natural way to handle such an exigency. For the basic techniques, see e.g. Efron (1982), Efron and Tibshirani (1993), or Davison and Hindley (1997). Bootstrapping has been widely applied to analyse sampling distributions in statistics and econometrics, and also to a number of substantive contexts in the social and economic sciences, such as privacy protection of economic data (Bowden 1992). In the present contribution, it is applied to a fresh substantive context, namely the problem of extending returns data to cover longer holding periods.

Turning to the domain of application, the distinction between longer and shorter holding periods is more than just a matter of the financial arithmetic of returns. Strategic modes of portfolio management commonly focus on longer period holding returns, with a number of associated recommendations. One such amounts to 'riding the risk premium'. A well known finding (starting with Banz 1981), is that certain dimensions such as size and book to market value, attract a pricing penalty in the market, leading to higher risk premiums; see also Fama and French (1992,1995), Shefrin and Statman (1995), Dennis at al (1995), or Kothari et al (1999). The size effect is the focus of the present paper, and the existence of such an effect is built into the recommendations associated with the Fama - French model of security returns now widely used by portfolio managers. According to this, if clients are interested only in the longer term, then their strategic portfolio manager can reap the rewards of patience, accepting the odd small corporate collapse, but nevertheless enjoying the long term higher average returns from investing in smallcap stocks, as at least a significant asset class in the portfolio. One might call such a strategy 'riding the risk premium'.

On the other hand, it is important to remember that long term investors have risk preferences as well. An investor with a 10 year horizon retains a utility function defined on the 10 year return, just as much as does an investor with a one period horizon. For instance, baby boomers might have a target retirement wealth and feel considerable anxiety at the prospect of a shortfall at the end of the proposed 10 year (etc.) holding period. So the empirical risk premium in and of itself cannot be the end of the story. Holding period returns must be considered jointly along with investor preferences. Taking the case of the historical U.S. small cap premium, the importance of risk profiling is addressed. Methodologically, risk profiling can be handled with the use of ordered mean difference methods (OMD). Originally conceived within a fund performance measurement framework (Bowden 2000), the applicability of OMD methodology has since been extended to capital market equilibrium testing (Bowden 2002), stochastic dominance and portfolio efficiency (Post 2002, Bowden 2003), and portfolio enhancement (Bowden 2002). One insight is that investors can be considered as a made up of a set of more elementary investors ('gnomes', say), each of whom has a simple sublinear two segment utility function. In the present paper, OMD methods are adapted to handle the issue of holding period returns under different attitudes towards risk. Because these methods do not make limiting assumptions about investor preferences (e.g. mean variance), or require specific distributional assumptions or pre testing, the data set will need to be extended for longer holding periods. This is where the bootstrapping is useful.

The scheme of the paper is as follows. Section II is methodological. It commences with a review of the elements of OMD theory and practice, as they apply to the present context, including the idea of investor surplus. Bootstrapping methods are then briefly exposited as a solution to the data exigency. Section II uses these methods to analyse the contribution of smallcap to large cap stocks, taking the latter as a benchmark. Section IV summarises the conclusions.

II Ordered mean difference and risk profiling

21 Portfolio efficiency via the OMD

Ordered mean difference techniques are in the first instance a diagnostic for portfolio efficiency. Suppose we have a proposed portfolio of return R (for brevity, henceforth just 'portfolio R'), and an asset of return *r* (or just 'security *r*') to be considered in conjunction with that portfolio, perhaps as a potential addition, or to see whether its existing allocation is correct. Imagine also, for the moment, that one knew exactly the manager's utility function for returns; denote this by U(R). Thus expected utility would be E[U(R)] and expected marginal utility is E[U'(R)]. The objective would be to choose a portfolio to maximize the expected utility. More extended justifications of the formulas used can be found in Bowden (2000,2002a).

Comparing the given security r with the benchmark R, consider the magnitude defined by

$$\tau_u = E[(r - R)\pi(R)]; \quad where\ \pi(R) = \frac{U'(R)}{E[U'(R))]}. \tag{1}$$

Evidently this consists in a weighting of the return differences by a weight function π(R), which gives extra weight to states of the world in which marginal utility of the benchmark - the value of an extra dollar from the base portfolio – is higher. It is a risk adjusted excess return, called the equivalent margin, or the investor surplus attached to *r* relative to benchmark RA value of 2%, for example, indicates that the investor would be willing to suffer a 2% penalty on the return of security *r* before it is offlisted from the benchmark portfolio. Asset *r* will be in equilibrium relative to portfolio *R* if $\tau_u = 0$. If $\tau_u > 0$, this will be a signal to increase one's holding of *r* and if $\tau_u < 0$, this should tell the investor to diminish holdings of *r*, or to go short if it is not already in the portfolio.

The usual problem is that one does not know the utility function U. Indeed only in very special cases would this ever be true. OMD analysis provides a way of overcoming such a difficulty. It breaks the problem down into the analysis of a special utility function portrayed in figure 1 below. This has just two linear segments, and depends on a single parameter P, which has the same dimensions as the portfolio return. We could write this utility function as

$$U_P(R) = \min(R - P, 0).$$

In spite of its odd appearance, the function $U_P(R)$ is a genuine risk averse utility function in its own right, associated with achieving a target return at R = P. The point P is called the node or focus. Increasing P (as in the diagram) indicates that the associated utility function $U_P(R)$ has diminishing risk aversion, for a fixed portfolio return distribution. The functions $U_P(R)$ are collectively called the utility generators, so named because for purposes of expected utility maximization, one can show that any utility function U(R) can always be taken to be a weighted sum of the utility generators with the weight assigned to the generator at P being proportional to the curvature ($-U''(P)$) of the given utility function U. This is called the 'spanning property'.

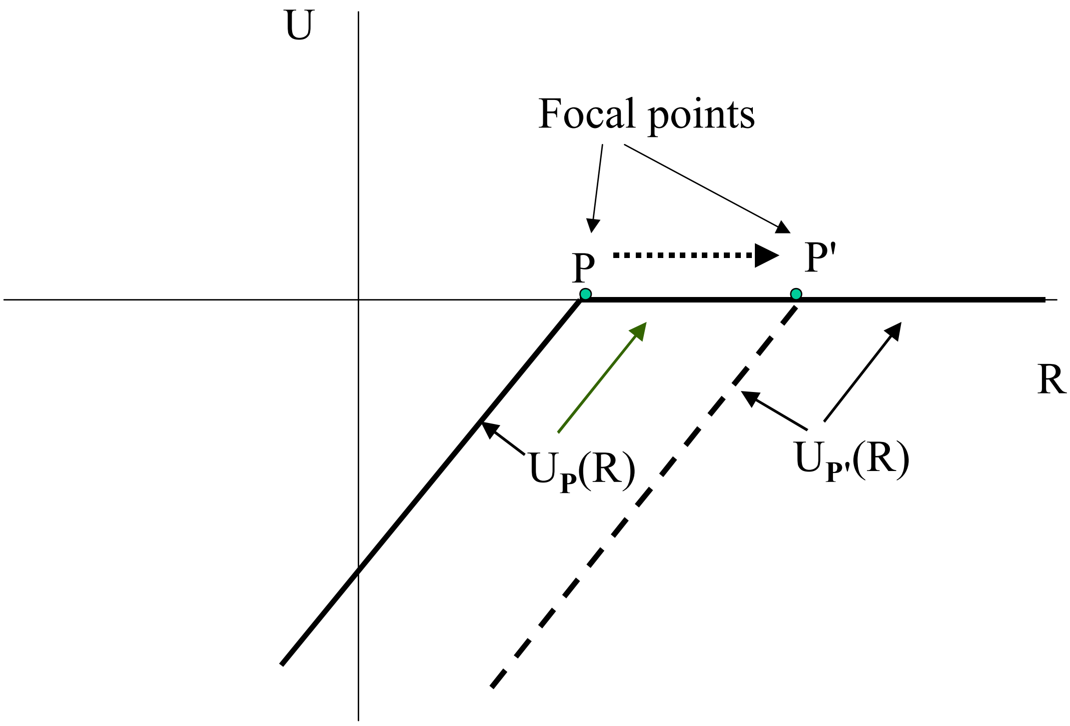

Figure 1: Utility generators

Although the generators are not strictly differentiable at R = P, an equivalent margin (or investor surplus) measure can be worked out for them as well, by replacing the U'(R) in formula (1) by U_P'(R), with a suitable interpretation of the latter derived from generalized distribution theory . The sample estimate for this is given by

$$\tau(P) = \frac{1}{\# R_t \leq P} \sum_{R_t \leq P} (R_t - r_t) \quad . \tag{2}$$

This is interpreted and computed as follows. First reorder the observations by ascending value of the benchmark return R. Then compute the running mean of the differences r-R for the reordered observations, up to and including R values that are less than or equal to the given number P.

Table 1 is an illustration of how this can be done with an Excel spreadsheet. Columns 1 and 2 are the natural returns data. Columns 3 and 4 are the data re-ordered according to ascending values of the benchmark (in Excel, use *tools/sort*). Column 5 is the difference r-R . Column 6 takes the running mean over increasing values of P. The schedule that plots the running means against P is called the ordered mean difference schedule.

r values	R values	reordered by R r	 R	r-R	running mean difference OMD	
0.022	0.011	-0.018	-0.045	0.027	0.027	
0.048	0.047	-0.021	-0.030	0.009	0.018	
0.032	0.018	0.004	-0.005	0.009	0.015	
0.004	-0.005	0.017	0.001	0.016	0.015	
0.027	0.008	0.027	0.008	0.019	0.016	
-0.021	-0.030	0.022	0.011	0.011	0.015	
-0.018	-0.045	0.032	0.018	0.014	0.015	
0.017	0.001	0.048	0.047	0.001	0.013	P

Table 1: Calculation of the OMD schedule

Risk profiling and the OMD

The OMD schedule has a number of uses and linkages. Among other things it offers an easy way of testing for stochastic dominance. However the interpretations of particular interest in the present context concern risk profiling. This arises from the spanning property of the utility generators. In making optimal portfolio choices, every investor acts as though he or she has operated off just one member of the generator family. In other words, there must exist some focal number P* such that the portfolio I choose by using $U_{P*}(R)$ is just the same as the one I would use from my true underlying utility function U(R).

From this it would follow that if my investor surplus is positive for every value of P - so that the OMD schedule lies wholly above the horizontal axis - then extra units of r should be invested, no matter who I am. More typically, one finds that the OMD schedule is monotonic either upwards or downwards, and intersects the horizontal axis. Suppose we have a situation such as figure 2 below, where the positive slope means that security r is aggressive relative to R. Individual 1 is more risk averse than individual 2, so that the representative generator focus $P^*_1 < P^*_2$ as indicated. It will pay individual 2 to go long(er) in security *r* as against the given benchmark R; but the opposite is true for individual 2, who should shed security *r*. Only for an investor whose representative P* is precisely equal to the crossing point A will security *r* be in personal portfolio equilibrium with the benchmark R.

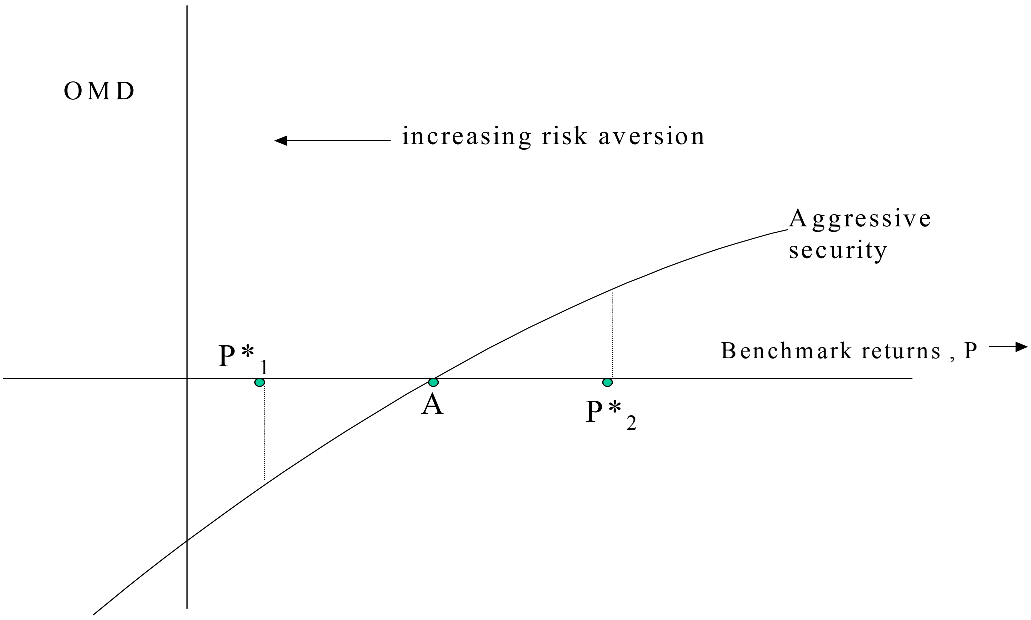

Figure 2: Risk profiling with the OMD schedule

For complete judgments about risk profiling one would like to locate the representative utility generator focus P* for a given investor.

(a) Suppose we are prepared to assume the existence of a market equilibrium. In this case even though investors may differ in their risk preferences, the market itself acts as though it was an investor, and has its own representative utility generator with node P^*_m say. Given a risk free rate this can be readily calculated. First we compute the running mean schedule of R itself. Analogous to formula (2) above this is given by:

$$\mu(P) = \frac{1}{\# R_t \leq P} \sum_{R_t \leq P} R_t \quad . \tag{3}$$

Note that $\mu(P) \to \mu_R$ as $P \to \infty$, where μ_R is the mean of R. Next, we locate the number P such that $\mu(P) = \rho$ where ρ is the given risk free rate. This is the required market P^*_m . Figure 3 illustrates. An equivalent procedure is to plot the running mean of the excess returns $R\text{-}\rho$ and locate the point where this crosses the horizontal axis. Figure 6 of section III does this. The latter procedure is better where the risk free rate can itself change over time, as of course does happen.

This approach is valuable because it enables us to separate out investors according to whether they would be classified as either more or less risk averse, than the market as a whole. If more risk averse, their representative generator focus P* will lie to the left of P^*_m and to the right if less risk averse.

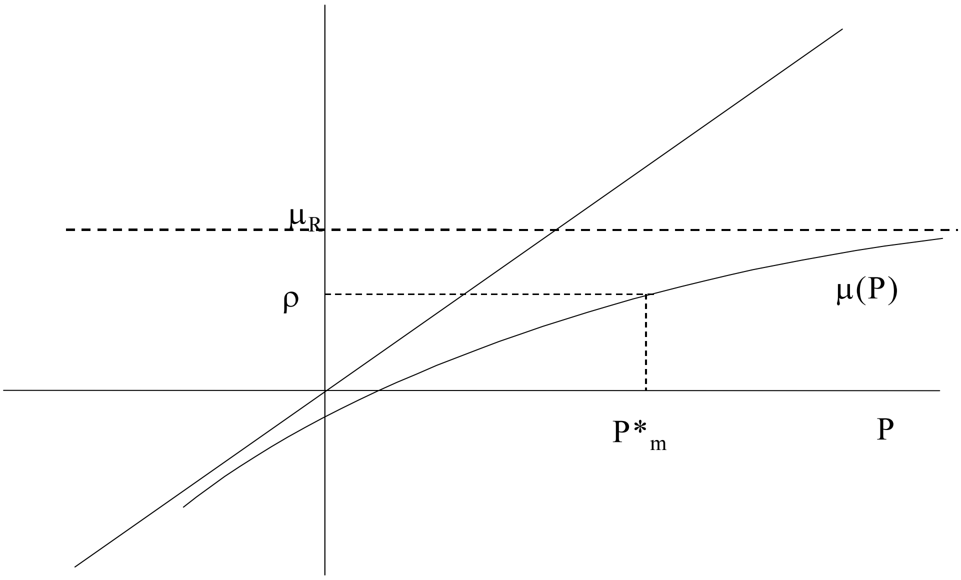

Figure 3: Locating the market representative generator focus.

(b) If there is no underlying market equilibrium, revelation experiments can in principle be designed that would enable the investor to decide just where he or she is located along the P axis. Short of such experiments, there is no way that the outsider can locate likely average values of P. All one can say in the case of figure 2 is that if the investor was very risk averse, security *r* should be divested relative to the benchmark portfolio. If the investor is strikingly non risk averse, e.g. approaching risk neutrality (infinite P), then he or she should add more units of security *r* .

2.2 Holding period returns and bootstrapping

For a holding horizon of T years, the holding period return is conventionally defined as the geometric average of the annual returns:

$$(1+r_g)^T = \Pi_{t=1}^{T}(1+r_t)$$

Longer holding periods are characteristic of the strategic portfolio approach. Portfolio planning based on annual returns is more characteristic of a fully dynamic approach, in which the current wealth state of the portfolio, as well as a more active informational approach, help to condition the portfolio to be chosen afresh in every period. If the returns are independently distributed, as in an efficient market, then the holding period return will be less than the arithmetic average of the annual returns, and this tendency increases with the length of the holding horizon T. The effect arises from the concavity of the geometric mean.

As earlier remarked, data availability become problematical with longer holding horizons. However, the original data set can be extended by bootstrapping the available observations. The simplest and most commonly used version (simple bootstrapping) draws further observations from the same empirical distribution as the original data. Technically, this is resampling without replacement from the original sample. In other words, given a sample r_1, r_2 ….r_{76} as in the annual Ibbotson data , assign each data point the probability 1/76, and draw a random number from the range 1-76. In this way one can arbitrarily extend the number of apparent observations; we used 10,000 in the work reported below. At first sight, this looks like getting something for nothing, and to be sure, the numerical values remain those of the original data set. However, sample statistics derived from the bootstrapped data will retain validity. For OMD schedules, all that will happen is that the same integer number t between 1-76, and hence the corresponding data pair (r_t,R_t), may be drawn repeatedly, giving rise to locally flat zones of the OMD schedule (see figure 5 below). The overall shape will nonetheless mirror that from a genuine larger sample.

III Application: Investor horizons and smallcap stocks

The object in the comparisons that follow is to see whether small caps add value to a base or benchmark portfolio of large cap stocks; how this might depend upon risk profiling; and whether the latter is dependent upon the holding period. The original data is the Ibbotson - Sinquefield (IS) annual return series, 1926-1971, sourced from the Ibbotson Associates Yearbook (2002). The data are generated by 10,000 random drawings from the sample set of integers, in effect generating a bootstrapped sample of 10,000 from the original 76 one period returns. Following this, 2,5, and 10 year holding period returns are computed. So one now has 10,000 one year returns, together with 5000, two year, 2000 five year and 1000 ten year. Computations were executed with Visual Fortran 6, on an ASUS laptop, and proved undemanding in terms of execution time. The same computations could alternatively have been done with a basic Excel spreadsheet - and a lot of scrolling.

Figure 4 summarises the respective one period descriptive statistics for each of the two series in isolation, i.e. the marginal data densities. The smallcap returns are noticeably leptokurtotic relative to the large cap, with both a higher mean and a higher variance. The latter underpin the usual strategic recommendation to 'ride the risk premium' over the longer term by overweighting in smallcap stocks, relative to a broad value weighted market index. It is assumed in doing so that one can overcome the survival problem, wherein smallcap stocks carry a higher risk premium simply because of the higher risk of bankruptcy. However, there are other indications that might give the manager pause, such as the differential kurtosis, and how all such descriptive statistics might be affected by a nominated holding period.

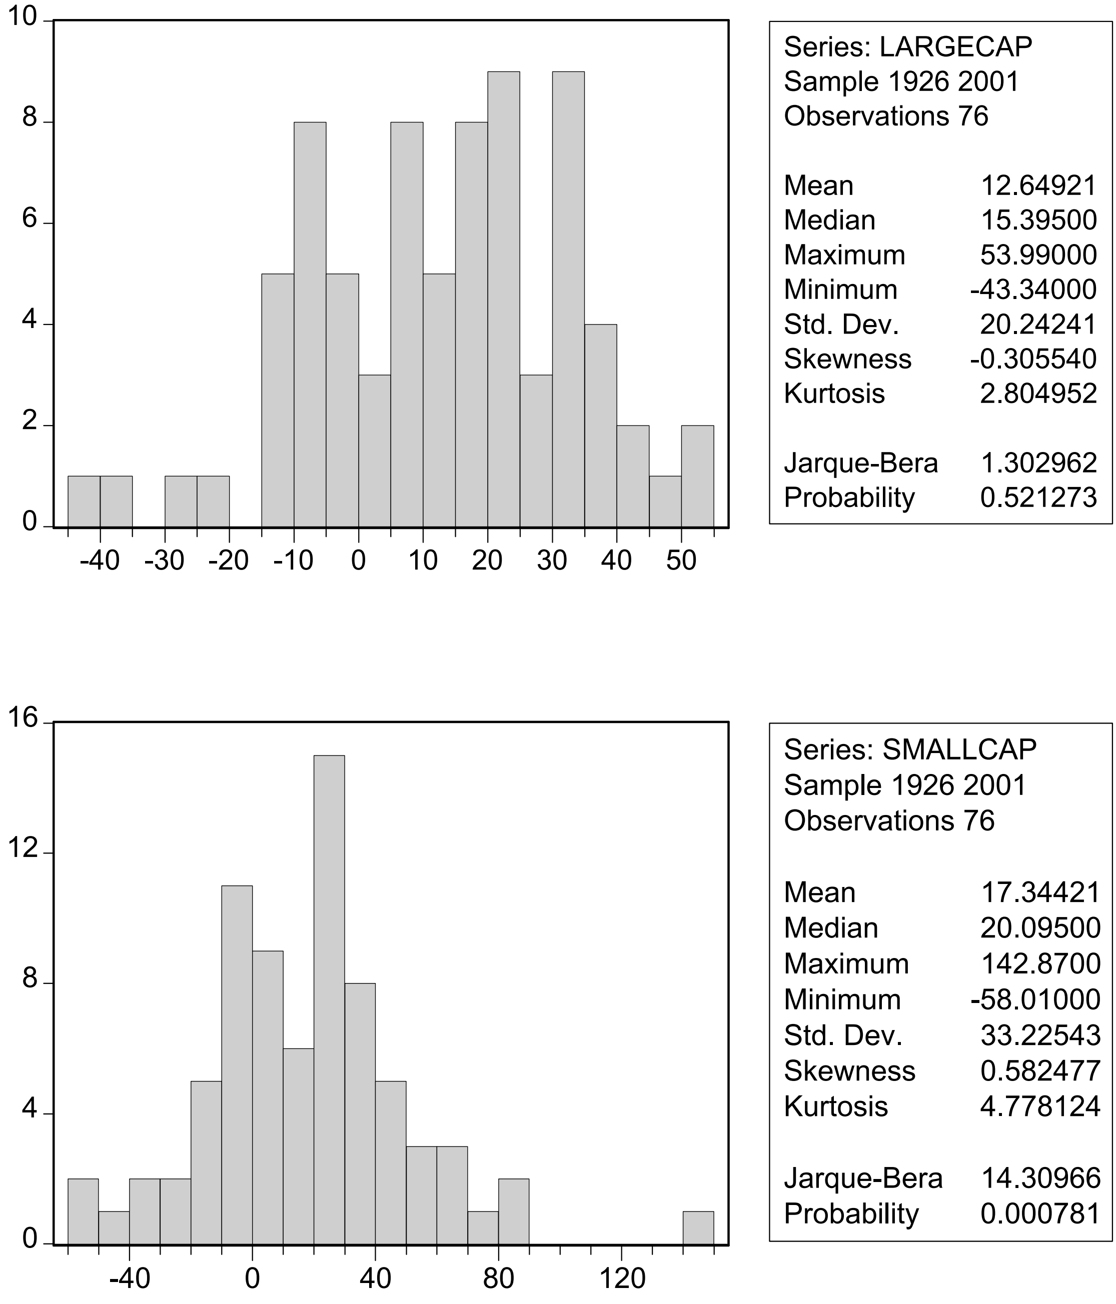

Figure 4: Descriptive statistics, marginal distributions

Figure 5 depicts on the same diagram the OMD schedules, smallcap against large cap, for all holding periods. As one would expect from the concavity properties of holding period returns, the plots for longer holding periods lie asymptotically below those for the one period returns. A risk neutral investor would therefore unequivocally prefer the one period returns. However there is a point at which the longer term holding periods cross over the shorter. Mildly risk averse investors might therefore derive greater investor surplus from a ten year holding period. On the other hand, very risk averse investors should not consider adding smallcap stocks at all; for any holding period, the OMD turns negative.

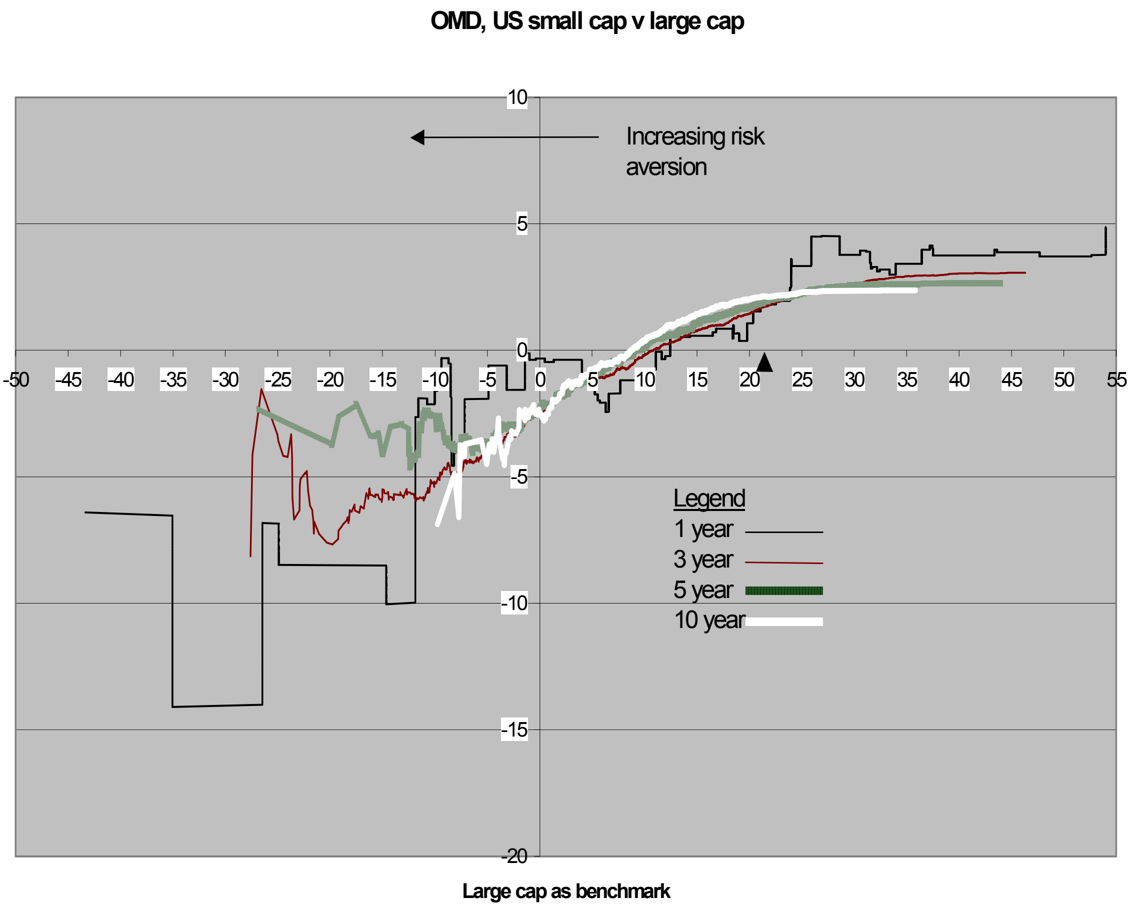

Figure 5: OMD schedules, small cap against large cap.

As mentioned in section II, we can calibrate the risk dimension by imagining a market equilibrium from the same set of data, and locating the value of $P=P^*_m$ at which the OMD of excess market returns crosses the horizontal axis. For this purpose we used the annual equity premium returns from the Ibbotson yearbook. Figure 6 illustrates.

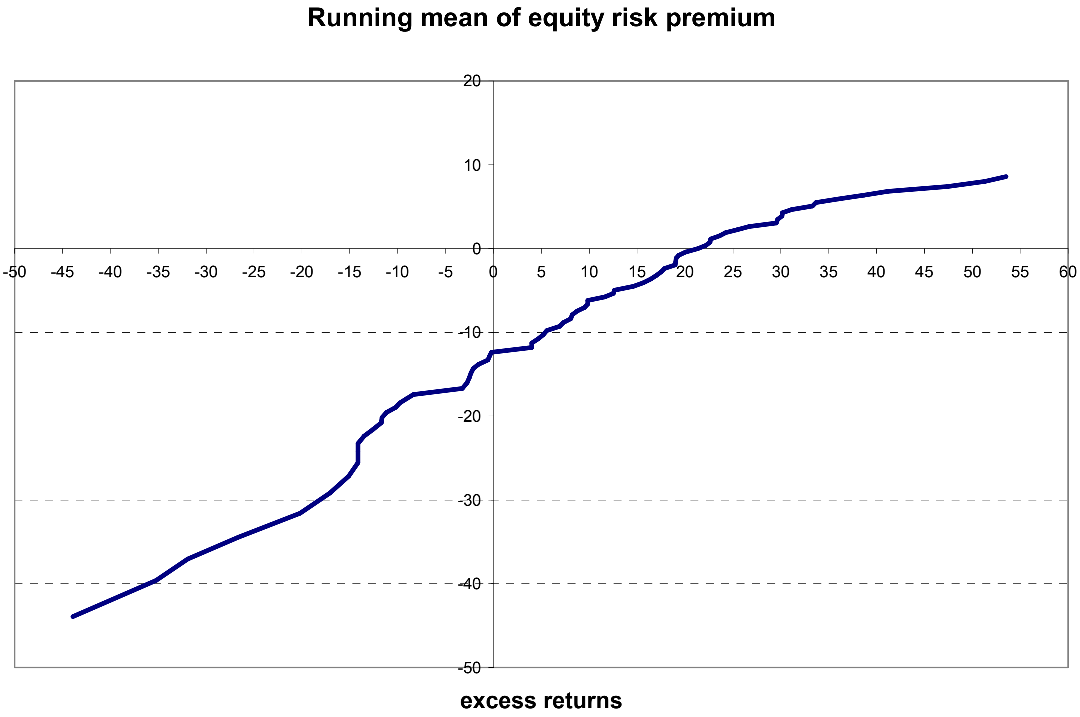

Figure 6: Locating the market P^*_m

The indicative market P^*_m is just over 20%, and this is located as the arrowhead in the previous figure 5. At this point the OMD schedules for all holding periods are roughly the same. An implication would be that if the investor is no more or no less risk averse than the market as a whole, then the holding period is not of concern. All yield the same result, namely that smallcap add significant[1] value to a base portfolio of largecap. On the other hand, if the investor is just a little more risk averse than the market, then the surplus from longer holding periods is greater. If the investor is markedly more risk averse, then he or she should not hold smallcap stocks at all. Indeed, such an investor might well be better off with a portfolio of T bills or high grade commercial bills.

[1] Formal significance bands for OMD schedules can be found in Bowden (2000), but have not been inserted here. For 10,000 observations at P=20, the bands would be extremely narrow.

IV Concluding remarks

The conclusions may be summarised as follows:

(a) *Methodological*. Ordered mean difference techniques deal with pairwise relationships, and are based on underlying portfolio considerations. As such, they are naturally adapted to issues of value added in an investment context. They can be demanding in terms of data requirements where extended holding periods are to be evaluated. However, this problem can be partly solved by bootstrapping on the original one period observation set. This can be expected to work whenever there are enough observations to make a one period OMD analysis sensible, say $\geq$ 50, on a monthly or annual basis. A broad enough range will then exist such that bootstrapping will little distort the true underlying joint distribution function.

(b) *Target holding periods do matter*. They affect the joint distribution of measured returns, extending therefore to their risk~return tradeoffs. Longer holding periods should not necessarily be regarded as an analgesic against risk. An investor who proposes to hold for ten years has a risk averse utility function defined over the ten year holding period, just as does an investor with a one year rebalancing period. Our results indicate that mildly risk averse investors, relative to the market as a whole, can indeed do better with longer term horizons, but this disappears when the degree of risk aversion rises, to the extent that smallcap stocks should be avoided. Smallcap stocks are at their very best with less risk averse investors, and in this case the horizon should be short.

References

Banz, R.W. (1981) 'The relationship between returns and market value of common stocks', *Journal of Financial Economics*, 9, 3-18.

Bowden, R.J. and A-B Sim (1992) 'The privacy bootstrap', *Journal of Business and Economic Statistics*, 10, 337-345.

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

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

Bowden, R.J., (2002b) 'The ordered mean difference approach to portfolio enhancement', School of Economics and Finance, Victoria University of Wellington

Bowden, R.J. (2003) 'Utility spanning and the ordered mean difference envelope', revised, School of Economics and Finance, Victoria University of Wellington.

Davison, A.C., and D.V. Hindley (1997) *Bootstrap methods and their applications*, New York: Cambridge University Press.

Dennis, P., S. Perfect, K. Snow, and K. Wiles (1995) 'The effect of rebalancing on size and book to market ratio portfolio returns', *Financial Analysts Journal*, 3, 47-57.

Efron,B. (1982) *The jakcknife, the bootstrap and other resampling plans*, Philadelphia: Society for Industrial and Applied Mathematics (SIAM).

Efron, B. and R. Tibshirani (1993) *An introduction to the bootstrap*, New York: Chapman and Hall.

Fama, E.F. and K.R. French (1992) ‘The cross section of expected stock returns’, *Journal of Finance*, 47, 427-65.

Fama, E.F. and K.R. French (1995) ‘Size and book to market factors in earnings and returns’, *Journal of Finance*, 50,131-85.

Ibbotson, R.G., and R.A. Sinquefield (1989) *Stocks, Bonds, Bills, and Inflation, Historical Returns (1926-1987)*, Homewood. Il: Dow Jones Irwin.

Ibbotson Associates (2002) *Stocks, Bonds, Bills, and Inflation Yearbook*, Chicago, Il: Ibbotson Associates.

Kothari, S.P., J. Shanken and R.G. Sloan (1999) ‘Another look at the cross section of expected stock returns’, *Journal of Finance*, 50,185-224.

Post, T. (2002) ‘Testing for stochastic dominance with diversification possibilities’, Erasmus University, Rotterdam; forthcoming, *Journal of Finance.*

Shefrin, S.P. and M. Statman (1995) ‘Making sense of beta, size and book-to-market’, *Journal of Portfolio Management*, 21, 26 34