OPTIMISATION – FACT OR FANTASY?

KEEPING INDEX FUNDS ON TARGET

by NORMAN SINCLAIR

Optimised index funds may benefit fund managers and investors — but the extent to which they are 'better' than 'conventional' funds may be illusory.

Since the stockmarket crash in October 1987 portfolio managers have shown a renewed awareness of index funds as a possible investment strategy. The adoption of a passive investment strategy such as an index fund requires the acceptance, on the part of the portfolio manager, of a number of presumptions about performance ability

First, an index fund manager accepts the proposition that he has a low or zero comparative advantage in security selection or "picking winners". Second, the passive portfolio manager accepts that he cannot "time the market" and initiate an early restructuring of the portfolio prior to a "bull" or "bear" trend. If both of these presumptions are true for a particular manager, then the best performance that he can aspire to is the average return on the market, assuming that the portfolio under management is a reasonable proxy for the market index.

Although the role of an index fund is simply to match the financial performance of some underlying target index, the construction of such a portfolio is not a simple exercise. However, it need not be as complex as some portfolio managers believe.

One obvious way to achieve average performance is to invest in all of the shares in the underlying index in the same proportions as in the index itself. For very large indexes, this fully replicative approach is impractical and excessively expensive. To achieve the objective of matching a target index, a portfolio manager need not explicitly mirror the target index. Even prohibitively large indexes such as the Australian Stock Exchange Index can be "mimicked" by smaller, statistically optimised portfolios. The major benefit of a passive investment strategy tied to an optimised index fund is that operating costs could be as low as one-ninth of those of operating conventionally managed funds, since there are fewer shares to monitor and lower transaction costs.[1]

However, before a portfolio manager commits the portfolio to an index fund, some basic questions need to be considered. The most obvious of these is: How close is the current portfolio to an "index fund"? Since an optimised portfolio may combine the same shares but in different proportions, does it really matter over a given investment horizon? In other words , can an *optimised* index fund improve on the existing portfolio, which is probably well-diversified and is virtually already an "index fund"?

In this paper, statistically optimised index funds are compared with randomly selected portfolios of equivalent size. The results do not unequivocally support the proposition that an optimised index fund is a better vehicle for tracking a target

Dr. N.A. Sinclair is a senior lecturer in the Department of Accounting and Business Law at the University of Melbourne and is a director of St Clair Consulting and Investment Pty Ltd, which specialises in financial analysis and research. *This research was undertaken with the financial assistance of Towers Perrin Australia. The author accepts sole responsibility for the opinions expressed in the article.*

index than a conventional, well-diversified portfolio and a "buy-and-hold" investment strategy. In practice, it appears that the optimisation simply addresses the already small differences between a conventional portfolio and a target index. However, optimisation may play a minor role in the rational reduction of the *scale* of a conventional portfolio.

Before examining the structure of an optimised index fund, it is worth focusing on the important principle of diversification which is embodied in portfolio construction. This principle states that as more shares are added to an existing portfolio, to the extent that the returns on those shares are not perfectly correlated with the returns on the existing portfolio, then the overall risk of the portfolio can be reduced by the inclusion of those shares, without necessarily sacrificing expected returns.

For the portfolio manager, the process of diversification involves the answers to two questions:

- Which shares ought to be added to the existing portfolio?
- What investment proportions ought to be assigned to those shares?

A naive approach to answering these questions merely spreads the investment dollars across economic sectors and across shares within each sector, resulting in a large number of shareholdings of arbitrary proportions. However, it is a well-known principle that, as the number of shareholdings approaches 30, even a random combination of shares with equal investment proportions will approximate, to a high degree of accuracy, a more scientifically constructed portfolio.[2] In other words, a portfolio manager who already holds a portfolio of at least this size is, to a large extent, already investing in an "index fund" irrespective of the arbitrary answers to the two questions above.

A statistically optimised index fund simply answers the two questions in a different way, so that the portfolio manager can ensure complete diversification, with an *optimal* number of selected shareholdings and *optimally* determined investment proportions.[3]

For a number of reasons, the construction of a statisically optimised index fund cannot be performed as an unconstrained exercise. Portfolio managers commonly impose a number of constraints on the construction and management of their portfolios. For example:

- A limit may be placed on the maximum percentage that can be invested in any one share issue.
- A limit may be placed upon the minimum percentage that can be invested in any one share issue to avoid unmarketable parcels.
- General restrictions may be placed upon investment in particular industries or shares, perhaps for non-financial reasons.
- The manager may want to pursue an active strategy which attempts to "bias" or "tilt" the portfolio towards shares with higher/lower risks, higher/lower dividend yields or higher/lower capital gains.

> ***Clearly, a portfolio manager would find full replication of this index to be prohibitively costly.***

The main implication of a constrained index fund is that as the portfolio manager places more and more constraints on the optimisation process, the resulting "optimised" portfolio will become less closely associated with its target than if the process had been completely unconstrained. Clearly, with enough realistic constraints it is questionable whether the "optimised" index fund will look any different from a conventionally constructed portfolio.

An empirical study

In this section of the paper, the following empirical issues are investigated:

- How closely associated with an underlying target index are optimised index funds and randomly selected portfolios.
  - ☐ over the historical time period used to construct the index fund and portfolios?
  - ☐ over an adjacent time period which was not used in the optimisation process?

These results are presented in Tables 1 and 2.

- What is the effect on a constrained optimised index fund as the re-optimisation interval increases? Rebalancing frequencies of 1, 3, 6, 12 and 24 months are examined to provide an insight into the transactions costs associated with rebalancing an optimised index fund.

These results are presented in Table 3.

A value-weighted index was compiled from all industrial and mining shares with the required share-price data on the AGSM-CRIF database in the period Feburary 1974 to December 1984 inclusive.[4] The minimum number of shares used in the index was 526 at September 1979 and the maximum was 841 at September 1981.

Clearly, a portfolio manager would find full replication of this index to be prohibitively costly. This value-weighted index was constructed to play the role of target for a smaller optimised index fund. Ninety-one shares in the index with complete share-price data over the time period provide the feasible set of shares for the optimisation process. These shares tend to be the larger blue-chip securities and so are likely to be very close to the actual shares that practising portfolio managers monitor. Of course, it is expected that the number of shares required to enable the optimised index fund to adequately mimic the target value-weighted index will be significantly less than 91.

The results in Tables 1, 2, and 3 contain a number of statistics on the performance of optimised and randomised index funds. In all three tables, there are two measures of the degree of association between optimised and randomised portfolios.

First, the mean absolute deviation is the average of the absolute differences between the returns on the optimised portfolio and the returns on the target index. That is, positive and negative deviations from the target are given equal weight when the objective is to mimic the target index. An ideal optimised portfolio would have a mean absolute deviation of zero.

Second, these tables also contain the correlation between the portfolio returns and the returns on the target index, which is a popular measure of association. An ideal optimised portfolio would have a correlation coefficient of 1.0.

Other statistics in Tables 1 and 2, such as the average return and standard deviation of return for the optimised portfolios, are statistically engineered to be as close to the target index as possible. The extent to which these statistics differ from the target also reflects upon the adequacy of the optimisation.

The results in Table 1 are obtained from the 107-month period (February 1974 to December 1982) over which the optimised portfolios were constructed. By construction, it is expected that the optimised portfolios will exhibit a closer association with the target index than any portfolio.

In Table 1, it is clear that the unconstrained optimised index fund has the lowest mean absolute deviation of 0.979 percent and the highest correlation coefficient of .978. Both of these statistics indicate a high level of association over the estimation period. Further, it requires only 28 of the 91 shares to achieve this level.

However, it is worth noting that the unconstrained optimised portfolio requires an unrealistic initial investment of 27 per cent in BHP and 11 per cent in CSR. When some realistic constraints are imposed on the investment proportions, such as a minimum of 0.5 per cent and a maximum holding of 10 per cent, then the mean absolute deviation increases to 1.246 per cent and the correlation drops to .965. Now, only 27 of the 91 shares are held. It is likely that the imposition of additional constraints would further reduce the association between the optimised portfolio and the target index.

Table 1 also shows results for two randomised portfolios. As expected, both measures of association for these portfolios are inferior to those of the optimised portfolios, moreso in the case of the equal-weighted portfolio. Obviously, a value-weighted portfolio with more emphasis on the larger shares represents a superior vehicle for tracking a large market index than an equal-weighted portfolio.

It is also interesting to note that over this time period, the average return per unit of standard deviation risk is higher for the randomised portfolios relative to both the optimised portfolios and the target index. To a large extent, this finding is ensured by the construction of the optimised index portfolio, in which the returns and risk of the portfolio are implicitly constrained to equal those of the target index over the estimation period. The randomised portfolios do not share this constraint and therefore gain/lose from their idiosyncratic returns over this time period. For example, in Table 1 the mean and standard deviation of the optimised portfolios bracket the target index within one-tenth of 1 per cent per month, which is closer than the randomised portfolios.

Table 1: Optimised and randomised index funds relative to a value-weighted market index (Feb. 1974 to Dec. 1982)

	Number of shares held	Average return (%/mth)	Standard deviation (%/mth)	Mean absolute deviation (%/mth)	Correlation of fund with target
Unconstrained optimised index fund	28	1.244	5.946	0.979	0.978
Constrained optimised index fund $.005 \leq w_i \leq .10$	27	1.041	6.015	1.246	0.965
Randomised index fund equally weighted	27	1.540	4.900	2.400	0.856
Randomised index fund value-weighted	27	1.350	6.416	1.616	0.945
Value-weighted target index	—	1.141	5.807	0.000	1.000

Table 2: Optimised and randomised index funds relative to a value-weighted market index (Jan. 1983 to Dec. 1984)

	Average return (%/mth)	Standard deviation (%/mth)	Mean absolute deviation (%/mth)	Correlation of fund with target	24-month value relative
Unconstrained optimised index fund	2.741	6.867	1.625	0.956	1.819
Constrained optimised index fund $.005 \leq w_i \leq .10$	2.459	6.042	1.33	0.958	1.722
Randomised index fund equally weighted	2.828	5.297	1.531	0.942	1.895
Randomised index fund value-weighted	3.081	6.917	1.946	0.941	1.968
Value-weighted target index	2.093	5.685	0.000	1.000	1.586

The results in Table 2 document the performance of optimised and randomised portfolios over a subsequent 24-month holding period and provide an "acid test" for the optimised index funds. That is, does an optimised portfolio track the target index any better than a randomised portfolio over an *independent time period*?

In Table 2, in terms of both mean and absolute deviations and correlation coefficients, the overall association with the target index remains high for all portfolios, although the absolute level of association has deteriorated. Both optimised portfolios have drifted away from the target index but remain relatively closer than the randomised portfolios. The mean absolute deviation statistic for the constrained optimised portfolio is the smallest at 1.333 per cent and the equal-weighted randomised portfolio is the next smallest, with a mean absolute deviation of 1.531 per cent. Also in Table 2 are 24-month value relatives which are analogous to $1 invested in each portfolio and held for 24 months. Again, the optimised portfolios are marginally closer to the target than the randomised portfolios.

It is clear that over the 24-month holding period, the optimised portfolios maintain their superior ability to track the target index, compared with their randomised equivalents. However, the performance margins between realistic alternatives, such as a constrained optimised portfolio and a value-weighted randomised portfolio, are not large and tend to be in the vicinity of 0.6 of 1 per cent per month. Whether this performance difference is worth the cost of optimisation is an empirical question.

The purpose of Table 3 is to monitor the statistics in Tables 1 and 2 as the frequency of rebalancing increases for the constrained optimised portfolio. In addition, by examining the nature of the changes in the investment proportions across the various rebalancing periods, some insight is gained into the likely impact of transaction costs on the construction and maintenance of optimised index funds.

It is clear from Table 3 that all

statistics reflecting the degree of association between the constrained optimised portfolio and the target index are unaffected by more frequent re-optimisation. For example, with monthly rebalancing the mean absolute deviation is 1.448 per cent and the correlation is .957, neither of which demonstrates any improvement in the ability of the optimised portfolio to track the target index. Therefore, it may be concluded that the benefits of more frequent rebalancing are small or non-existent, relative to a simple buy-hold policy with initial optimisation.

The major result in Table 3 concerns the distribution of changes in investment proportions as the frequency of re-optimisation increases. Both the number of trades falling in a defined interval, and the ratio of those trades to all possible trades, are shown in the table. The latter percentage (in parentheses) may be taken as an indicator of the proportion of the portfolio to be turned over for the complete holding period. For example, with quarterly rebalancing there are 819 (9 times 91) potential trades, of which 56 (6.8 per cent) occurred and involved changes in investment proportions of between 1 per cent and 5 per cent. Similarly, 102 (12.5 per cent) involved buying additional shares whereas 98 (12 per cent) involved the sale of shares.

It is clear that as the frequency of rebalancing increases from bi-annully to monthly, greater trading volumes are required to optimise the portfolio yet there is ostensibly no improvement in the association between the portfolio and the target index.

A second important insight in Table 3 is that as the rebalancing frequency increases, a higher proportion of trades falls outside the minimum transaction constraint of 0.5 per cent, and therefore would not be initiated. Inevitably, this would further disassociate the optimised portfolio from the target index and reduce the superiority over the randomised portfolios noted in Table 2. As the rebalancing frequency decreases to bi-annually, the proportion of inadmissable trades drops to 3.8 per cent and the trades tend to be larger.

The overall conclusion to be drawn from Table 3 is that reoptimising the portfolio more frequently will generate considerable transaction costs as a greater proportion of the portfolio is turned over. However, more frequent rebalancing does not appear to enhance the ability of the portfolio to track the target index and would more likely impede it when inadmissable transactions are indicated by the optimisation.

On the strength of this data, bi-annual rebalancing performs as well as more frequent rebalancing, with only a slight superiority over a randomised, value-weighted portfolio and a buy-and-hold policy.

The results indicated that within the period of the optimisation, optimised portfolios will always, by construction, mimic their target indexes better than any other portfolios. However, when the tracking ability of the optimised and randomised portfolios is assessed using "out-of-sample" data over an adjacent holding period, the superiority of optimised index funds is largely diminished and an equal-weighted randomised portfolio does almost as well.

More frequent rebalancing did not improve the performance of the optimised portfolio, but did tend to identify trades which would not be consummated and would therefore lead to further deterioration in the ability of the optimised portfolio to track the target index. In turn, the marginal difference in tracking performance between optimised and randomised portfolios would diminish.

Table 3: Effect of periodic rebalancing on a constrained optimised index fund relative to a value-weighted market index (Jan. 1983 to Dec. 1984)

	Rebalancing Period					
	Month	Quarter	Half-yearly	Annual	Bi-annual	Target index
Average return (%/month)	2.624	2.487	2.544	2.547	2.459	2.093
Standard deviation (%/month)	6.238	6.231	6.200	6.169	6.042	5.685
Mean absolute deviation from target	1.448	1.498	1.507	1.480	1.333	0.000
Correlation with target index	0.957	0.953	0.952	0.949	0.958	1.000
24-month value relative	1.784	1.728	1.752	1.754	1.722	1.586
Distribution of changes in investment proportions*						
.00 < [Δw_i] ≤ .005	348 (15.3)	84 (10.3)	34 (7.5)	12 (4.4)	7 (3.8)	
.005 < [Δw_i] ≤ .01	135 (5.9)	60 (7.3)	34 (7.5)	12 (4.4)	6 (3.3)	
.01 < [Δw_i] ≤ .05	68 (3.0)	56 (6.8)	35 (7.7)	29 (10.6)	16 (8.8)	
.05 < [Δw_i] ≤ .10	0	0	0	0	0	
Number of Buy transactions	278 (12.2)	102 (12.5)	54 (11.9)	30 (11.0)	14 (7.7)	
Number of Sell transactions	273 (12.0)	98 (12.0)	49 (10.8)	23 (8.4)	15 (8.2)	

The figures in parentheses represent a percentage of total trades possible..For example, the total number of possible trades with half-yearly rebalancing is 91 x 5 = 455 since the portfolio is also rebalanced at the end.

REFERENCES

Brown, P. and Tippett, 1981, "The Effect of Diversification on Australian Portfolios: An Analysis", *JASSA*, 3, 15-16.

Ball, R.,1977"Will Australian Investors Get Index Funds", *JASSA*, 3-10.

Hayes, R.E.,1977, "What Happened to Index Funds? — 1980"*JASSA*, 11-15.

NOTES

1. See Ball (1977) for an enlightening discussion of the rationale of index funds as a vehicle for investment. Hayes (1977) provides a somewhat erroneous contrary viewpoint. Optimised index funds are a subset of index funds generally due to their reliance on statistical techniques to select both securities and investment proportions.

2.Brown and Tippett (1981) formally demonstrate that only 1/Nth of the diversifiable risk remains in a randomly selected equally-weighted portfolio containing N securities. That is, for a portfolio of 30 shares, only 3.33 per cent of potentially diversifiable risk remains. A market index contains zero diversifiable risk.

3. A common way to construct optimised index funds is to use a quadratic programming algorithm in which the diversifiable risk is constrained to zero whilst the systematic risk is set equal to the target index. Quadratic programming algorithms are flexible enough to allow additional constraints such as those introduced later in the paper.

4. The Centre for Research in Finance (CRIF) is based at the Australian Graduate School for Management (AGSM) in the University of New South Wales and is a major supplier of research databases. The portfolio construction issues dealt with in this paper are general enough to be unaffected by the time period covered by the database. There is no reason to believe that similar results would not be obtained from more contemporaneous data. □