Optimal Selection of Passive Portfolios

Andrew Rudd

Andrew Rudd is Assistant Professor of Finance and Operations Research at the Graduate School of Business and Public Administration, Cornell University. *He wishes to thank Ron Lanstein and Barr Rosenberg for helpful discussions.*

Introduction

In the last few years, the use of "index funds" as one part of a total investment strategy has gained acceptance among investors, particularly among corporate pension sponsors. Use has been spurred partly by the passage in 1974 of the Employee Retirement Income Security Act (ERISA), which caused many pension sponsors to re-evaluate the management of their pension portfolios.* Index funds are appealing both on economic and fiduciary grounds, because typically they have lower management fees and higher diversification than managed portfolios. This has led to considerable growth in indexed pension assets, as is demonstrated by the figures in Exhibit 1. The majority of these assets are managed by external money managers, but there is a distinct trend for corporations to manage a proportion of their pension assets in an in-house index fund. This trend is also true for public funds, as evidenced by the recent formation of a $500 million index fund run by the New York State Common Retirement Fund [5].

An index fund is a "passive" (as opposed to "active") portfolio that is designed to track closely a visible index. In general, passive portfolios need not be indexed, but could have distinct sectoral or asset emphasis depending on the investor's attitudes toward risk and the economic environment. However, all passive portfolios, whether indexed or not, are designed to be stable and to match the long-term performance of one segment of the capital markets rather than to be managed with the aim of earning incremental rewards from transient asset or market behavior.

The method of selecting and revising passive portfolios has received virtually no attention in the literature. In the case of index funds, Black and Scholes [1], Shapiro [15], and Good, Ferguson, and

Exhibit 1. Growth of Indexed Assets ($ million)

	Dec. 1978	Dec. 1977	Dec. 1976
Assets managed externally	$6,546	$4,236	$1,496
Assets managed internally	1,021	265	95
Total	7,567	4,501	1,591

Source: *Pensions & Investments*, December 18, 1978

*See Langbein and Posner [7] for more discussion of the legal environment of portfolio management.

Treynor [4] have briefly commented on the problem without suggesting solutions. Of course, if the fund were exactly indexed, its composition would duplicate that of the index, and in the absence of cash flows the selection process would be purely mechanical; the portfolio would require revision only when the index was changed. However, cash flows from dividend income and additional contributions must be invested. It is impractical to allocate this incremental flow perfectly across all assets in the index in order to maintain exact duplication. Instead, larger investments must be made in a smaller number of assets which are dynamically balanced through time as additional flows are received. Hence, in order to achieve minimal tracking error, the rebalancing of the fund must be as efficient as possible.

Passive Portfolio Strategies

Treynor and Black [16] first introduced the idea of passive versus active management strategies. The distinction arises from the separation of the investment decision into two components. The first component involves an analysis of the long-term risk and rewards in the capital markets, leading to the resolution of the "normal" exposure of the portfolio to economic events. The normal exposure may be expressed in terms of an asset mix (for instance, an allocation of funds to cash and equivalents, bonds or equities) or the risk level of the portfolio (for instance, specifying the portfolio beta).

The second component entails the active management of the portfolio to take advantage of perceived mis-pricing and consequent revaluation in the future. This adjustment is a matter of judgment; in the absence of any subjective information, there is no reason to shift the portfolio from its normal position. Hence, a managed portfolio can be thought of as the sum of two sub-portfolios. These are the active and passive parts; the passive sub-portfolio is the central core around which the active manager places bets by underweighting some stocks and overweighting others.

This distinction between active and passive strategies has important implications for investment decision-making. First, it lends itself to a natural organizational structure of an investment management department, as personnel are either forming judgments on security valuation or are occupied with the more mechanical tasks of passive management. Second, it permits greater control of active decisions and transactions costs. This follows since there is no requirement to make transactions to ensure portfolio diversification, which is obtained from the passive component. All transactions are explicitly made in response to judgment and to increase portfolio return.

The composition of the appropriate passive portfolio should not be the result of disequilibrium analysis (that is, hypothesizing some perpetually undervalued sector), because this implies that investors are consistently stupid in not taking advantage of the continual mis-pricing. Rather, the decision should be based on equilibrium arguments, so that in the absence of special circumstances the appropriate passive portfolio should be the market portfolio, which includes all assets with portfolio weights proportional to their respective capitalizations.

The special circumstances that may induce non-market passive portfolios include 1) differential taxation on capital gains and dividends implying that tax-exempt portfolios should be yield-biased (see [12]); 2) high-beta portfolios for institutions with a long-term investment horizon and no intermediate claims on portfolio assets; 3) liability covariance adjustments to take advantage of the correlation between portfolio liabilities and assets (for instance, the liabilities of life insurance companies are expressed in nominal terms; hence it is reasonable for their assets to be mainly bonds that have nominal pay-offs); and 4) legal indentures or social responsibility issues that force exclusion of certain assets; for example, some funds are legally obliged not to hold any "sin" stocks or stocks with South African involvement (see [13]). Notice that the cost of the bias is implicitly offset by a greater gain in these examples; for instance, in the last case, the increase in "moral" utility must be larger than the increase in "financial" disutility.

In the past, however, few sponsors or money managers have designed passive portfolios to capture the increase in utility to be gained in these situations. Typically, managers have implemented the single, simple passive strategy of using the Standard & Poor's 500 stock index (S&P 500) as a surrogate for the market portfolio, and then matching portfolio performance with this benchmark. In order to study the implications of this decision, it is useful to classify passive strategies into three types according to the portfolio goals and formation method. These are:

1. Matching a portfolio to an explicit index with no constraints on formation (*i.e.*, all assets in the index are available for purchase, and there are no bounds on asset holdings).

2. Matching a portfolio to an explicit index with constraints on formation (*i.e.*, some assets in the index

may be unavailable, or there may be bounds on asset holdings).

3. Matching a portfolio to a set of attributes rather than to an explicit index (*e.g.*, the yield-biased portfolio is a good example, where the attributes may include maintaining the most diversified portfolio with a specified yield).

In the last classification, the majority of marketed passive portfolios belong to the first type. Type two and type three portfolios are considered more difficult to select, and they have only recently been endorsed by pension sponsors.

Almost certainly, much of the reluctance to follow these more difficult — yet more rewarding (in terms of utility) — passive strategies has been due to the lack of any acceptable selection procedure. This is underscored by information in Exhibit 2, which shows the number of stocks in the S&P 500 index funds of the leading managers at the end of 1977. The exhibit indicates that the methodology used to form the funds varies widely, as the number of stocks held ranges from approximately 50 to 500. Notice that only one manager invests in all the stocks in the index. Chase Manhattan Bank invests in 499 companies, leaving out its own stock for legal reasons. The argument of fiduciary responsibility led Wells Fargo Bank to screen out the few companies from the S&P 500 believed likely to go bankrupt. The remaining index funds show that specific selection methods have been used, in that the number of stocks held is much less than the 500 that constitute the index.

For any passive strategy, there are two significant advantages in holding as few stocks as possible. The first is the decreased transaction costs on forming the portfolio. However, this advantage may be smaller than for normal transactions since passive trades are "informationless" (*i.e.*, they are only made to invest cash flow or to rebalance the portfolio). Transaction costs per unit traded are smaller than in trades where the broker or dealer believes a spread is required for protection against an information disadvantage. (See Black and Scholes [1] for further discussion of trading policy in the context of index funds.)

Forming a portfolio with few stocks implies that more frequent rebalancing may be required, which increases transaction costs. If the cash flow is sufficient, though, rebalancing may be achieved by judicious new investment, avoiding additional costs. The second, and probably larger, advantage is the decrease in administrative overhead; administration costs (particularly custodial and accounting) are reduced substantially by holding fewer companies.

Exhibit 2. Index Funds, Managers and Number of Stocks Held

Manager	No. of stocks in index fund
American National Bank	240–340
Wells Fargo	496
Bankers Trust	250
Batterymarch Financial	250
Harris	250
Manufacturers National (Detroit)	400
Chase Manhattan Bank	499
Analytic Investment Management	35–50
First Index Investment Trust	500
Continental Illinois	300–350
First National of Chicago	485
Marine National Exchange Bank	350
St. Louis Union Trust	240–280
State Street Bank & Trust	250
Metropolitan Life Insurance Company	350
First National Bank of Minneapolis	270
Aetna	120–150
Union Bank (California)	250
Fidelity Bank (Pennsylvania)	60–70
Girard Bank	475
Trust Company of Georgia	250
Birmingham Trust	159
First National of Akron	100

Source: *Pensions & Investments*, December 19, 1977

The major disadvantage of holding only a few stocks is related to the tracking error of the portfolio. There are two important criteria related to risk that are useful in measuring the performance of a passive strategy.

First, the mean difference between the return on the portfolio and bogey should be zero, where the bogey represents the characteristics which the passive strategy is attempting to match. The bogey could be an index, in which case the characteristics would be the set of assets (and their holdings) comprising the index, or, for a type three portfolio, the set of attributes defining the strategy. According to the Capital Asset Pricing Model (CAPM), the proper control of the systematic risk level or beta will ensure minimal mean error in the absence of transactions costs. The portfolio and the bogey must have the same exposure to aggregate economic events, which obtains when they both have the same systematic risk; that is, the portfolio must have a beta of unity relative to the bogey. Hence, in the context of the CAPM, minimizing the mean difference between the returns on the portfolio and bogey requires decisions concerning beta and the trading strategy used during rebalancing. This must be established so that transactions costs are as close to zero as possible.

Second, the random variability in the return of the portfolio about the bogey should be as small as possible. This random variability is the result of the unsystematic or residual risk in the portfolio. Residual risk, usually reported in units of annual standard deviation, measures the tracking error of the portfolio that results from imperfect diversification compared to the bogey.

For example, in the case of the S&P 500 index funds, an index fund not holding all 500 stocks necessarily exhibits residual risk. As the number of stocks held in the portfolio decreases, the level of portfolio residual risk typically rises, increasing the probability for significant tracking error. The level of residual risk in currently-marketed S&P 500 index funds varies from zero to about 1.5% annual standard deviation. Those funds that hold virtually all the stocks in the index are at the lower end of this range, while those composed of fewer stocks are at the upper end. For contrast, actively-managed portfolios may exhibit from 2% to 15% annual residual standard deviation; at the lower end of this range are the "closet indexers," while at the upper end are the special equity funds. The median for institutional portfolios is about 6%.

Finally, there is the question of practical implementation. If each client requires its assets to be separately held, or if the dollar value of funds invested passively is small, it will be impossible to purchase economic shareholdings in many companies. In this case, a portfolio with a small number of stocks will be forced upon the manager. Conversely, if all client accounts are to be commingled, forming a portfolio with a large number of stocks is a practical possibility.

The central question remains: Given a certain pool of money, how should the assets and their portfolio weights be determined to most efficiently form a passive portfolio?

Portfolio Selection By Universe Stratification

The most common method of index fund selection is stratifying the universe of assets. This procedure essentially divides the universe into "cells." The portfolio is then selected by investing in assets from the cells. The measure of how closely the portfolio will track the index is approximated by determining the difference between portfolio and index cell holdings. If the difference between these two holdings is small, intuition suggests that the portfolio will track the index closely, or that the index fund will exhibit little residual risk about the index.

Most index fund selection programs currently in use incorporate the following simple heuristic:

1. Assemble the universe of assets from which the fund is to be selected and the index to which the fund will be matched.
2. Specify the dollar value, $D, of the index fund.
3. Select the minimum investment size, h, as a proportion of the total value of the fund. The reciprocal of this size, 1/h, determines the number of "units" in the fund. The dollar value of each unit is $hD.
4. Form a capitalization-weighted portfolio with total value $D from all the assets in the universe, and rank holdings from largest to smallest. The index fund is formed by purchasing assets that have holdings greater than minimum position size h. The size of the transaction is the integer number of units such that the dollar value of the transaction is closest to the dollar value of the position in the capitalization-weighted portfolio. Purchases begin with the largest capitalization asset and stop when all units are used. If all purchases are completed, and there are still units remaining, continue with Step 5.
5. Assign each asset in the universe to one of N mutually exclusive and exhaustive cells. In practice, this amounts to assigning each asset to one industry group. Calculate the investment holdings of the index and index fund in each of these cells.
6. Rank the cells in order of the difference between the fund and index holdings, from the most to the least deficient. Starting with the most deficient cell, and with assets in that cell that have not yet been purchased in step 4, purchase single units until the index fund cell holding matches the index cell holding to a discrepancy of one unit or less. Continue until all units have been invested.

In summary, the heuristic selects the fund first by matching the fund and the index by company size and then by one other dimension (usually industry groups).

An interesting variation of this heuristic used to control transaction costs involves sampling at Steps 4 and 6. Instead of specifying the exact company that must be purchased, a group of similar companies is isolated. This list is then given to the trading desk, which "samples" from the group depending on the depth of the market for each company.

The selection procedure is unsophisticated in that the only control of tracking error is by minimizing the deviations of portfolio holdings from index holdings along the two dimensions of, usually, capitalization and industry groups. Unfortunately, the intuition that keeping the differences in holdings small will cause the tracking error also to be small is frequently erroneous. In fact, of two 250-stock index funds given in Exhibit 2, one has almost three times the residual standard deviation of the other.

The reason is that residual risk is composed of two components, specific risk and extra-market covariance (XMC). Specific risk is that part that is specific to an individual company and that is independent of the specific risk of other companies. To immunize a portfolio from each company's specific risk, the portfolio holding should equal the index holding. The specific risk contribution of an asset to portfolio risk is proportional to the squared difference between these holdings. Since large differences in holdings are proportionally more serious than small ones, the heuristic, by matching portfolio holdings with the index holdings, starting with the largest companies, probably does a good job of controlling specific risk.

It is with the other component of residual risk, XMC, that the problem lies. XMC is the risk due to common factors in the economy. It is present in the portfolio whenever the portfolio has a different exposure to these common factors than the index exposure. A recent example is that of small capitalization companies; if the average sizes of the companies in the portfolio and in the index are different, the portfolio and index performance may well be different. Money managers would refer to this as the effect of "second tier" stocks. In other words, XMC is the risk that results from groups of companies moving together in a manner unrelated to the overall movement of the index.

What are these common factors? In his empirical research, King [6] suggested that they may be associated with industry groups. Recently, Rosenberg and Marathe [10, 11] isolated 45 common factors: 39 related to industry groups and 6 to other company attributes (price and earnings variability, size, financial risk, growth orientation, and past performance). Hence, whenever the portfolio and index differ on any one of these 45 dimensions, the portfolio will exhibit residual risk relative to the index.

It is here that the stratification heuristic fails. By starting with the larger capitalization stocks, the index fund may well be underweighted with small stocks. Since the 45 dimensions are related (for instance, small stocks tend to be growth-oriented), there is considerable opportunity that use of this selection procedure causes exposure to several common factors. In other words, stratification makes little attempt to control XMC.

Similarly, the procedure does not control the systematic risk level of the index fund. Again, depending on the position size, h, and the actual universe used, the procedure may overweight large capitalization stocks that tend to have lower systematic risk, causing the index fund and the index to respond with different magnitudes to economic events. Unless the portfolio risk is analyzed subsequent to selection (by performing a historical simulation, for instance), there is no knowledge of the magnitude and direction of the bias, just knowledge that bias exists.

Finally, there is no indication of the relative benefit or cost of adding or deleting a stock. This is particularly important when revising the portfolio, because the cost of the transaction must be weighted directly against the benefit, in this case of more closely tracking the index. Since there is no quantitative measure of the tracking error, it is not known whether the transaction cost is greater than the disutility arising from the portfolio risk. In other words, there is no indication whether the transactions will be beneficial or not.

The major advantage of the procedure is that it is simple to use and to understand. No sophisticated modeling or estimation is required. It is also an advantage to be able to control the asset position size, as it gives an indication of the approximate number of securities in the index fund.

Finally, it is worth noting that stratification is not an optimizing procedure. The portfolio is deemed selected when the heuristic stops (that is, all funds are invested), and there is no guarantee that the index fund is in any sense optimal.

The Optimization Approach

An alternative approach to passive portfolio selection is optimization. The problem of selecting a portfolio to track an index is naturally stated as an optimization problem. As discussed above, the requirements for an index fund are:

1. A beta of unity relative to the index;
2. Minimum differential exposure to factors common to the index and the portfolio; that is, the fund should have minimum XMC relative to the index; and
3. Minimum specific risk.

These latter two requirements define the objective of minimizing residual risk. Further, rebalancing of the fund should only take place when transactions costs are not greater than the increase in utility arising from either maintaining a beta closer to unity or decreasing residual risk.

Of course, this approach is far more complex, requiring more sophisticated computer techniques than the stratification procedure. At the same time, it is far more flexible, with none of the disadvantages mentioned in the previous section.

The use of quantitative portfolio optimization methods was originally suggested by Markowitz [9]. Since the original work, there have been many formulations of portfolio optimization in order to promote its application to practical money management. Actual implementation of optimization techniques, though, has not been very widespread until recently, when realistic optimization programs became available.

The optimization approach results in a quadratic program, since residual risk is inescapably a quadratic function of the asset holdings. The control of systematic risk is effected as a single linear constraint requiring the portfolio beta to be unity, while the control of residual risk is implemented directly from a model of the residual covariance among assets. The quadratic program to capture the goals above can be specified as follows:

$$\min_{x_i} \omega_P^2$$

subject to:

$$\beta_P = \sum_{i=1}^{N} x_i \beta_i = 1$$

$$\sum_{i=1}^{N} x_i = 1 \; ; \; x_i \geq 0$$

where ω_P^2 is portfolio residual variance;
β_P is the portfolio beta, defined as the weighted sum of the asset betas, where the weights are the asset holdings; and
x_i is the revised holding of the i^{th} asset in the portfolio, assumed non-negative.

The first constraint defines the portfolio beta, while the second requires that all funds should be fully invested. In other words, index fund selection in the absence of transaction costs reduces to locating a portfolio with unit beta and minimal residual risk (where these latter measures are defined relative to the index).

One difficulty is the realistic modeling of transaction costs, since the penalty is typically a non-linear function of the transaction, including both a set-up cost and cost dependent on the amount traded. Since this specification makes the problem computationally infeasible, it is assumed that costs are proportional to the transaction. This treatment of transaction costs transforms the objective function of the problem to:

$$\text{Max } \lambda_T T + \lambda_V \omega_P^2$$

$$\text{where } T = \sum_{i:x_i > x_i^0} t_{pi}(x_i - x_i^0) + \sum_{j:x_j^0 > x_j} t_{sj}(x_j^0 - x_j),$$

t_{pi} and t_{sj} are the proportional purchase and sale transaction costs, λ_T and λ_V are negative coefficients defining the after-transaction cost risk-return tradeoff, and x_i^0 is the initial holding in the i^{th} asset. T represents the total transaction cost incurred during revision. Incorporating transaction costs directly into the objective function permits the effect of adding and/or deleting a stock to be quantified so that the cost of transaction can be compared directly with the benefit of the trade. The only trades undertaken are where the benefit is greater than the cost.

The following case study contrasts the two methods.

Index Fund Selection from a Non-Standard Universe

A major bank trust department wishes to offer an S&P 500 index fund service, the only requirement being that the assets appear on its approved list. The approved list contains about 350 companies approved for purchase by the Investment Committee. Approximately 250 of them belong to the S&P 500. The initial investment in the index fund is set at $40 million.

The specification requires that a portfolio be formed that matches the index performance without using all the constituent assets. This problem belongs to the second type in the classification given earlier. It is inherently difficult for the stratification procedure since there are assets in the universe (the 350-stock approved list) that are not in the index (the S&P 500). More complex heuristics than those described above may select (at greater cost) better index funds than those described below, although it is doubtful that one specialized stratification program could perform adequately on all type two problems. This situation though, describes a real study, and modern investment technology should be sufficiently robust to be applicable to real-world situations.

Stratification Results

The stratification analysis was run several times for different variations of this problem. In each of the runs, the selected portfolio had a beta different from unity within the range being from .97 to 1.02.

The first two analyses tried to establish the best index fund that could be formed from the entire 350 stocks. The minimum position size was set to 0.001 so that the fund was formed from 1,000 units each worth $40,000. At the end of step 4 of the stratification heuristic, approximately 70% of the index fund had been invested. The 350-stock approved list was then divided into 43 industry groups and the remaining 30% of the index fund invested.

A risk analysis of the resulting portfolio was then performed. It was determined that the portfolio beta was 1.01 and the annual residual standard deviation something over 1.06%. This level of risk is somewhat surprising for a large portfolio (350 stocks), so the stratification program was re-run to establish the level of residual risk in the best 350-stock index fund where all 350 stocks belong to the S&P 500. A risk analysis of this 350-stock index fund revealed it had a beta of .997 and an annual residual standard deviation of 0.23%.

The difference in risk levels is essentially due to the "bias" in the 350-stock approved list and an inconsistency in the stratification procedure. The bias exists because the trust department (at the time the analysis was performed) favored small growth companies, so the approved list was not representative of the index. The inconsistency is more obvious with the next three analyses, which were aimed at establishing the risk levels for (approximately) 250-stock index funds. In this case, the minimum position size was set to .002 so that the fund was formed from 500 units each worth $80,000.

The first run formed a 250-stock index fund, where all 250 companies were drawn from the index. This fund had a residual standard deviation of 0.50%. Stratification of the approved list aiming for a 250-stock index fund produced a 231-stock portfolio with a residual standard deviation of 1.122%. However, when the universe was restricted to those stocks in the approved list *and* the S&P 500, the stratification program produced a 214-stock portfolio with a residual standard deviation of .958%; in other words, further restriction of the universe produces a portfolio that meets the desired goals more closely.

This unhappy result points out a major problem with the procedure. When selecting from the 350-stock universe, one should not treat assets that are not in the S&P 500 the same way as those that are. This follows from the fact that the entire holding of non-S&P 500 companies contributes to specific risk, whereas it is the difference between portfolio and S&P 500 holdings that contributes to specific risk for S&P 500 companies. (Recall that, to immunize a portfolio from a company's specific risk, the portfolio and index holding should be the same.) In short, adding non-S&P 500 companies always increases specific risk, so it is no surprise that the stratification procedure produces a more diversified portfolio with fewer (but all S&P 500) securities.

Finally, to gauge the effect of the bias in the approved list, ten major stocks in the S&P 500 but not in the approved list were added to the universe. The minimum position size was reset to 0.001, and the resulting residual standard deviation was 0.710%.

Optimization Results

The betas of individual assets and the variances and covariances of residual returns of all assets were drawn from the predictive models of investment risk in [10] and [11], updated by the Fundamental Risk Measurement Service. The index fund was formed using the optimization approach described in [14] and outlined in the appendix. The instructions given to the optimization program were the following: Form a portfolio with 1) a systematic risk level equal to that of the S&P 500, that is, with a beta equal to unity relative to the index, and 2) minimum residual risk relative to the S&P 500.

Notice that, since residual risk incorporates specific risk, the optimization approach will treat the non-S&P 500 companies differently from S&P 500 companies. Adding a non-S&P 500 company (and hence increasing specific risk) can only be justified if portfolio XMC is reduced more than specific risk is increased, or if it facilitates meeting the beta constraint. This effect is borne out by the results of the study.

Using the 350-stock approved list, an index fund was formed with 297 stocks with a portfolio beta of exactly unity and an annual residual standard deviation of 0.54%, or, about half that of the fund produced by stratification. Omitting the 100 non-S&P 500 stocks causes the level of residual risk to increase, as can be seen in Exhibit 3 (which gives the detailed comparison of the two methods).

In order to verify that the cross-sectional results obtained from the optimization were realistic, the actual performance of the optimized index fund was traced over an 18-month period. The experimental design was as follows: At the end of December 1976, an index fund was formed by the optimization from the approved list. This portfolio is the one referred to above as having a beta of unity and an annual residual standard deviation of 0.54%. The monthly performance of the fund relative to the index was then computed for January, February, and March 1977. At the end of March, the fund was revised using the optimization, including proportional transactions costs of 0.10% for both purchases and sales. (This penalty was estimated as a cost of 1% amortized over a ten-year period, which conforms to the magnitudes reported by index fund managers.) The revised portfolio was then tracked over the next three months and rebalanced at the end of June 1977. Continuing in this manner, quarterly revisions were performed in September and December 1977 and March 1978.

Exhibit 4 shows the actual monthly performance (expressed as a percentage rate of return) of the index fund relative to the index for the 18 months. The right column shows the total tracking error that arises from exposure to systematic and residual risk different from that of the index. The effect of the residual risk is further attributed to XMC and specific risk. As is expected, the tracking error arises predominately from the specific risk of the index fund.

The standard deviation of the experienced monthly tracking error is 0.19%, or 0.65% annually. The predicted annual residual standard deviation computed at the end of December 1976 was 0.54%. The difference between the experienced and predicted risk is largely due to the quarterly revision period and the inclusion of transaction costs. Thus the levels of predicted risk that were obtained from the cross sectional analysis can be interpreted as the typical levels of risk that would be experienced in a real-world application.

Exhibit 3. Comparison of Stratification and Optimization Selection Methods

Analysis	Portfolio Annual Residual Standard Deviation %: Stratification	Optimization
Index fund from S&P 500 (max. 350 stocks)	0.23	0.13
Index fund from approved list (max. 350 stocks)	1.06	0.54
Index fund from S&P 500 (max. 250 stocks)	0.50	0.49
Index fund from approved list (max. 250 stocks)	1.12	0.66
Index fund from S&P 500 stocks in approved list (subset of approved list)	0.96	0.76
Index fund from approved list plus 10 specified S&P 500 stocks	0.71	0.15

Exhibit 4. Performance of the Optimized Index Fund Relative to the Index in Percent Return

Month		Systematic	XMC	Specific	Total Residual	Total
January	1977	0.0	-0.11	-0.25	-0.36	-0.36
February	1977	0.00	-0.00	-0.06	-0.06	-0.06
March	1977	0.00	0.00	-0.11	-0.10	-0.10
April	1977	0.0	-0.02	0.04	0.02	0.02
May	1977	0.00	0.14	-0.12	0.02	0.02
June	1977	0.0	-0.01	-0.07	-0.08	-0.08
July	1977	0.0	-0.07	-0.07	-0.14	-0.14
August	1977	-0.00	0.14	0.01	0.14	0.14
September	1977	-0.00	-0.10	0.41	0.30	0.30
October	1977	-0.00	-0.02	0.02	-0.01	-0.01
November	1977	0.00	0.09	-0.04	0.05	0.06
December	1977	0.00	-0.01	-0.06	-0.07	-0.07
January	1978	0.0	-0.11	-0.23	-0.35	-0.35
February	1978	0.0	-0.05	0.08	0.04	0.04
March	1978	0.00	-0.03	-0.04	-0.07	-0.07
April	1978	-0.01	0.20	0.25	0.45	0.44
May	1978	-0.00	-0.01	-0.07	-0.08	-0.08
June	1978	0.00	-0.01	-0.09	-0.09	-0.09
Mean Value		-0.00	0.00	-0.02	-0.02	-0.02
Median Value		0.00	-0.01	-0.06	-0.07	-0.06
Standard Deviation		0.00	0.09	0.15	0.19	0.19

General Passive Strategies

There is no doubt that passive investment strategies are here to stay. What will almost certainly change is the simple strategy of indexing to the S&P 500. Bond index funds, international index funds, NYSE index funds, and more general passive strategies, including a yield-biased passive portfolio ([3]) and a "sin-free index" ([2]), are currently being offered to pension sponsors. LeBaron [8] mentions other passive investment devices that he believes will be marketed in the future.

These more general passive strategies require forming a portfolio matched to specific attributes, instead of a widely visible index with known holding proportions (type three passive portfolios). A high-yield or a high-beta strategy are good examples of these general strategies that are clearly more complex to select. Another timely example concerns multiple-management of portfolios. Large pension funds are typically managed by several asset managers, each managing a part of the fund with (perhaps) a different style. It makes no sense to measure the performance of each of these different styles against a neutral market portfolio or bogey. The sponsor should define the attributes of the "normal" passive portfolios for each of the different styles to act as bogeys for performance measurement. These passive portfolios should be chosen such that the value-weighted aggregate coincides with the optimal passive portfolio of the sponsor. The problem now is to select the passive portfolios that represent the normal position of each of the styles.

A related situation occurs when asset managers tend to have the same style. In this case the aggregate portfolio may be highly exposed to one common factor. To counter this undesirable bias, a passive portfolio could be formed to offset or hedge the common bias from the individual active managers. This portfolio, sometimes called a compensating core, provides a good example of the use of a passive strategy.

In these cases, the stratification procedure will give uncertain results, as there is no clearly defined universe from which to stratify. The optimization approach will not suffer from this problem, since the optimization goals are defined directly in terms of the desired attributes. For instance, the yield-biased passive strategy can be calculated by finding the minimum residual risk portfolio with desired yield and, perhaps, unit beta.

In the past, when the only common passive strategy was the type one S&P 500 index fund, the use of stratification could be justified on the grounds of simplicity. In the future, the selection of general passive strategies of types two and three will require a more sophisticated and (unfortunately) more complex procedure such as the optimization approach.

References

1. Fischer Black and Myron Scholes, "From Theory to a New Financial Product," *Journal of Finance* (May 1974), pp. 309–412.
2. Phyllis Feinberg, "The 1978 Pension Olympics," *Institutional Investor* (February 1979), pp. 55–56.
3. Joanne Gamlin, "Yield Tilt Concept Slides Assets to Wells," *Pensions & Investments* (June 4, 1979), p. 2.
4. Walter Good, Robert Ferguson, and Jack Treynor, "An Investor's Guide to the Index Fund Controversy," *Financial Analysts Journal* (November-December 1976), pp. 27–36.
5. "Index Funding is Completed," *Pensions & Investments* (January 1, 1979), p. 2.
6. Benjamin King, "Market and Industry Factors in Stock Price Behavior," *Journal of Business* (January 1966), pp. 139–190.
7. John Langbein and Richard Posner, "Market Funds and Trust Investment Law," *American Bar Foundation Research Journal* (Winter 1976), pp. 1–34.
8. Dean LeBaron, "The Future of Indexing," paper presented at the Eighth Annual Meeting, Financial Management Association, Minneapolis, October 12–14, 1978.
9. Harry Markowitz, "Portfolio Selection," *Journal of Finance* (March 1952), pp. 77–91.
10. Barr Rosenberg and Vinay Marathe, "Common Factors in Security Returns: Microeconomic Determinants and Macroeconomic Correlates," *Proceedings of the Seminar on the Analysis of Security Prices*, University of Chicago (May 1976), pp. 61–116.
11. Barr Rosenberg and Vinay Marathe, "Prediction of Investment Risk: Systematic and Residual Risk," *Proceedings of the Seminar on the Analysis of Security Prices*, University of Chicago (November 1975), pp. 85–225.
12. Barr Rosenberg and Andrew Rudd, "The Yield/Beta/Residual Risk Trade-Off," Working Paper 66, Research Program in Finance, Berkeley Institute for Business and Economic Research, University of California, 1978.
13. Andrew Rudd, "Divestment of South African Equities: How Risky?" *Journal of Portfolio Management* (Spring 1979), pp. 5–10.
14. Andrew Rudd and Barr Rosenberg, "Realistic Portfolio Optimization," in *Portfolio Theory-Studies in Management Sciences*, Vol. II, E. J. Elton and M. J. Gruber, eds., Amsterdam, North-Holland, 1979.
15. Harvey Shapiro, "How Do You Really Run One of those Index Funds," *Institutional Investor* (February 1976), pp. 24–25.

16. Jack Treynor and Fischer Black, "How to Use Security Analysis to Improve Portfolio Selection," *Journal of Business* (January 1973), pp. 66–86.
17. Balder Von Hohenbalken, "A Finite Algorithm to Maximize Certain Pseudoconcave Functions on Polytopes," *Mathematical Programming* (October 1975), pp. 189–206.

Appendix. Formulation of the Quadratic Program For Portfolio Revision

It is shown in [14] that a realistic objective for optimization is:

Max $\lambda_\alpha \alpha_P + \lambda_Y Y_P + \lambda_{YQ} (Y_P - Y_T)^2 + \lambda_T T + \lambda_\beta (\beta_P - \beta_T)^2 + \lambda_V \omega_P^2$

subject to:

$$\sum_{i=1}^{N} x_i = 1$$

$$\beta_{min} \leq \sum_{i=1}^{N} \beta_i x_i \leq \beta_{max}$$

$$Y_{min} \leq \sum^{N} y_i x_i \leq Y_{max}$$

$$L_i \leq x_i \leq U_i \quad \text{for } i=1, \ldots, N$$

where $T = \sum_{i:x_i > x_i^0} t_{pi} (x_i - x_i^0) + \sum_{j:x_j^0 > x_j} t_{sj} (x_j^0 - x_j)$,

$$\omega_P^2 = \sum_{i=1}^{K} \sum_{j=1}^{K} \gamma_i \gamma_j F_{ij} + \sum_{i=1}^{N} (x_i - \beta_P x_{Mi})^2 \sigma_i^2$$

and N is the number of assets.

The notation used is as follows: λ_α, λ_Y, λ_{YQ}, λ_T, λ_β and λ_V are parameters specifying the tradeoff between the respective components in the investor's objective function; α_P is the expected portfolio residual return or "alpha;" Y_P is the portfolio yield; β_P is the portfolio beta; ω_P^2 is the portfolio residual variance; Y_T and β_T are the portfolio target yield and beta, respectively; Y_i and β_i are asset i yield and beta, respectively; β_{min}, Y_{min}, L_i, β_{max}, Y_{max}, and U_i are minimum and maximum bounds on portfolio beta, yield, and the asset i holding, respectively; t_{pi} and t_{sj} are the proportional purchase and sale penalty, respectively; and x_i^0, x_i and x_{Mi} are the initial and revised proportions of asset i in the portfolio and market, respectively.

Extra-market covariance is explained by a K-factor model where γ_i is the exposure of the portfolio to the i^{th} factor and F_{ij} is the covariance between the i^{th} and j^{th} factors, while the other component of residual risk, specific risk, is a weighted sum of the individual asset i specific risks, σ_i^2. The weights are the squared differences between the portfolio and an equal-beta-levered market portfolio.

In the passive portfolio revision problem, the formulation simplifies by setting:

$$\begin{aligned} \lambda_\alpha &= \lambda_Y = \lambda_{YQ} = 0 \\ \beta_{min} &= \beta_{max} = 1 \\ Y_{min} &= L_i = 0; i = 1, \ldots, N \\ Y_{max} &= U_i = 100\%; i = 1, \ldots, N \end{aligned}$$

resulting in a problem that requires:

$$\max_{x_i} \lambda_T T + \lambda_V \omega_P^2$$

subject to:

$$\sum_{i=1}^{N} x_i = 1$$

$$\beta_P = \sum_{i=1}^{N} \beta_i x_i = 1$$

$$x_i \geq 0$$

where the quadratic term for portfolio beta is a constant, since portfolio beta is constrained to unity and hence drops out of the objective function.

There are numerous algorithms that may be used for this problem. A particularly efficient code is the quadratic programming variant of Von Hohenbalken's general mathematical programming code [17]. For further discussion of this approach, see [14].