Ex-Post Evaluation of Dynamic Portfolio Strategies (or How to Tell Whether a Million Dollars Has Been Thrown Away)

Stewart D Hodges

Director
Financial Options Research Centre
University of Warwick

September 1991

Financial Options Research Centre
School of Industrial and Busines Studies
University of Warwick
Coventry
CV4 7AL
Phone: 0203 523606

FORC Preprint: 91/23

Ex-Post Evaluation of Dynamic Portfolio Strategies (or How to Tell Whether a Million Dollars Has Been Thrown Away)

1. Introduction

It has been demonstrated, in a paper by Dybvig [1988a] with a related title to this, that a number of portfolio (asset mix) strategies followed by practitioners are significantly inefficient. The inefficiencies arise not because of imperfect diversification across stocks, but because of poor diversification over time. Dybvig showed how the costs of following such strategies may be calculated, provided we know the parameters of the market return generating process and full details of the strategies concerned. The costs of following policies such as stop-loss, lock-in, random market timing or repeated portfolio insurance were shown to be substantial. However, in terms of applying this approach to practical problems of measuring the performance of investment portfolios, there remains the difficulty that usually we can only observe a single time series of what the fund manager actually did, and we are unlikely to know what policy would have been pursued along other possible unrealised paths of the probability tree. The purpose of this paper is to develop a powerful and robust technique which will enable Dybvig's payoff distribution pricing approach to be applied for performance measurement purposes with realistically limited amounts of observational data.

The advent of active approaches to asset allocation, of Futures and Options Funds, and of institutions carrying substantial portfolios of derivative instruments, has posed important unsolved problems for those involved with performance measurement. The LIFFE/LTOM [1990] Consultative Document reflects these concerns. The main difficulties are, first, to obtain accurate and relevant information about both value and risk exposure and, second, to find meaningful ways to interpret what has been achieved, taking due account of the risks that have been run. This paper is concerned with some aspects of the latter problem. It is motivated specifically by the question: "What is an appropriate yardstick to use when evaluating the performance of a portfolio in which the asset mix is actively varied through time?" We should not under-estimate either the academic or the practical difficulties of this question. The literature on performance measurement is both extensive and sophisticated, yet fails to provide any real guidance on this issue[1]. It should therefore be no surprise that the LIFFE/LTOM document states "We would welcome suggestions on practical ways in which risk adjusted performance figures for futures and options might be constructed". In contrast, the more conservative National Association of Pension Funds [1990] "Report Into Investment Performance Measurement" is altogether more suspicious of even the principle of risk adjustment but "does not close the door upon the search for generally accepted 'risk' measurements in the future."

The early works on risk adjusted performance measures (such as Treynor [1965], Sharpe [1966], and Jensen [1969]) focus on risk adjustment in an essentially static context in which the Capital Asset Pricing Model (CAPM) is assumed to hold. There is a now a considerable literature on the problems associated with such CAPM based measures including Dybvig and Ross [1985a, 1985b], and Grinblatt and Titman [1989]. Fama's [1972] approach at decomposing performance into components attributable to different sources provides further valuable insights, and includes a way of looking at market timing. Market timing and its evaluation has been specifically addressed in a number of papers including Sharpe [1975], Merton [1981], Hendrickson and Merton [1981], and more recently Jagannathan and Korajczyk [1986]. Many of the advances here are of an econometric nature. A number of other papers, such as Bookstaber and Clarke [1984, 1985], have discussed the effects of options in complicating the measurement of portfolio performance. Although this literature recognises the distributional problems arising from an actively managed asset mix, to date no previous work on performance measurement seems to have taken advantage of the opportunities provided by Dybvig's [1988a, 1988b] work on payoff distribution pricing.

The key result that we exploit in this paper is the insight that a portfolio management strategy is efficient (ie. not dominated by other strategies) if and only if it leads to wealth at the investment horizon being monotonic non-increasing in the state-price density. In a Black-Scholes type of world (in which there is a positive risk premium) this means that final wealth must be a monotonic non-decreasing function of the value of the risky asset at the horizon date. An appropriate way to measure the efficiency of a portfolio's observed sample path is therefore to estimate how closely the observed sequence of decisions may be approximated by a strategy which generates wealth that is suitably monotonic. A suitable basis from which to approximate the space of efficient portfolio strategies is formed by the strategies which replicate option spreads. We can estimate the approximating efficient strategy as a non-negative linear combination of the basis strategies. The estimates we need can be obtained in a straightforward way using least squares subject to non-negativity constraints. Efficient policies have a characteristic of being in some sense rather smooth, and this gives our estimates considerable robustness. The approximating strategy provides an important benchmark for performance comparisons, and does so even under much weaker conditions than those initially assumed. The extent to which the actual strategy departs from the approximating strategy enables us to obtain a measure of the opportunity cost of the strategy followed. The absolute magnitude of this cost depends on the level of the market risk premium. The technique is illustrated by means of a number of simulated examples, and issues concerning its application and extension are discussed.

The structure of the remainder of the paper is as follows. Section 2 provides the notation and formal results that are used from Dybvig's theory of payoff distribution pricing. Section 3 develops the methodology already outlined for measuring the characteristics and performance of dynamically managed portfolios. For this section and section 4 we assume a world in which Black-Scholes option pricing is valid. Section 4 gives the results of some simulations of measuring the efficiency of single sampled time series drawn from a variety of strategies. These include both efficient and inefficient strategies, including some of those previously described by Dybvig [1988a]. These simulations indicate that the technique has considerable power and they confirm the important role of the investment horizon of the fund. In section 5, we discuss the validity of the technique under more general assumptions. The concluding section, summarizes the results which have been obtained, and suggests areas in which the approach may be extended further.

2. The Payoff Distribution Pricing Model

Apparently little known but important work by Cox and Leland [1982] showed that when the riskless rate is constant and the risky asset follows either Geometric Brownian Motion or a geometric binomial process, portfolio strategies which fail to be path independent are stochastically dominated by other strategies and are thus unambiguously inefficient. In other words, within this framework, if the outcome from a particular strategy depends not just on where the value of the risky asset ends up, but also on what path it took to get there, then the strategy is inefficient. This immediately implies that a number of commonly used strategies such as lock-in strategies, stop loss strategies, and strategies involving rolling over portfolio insurance are inefficient, but does not tell us the magnitude of the inefficiency. However, subsequent work by Dybvig [1988a, 1988b] has succeeded in measuring the inefficiency costs of such strategies, and has concluded that they are substantial.

We begin by recalling some key results from Dybvig's analysis. Consider first a complete market which is frictionless and has a finite number of equally probable states s=1,..,S. for which the Arrow-Debreu pure securities have prices p_s, and probabilities π_s. An agent with utility function u(c) will choose state contingent consumption c_s so as to solve the following problem:

$$\text{Maximise} \quad \Sigma_s \ \pi_s \ u(c_s) \text{ subject to}$$

$$\Sigma_s \ p_s \ c_s = w_0.$$

We assume the usual differentiability conditions for u(.), and obtain the usual first order conditions that

$$u'(c_s) = \lambda \ p_s \ / \ \pi_s$$

for some λ and for all s. In other words the agent's marginal utility of wealth in state s is proportional to the terminal state-price density $\xi_s = p_s \ / \ \pi_s$. This optimization characterizes the probability distribution of the agent's terminal consumption c_s. We are accustomed to think of a risk averse agent for whom marginal utility is clearly a non-increasing function of consumption. Slightly less obviously, any investor who maximises expected utility and prefers more to less will choose the c_s to be a non-increasing function of the state-price density. Any other allocation could be improved on to give the same distribution for c_s at a lower cost, and so would not be optimal. Dybvig defines his payoff distribution pricing model (PDPM) by a pricing function PD(F) for distributions F of terminal consumption, such that PD(F) repesents the cheapest cost by which terminal consumption can be obtained that will have distribution F. To calculate this minimum cost we would need to sort c_s on the state-price density and allocate consumption to states on that basis. These ideas provide Dybvig's main theorem as stated in [1988a] (the formal proof is in Dybvig [1988b]):

Theorem 1. Payoff Distribution Pricing

The following are equivalent:

1. The consumption pattern c is chosen by some agent who has strictly increasing von Neumann-Morgenstern preferences over terminal wealth.
2. The consumption pattern has an asset price equal to the distributional price of its distribution function, that is,
   $PD(F) = \Sigma_s \; p_s c_s$ where F is the distribution of c_s.
3. Consumption is non-increasing[2] in the terminal state-price density, ξ_s.

We may note that in the kinds of binomial tree commonly used for option evaluation, the risk neutral probabilies which are used for weighting the various outcomes represent (undiscounted) state prices. For option valuation the objective probabilities usually serve no role at all. However this is certainly not the case for portfolio selection, where both are relevant. Banz and Miller [1978] and Breeden and Litzenberger [1978] were the first to demonstrate how such state prices might be recovered from the prices of traded options.

Our treatment of these ideas has so far been in the context of an economy with a finite number of states. They apply equally to continuous time formulations involving a continuum of states, provided that the assumptions of completeness and no arbitrage opportunities are satisfied[3] (plus some other essentially technical conditions for things to be well behaved). Thus the theorem applies to a Black-Scholes kind of world in which the risk free rate is deterministic and the value of the market portfolio follows a Geometric Brownian Motion. In such a Black-Scholes type of world with a positive risk premium the state-price density is a decreasing function of the value of the risky asset. In this framework efficient portfolios must provide payoffs which are monotonic non-decreasing in the value of the risky asset at the horizon date. This squares quite well with intuitive notions concerning risk sharing within an economy.

Dybvig [1988a] goes on to describe and measure the inefficiency of a number of path dependent portfolio strategies. Under plausible assumptions for the market risk premium, he finds that a stop-loss strategy employed over a year may incur an efficiency loss of 80 or even 90 basis points. These figures were calculated on the assumption that the market index provides an expected risk premium of 8 percent. Since such a strategy may be expected to be in the market only about 50% of the time, the strategy is only earning a risk premium of about 4%. The efficiency loss can thus be of the order of 20 percent of the risk premium earned from being exposed to the market. We shall describe this kind of strategy as being 80 percent efficent. Lock-in strategies may incur an efficiency loss of up to 80 basis points a year, while a random timing strategy employed over a year (eg. by a would be market timer with no ability to forecast the market) can be inefficient by almost two hundred basis points a year. Finally, it is shown that although one year portfolio insurance is efficient for a one-year horizon, if such insurance is rolled over for several years it is soon inefficient to the extent of 50 basis points a year.

3. Measurement of the Inefficiency of Observed Portfolios

The purpose of this section is to show how we may extend Dybvig's results so as to make it possible to measure portfolio efficiency without needing to know the full details of the portfolio manager's strategy. In practice we are unlikely to know what the portfolio manager would have done along paths of the tree which did not occur. We consider what can be inferred if we simply observe the dollar funds invested in the risky asset at a limited number of discrete points in time. We assume that we observe a single implemented strategy y, as revealed by the dollars y_0, y_1,...,y_T invested in the economy's risky asset at the discrete time points 0,1,...,T. We further assume that the fund is intended to have an investment horizon of at least H, with H > T. For the purposes of this section and the next we shall assume that the economy provides a deterministic riskless interest rate, and that the price of the single risky asset, S, follows a Geometric Brownian Motion with a constant risk premium. (Alternatively you may legitimately visualise a geometric binomial process). These assumptions are very much stronger than we really require but assist with parts of the exposition. We shall consider what is involved in relaxing them in section 5.

We will use the symbol $\mathscr{E}_H$ to denote the set of all portfolio strategies which are efficient in the sense that they maximise expected utility for some non-decreseasing utility function defined at the horizon date H. This is the same notion of efficiency as employed by Dybvig.[4] What we want to learn is whether an observed strategy y belongs to $\mathscr{E}_H$, and if not we would like to measure the magnitude of the opportunity cost involved. We will define a norm $\|.\|$ so that we can use $\|x-y\|$ to measure how far one strategy x lies from another one y. What we have in mind is that x will be an efficient strategy, y will be the observed strategy, and we would like to choose our norm so that $\|x-y\|$ provides a measure of the opportunity cost of departing from an efficient strategy. To measure efficiency all we need to do is to find the member $x^* \in \mathscr{E}_H$ which lies closest to y, and $\|y-x^*\|$ will then tell us the extent to which the revealed policy y is inefficient. We therefore solve the optimization problem:

$$\underset{x^* \in \mathscr{E}_H}{\text{Minimize}} \qquad \|x^*-y\|.$$

It remains to show how we can conveniently identify which strategies belong to our set $\mathscr{E}_H$ of efficient strategies, how we may define a reasonably appropriate norm $\|.\|$, and how we can complete the calculations in a practical manner. We will deal with these technical aspects next.

Efficient Portfolio Strategies

The assumptions of our economy mean that our world is consistent with using the Black-Scholes model for evaluating options. Indeed any contingent claim based on the horizon value of the risky asset can be priced quite easily. In this economy efficient portfolio strategies are those which will result in horizon payoffs that are monotonic non-decreasing in the horizon value of the risky asset. For any such payoff function $f(S_H)$ we can use option pricing theory to determine what position it is necessary to have in the market risky asset for every (S,t) combination (with $t<H$). Thus we have a simple method for defining efficient strategies in terms of the payoffs they result in, and for calculating the holdings that they imply should be held in the risky asset through time.

Norms and Opportunity Costs

Without loss of generality we may assume that the risk-free rate is not simply deterministic but is actually zero. This is because we can always work with discounted values without affecting any aspects of the structure involved. We shall let μ denote the expected rate of return on the risky asset, and σ its standard deviation. We shall consider the risks and rewards associated with a strategy that consists of an efficient strategy plus an orthogonal disturbance. Consider a strategy which invests dollar amounts in the risky asset:

$$y_t = x_t + \varepsilon_t$$

where x_t is drawn from an efficient strategy and ε_t is a noise term with zero mean over t, variance$[\varepsilon_t] = \upsilon^2$, and ε_t is uncorrelated with x_t. We know, by assumption, that the strategy x_t provides outcomes which are suitably monotonic in the state-price density. However, the presence of the noise term ε_t creates a scattering of outcomes which destroys that monotonicity. Instant by instant we can apply the Payoff Distribution Pricing Theorem to find the cost imposed by ε_t. At any time t, the expected rate of return and variance rate of return obtained from the strategy y (conditional on ε_t) are:

$$E[\text{return}(y|\varepsilon_t)] \quad = \mu\,(x_t + \varepsilon_t)$$

$$Var[\text{return}(y|\varepsilon_t)] \quad = \sigma^2\,(x_t + \varepsilon_t)^2.$$

When we take account of expectations over possible realizations of ε_t these become:

$$E[\text{return}(y)] \quad = \mu\, x_t$$

$$Var[\text{return}(y)] \quad = E[\sigma^2\,(x_t + \varepsilon_t)^2]$$

$$= \sigma^2\, E[x_t^2 + 2\, x_t\, \varepsilon_t + \varepsilon_t^2]$$

$$= \sigma^2\,(x_t^2 + \upsilon^2).$$

Note that, there has been no increase in the expected return compared to the efficient strategy x_t, but that the variance of return has increased by $\sigma^2\,\upsilon^2$. For this level of risk the strategy should be earning an expected return of

$$\mu \sqrt{x_t^2 + \upsilon^2}$$

instead of $\mu\, x_t$. We may interpret

$$\mu\,(\sqrt{x_t^2 + \upsilon^2} - x_t)$$

as the instantaneous inefficiency loss associated with y_t. The average loss over time is

$$L = \Sigma_t\; \mu\,(\sqrt{x_t^2 + \upsilon^2} - x_t)\; \Delta t/T.$$

The average risk premium gained is

$$G = \Sigma_t\; \mu\, x_t\; \Delta t/T,$$

and we may calculate the percentage efficiency loss as

$$\text{Percentage Efficiency Loss} \quad = \quad L/G\ \%.$$

This approximates to

$$\text{Percentage Efficiency Loss} \quad = \quad ½\ \upsilon^2/(E[x_t])^2\ \%.$$

One major advantage of this form is that it does not involve the risk premium μ. We may also note that if we fit an efficient strategy x_t econometrically to y_t it is no great problem for us that the residual ε_t are likely to display considerable autocorrelation. Since $E[x_t] = E[y_t]$, we conclude that we can fit x_t to y_t using a least-squares criterion, and then obtain from the residual variuance υ^2 a good approximation to the opportunity cost implied by the departures from an efficient strategy.

The Estimation Procedure

In order to measure efficiency, we form a very natural linear approximation to efficient strategies, and thus use a linear estimation technique to estimate the nearest efficient strategy to the one observed. The whole process is illustrated graphically in Figure 1. First, define a grid consisting of different values $P_1 < P_2 <....< P_{k-1}$ that S_H might attain[5] at the horizon H. Define the following k contingent claims based on S_H:

$$
\begin{aligned}
c_1(S_H) &= S_H - P_1 && \text{for } S_H < P_1\\
&= 0 && \text{for } S_H \geq P_1\\
c_j(S_H) &= 0 && \text{for } S_H < P_{j-1}\\
&= S_H - P_{j-1} && \text{for } S_H \in [P_{j-1}, P_j]\\
&= P_j - P_{j-1} && \text{for } S_H > P_j\\
& \quad \text{for } j = 2,....,k\text{-}1. &&\\
c_k(S_H) &= 0 && \text{for } S_H < P_{k-1}\\
&= S_H - P_{k-1} && \text{for } S_H \geq P_{k-1}.
\end{aligned}
$$

These are easily recognised as the payoffs from conventional options positions: c_1 is a short position in a put at P_1 (ie. with exercise price P_1). The claims c_2 to c_{k-1} are bull spreads (eg. c_2 is long a call at P_1 and short a call at P_2), and c_{k-1} is a long position in a call at P_{k-1}. Their payoffs as a function of S_H are depicted on the right hand side of Figure 1.

By choosing a suitably fine grid it is possible to ensure that functions of the form

$$c(S_H) \quad = \quad \Sigma_j \; \alpha_j \; c_j(S_H)$$

with $\alpha_j \geq 0$ can approximate arbitrarily closely[6] to any continuous monotonic non-decreasing function $f(S_H)$ which might characterize an efficient portfolio strategy in $\mathscr{E}_H$. This summation is illustrated at the bottom of the right hand side of Figure 1. Let $x_{tj}(S_t)$ denote the number of dollars to be invested in the risky asset at date t (t=0,1,...,T) in order to achieve the claim c_j, where S_t is the price of the asset at date t. These sequences through time of the holdings (conditional on S_t) required to synthesize each of the option spreads are illustrated on the left hand side of Figure 1. The investment required in the risky asset at date t in order to achieve the payoffs f(S) at the horizon date is therefore:

$$x_t(S_t) \quad = \quad \Sigma_j \; x_{tj}(S_t) \; \alpha_j.$$

This is the time sequence of the total holding required to synthesize the total payoff $c(S_H)$, and it is illustrated at the bottom left hand side of Figure 1. Thus the estimation is completed by solving the following conventional linear least squares estimation for the α_j subject to non-negativity constraints:

$$\text{Minimize} \quad \Sigma_t \; \varepsilon_t^2, \text{ where}$$

$$y_t \quad = \quad \Sigma_j \; x_{tj}(S_t) \; \alpha_j \; + \; \varepsilon_t, \quad (t = 0, 1,..., T),$$

and $\qquad \alpha_j \geq 0, \quad (j = 1,....,k).$

Writing this in matrix notation as

$$y = X\alpha + \varepsilon,$$

the first order conditions are

$$X'X\alpha - z = X'y$$

$$\alpha' z = 0, \qquad \alpha, z \geq 0.$$

These equations[7] are readily solved using a computer package (or otherwise). Note that, particularly for large k, the columns of X will exhibit multicollinearity. The non-negativity constraints on α nevertheless make the optimization well behaved, though it is possible that α may not be uniquely determined.[8]

Once the α have been estimated, $\mathbf{X\alpha}$ provides a T+1 vector of the market exposure for the nearest efficient strategy to the observed $y_0, y_1, \ldots, y_T$. Note that this can also be extended to give fitted values for dates t at which y_t was not observed. We also get from $\alpha'\mathbf{c}(\mathbf{S_H})$ an objective function for the fund defined as wealth contingent on the value of the index at the horizon date H. The interpretaton of this function is that it gives the objective contingent on S_H of that investor with horizon H, for for whom the observed sequence $y_0, y_1, \ldots, y_T$ gives the minimum inefficiency. We can think of this informally as the objective of the investor who (with hindsight) would have been most likely to wish to invest in the observed portfolio strategy.

4. Examples of Implementing the Technique

We now show the results of some simulations of the technique just described. In each case we have used the following parameter values, which are the same as in Dybvig [1988a]:

risk free rate, $r = 0.08$
rate of drift of risky asset, $\mu = 0.16$
annualised standard deviation = 0.20

The simulations have been performed to look at a single sample path over a period of two years with monthly observations, ie, just 24 observations in all. The approximation is made with $k = 7$ (with more observations a larger value would have been chosen). These low numbers have been deliberately chosen to demonstrate the power and robustness of the technique - prettier pictures might have been obtained for larger parameter values. In this regard we may also note that if we were in Dybvig's position of knowing in advance the nature of the management strategy we could apply our method to multiple simulations, N (say), so that **X** becomes (T+1)N x k etc.

Figure 2 shows a time series for the simulated market index, observed monthly over two years. It starts at a base of 100.0. The simulated portfolio starts with initial capital of $1,000,000. The remaining figures and tables refer to estimations performed for a number of portfolio strategies. The strategies reported are as follows:

Strategy 1 Invest in bonds plus 10,000 two year calls at 130. This is a simple static portfolio insurance policy with a two year horizon. When we evaluate this strategy with a two year horizon it is efficient. When we evaluate it using a five year horizon it is clearly not. The estimated ex-post inefficiency under the five year horizon is actually rather modest at 0.9%. For other realizations of the sample path of the market it might have been higher (or lower) and there is no reason to expect the ex-post inefficiency revealed by a single sample path to bear any relationship to the ex-ante inefficiency (eg. of Dybvig) calculated over all sample paths. Figure 3 shows the estimated objective under the five year horizon, and Figure 4 plots the actual (dotted line) and fitted efficient (solid line) strategies against time.

Strategy 2 Lock-in strategy at $1.2 million. This strategy has an ex-post inefficiency of 5.7% with respect to a three year investment horizon. Figure 5 shows the corresponding estimated objective, and Figure 6 plots the actual and efficient strategies, but this time against the market index level instead of time.

Strategy 3 Repeated portfolio insurance. The fund has been simulated as purchasing nine-month portfolio insurance (ie. bonds plus nine month calls) and this is renewed without changing the market exposure every six months (ie. when the calls have three months to run). This seems a fairly mild strategy. Nevertheless the analysis shows that from the perspective of a three year horizon it is inefficient to the tune of 2.5%. Figures 7 and 8 again show the estimated objective, and the actual and fitted strategies.

Some other simulations have also been performed which there is not space here to report in any detail. In particular constant proportional portfolio insurance of the kind proposed by Black and Jones [1987], and Black and Perold [1990] has been simulated. It is a special case of the stationary policies described by Brennan and Schwartz [1987] as being efficient for any horizon. The simulations confirm the efficiency of such policies, and also are able to detect the inefficiency which occurs if the growing floor for the fund $F(t) = F\ e^{rt}$ is replaced by a floor which grows at a rate different from the riskless interest rate.

5. Generalizations and Extensions

The results described above are most encouraging. We might not have expected to be able to recover so much information about objectives and efficiency from a handful of portfolio position data along a single realised path of outcomes[9]. Nevertheless, the assumptions made so far are uncomfortably strong: an economy with a deterministic interest rate and a single risky asset following Geometric Brownian Motion with a constant risk premium. For practical work we would certainly need to relax these considerably. In this section we shall see that there are no conceptual hurdles to extending the method to encompass economies where there are several state variables and also multiple risk premia which vary through time.

We shall discuss first how the model may be employed if the risk premium in our original simple economy varies through time. Recent empirical studies (for example those of Fama and French [1988], and Poterba and Summers [1988]) suggest that a constant risk premium is unlikely, so this kind of generalization becomes very important. We now have the following process for the risky asset:

$$dS = \mu(S,t)\, S\, dt + \sigma\, S\, dz.$$

First we note that the Black-Scholes option valuation model is still valid as it does not depend on $\mu(S,t)$. We need to assume specifically that $\mu(S,t)$ is consistent with an equilibrium characterized by investors who maximize the expected utility of their consumption[10] at a horizon date at least as distant as H. This then implies that efficient portfolio strategies can be spanned as before and we can certainly test whether a particular strategy is efficient. The main components of our methodology still apply. The only problem is the measurement of the degree of the efficiency loss. Even here our percentage efficiency loss measure (which does not involve μ) should provide a robust though inexact measure of inefficiency.

We describe next how our technique can be generalized to models of the economy which contain several state variables and several investment instruments. The motivation for this is to be able to encompass such realities as stochastic interest rates, inflation and volatility. We will retain the assumption that the market is complete. Consider two state variables r and s, so that there exist state prices p_{rs} and state probabilities π_{rs}. As before, for an efficient strategy, horizon wealth must be monotonic nonincreasing in the state-price density function, $\xi_{rs} = p_{rs} / \pi_{rs}$. We may construct efficient portfolios as non-negative linear combinations of bear spreads on the state-price density itself. An appropriate norm must also be found. The difficulty of this step depends somewhat on the details and complexity of the model. For many models a least squares approach will still work. It is clear that in principle the number of state variables is no obstacle.

We have so far ignored a difficulty which arises depending on how far the portfolio wealth departs from the value predicted by the fitted strategy. This can occur through the cumulative effects of not employing an "efficient" strategy. A resulting difficulty is that deviations between the value of the actual portfolio and that of the fitted one will, if sufficiently large, justify further departures in the subsequent strategy. The issue has been ignored in this paper as likely to be a second order problem. However, it seems to lead to an iterative estimation and deserves further investigation.

In this paper we have also ignored the important issue of forecasting ability. Market timing behaviour based on forecasts of the future may justifiably lead a fund to depart from the kinds of strategies we have labeled as efficient. These strategies only remain efficient in a world which is characterized by the kind of equilibrium we have briefly described above. If instead we wish to consider more realistic assumptions in which assymetries of information provide some participants with differential forecasting ability, we shall have to re-intepret our inefficiency measures as the opportunity costs associated with trading on information.

One final reason why managers may not follow strategies like our "efficient" ones in practice concerns transactions costs. There has been significant recent progress on the topic of intemporal portfolio management under transactions costs (eg. Taksar, Klass and Assaf [1988], Dumas and Luciano [1989], and Davis and Norman [1990]) and in the replication of options positions under costs (eg. Hodges and Neuberger [1989]). It would be interesting to see what implications that work carries for the technique we have described here.

6. Conclusions

The paper has demonstrated a practical procedure for estimating (from minimal data) the opportunity costs implied by dynamic portfolio strategies which are inefficient. The measure is based on observing what a manager actually did, rather than on knowing what he would do under every eventuality. The technique provides a measure of the extent to which the strategy used was imperfectly diversified in an intertemporal sense. It is analogous to the opportunity cost imputed by Fama [1972] for unsystematic risk stemming from imperfect cross-sectional diversification. The analysis builds on earlier work by Dybvig [1988a], and the type of estimation proposed is not dissimilar to that in Admati and Ross [1985]. Simulations demonstrate that the technique is remarkably robust.

Footnotes

1. It is both significant and typical, for instance, that Kritzman [1990] which covers a great deal of material on various types of dynamic strategies, contains no material at all on performance measurement.

2. It may be noted that Dybvig [1988a] contains a misprint, with this part printed incorrectly as "Consumption is nondecreasing in the terminal state-price density".

3. In this framework the state-price density is essentially the Radon-Nikodym derivative of the risk neutral probability measure with respect to the objective one.

4. Note that this concept of efficiency is far more general than that of mean-variance efficiency, which would be quite inappropriate in this context. See Dybvig and Ingersoll [1982] for a discussion of the role of options in an economy in which tyhe underlying assets are priced in a mean-variance framework.

5. Note that these subscripts on P do not refer to dates as does that on S.

6. See, for example, Powell [1981] for related convergence theorems for spline functions such as these. The extension of these results to our specific problem is straightforward.

7. As stated, the equations ensure that $E[x_t\ \varepsilon_t] = 0$, but not that $E[\varepsilon_t] = 0$. To ensure that this latter condition always holds exactly we need to append an additional constraint.

8. The use of Singular Value Decomposition (see, for example, Press et al [1989], Golub and Van Loan [1983]) provides a robust way of performing the computations.

9. Part of the reason for this lies in the relation to the following recoverability result described by Dybvig [1986]. The key assumptions are that the portfolio is managed under an efficient strategy and that it is possible to calculate the state-price density function. Under these assumptions if we can observe the asset mix along a single realised sample path, we can recover (and actually construct) the entire portfolio strategy over the whole probability tree. It turns out that small disturbances in the strategy employed along the observed path result in large perturbations to the terminal wealth profile (as a function of the value of the market portfolio) that would be derived under the path independence assumption. (This casts some doubt on the conjecture in that paper that the utility function is recoreable in the continuous limit). Conversely, monotonicity of terminal consumption implies considerable smoothness in the corresponding efficient investment strategy. It is this smoothness which makes our approach well behaved, and also means that the grid of horizon date prices can be rather coarse.

10. The papers of Merton [1971], [1973], Breeden [1979] and Cox, Ingersoll and Ross [1985] treat the issues of optimal consumption, optimal portfolio selection and equilibrium within continuous time economies. Further properties are developed by Cox and Huang [1987]. Bick [1987a, 1987b] deals particularly with a simpler Black-Scholes like economy. The characterisation of equilibria of the kind described is the subject of a further paper by the author.

References

Admati, A R and S A Ross (1985), "Measuring Performance in a Rational Expectations Equilibrium Model", *Journal of Business*, 58, 1-26.

Banz, R and M Miller (1978), "Prices for State Contingent Claims: Some Estimates and Applications", *Journal of Business*, 51, 653-672.

Bick, A (1987a) "On the Consistency of the Black-Scholes Model with a General Equilibrium Framework", *Journal of Financial and Quantative Analysis* 22, 259-275.

Bick, A (1987b), "On Viable Diffusion Price Processes", Working Paper No 449, Salomon Brothers Center for the Study of Financial Institutions, New York University.

Black, F and R Jones (1987), "Simplifying Portfolio Insurance", *Journal of Portfolio Management*, Fall 1987, 48-51.

Black, F and A F Perold (1990), "Theory of Constant Proportion Portfolio Insurance", WP90-028, Harvard Business School.

Black, F and M Scholes (1973), "The Pricing of Options and Corporate Liabilities", *Journal of Political Economy*, 81, 637-654.

Bookstaber R and R Clarke (1984), "Options Portfolio Strategies: Measurement and Evaluation", *Journal of Business*, 57, 469-492.

Bookstaber R and R Clarke (1985), "Problems in Evaluating the Performance of Portfolios with Options", *Financial Analysts' Journal*, Jan-Feb 1985, 48-62.

Breeden, D T (1979), "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities", *Journal of Financial Economics*, 7, 265-296.

Breeden, D T and R H Litzenberger (1978), "Prices of State Contingent Claims Implicit in Option Prices", *Journal of Business*, 51, 621-651.

Brennan, M J and E S Schwartz (1987), "Time Invariant Portfolio Insurance Strategies", *Journal of Finance*, 43, 283-299.

Cox and Huang (1987), "Optimal Consumption and Portfolio Policies when Asset Prices Follow A Diffusion Process", WP# 1926-87, AP Sloan School of Management, MIT.

Cox, J C, Ingersoll, J E and S A Ross (1985), "An Intertemporal Equilibrium Model of Asset Prices", *Econometrica*, 53, 363-385.

Cox, J C and H E Leland (1982), "On Dynamic Investment Strategies", Proceedings of the Seminar on the Analysis of Security Prices", Centre for Research in Security Prices, University of Chicago.

Davis, M H A and A R Norman (1990), "Portfolio Selection With Transactions Costs", *Mathematics of Operations Research*, 15, 676-713.

Dumas, B and E Luciano (1989), "An Exact Solution to the Portfolio Choice Problem Under Transactions Costs", Working Paper, Wharton School, University of Pennsylvania.

Dybvig, P H (1981), "A Multiperiod Recoverability Result", Working Paper #16, Yale School of Management, (1981, revised 1982 and July 1986).

Dybvig, P H (1988a), "Inefficient Dynamic Portfolio Strategies, or How to Throw Away a Million Dollars", *Review of Financial Studies*, 1, 67-88.

Dybvig, P H (1988b), "Distributional Analysis of Portfolio Choice", *Journal of Business*, 61, 369-393.

Dybvig, P H and J E Ingersoll (1982), "Mean-Variance Theory in Complete Markets", *Journal of Business*, 55, 233-51.

Dybvig, P H and S A Ross (1985a), "Performance Measurement Using Excess Returns Relative to a Security Market Line: Differential Information", *Journal of Finance*, 40, 384-399.

Dybvig, P H and S A Ross (1985b), "The Analytics of Performance Measurement Using a Security Market Line", *Journal of Finance*, 40, 401-416.

Fama, E F (1972), "Components of Investment Performance", *Journal of Finance*, 27, 551-567.

Fama, E F and French K R (1988), "Dividend Yields and Expected Stock Returns", *Journal of Financial Economics*, 22, 3-25.

Golub, G H and C F Van Loan (1983), *Matrix Computations*, John Hopkins University Press.

Grinblatt, M and S Titman (1989), "Portfolio Performance Evaluation: Old Issues and New Insigths", *Review of Financial Studies*, 2, 393-421.

Henrikson R D and R C Merton (1981), "On Market Timing and Investment Performance II: Statistical Procedures for Evaluating Forecasting Skills", *Journal of Business*, 54, 513-533.

Hodges, S D and A Neuberger (1989), "Optimal Replication of Contingent Claims Under Transactions Costs", *The Review of Futures Markets*, 8,222-239.

Jagannathan, R and R A Korajczyk (1986), "Assessing the Market Timing Performance of Managed Portfolios", *Journal of Business*, 59, 217-235.

Jensen, M (1969), "Risk, the Pricing of Capital Assets, and the Evaluation of Investment Portfolios", *Journal of Business*, 42, 167-247.

Kritzman M P (1990), *Asset Allocation for Institution Portfolios*, R D Irwin.

LIFFE/LTOM (1990), The Reporting and Performance Measurement of Futures and Options, LIFFE/LTOM Consultative Document, May 1990.

Merton R C (1971), "Optimum Consumption and Portfolio Rules in a Continuous-Time Model", *Journal of Economic Theory* 3, 373-413.

Merton R C (1973), "An Intertemporal Asset Pricing Model", *Econometrica*, 41, 867-887.

Merton R C (1981), "On Market Timing and Investment Performance. 1. An Equilibrium Theory of Value for Market Forecasts", *Journal of Business*, 54, 363-406.

National Association of Pension Funds "Committee of Enquiry Report into Investment Performance Measurement", December 1990.

Poterba, J M and L H Summers (1988), "Mean Reversion in Stock Prices: Evidence and Implications", *Journal of Financial Economics*, 22, 27-59.

Powell, M J D (1981), *Approximation Theory and Methods*, Cambridge University Press.

Press, W H, B P Flannery, S A Teukolsky and W T Vetterling (1989), *Numerical Recipes: The Art of Scientific Computing*, Cambridge University Press.

Sharpe, W F (1966), "Mutual Fund Performance", *Journal of Business*, 39, 119-138.

Sharpe, W F (1975), "Likely Gains from Market Timing", *Financial Analyst's Journal*, March-April, 60-69.

Taksar, M, M J Klass and D Assaf (1988), "A Diffusion Model for Optimal Portfolio Selection in the Presence of Brokerage Fees", *Mathematics of Operations Research*, 13, 277-94.

Treynor, J (1965), "How to Rate Management of Investment Funds", *Harvard Business Review*, 44, 131-136.

Figure 1

THE APPROXIMATION PROCEDURE

Approximating Strategies of $\mathcal{E}_H$ Approximating Monotonic Payoffs

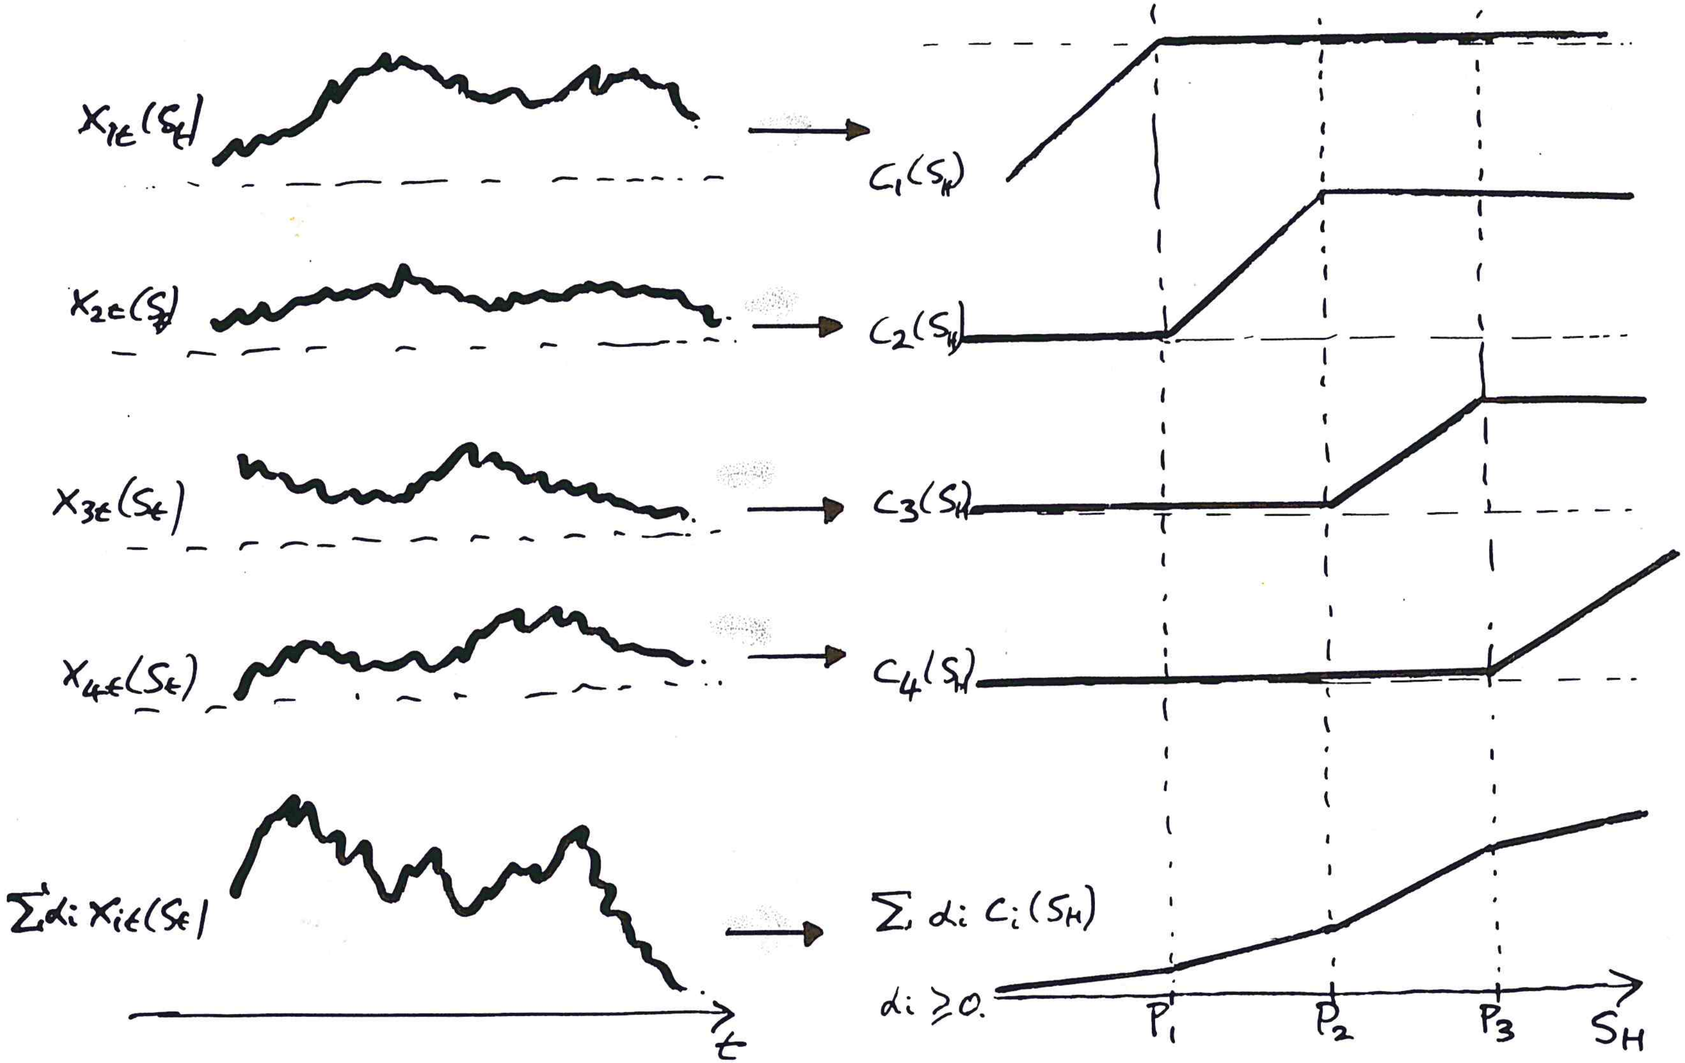

Figure 2
Market Index

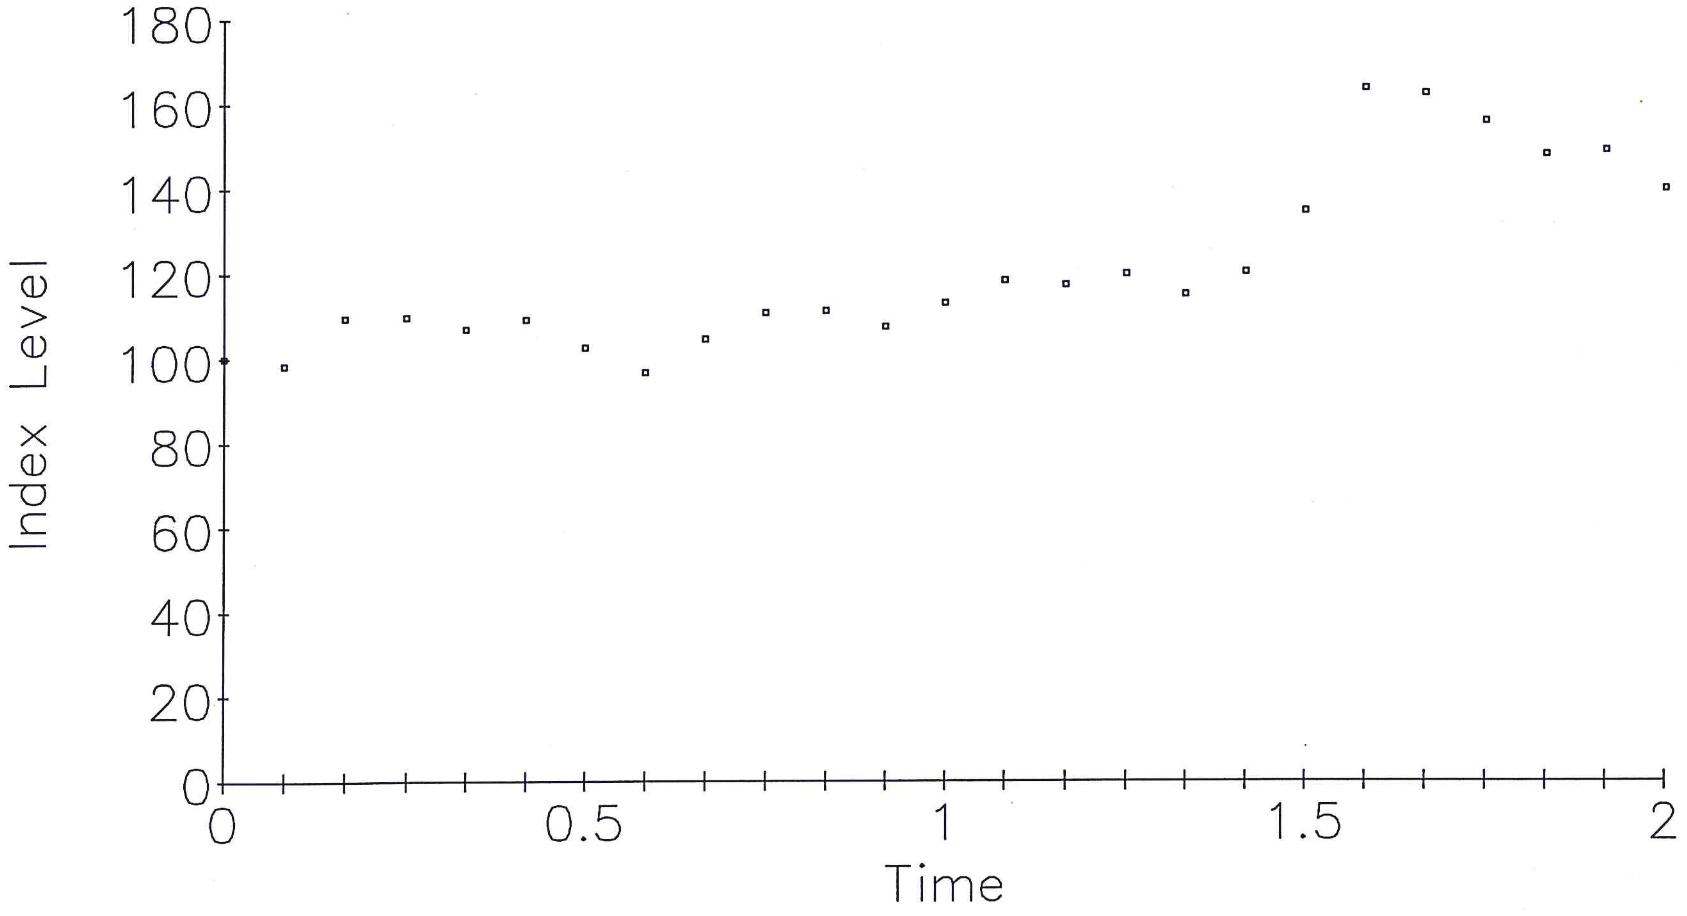

Figure 3: Estimated Objective
Call Option Strategy

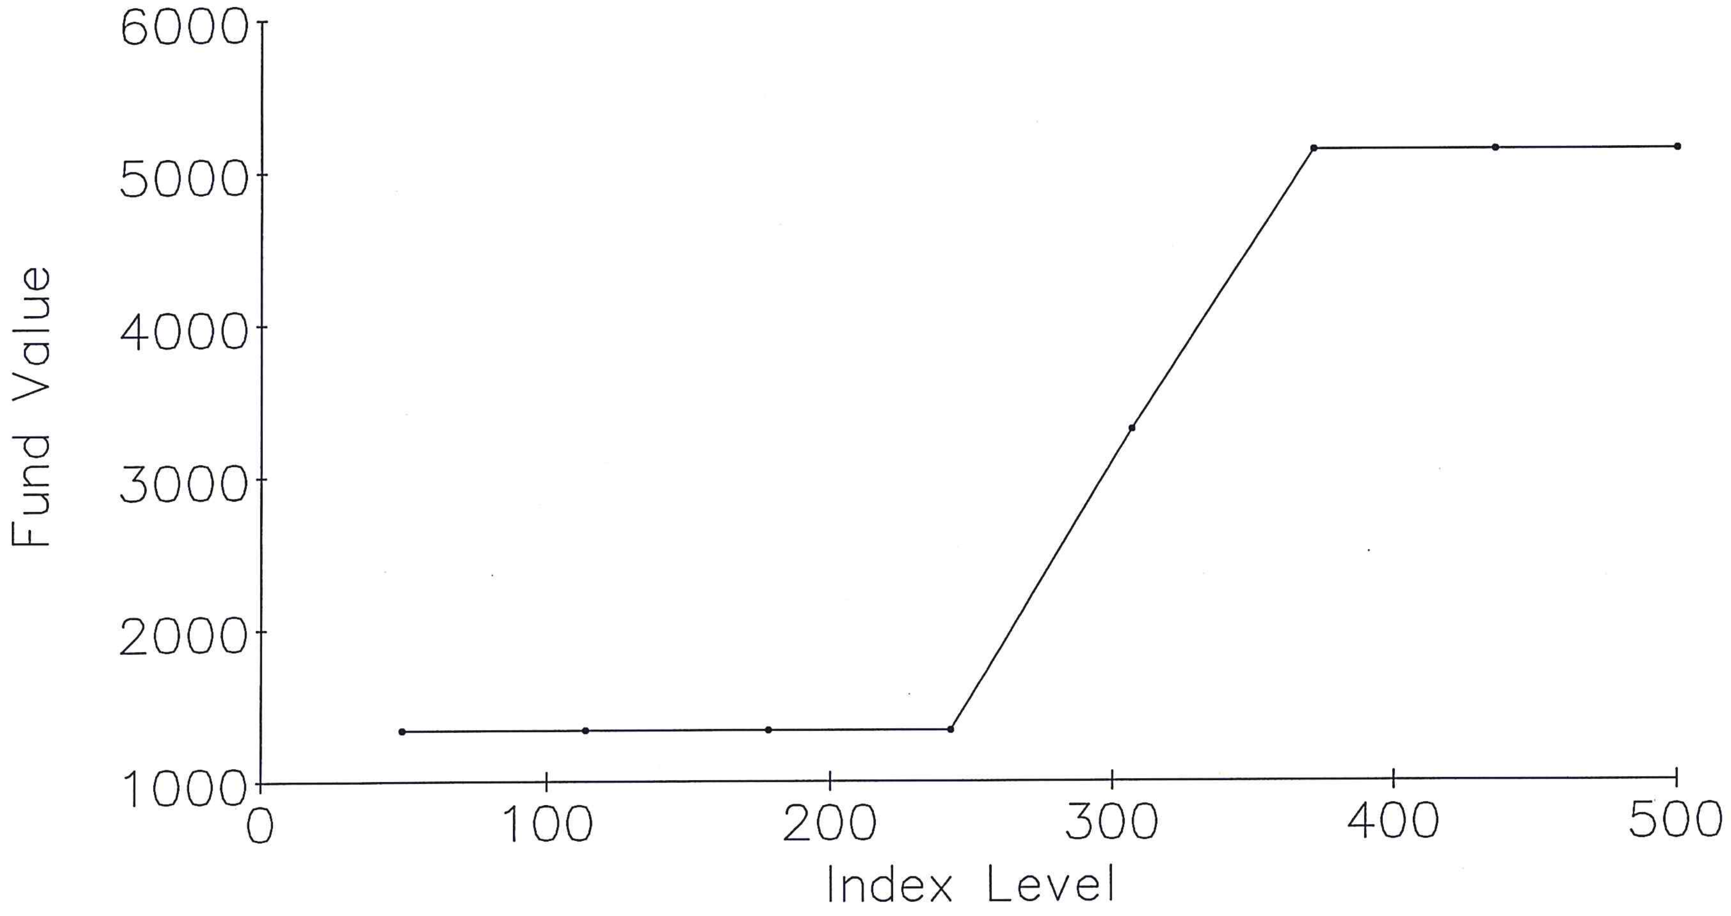

Figure 4
Call Option Strategy

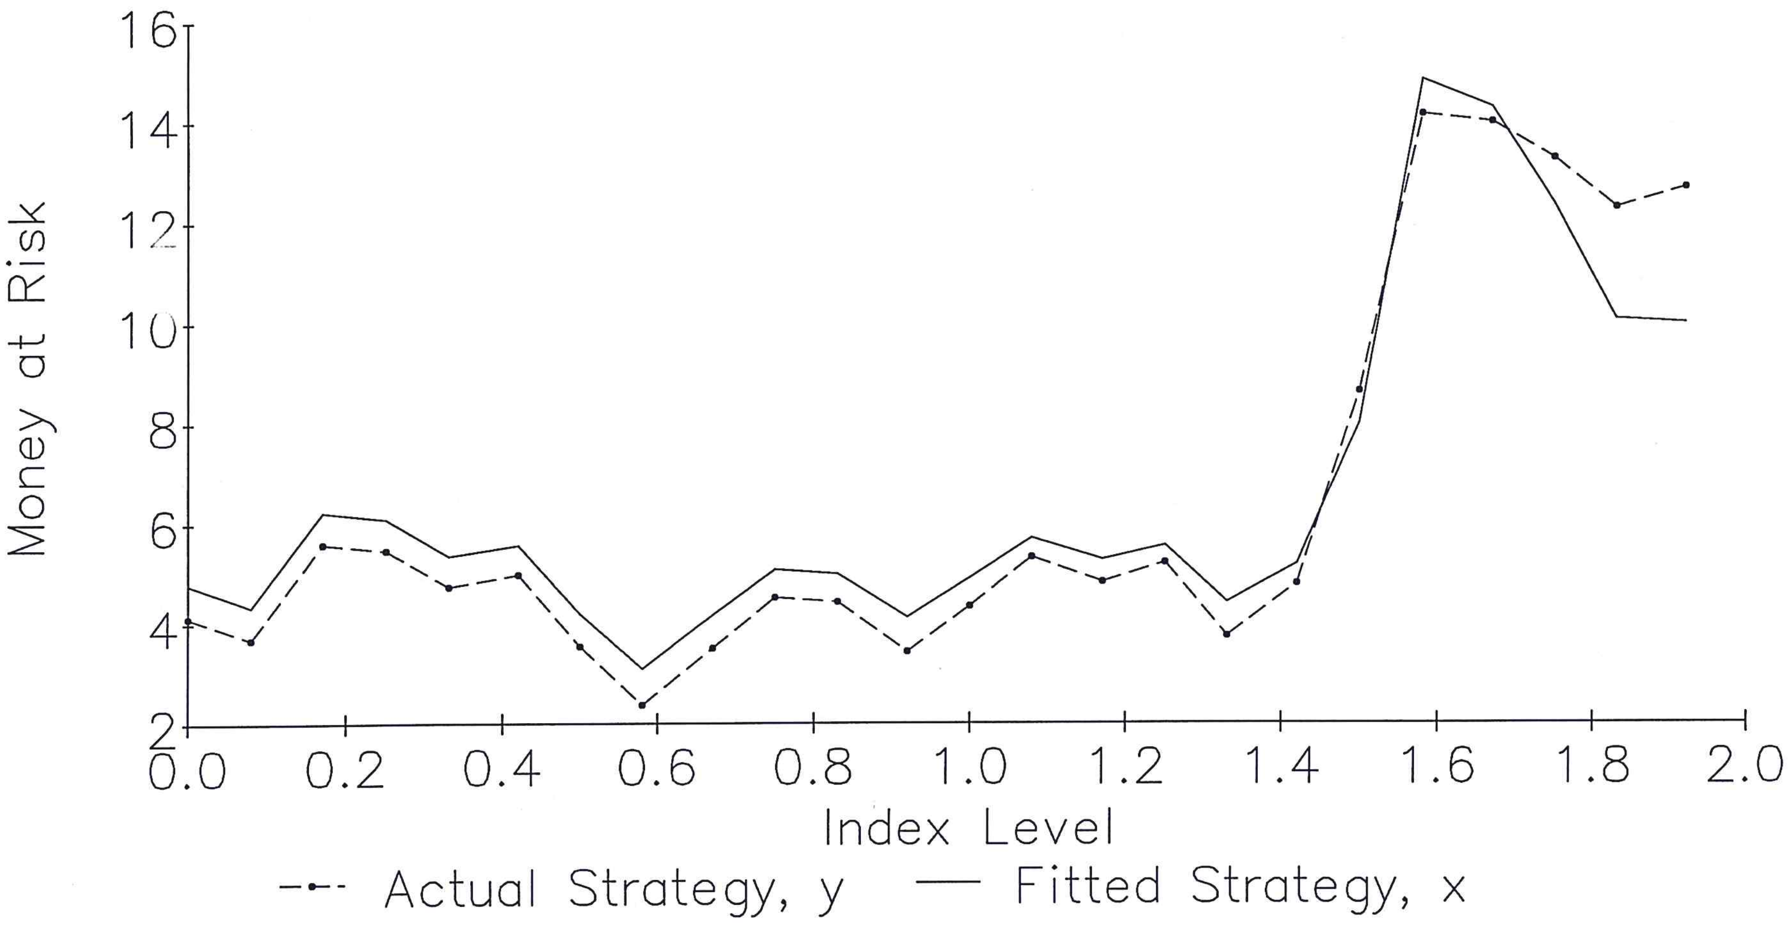

Figure 5: Estimated Objective
Lock-In Strategy

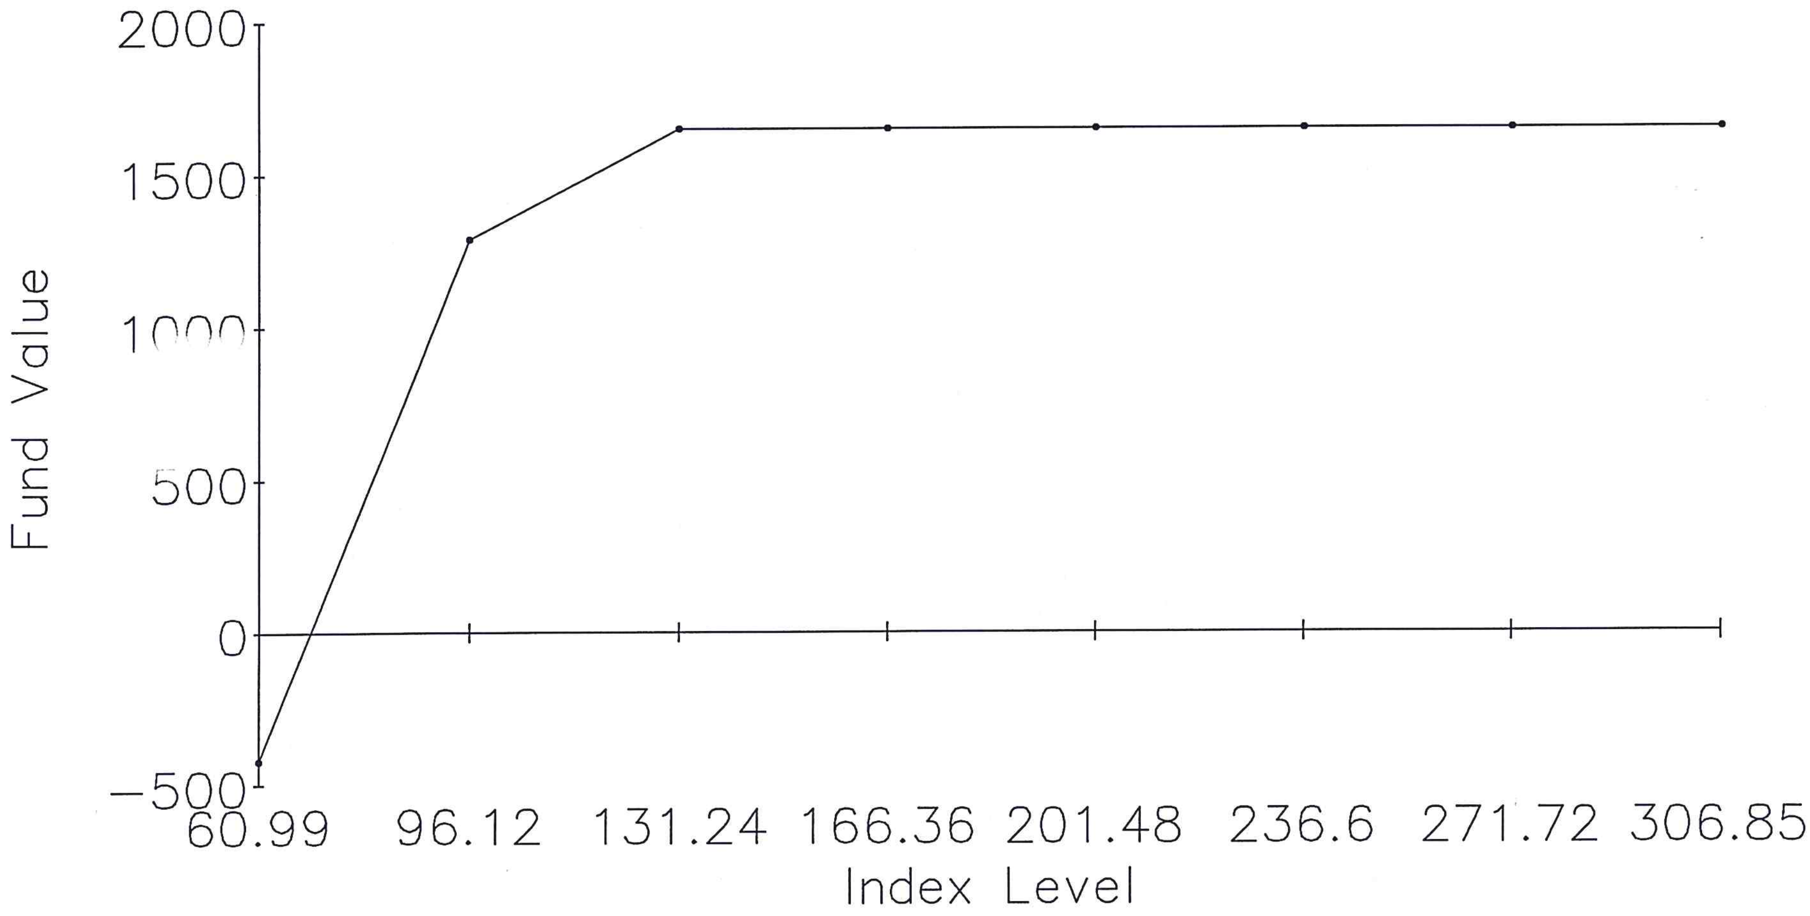

Figure 6
Lock-in Strategy

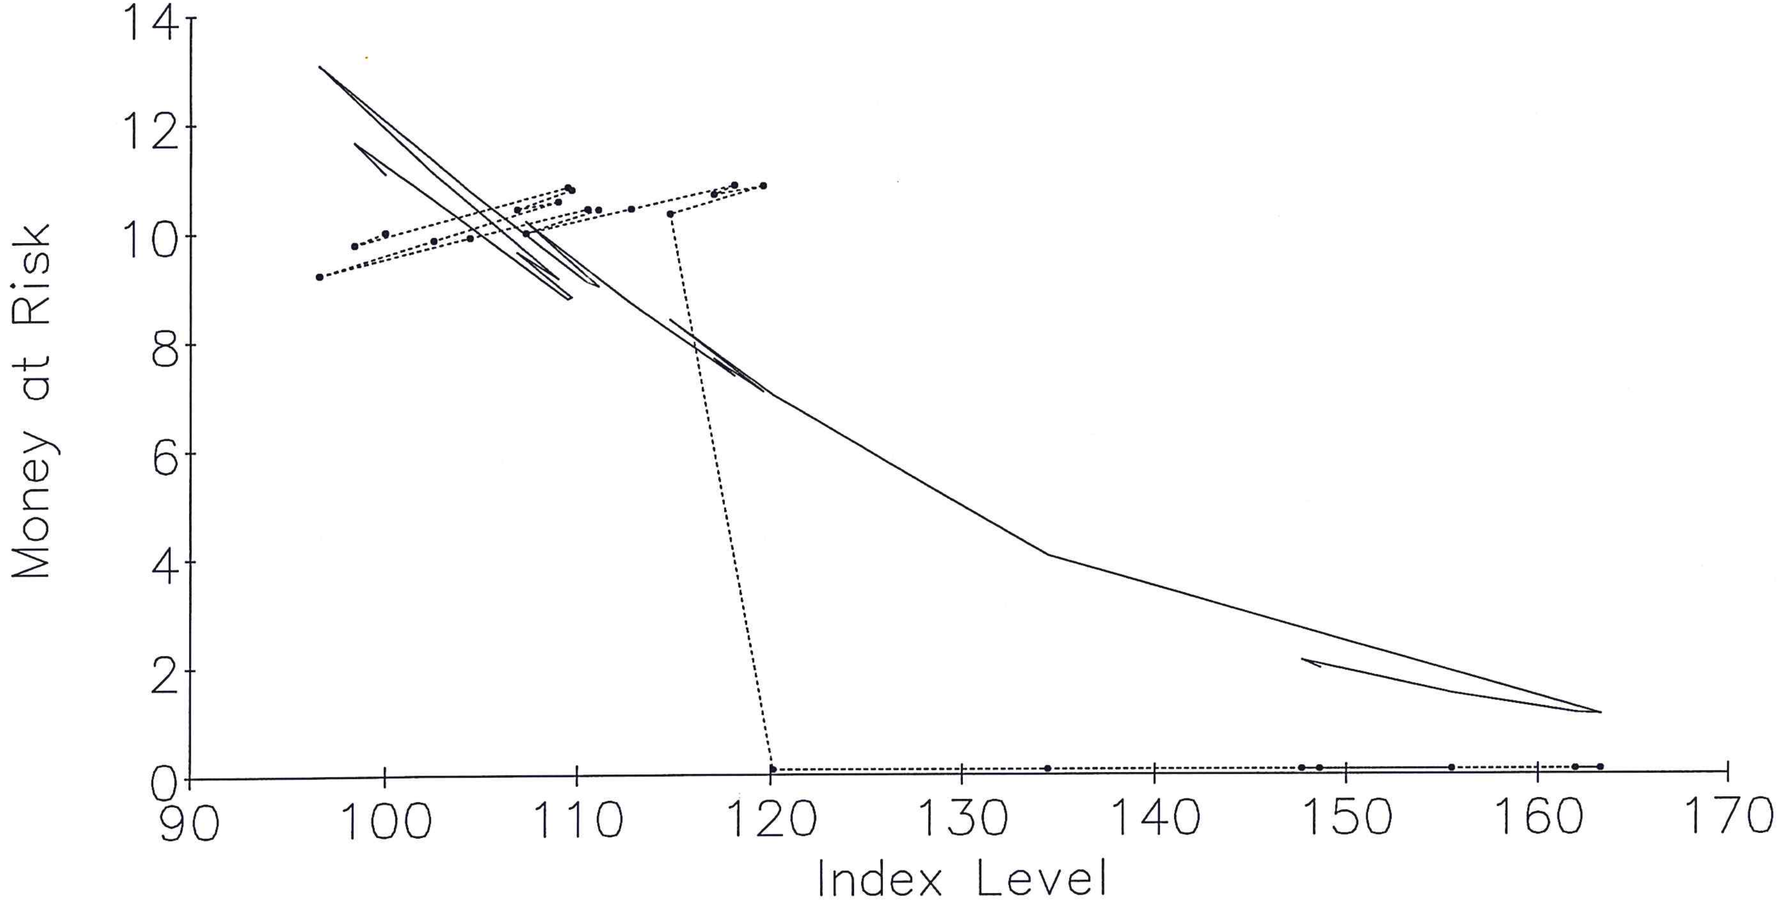

Figure 7: Estimated Objective:
Repeated Portfolio Insurance

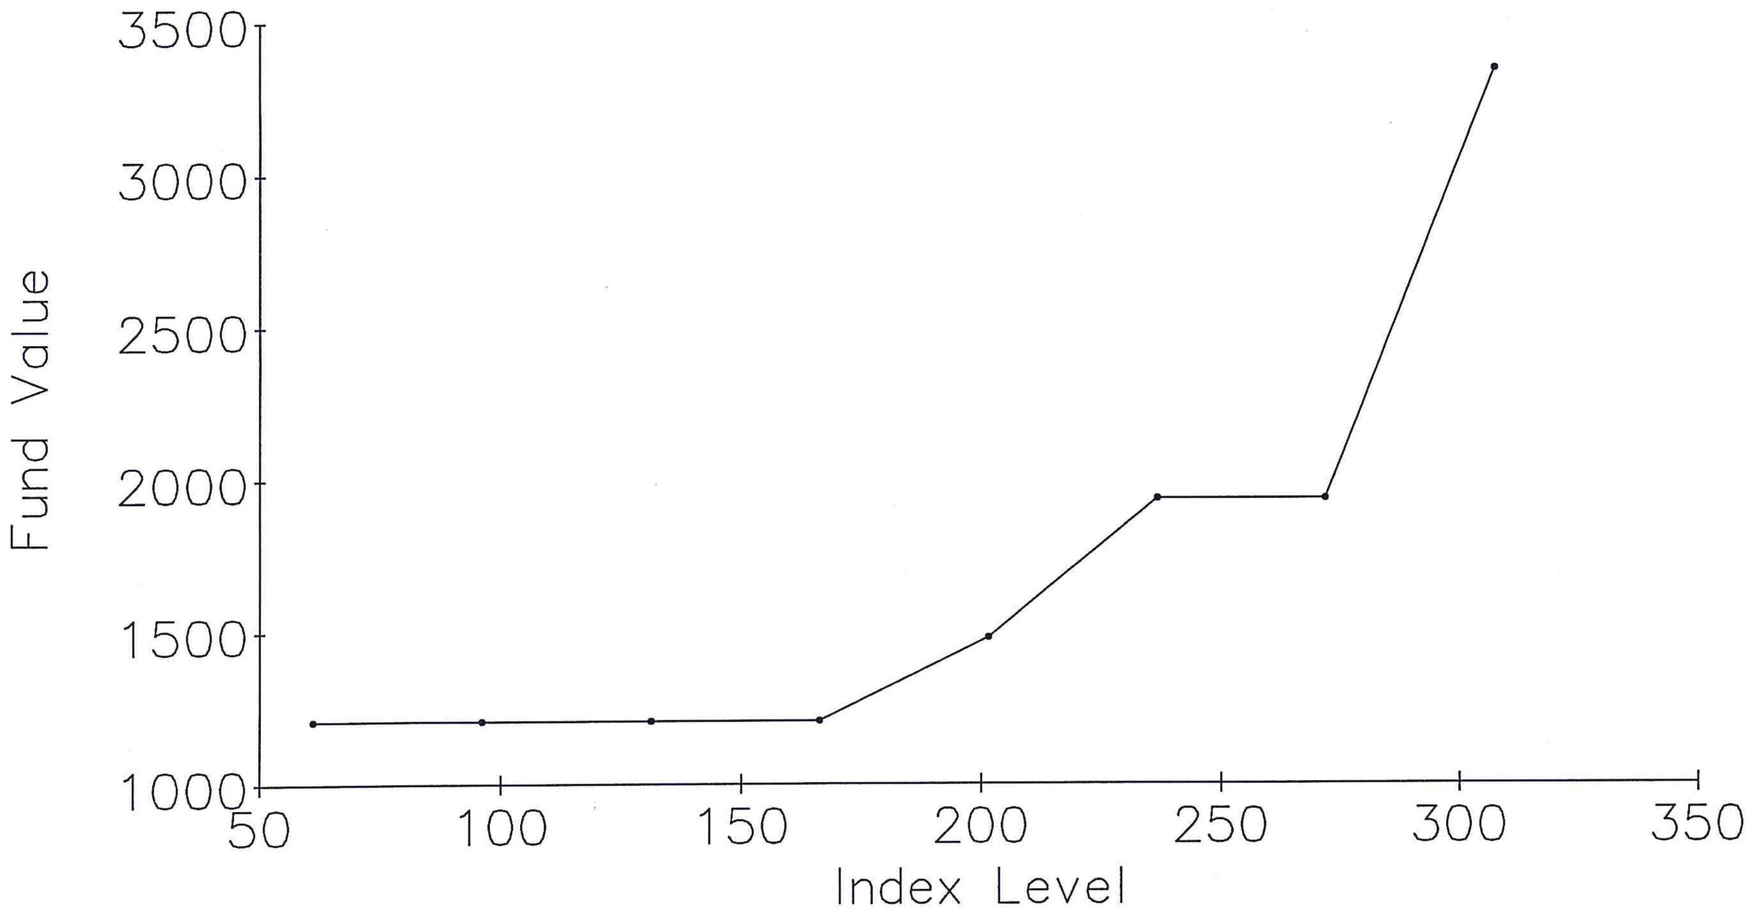

Figure 8
Repeated Portfolio Insurance