COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

AT YALE UNIVERSITY

Box 2125, Yale Station

New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 826-R

INEFFICIENT DYNAMIC PORTFOLIO STRATEGIES

or

HOW TO THROW AWAY A MILLION DOLLARS IN THE STOCK MARKET

Philip H. Dybvig

Revised January 1988

Inefficient dynamic portfolio strategies

or

How to throw away a million dollars in the stock market

Philip H. Dybvig [†]

Send correspondence (including requests for reprints) to:

Philip H. Dybvig
Yale School of Management
Box 1A
New Haven, CT 06517
203-432-6020 (O)
203-248-0051 (H)

November, 1985

Revised January, 1988

[†]School of Management, Economics Department, and Cowles Foundation at Yale University. I am grateful for helpful discussions with Michael Brennan, Stephen Brown, Kent Dybvig, David Feldman, Mike Granito, Roger Ibbotson, Jon Ingersoll, Alan Kraus, Steve Ross, Eduardo Schwartz, and participants in various seminars. I am also grateful for financial support from the Sloan Research Fellowship program.

Abstract

A number of portfolio strategies followed by practitioners are dominated because they are incompletely diversified over time. The Payoff Distribution Pricing Model is used to compute the cost of following undiversified strategies. Simple numerical examples illustrate the technique, and computer-generated examples provide realistic estimates of the cost of some typical policies using reasonable parameter values. The cost can be substantial and should not be ignored by practitioners. A section on generalizations shows how to extend the analysis to term structure models and other general models of returns.

1. Introduction

Portfolio managers regularly use a number of dynamic portfolio strategies that have not received careful theoretical analysis. Some examples are lock-in strategies, stop-loss strategies, rolling over portfolio insurance, and contingent immunization. The lack of analysis has been due largely to the inadequacy of the traditional theoretical tools. Specifically, mean-variance analysis is not valid when the portfolio return is non-linearly related to market returns, as it will be under these strategies.[1] Cox and Leland [1982] have shown that when the riskless rate is constant and the risky asset follows geometric Brownian motion or a geometric binomial process, strategies such as these are inefficient. Unfortunately, the Cox-Leland approach, while elegant and insightful, does not tell us the *magnitude* of the inefficiency. The purpose of this paper is to use the Payoff Distribution Pricing Model (Dybvig [1980, *in press*]) to compute directly the cost of the inefficiency. As a result, we can now compare the importance of general lack of diversification with non-modeled costs, such as trading commissions. The results indicate that the inefficiency costs of the strategies are substantial and should not be ignored by practitioners.

A common misconception among students first learning about the efficient markets hypothesis is that portfolio managers can do no damage. Of course, this is not true, because managers choosing random or poorly diversified portfolios throw away investors' money by obtaining them less return than is justified for the amount of risk taken on. For example, in

---

1. See Dybvig and Ingersoll [1982] for a discussion of the difficulty of using mean-variance analysis for evaluating options and other non-linear claims, and Dybvig and Ross [1985a,b] for a general discussion of why mean-variance performance measures may not be valid even in the absence of measurement error.

the mean-variance world an efficient portfolio choice could have given the investors the same mean and variance of terminal wealth at a lower cost. In an intertemporal context, things get a bit more complicated. Besides the importance of diversification across assets, an efficient portfolio choice must also be diversified across time. Furthermore, a nonconstant portfolio choice over time may be optimal, but such a portfolio choice must react appropriately to information arrival.

Fortunately, there is a simpler way of viewing the multiperiod problem. As Ross [1978] has emphasized, the space of feasible consumption bundles is quite generally a linear space. Therefore, if all consumption takes place at the end, we can replace the original dynamic problem with an equivalent one-period problem with the appropriate terminal state prices.[2] Use of state prices to reduce a multi-period problem to a one-period problem is the basis of Cox and Leland [1982], and has been emphasized by many others starting perhaps with Ross [1976] and Rubinstein [1976].[3]

Once we assume that all consumption takes place at the end, we apply the Payoff Distribution Pricing Model (PDPM), which allows us to calculate a lower bound on the cost of the efficiency loss. Here are the assumptions of the PDPM. (See Dybvig [*in press*] for a formal development of the Payoff Distribution Pricing Model.)

1. Agents' preferences depend only on the probability distribution of terminal wealth.

---

2. We will always take consumption to occur at the end. More generally, if preferences are time-separable, the analysis is unchanged if we treat consumption at each date separately.

3. Other papers emphasizing state prices and reduction of a multiperiod problem to one period include Banz and Miller [1978], Brennan and Solanki [1981], Cox and Leland [1982], Cox, Ross, and Rubinstein [1979], Cox and Huang [1985], and Pliska [1986].

2. Agents prefer more to less, i.e., given a choice between two ordered random terminal wealths, an agent will always choose the larger.

3. The market faced by an individual comes from our standard model of a perfect market (no taxes, transaction costs, or information asymmetries) that is complete over finitely many equally probable terminal states or some atomless continuum of states. Such a market allows short sales without penalty.

Informally, the assumptions are state independence of preferences, preference of more to less, and completeness of frictionless complete markets with equally probable states.[4]

The first assumption says that preferences depend only on the probability distribution of terminal consumption. This assumption allows von Neumann-Morgenstern preferences or more generally Machina [1982] preferences over wealth, but precludes state-dependent preferences (including those induced by non-traded wealth). The second assumption, preference of more to less, would not be reasonable for ice cream but is certainly reasonable for wealth. The third assumption, completeness of markets over equally probable or continuous terminal states, is a natural assumption in the presence of continuous trading or a complete set of options. The assumption of equally probable terminal states is for convenience; it allows us to use first-order stochastic dominance. If we allow terminal state probabilities to be unequal and assume concavity of preferences, the analysis is messier, but exactly the same numerical results are valid. (See Appendix I of Dybvig [*in press*].)

These assumptions imply that any optimal strategy purchases more consumption in terminal states in which consumption is cheaper. What is new

---

4. By definition, any atomless distribution has equally probable states each having probability zero. (An atom is an indivisible state with positive probability.) Whenever we refer to a continuum we will implicitly mean a nonatomic continuum.

to the PDPM is the idea of computing how much the cheapest portfolio generating a given *distribution function* of consumption should cost, and the development of simple machinery for doing so. This cost is given by the change in price in response to swapping consumption across terminal states to make the consumption a decreasing function of the state price density while maintaining the same marginal distribution.

Section 2 contains simple numerical examples. Section 3 presents some computer-generated numerical results for reasonable parameter values. Section 4 discusses generalizations, particularly to term structure models. The paper is intended to be self-contained in the sense that it does not require any prior knowledge of the PDPM.

2. Some numerical examples

Here are some simple examples designed to illustrate the principle behind applying the Payoff Distribution Pricing Model to measuring inefficiency. In these examples, we will use the binomial model of stock returns introduced by Cox, Ross, and Rubinstein [1979]. For convenience, we will assume numerically simple parameters: the initial wealth level and initial stock price are both 16, the riskless rate is always zero, and in each period the stock doubles in price or halves in price, each with probability 1/2. We analyze a four-period model since, given our other assumptions, this is the shortest time span over which the analysis does not degenerate.[5] Obviously, these examples are for illustration only; we will

5. Readers who are familiar with the path independence results of Cox and Leland [1982] may find this confusing, since the strategies we consider will have path-dependent strategies in three or even two periods. However, these strategies will not be inefficient for agents with concave preferences that are not necessarily strictly concave. To get inefficiency for these general agents, we need something slightly stronger than path dependence, which is that a path with strictly higher state price should have strictly higher

analyze more realistic examples in Section 4, using the general form of the Payoff Distribution Pricing Model.

Before moving to the examples, we first summarize some important properties of the binomial model. (All of these properties have appeared in the literature in one form or another.) Stock and bond returns are shown graphically in Table 1. For binomial models, it is most common to represent the stock price movements by an *ingrown tree*. An expanded tree in which all possible stock price paths are distinguished will also be useful, since we will be studying portfolio strategies for which the terminal portfolio value will depend on the whole path of stock prices and not just on the final stock price. The bond price is constant over time and in all states and is represented by a line segment.

We can see from Table 1 that the usual convention of representing the stock price in terms of an ingrown tree is simply a shorthand that combines all the states in which the stock price is the same. In the expanded tree, each state has the same probability, $1/16 = (1/2)^4$, because at each node the up and down probabilities are both 1/2. We could write the bond in an expanded tree in the same way, but the result would be a boring tree with 16 at each node.

From option pricing theory (and explicitly Cox, Ross and Rubinstein [1979]), we know that every contingent claim paying off various amounts in the last period can be priced, because each contingent claim can be duplicated by some hedging strategy. In particular, we can price a claim that pays 1 in a given state and 0 in all other states. By definition, the

---

consumption. For a general discussion of the relation between the amount of regularity assumed of utility functions and the first order conditions in terms of the state price density, see Dybvig and Ross [1982], especially in Table 1 and the related discussion.

price of this claim is called the *state price* of the given state. State prices are useful because the value of any security can be written as the sum across states of the state price times the value of the security in the state.

To compute the state price for the binomial model, look first to a single period. Suppose the value of an asset next period is v_1 if the stock goes up and the value next period is v_2 if the stock goes down. We want to duplicate holding the asset. If we invest an amount v_S in stock and an amount v_B in bond today, tomorrow we will have $2v_S + v_B$ if the stock goes up, and $v_S/2 + v_B$ if the stock goes down. If this investment duplicates the asset's value, then we have that

$$v_1 = 2v_S + v_B$$

and

$$v_2 = v_S/2 + v_B.$$

Solving for v_S and v_B, we get that $v_S = 2(v_1-v_2)/3$ and $v_B = (4v_2-v_1)/3$, which is the hedging strategy. Note that $v_S + v_B = v_1/3 + 2v_2/3$, which is the one-period pricing relation. In other words, the up state has price $1/3$ and the down state has price $2/3$.

Of course, we can use this procedure to obtain the state price of any node, and by extension the value of any claim. In particular, the state price of any node equals the price of a claim that pays 1 in that state at that time and zero otherwise. By folding back, we conclude that the state price of any node is $(1/3)^u(2/3)^d$, where u is the number of times the stock price goes up and d is the number of times the stock price goes down. This formula applies to all time intervals. For example, the value of a security at any point is equal to $1/3$ times the value one period later if the stock goes up plus *2/3* times the value one period later if the stock goes down. Working backwards a period at a time using state prices is analogous to solving the Black and Scholes [1973] differential equation, while valuing a claim directly by summing over the four-period state prices is analogous to using the Rubinstein [1976] integral approach to option pricing. From now on, we will focus on the approach using state prices. The reader should keep in mind, however, that the derivation of the state prices tells us explicitly how to compute the amounts of stock and bond held at each point in time in the dominating strategy.

The one aspect of Table 1 we have yet to discuss is the *state price density* (or *state price per unit probability*), which is simply the state price divided by the probability.[6] It is useful to think in terms of this ratio, which plays a central role in the PDPM. For one thing, maximizing a von Neumann-Morgenstern utility function gives you a first order condition that the marginal utility is proportional to the terminal state price density. Suppose an agent solves the following problem:

Choose c_i*'s to*

maximize $\sum \pi_i u(c_i)$

subject to $\sum p_i c_i = w_0$,

6. One special feature of Table 1 is that the terminal state price density is a function only of the terminal stock price. This is a very special feature of this particular example and certain other examples including economies with geometric i.i.d. stock price movements (see Cox and Leland [1982]). Especially in models of the term structure (with random interest rate movements) it is not reasonable to assume that the state price density is a function only of the natural state variables. Fortunately, as we will see in Section 4, the approach in this paper does not require the state price to be a function only of the state variables driving asset returns.

where c_i is consumption in terminal state i, π_i is the probability of terminal state i, $u(\cdot)$ is the agent's utility function, and p_i is the state price of terminal state i. If $u(\cdot)$ is differentiable, then the first-order condition is that for some λ,

$$\pi_i u'(c_i) = \lambda p_i$$

or

$$u'(c_i) = \lambda p_i/\pi_i \equiv \lambda \rho_i, \tag{1}$$

i.e., the agent's marginal utility of wealth in terminal state i is proportional to the terminal state price density $\rho_i \equiv p_i/\pi_i$.[7] A second important feature of the state price density is that if we combine states with the same state price density, the combined aggregate state will also have the same state price density. Perhaps more importantly, we can define the state price density even if there is a nonatomic continuum of states (in which case both the state price and the probability are zero), as in the diffusion models. The state price density is defined at each node as the ratio of the state price to the probability of the node. The state price density follows a multiplicative process whose movements locally price all assets correctly. In our specialized binomial model, the state price density at a node following u ups and d downs is given by $\rho_n = p_n/\pi_n = ((1/3)^u(2/3)^d)/((1/2)^u(1/2)^d) = (2/3)^u(3/2)^d$.

We will need a few concepts and results of the Payoff Distribution Pricing Model (PDPM). An *asset pricing model* (such as the CAPM, APT, or the

---

7. If $u(\cdot)$ is concave but not everywhere differentiable, $u'(\cdot)$ should be interpreted as some element of the marginal utility correspondence which is the closed interval bounded by the right and left derivatives. See Dybvig and Ross [1982].

Black-Scholes model) gives us the price of a random cash flow, as in the budget constraint to the agent's maximization problem above. In our context, we can write the asset pricing model as $P_A(c) = \sum_i p_i c_i$. The Payoff Distribution Pricing Model assigns a price to a distribution function of consumption by assigning to it the price of the least expensive consumption pattern having that payoff. In other words, we can write the distributional pricing function as $P_D(F) = min(P_A(c)|c{\sim}F)$, where ~ means "is distributed as." For the extensions in Section 4, we will refer to a general formula for this minimum cost in terms of the distribution functions of c and ρ, but for now all we need is the following Theorem that combines several results from Dybvig [*in press*].

Theorem: The following are equivalent.

1. The consumption pattern c is chosen by some agent with 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, $P_A(c) = P_D(F_c)$.

3. Consumption is nondecreasing in the terminal state price density.

Proof: See Dybvig [*in press*], Theorems 1 and 2. □

This Theorem is useful to us for two different reasons. First, it says that $P_A(c) - P_D(F_c)$ is a tight lower bound on the amount of initial wealth an agent would pay to switch from c to an optimal strategy, given that we do not know the agent's actual preferences. (This is a bound because all agents are indifferent between c and the strategy underlying $P_D(c)$, and the bound is tight because the theorem tells us that there is some agent who would follow that underlying strategy, implying the bound is achieved for that agent.) Second, it tells us how to compute the bound, namely, by swapping consumption across terminal states, leaving the distribution function unchanged, until consumption is nondecreasing in the terminal state price density.

Now that our reviews of the binomial model and the PDPM are out of the way, we are ready to proceed to our examples. All of our examples use the concepts and tools of the PDPM to quantify the amount of damage done by following an inefficient policy, that is, a policy for which consumption is not nonincreasing in the terminal state price density. Our first example examines a policy of holding stock initially but limiting potential losses by switching from the stock to the bond if ever the portfolio value falls too much. We refer to this policy as a *stop-loss* strategy.

Example 1 Stop-loss strategy

The rule under this strategy is to invest in the stock until the portfolio value falls to 8, and to stay in the bond from then on. The value of the portfolio under this strategy is given in the ingrown tree in Table 2. The probabilities are computed by adding up the number of paths to the terminal node and multiplying by 1/16. For example, there are three paths (up-up-up-down, up-up-down-up, and up-down-up-up) having terminal wealth of 64 and .two paths (up-up-down-down and up-down-up-down) having terminal wealth of 16. Horizontal paths corresponding to holding the bond have to be counted twice *per period the bond is held*, since the horizontal line captures both states. Therefore, there are ten paths with terminal wealth of 8 (up-down-down-up, up-down-down-down, down-up-up-up, down-up-up-down, down-up-down-up, down-up-down-down, down-down-up-up, down-down-up-down, down-down-down-up, and down-down-down-down).

The second strategy in Table 2 was chosen to obtain the same distribution of terminal wealth (allocated differently across states) but with consumption ordered opposite of the terminal state price density which, from Table 1, is ordered opposite of the stock price. To do this, we walk down the two probability distributions together. First, we assign the 1/16 probability of 256 to the terminal state in which the stock reaches 256. Next, we assign the 3/16 probability of 64 to three of the four terminal states in which the stock price reaches 64. Now, we assign 1/16 of the 2/16 probability of getting 16 to the remaining terminal state in which the stock is 64, and the remaining 1/16 to one of the terminal states in which the stock is 16. In all the remaining terminal states (10 of them) the amount we get is 8. Because this selection makes consumption nonincreasing in the terminal state price density, by Theorem 1 the resulting portfolio strategy is efficient. The values earlier on in the tree are computed by walking back period by period using the 1/3–2/3 weighting rule. We find that this portfolio strategy, while giving exactly the same probability distribution of terminal wealth as the stop-loss strategy, costs only 15 65/81 (as compared to 16). □

What is really going on here? If we compare the two strategies' terminal wealth state-by-state, we find that they differ only in the states up-down-up-down and down-up-up-up. Since the latter state has more up's and fewer down's, the terminal state price density is lower. The efficient dominating strategy has its higher consumption in that state (16 versus 8) while the stop-loss strategy has it reversed (8 versus 16). The savings is the difference in cost of the two strategies (16 - 1280/81 = 16/81), which is the probability (1/16) times the difference in terminal state price density (64/81 − 32/81 = 32/81) times the amount of consumption moved (16 − 8 = 8). In richer examples with more periods or more elaborate strategies, there would be more terminal states in which the inefficient and dominating strategies disagree. Nonetheless, the concept would be the same: the dominating strategy would move consumption from expensive terminal states to cheaper states.

The second example is in the same spirit as the first, but in reverse. The policy is to hold stock initially, but to switch into bonds (to lock in the gain) if there is sufficient improvement in the portfolio value. We refer to this policy as a *lock-in strategy*.

Example 2 Lock-in strategy

The rule under this strategy is to invest in the stock until the portfolio value rises to 32, and to stay in the bond from then on. The value of the portfolio is given by the ingrown tree in Table 3. In terms of which paths can occur (and therefore the probabilities of the outcomes), the ingrown tree is just the same as the stop-loss tree of Example 1 in Table 2, only upside down. The efficiency loss is different, however, since the quantities and terminal state prices are different when we turn the tree upside down.

The second strategy in Table 3 was chosen to obtain the same terminal distribution of terminal wealth as the lock-in strategy but with consumption ordered opposite of the terminal state price density. This process is just as in Example 1, except starting from the opposite side. Computing the initial investment required for this strategy, we find that it costs only 15 17/81 (as compared to 16).

□

As in Example 1, if we compare the terminal wealth of the lock-in strategy with its dominating strategy state-by-state, we find that the two differ in only two states (up-down-down-down and down-up-down-up). The improvement made by the dominating strategy is to move the larger consumption from the more expensive of the two states to the cheaper one.

In our next example, we compute the potential cost of hiring someone who claims to have timing ability but actually may not.

Example 3 Random market timing strategy (market timer who can't)

The rule under this strategy is to invest in the stock in some two of the four periods (half the time) and to invest in the bond in the other two. The timing is based on any random rule that is independent of market returns. The distribution of terminal wealth is the same whatever the timing; two examples ("A" and "B") are illustrated in table 4. Strategy "A" has the stock investment in the first two periods. Strategy "B" has the stock investment in the first and last periods. Since the terminal distribution is the same independent of the random choice, the unconditional distribution is the same under each choice.

As before, the way to dominate the strategy is to move the large amounts of consumption to terminal states in which consumption is cheaper. From Table 4, we can see that the move to the dominating strategy requires two switches, one between up-up-down-down and up-down-up-up, and the other between down-up-down-down and down-down-up-up. The first switch reduces cost by the product of the probability 1/16, the amount 48 (= 64 - 16) of consumption moved, and the difference 32/81 (= 64/81 - 32/81) in terminal state price density, for a cost reduction of 1 5/27 (= 32/27). The second switch reduces cost by the product of the probability 1/16, the amount 12 (= 16 - 4) of consumption moved, and the difference 64/81 (= 128/81 - 64/81) in terminal state price density, for a cost reduction of 16/27. Combining these two changes we have a total cost reduction of 1 7/9 (= 1 21/27), which reduces the initial cost from 16 to 14 2/9. □

This concludes our section of simple numerical examples. In Section 3, we report computer-based calculations of the loss under more reasonable parameter values.

3. Realistically, how large is the cost?

In Section 2, we looked at three numerical examples that showed the theoretical principle behind measuring the cost of following a dominated strategy. Now we compute the cost in more realistic situations. The calculations approximate continuous lognormal stock movements using a binomial process with a daily grid. To approximate current conditions with a round number, the short riskless rate is assumed to be 8%. To approximate the Ibbotson-Sinquefield historical returns on a well-diversified portfolio, we assume that the stock has an expected return of 16% (i.e., an excess return of 8% annually) and an annual standard deviation of about 20% (in logs).

Table 5 summarizes how these parameter assumptions map into per-period returns. As the time increment Δt gets smaller and smaller, the stochastic process described in Table 5 and the corresponding pricing converges to a standard lognormal diffusion model for the stock price (as is consistent with Black and Scholes [1973] with a constant mean return). (For related analyses see Banz and Miller [1978], Brennan and Solanki [1981], Cox and

Leland [1982], Cox, Ross, and Rubinstein [1979], Garman [1976], Ross [1976], and Rubinstein [1976].) Therefore, we can consider our numerical results to be an approximation to what would be obtained in continuous time.

Now we are ready to look at some numerical results. Figures 1, 2, and 3 are plots of numerical estimates of the cost of following the three inefficient strategies from Section 2. The estimates were made using a set of routines for analyzing probability distributions developed using Scheme, which is a dialect of Lisp (see R. K. Dybvig [1987]). The program computes the minimum cost by matching consumption levels in reverse order of the terminal state price density as in Section 2. In computing the terminal distribution, the routines manage the size of the problem by combining indistinguishable states along the way.

Figure 1 shows the cost of following a stop-loss strategy, as a function of the limit value. (The "jaggedness" of the plot comes from the coarseness of the daily binomial approximation to the diffusion, and the plot becomes smoother when we move to a half-day interval.) For example, assume the current portfolio value is *$2* billion and we plan to switch into stock if the value falls to *$1.8* billion or below. Then the limit value is *90%* ($= 1.8/2.0$). Figure 1 says that we will be throwing away about *60* basis points or *$12* million by following the stop-loss strategy for a year, as compared to following the efficient strategy giving the same distribution of terminal wealth. While this ignores transaction costs for the two strategies, *60* basis points over a year is a large number, and we can surely do better than a stop-loss strategy. When the limit value is small, the efficiency loss is small, since the limit is rarely achieved and the portfolio strategy is nearly the same as holding the stock (which is efficient). Similarly, as the limit value approaches *100%* from below, the probability of hitting the limit close to the starting time increases to one, and the strategy looks more and more like holding the bond (which is also efficient). When the limit value is *100%* or more, the strategy switches immediately to the bond and the strategy is precisely holding the bond (which is efficient). Intuitively, the loss is largest when there is a large chance both of hitting and of missing the limit soon after the start. This is consistent with Figure 1, which shows the largest loss at a limit of about *95%* (with a one year horizon), which is *1/4* of the one-year standard deviation of the stock.

Figure 2 shows the efficiency loss of a lock-in strategy, which is in some sense a mirror image of a stop-loss strategy. The loss is not exactly symmetric since the terminal state price density is higher for lower stock prices. This means that the damage done by the stop-loss strategy (which usually occurs when the stock has gone down) is more costly than damage done by the lock-in strategy. Nonetheless, the cost of following a lock-in strategy is substantial and should not be ignored by practitioners.

Figure 3 shows the efficiency loss of a random timing strategy (a "timer who can't"), as a function of the fraction of time the timer holds the stock. The efficiency loss of this strategy can be as high as *200* basis points over a year, or *$40* million for a *$2* billion dollar portfolio! While Figure 3 contains simulation results, this is one of the few cases we can actually solve analytically for the diffusion model. The random timer follows a strategy which holds the stock a fixed fraction f of the time and the bond a fraction $(1-f)$ of the time. (Since we are not including transaction costs, the exact allocation doesn't matter — it only matters that the stock is held exactly a fraction f of the time.) If w_0 is the initial wealth, then the wealth $\tilde{w}_T$ at the end (time T) is lognormally distributed, as

$$log(\tilde{w}_T) \sim N(log(w_0) + \mu fT + r(1-f)T + \sigma^2 fT/2,\ \sigma^2 fT). \tag{2}$$

An alternative strategy with initial wealth x_0 and a fixed portfolio weight α is also lognormally distributed, and is efficient for $\alpha > 0$. Its terminal distribution is

$$log(\tilde{x}_T) \sim N(log(x_0) + \mu\alpha T + r(1-\alpha)T + \sigma^2\alpha^2 T/2,\ \sigma^2\alpha^2 T). \tag{3}$$

To give these two the same distribution (for $\alpha > 0$), we must choose $\alpha = \sqrt{f}$ to match variances and then

$$log(x_0) = log(w_0) - (\sqrt{f}-f)(\mu-r)T \tag{4}$$

to match means. In log terms, $(\sqrt{f}-f)(\mu-r)T$ is the loss, and up to scaling this is essentially what is plotted in Figure 3.[8]

---

8. The formula for the loss in the random timer case is linear in T. Another way of saying this is that if you split the time interval into two parts, the value as a percentage of potential on the whole is the product of the value as a percentage of potential on each half. This is a special case of a general result. Suppose security returns are independent over time and that the return to the inefficient portfolio strategy in the two sub-periods is independent. Then the value as a percentage of potential on the whole period is less than or equal to the product of the values as a percentage of potential on the subperiods. To prove this, consider making the dominated strategies on the subperiods your strategy over the whole period. This may not be optimal over the whole period, but achieves a value that proves the bound. This result says that when stock returns are independent, you cannot recover from past inefficient policies. As an example of applying this, rolling over portfolio insurance each year is efficient in each period but inefficient over two years (an example of inequality). Rolling over portfolio insurance each year over four years is at least twice as bad

Figure 4 shows the efficiency loss to a strategy not analyzed in Section 2, which used portfolio insurance repeatedly. (Because of technical limitations of my computer program, I have chosen slightly different parameters and weekly rebalancing for this example. Both choices reduce the number of terminal nodes and help to keep the size manageable in spite of exponential growth.) Since many managers create synthetic portfolio insurance with a one-year horizon repeated annually, the large size of the efficiency loss shown here (over *5%* in *10* years) is especially troubling.

4. Generalizations

In this section, we derive a formula for the state price density in the general case when asset prices follow general Itô processes. The main economic assumptions we need are completeness of markets and the absence of arbitrage.[9] We then turn to applications involving the term structure of interest rates. While numerical analysis like that in Section 3 has not been performed for the term structure applications, doing such computations is a promising avenue for future research.

The state price density ρ gives a representation of the linear pricing rule of Ross [1978]. Letting P be the vector of re-invested price series, then we have that the price at time s can be written in terms of ρ and the price at a later time t as

---

(measured in logs) as rolling over portfolio insurance each year for two years (by the result — in fact it is even worse).

9. There are some additional technical assumptions that would be required in a more formal analysis. See, for example, Harrison and Pliska [1981], Cox and Huang [1986], or Dybvig and Huang [1987] for related results. In those papers, the emphasis is on the risk neutral probabilities (called martingale probabilities). The state price density is equal to a discount factor times the Radon-Nikodyn derivative of the risk-neutral probabilities applied to indicator sets, with respect to the probability measure.

$$\rho_s P_s = E_s[\tilde{\rho}_t \tilde{P}_t], \tag{5}$$

or in particular if we take $\rho = 1$ at $s = 0$, we have that

$$P_0 = E_0[\tilde{\rho}_t \tilde{P}_t]. \tag{6}$$

Equation (5) says that ρP is a martingale, which implies that ρP has no drift.

We assume that prices follow an Itô process. Specifically, the n-vector of risky asset price changes is given by

$$\frac{d\tilde{P}}{P} = \mu dt + \sigma d\tilde{Z}, \tag{7}$$

where the division is componentwise, μ is the n-vector of expected returns, $\tilde{Z}$ is a k-dimensional Wiener process, and σ is a $k \times n$ matrix of risk exposures. Under reasonable assumptions, ρ itself follows an Itô process involving only $\tilde{Z}$. We will take this as given. Then we have that

$$\frac{d\tilde{\rho}}{\rho} = adt + b' d\tilde{Z} \tag{8}$$

where $\rho_0 = 1$ or equivalently

$$\tilde{\rho}_t = \exp\left(\int_{s=0}^{t} a \; ds + \int_{s=0}^{t} b' d\tilde{Z}_s - \frac{1}{2} \int_{s=0}^{t} b'b \; ds \right) \tag{9}$$

for some random 1-dimensional process a and some k-dimensional process b. Since (5) holds for all assets,

$$\frac{d(\tilde{\rho}\tilde{P})}{\rho P} = \frac{d\tilde{\rho}}{\rho} + \frac{d\tilde{P}}{P} + \frac{d\tilde{\rho}}{\rho}\,\frac{d\tilde{P}}{P} \tag{10}$$

has no drift, or by Itô's lemma, this implies that

$$0 = drift\left(\frac{d(\tilde{\rho}\tilde{P})}{\rho P}\right) = ae + \mu + \sigma b, \tag{11}$$

where e is an n-vector of ones. If we eliminate locally redundant assets, completeness implies that σ is square and nonsingular, and therefore

$$a = -r \tag{12}$$

and

$$b = -\sigma^{-1}(\mu - re) \tag{13}$$

where r is the local riskless rate. Generally, a and b are solutions to (10), even if we have not eliminated locally redundant assets. Of course, completeness of markets implies that there is a locally riskless portfolio.

To use (8), (11), and (12) with the payoff distribution pricing model, we have to use the continuous state-space analogue of valuation using consumption ordered in reverse of the terminal state price density. This analogue implies that the cost of the dominating portfolio (the distributional price) is

$$P_D = \int_{\gamma=0}^{1} F_{\rho}^{-1}(\gamma)F_{c}^{-1}(1-\gamma)\,d\gamma, \tag{14}$$

which is equation 3 of Dybvig [in press]. This is a general expression for the expectation of the product of two variables ρ and c that are perfectly inversely related. F_{ρ}^{-1} has the same units as ρ (state price over probability), F_c^{-1} has units of consumption, and γ has units of probability. The arguments γ and $1-\gamma$ signify inverse ordering, and the integral corresponds to summing across states in the finite model.[10]

Term Structure Models

To illustrate the evolution of the state price density in continuous time models, this section computes the state price density in closed form for a class of models with interest rate uncertainty. Because interest rates can move randomly, there is a nontrivial term structure of interest rates in these models. Throughout the rest of this section, we will assume either that preferences are over nominal payoffs, or that we are expressing all returns in real terms (which is formally equivalent).

To illustrate the computation of ρ, we assume the vector b of asset risk premia is constant. Then we can write the state price density (9) as

$$\tilde{\rho}_t = exp\left(-\int_{s=0}^{t} r_s ds + \int_{s=0}^{t} b' d\tilde{Z}_s - \frac{t}{2} b'b\right). \tag{15}$$

In particular, if $b = 0$, we have the "local expectations hypothesis" (see Cox, Ingersoll, and Ross [1981]), which is a reasonable assumption if all the assets in our list are bonds or derivative of bonds. In this case, every efficient portfolio has a terminal value that is a nondecreasing function of the compounded return on rolling over shorts. One interesting implication of this result is that contingent immunization is not an

10. For details, see Dybvig [in press]. To define the inverse distribution function for discrete variables (or more generally at mass points), we put "risers" on the step function. For example, suppose a random variable is either 1 or 2, each with probability $1/2$. Then the inverse distribution function is defined to be 1 on $(0,1/2]$ and 2 on $(1/2,1)$. The value assigned to the endpoints 0 and 1 don't matter, because they don't affect the integral in (14).

efficient strategy.[11] This is formally true from (15) whenever the local expectations hypothesis holds; more generally (9) and (12) tell us that we would have to make a bizarre assumption about the movement of the vector b of risk premia to make contingent immunization efficient.

To apply our analysis to a term structure model, we would want to use (14), for which we need to know the distribution of the state price density following the process as given by (15). In general, we can do this numerically, but we want to consider a special case in which we can compute the distribution analytically. We will use a special case of Vasicek [1977]. Loosely speaking, Vasicek showed that if interest rates follow a Gaussian process, then we can compute bond prices. (One has to make an assumption about risk premia as well, our assumption that the vector b of risk premia is a constant is sufficient.) Vasicek's model is attractive analytically because the normality makes it tractable. (Unfortunately, however, it is not a good approximation to the actual movement of interest rates, except perhaps over very short periods of time.[12]) From (6) and (15), we can see how to compute bond prices in the Vasicek models. By (6), the price at 0 of a bond paying 1 at t is $E_0[\tilde{\rho}_t/\rho_0] = E_0[\tilde{\rho}_t]$. If r and Z

11. Intuitively, a contingent immunization strategy switches from one risky portfolio into an immunized portfolio using a cut-off rule that is qualitatively similar to the stop-loss strategy, if we consider the initial portfolio as the stock and the immunized portfolio as the bond. (Using the immunized portfolio as *numeraire* makes the analogy almost exact.) Therefore, while the qualitative properties of the efficiency loss should be as in Figure 1, without further analysis we cannot be sure of the magnitude of the loss. It does seem, however, that if the size of the loss we are insuring is significant, the efficiency loss will be significant, too.

12. For example, interest rates can go arbitrarily negative in Vasicek's model, and will go negative frequently under reasonable variance assumptions. Also, it has been shown empirically that the variance of interest rates changes over time in a way that can be predicted by looking at yield curves, contradicting the assumption of constant variance — see Brown and Dybvig [1986].)

are jointly normal (as they are in Vasicek's model), we can compute this expectation using the normal moment generating function.

For our extended example, assume that r follows the following simple mean reverting process.

$$dr = \kappa(\bar{r}-r)dt + \Sigma' d\tilde{Z}, \tag{16}$$

where κ, $\bar{r}$, and Σ are known and constant, and $\tilde{Z}$ is the k-dimensional Wiener process that drives security prices. Then we have that

$$r_t = \bar{r} + (r_0-\bar{r})e^{-\kappa t} + \int_{\tau=0}^{t} e^{-\kappa(t-\tau)}\Sigma' d\tilde{Z}_\tau, \tag{17}$$

and that

$$\int_{s=0}^{t} r_\tau d\tau = \bar{r}t + (r_0-\bar{r})\frac{1-e^{-\kappa t}}{\kappa} + \int_{\tau=0}^{t} \left(\frac{1-e^{-\kappa(t-\tau)}}{\kappa}\right)\Sigma' d\tilde{Z}_\tau. \tag{18}$$

From (15) and (18), we have that

$$log(\rho_t) = -\left(\frac{b'b}{2} + \bar{r}\right)t - (r_0-\bar{r})\frac{1-e^{-\kappa t}}{\kappa} + \int_{\tau=0}^{t} \left(b' - \frac{1-e^{-\kappa(t-\tau)}}{\kappa}\Sigma'\right)d\tilde{Z}_\tau. \tag{19}$$

Therefore, ρ_t is normally distributed with mean

$$M = -\left(\frac{b'b}{2} + \bar{r}\right)t - (r_0-\bar{r})\frac{1-e^{-\kappa t}}{\kappa} \tag{20}$$

and variance

$$V = \int_{\tau=0}^{t} \left(b - \frac{1-e^{-\kappa(t-\tau)}}{\kappa}\Sigma\right)'\left(b - \frac{1-e^{-\kappa(t-\tau)}}{\kappa}\Sigma\right) d\tau$$

$$= \left(b'b - \frac{2b'\Sigma}{\kappa} + \frac{\Sigma'\Sigma}{\kappa^2}\right)t + \frac{2}{\kappa}\left(b'\Sigma - \frac{\Sigma'\Sigma}{\kappa}\right)\frac{1-e^{-\kappa t}}{\kappa} + \frac{\Sigma'\Sigma}{\kappa^2}\,\frac{1-e^{-2\kappa t}}{2\kappa}. \qquad (21)$$

Therefore, ρ_t is distributed lognormally, and $log(\rho_t)$ has mean M and variance V. By the normal moment generating function, then, the bond price is given by $exp(M+\frac{1}{2}V)$.

The lognormality of ρ_t implies it is possible (although we do not do so here) to compute analytically the cost of some types of random timing strategies. Of course, more numerical work is required to compute the cost of following other dominated strategies. In numerical work, having a closed-form expression for the distribution of the state price density is very useful, because computing it numerically requires us to keep track of two state variables, r and ρ. In some sense, this is why term structure models are difficult to solve analytically: the state price density is not a function of the natural state variable (the interest rate).

5. Conclusion

Our numerical results show that the efficiency loss to inefficient strategies may in fact be very large, even given very realistic assumptions. The strategies we have considered, stop-loss, lock-in, random timer, and repeated portfolio insurance, are very similar to strategies used in practice. It is interesting to note that the efficiency loss is the same whether or not the strategy was "planned" in advance; in other words, a manager deciding to lock in the gains at the time a boundary is reached has the same terminal distribution of wealth as a manager who planned from the start to follow this strategy.

Much work remains. In one direction, it would be nice to extend the analysis to include transaction costs explicitly. Short of that, we can add the transaction cost to the cost described here to get an overall measure of the cost of a given policy, and it would be useful to have a collection of examples of this sort to aid our understanding. Along other lines, it is possible to measure the efficiency loss of other strategies. For example, it would be nice to know the magnitude of loss from contingent immunization and other fixed-income strategies.

References

Banz, Rolf, and Merton Miller, 1978, Prices for State-Contingent Claims: Some Estimates and Applications, *Journal of Business*, 51, 653–672.

Brennan, M. and R. Solanki, 1981, Optimal Portfolio Insurance, *Journal of Financial and Quantitative Analysis*, 16, 279–300.

Black, Fischer, and Myron Scholes, 1973, The Pricing of Options and Corporate Liabilities, *Journal of Political Economy*, 81, 637–654.

Breeden, Douglas, and Robert Litzenberger, 1978, Prices of State Contingent Claims Implicit in Options Prices, *Journal of Business*, 51, 621–652.

Brown, Stephen J. and Philip H. Dybvig, 1986, The Empirical Implications of the Cox, Ingersoll, Ross Theory of the Term Structure of Interest Rates, *Journal of Finance*, 41, 617–632.

Cox, John C., and Hayne Leland, 1982, On Dynamic Investment Strategies, *Proceedings, Seminar on the Analysis of Security Prices*, 26, no 2, Center for Research in Security Prices, University of Chicago .

Cox, John C., and Chi-fu Huang, 1985, A variational problem arising in financial economics with an application to a portfolio turnpike theorem, working paper, MIT.

Cox, John C., Jonathan E. Ingersoll, Jr., and Stephen A. Ross, 1981, A Re-examination of Traditional Hypotheses about the Term Structure of Interest Rates, *Journal of Finance*, 36, 769–799.

Cox, John C., Stephen A. Ross, and Mark Rubinstein, 1979, Option Pricing: A Simplified Approach, *Journal of Financial Economics*, 7, 229–263.

Dybvig, Philip H., 1980, Some New Tools for Testing Market Efficiency and Measuring Mutual Fund Performance, working paper, Yale University.

Dybvig, Philip H., *in press*, Distributional Analysis of Portfolio Choice, *Journal of Business*.

Dybvig, Philip H., and Chi-fu Huang, 1986, Nonnegative Wealth, Absence of Arbitrage, and Feasible Consumption Plans, working paper, Yale University.

Dybvig, Philip H., and Jonathan E. Ingersoll, Jr., 1982, Mean-Variance Theory in Complete Markets, *Journal of Business*, 55, 233-251.

Dybvig, Philip H., and Stephen A. Ross, 1982, Portfolio Efficient Sets, *Econometrica*, 50, 1525-1546.

Dybvig, Philip H., and Stephen A. Ross, 1985a, The Analytics of Performance Measurement Using a Security Market Line, *Journal of Finance*, 40, 401-416.

Dybvig, Philip H., and Stephen A. Ross, 1985b, Differential Information and Performance Measurement Using a Security Market Line, *Journal of Finance*, 40, 384-399.

Dybvig, R. Kent, 1987, *The Scheme Programming Language*, Prentice-Hall.

Garman, Mark, 1976, A General Theory of Asset Valuation Under Diffusion State Processes, working paper, University of California at Berkeley.

Harrison, J. Michael, and Stanley Pliska, 1981, Martingales and Stochastic Integrals in the Theory of Continuous Trading, *Stochastic Processes and Their Applications*, 11, 215-260.

Machina, Michael, 1982, "Expected Utility" Analysis Without the Independence Axiom, *Econometrica*, 50, 277-323.

Pliska, Stanley, 1981, A Stochastic Calculus Model of Continuous Time Trading: Optimal Portfolios, *Mathematics of Operations Research*, 11.

Ross, Stephen A., 1976, Options and Efficiency, *Quarterly Journal of Economics*, 90, 75-89.

Ross, Stephen A., 1978, A Simple Approach to the Valuation of Risky Streams, *Journal of Business*, 51, 453-475.

Rubinstein, Mark, 1976, The Valuation of Uncertain Income Streams and the Pricing of Options, *Bell Journal of Economics*, 7, 407–425.

Vasicek, Oldrich, 1977, An Equilibrium Characterization of the Term Structure, *Journal of Financial Economics*, 5, 177–188.

Table 1 Security Returns, state probabilities, and state prices

Bond:

16 — 16 — 16 — 16 — 16

Stock:

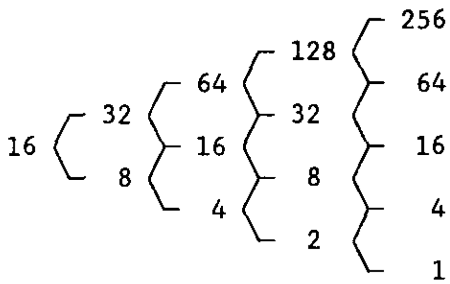

state probability	state price	state price density
1/16	1/81	16/81
4/16	8/81	32/81
6/16	24/81	64/81
4/16	32/81	128/81
1/16	16/81	256/81

Stock (expanded):

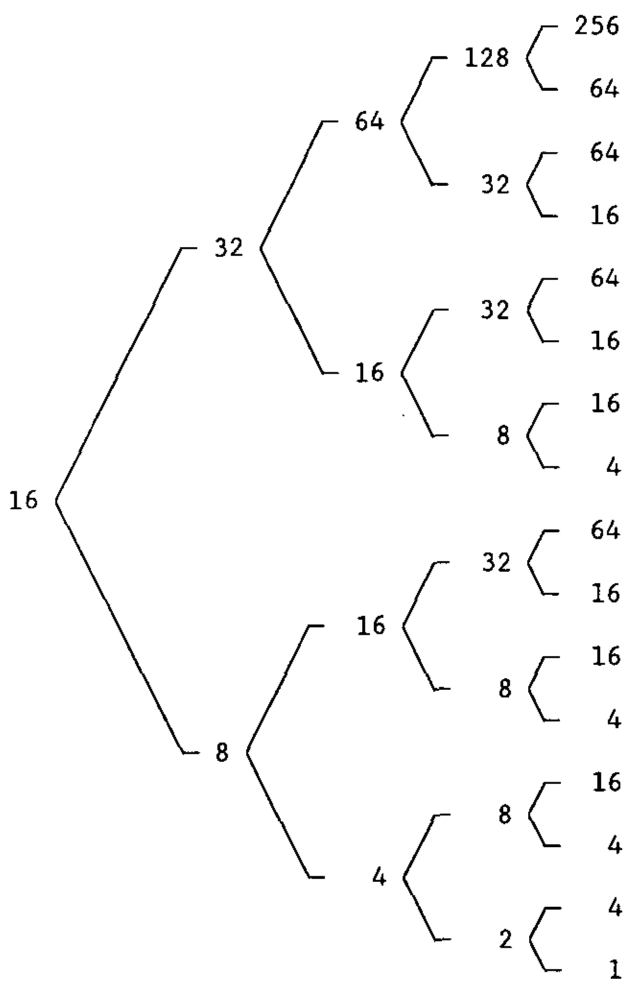

state probability	state price	state price density
1/16	1/81	16/81
1/16	2/81	32/81
1/16	2/81	32/81
1/16	4/81	64/81
1/16	2/81	32/81
1/16	4/81	64/81
1/16	4/81	64/81
1/16	8/81	128/81
1/16	2/81	32/81
1/16	4/81	64/81
1/16	4/81	64/81
1/16	8/81	128/81
1/16	4/81	64/81
1/16	8/81	128/81
1/16	8/81	128/81
1/16	16/81	256/81

Table 2 The stop-loss strategy (limit = 8) and a dominating strategy

Here is the stop-loss strategy.

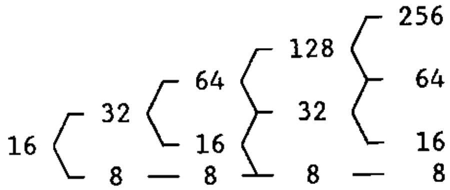

terminal wealth	probability
256	1/16
64	3/16
16	2/16
8	10/16

Here is a dominating strategy (which is itself undominated). This strategy costs only 15 65/81 (= 1280/81) but gives the same terminal probability distribution of wealth (in different states).

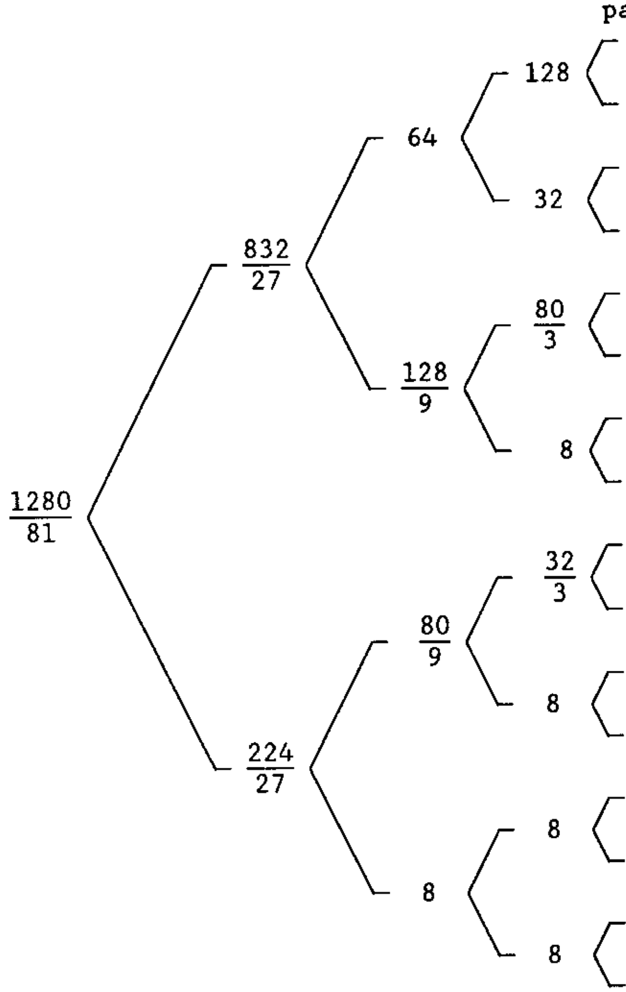

dominating payoff	stop-loss payoff	state price density
256	256	16/81
64	64	32/81
64	64	32/81
16	16	64/81
64	64	32/81
8	16	64/81
8	8	64/81
8	8	128/81
16	8	32/81
8	8	64/81
8	8	64/81
8	8	128/81
8	8	64/81
8	8	128/81
8	8	128/81
8	8	256/81

Table 3 The lock-in strategy (limit = 32) and a dominating strategy

Here is the lock-in strategy.

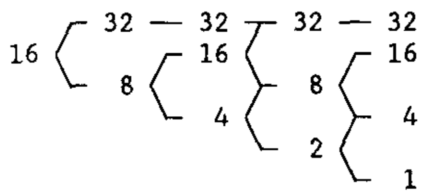

terminal wealth	probability
32	10/16
16	2/16
4	3/16
1	1/16

Here is a dominating strategy (which is itself undominated). This strategy costs only 15 17/81 (= 1232/81) but gives the same terminal probability distribution of wealth (in different states).

				dominating payoff	lock-in payoff	state price density
$\frac{1232}{81}$	$\frac{736}{27}$	32	32	32	32	16/81
				32	32	32/81
			32	32	32	32/81
				32	32	64/81
		$\frac{224}{9}$	32	32	32	32/81
				32	32	64/81
			$\frac{64}{3}$	32	32	64/81
				16	32	128/81
	$\frac{248}{27}$	$\frac{176}{9}$	32	32	32	32/81
				32	32	64/81
			$\frac{40}{3}$	32	16	64/81
				4	4	128/81
		4	8	16	16	64/81
				4	4	128/81
			2	4	4	128/81
				1	1	256/81

Table 4 Random market timing strategy (50% stock) and a dominating strategy

Here is random market timing strategy "A."

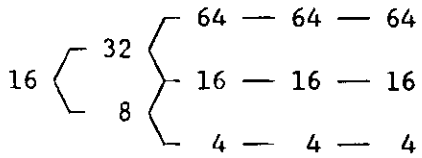

probability
4/16
8/16
4/16

Here is random market timing strategy "B."

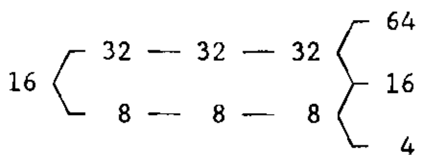

probability
4/16
8/16
4/16

Here is a dominating strategy (which is itself undominated). This strategy costs only 14 2/9 (= 129/9) but gives the same terminal probability distribution of wealth (in different states).

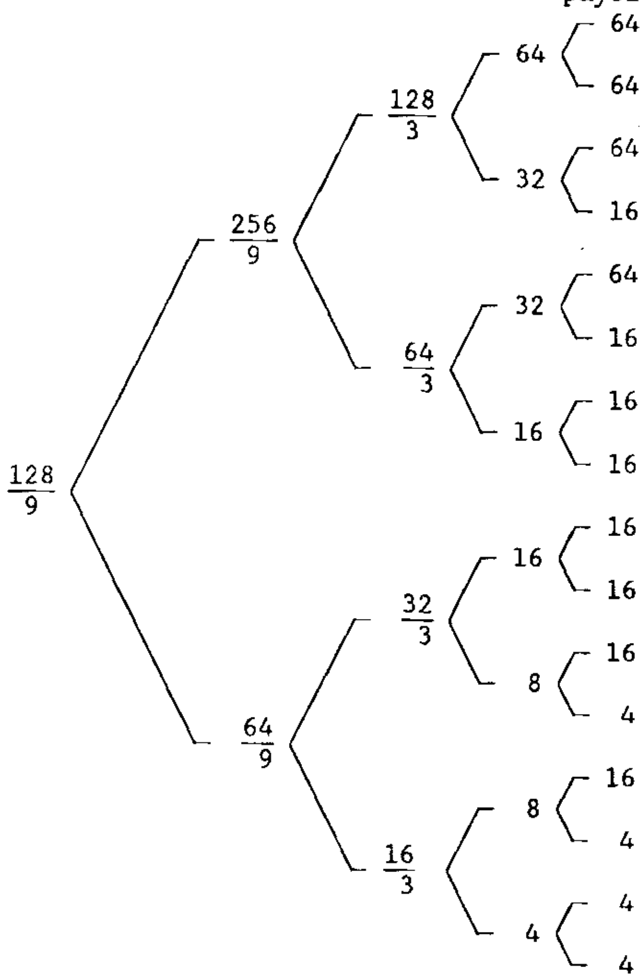

dominating payoff	timing "A"	state price density
64	64	16/81
64	64	32/81
64	64	32/81
16	64	64/81
64	16	32/81
16	16	64/81
16	16	64/81
16	16	128/81
16	16	32/81
16	16	64/81
16	16	64/81
4	16	64/81
16	4	128/81
4	4	128/81
4	4	128/81
4	4	256/81

Table 5 One-period returns in terms of the underlying parameters

Parameters used in numerical work:

$r = 0.08$ (annual interest rate = 8%, continuous compounding)

$\mu = 0.16$ (annual expected return = 16%, for an 8% risk premium)

$\sigma = 0.2$ (annual proportional standard deviation = 20%)

$\Delta t = 1/360 \approx .0027778$ (daily)

$\sqrt{\Delta t} \approx .0527046$

Here is the one-period Bond return:

$$1 \longrightarrow 1 + r\Delta t \approx 1.0002222$$

Here is the one-period Stock return:

$$1 \left\langle \begin{array}{lr} 1 + \mu\Delta t + \sigma\sqrt{\Delta t} \approx 1.0109854 & \text{probability } 0.5 \\ 1 + \mu\Delta t - \sigma\sqrt{\Delta t} \approx 0.9899035 & 0.5 \end{array} \right.$$

State price density:

$$1 \left\langle \begin{array}{l} \dfrac{1}{(1+r\Delta t)}\left[1 - \dfrac{(\mu-r)\Delta t}{\sigma\sqrt{\Delta t}}\right] \approx 0.9787007 \\ \dfrac{1}{(1+r\Delta t)}\left[1 + \dfrac{(\mu-r)\Delta t}{\sigma\sqrt{\Delta t}}\right] \approx 1.0208550 \end{array} \right.$$

Legends for Figures

Figure 1 Efficiency loss of a stop-loss strategy

This figure gives the efficiency loss of a stop-loss strategy, in basis points (hundredths of one percent of the initial investment). Under the stop-loss strategy, the manager invests the entire portfolio in stocks until the portfolio value reaches or falls below the limit value. When the limit value is at or above 100% of the initial wealth, the switch takes place immediately and the strategy is the same as just holding the bond (and is therefore efficient). The size of the efficiency loss can be dramatic: at its worse it is nearly 1% of the portfolio value in only a year!

Figure 2 Efficiency loss of a lock-in strategy

This figure gives the efficiency loss of a lock-in strategy, in basis points (hundredths of one percent of the initial investment). Under the lock-in strategy, the manager invests the entire portfolio in stocks until the portfolio value reaches or exceeds the limit value. When the limit value is at or below 100% of the initial wealth, the switch takes place immediately and the strategy is the same as just holding the bond (and is therefore efficient). Again the size of the efficiency loss can be dramatic: at its worse it is roughly 0.8% of the portfolio value in only a year. (It is not exactly the same as for the very similar stop-loss strategy, since the state prices are not symmetrical for increases and decreases.)

Figure 3 Efficiency loss of a random timing strategy

This figure gives the efficiency loss of a random timing strategy. A random timing strategy is a strategy followed by an agent who claims to have market timing ability but really does not. By assumption, such a manager spends a fixed fraction of the time fully invested in stocks and a fixed fraction of the time fully invested in bonds, using a rule that is independent of security returns. For the limits with 0% or 100% of the time spent invested in the stock, the strategy is efficient, since these limits correspond to buying and holding the bond or stocks, respectively. For other cases, the efficiency loss is even larger than for the stop-loss and lock-in strategies: at its worse it is nearly 2% of the portfolio value in only a year!

Figure 4 Efficiency loss of repeated portfolio insurance

This figure shows the efficiency loss of using portfolio insurance repeatedly. Under portfolio insurance, a dynamic strategy based on option pricing theory varies the portfolio mix between stocks and bonds to create a payoff at the end of the insurance horizon which is proportional to the larger of payoff to holding the stock and the initial investment. This plot shows that while following this strategy for one year is efficient, following it repeatedly with a one-year horizon is poorly diversified over time and is very costly: in 10 years, the strategy throws away over 5% of the initial investment!

FIGURE 1
Efficiency loss of a stop-loss strategy

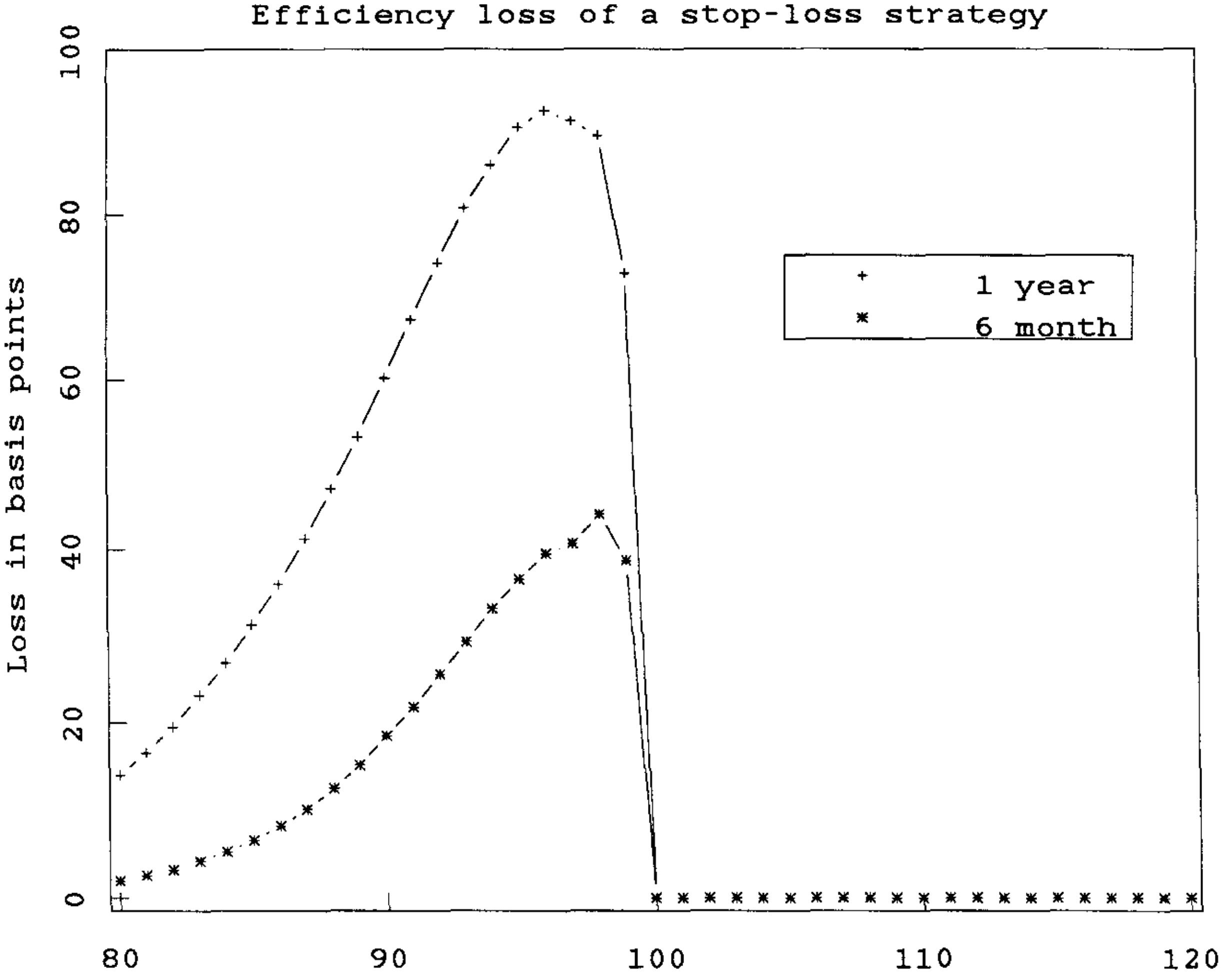

Limit value, percentage of initial investment
r=.08 mu=.16 sigma=.2 dt=1/360

FIGURE 2
Efficiency loss of a lock-in strategy

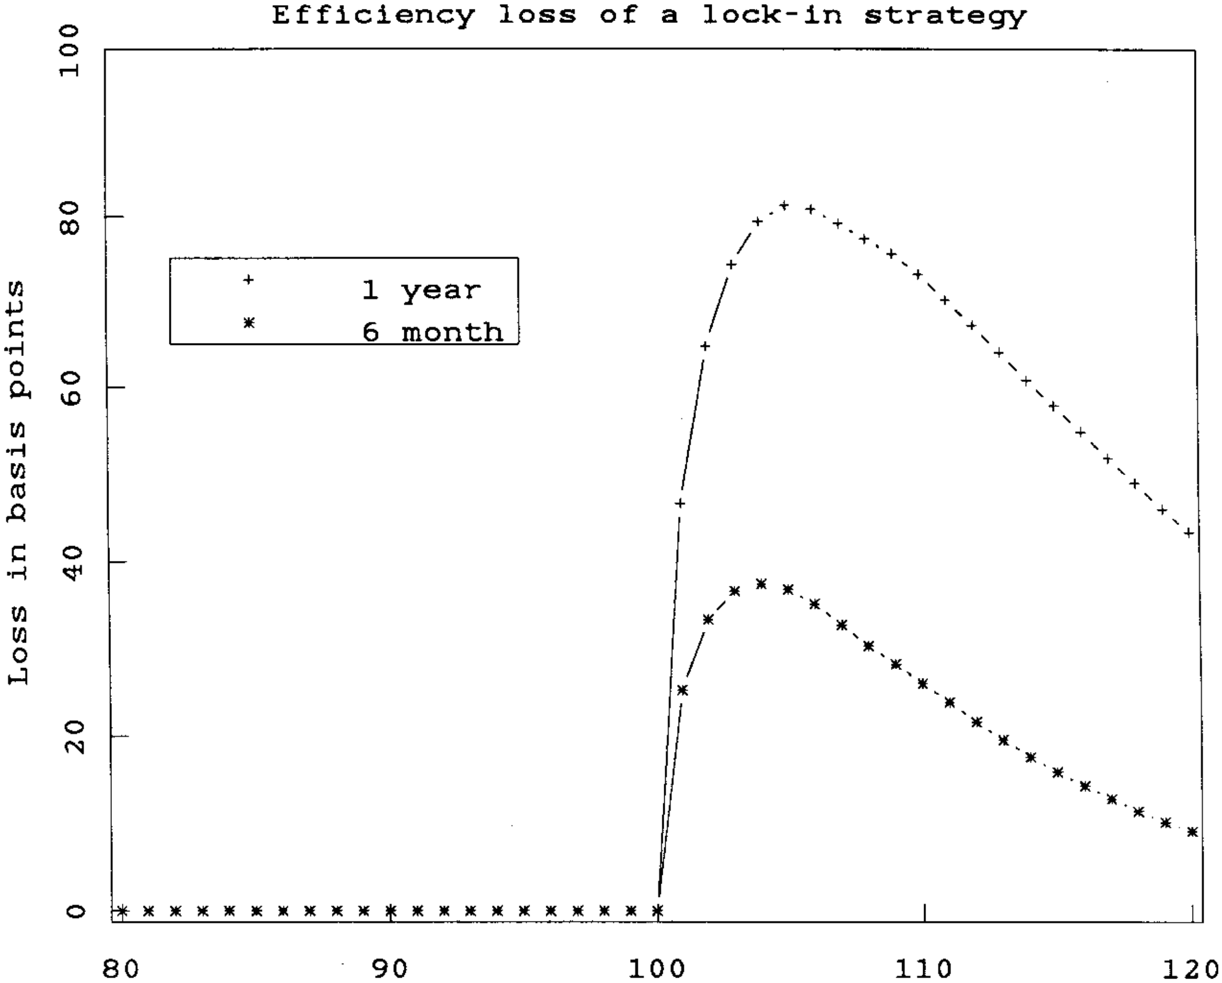

Limit value, percentage of initial investment
r=.08 mu=.16 sigma=.2 dt=1/360

FIGURE 3
Efficiency loss of a random timing strategy

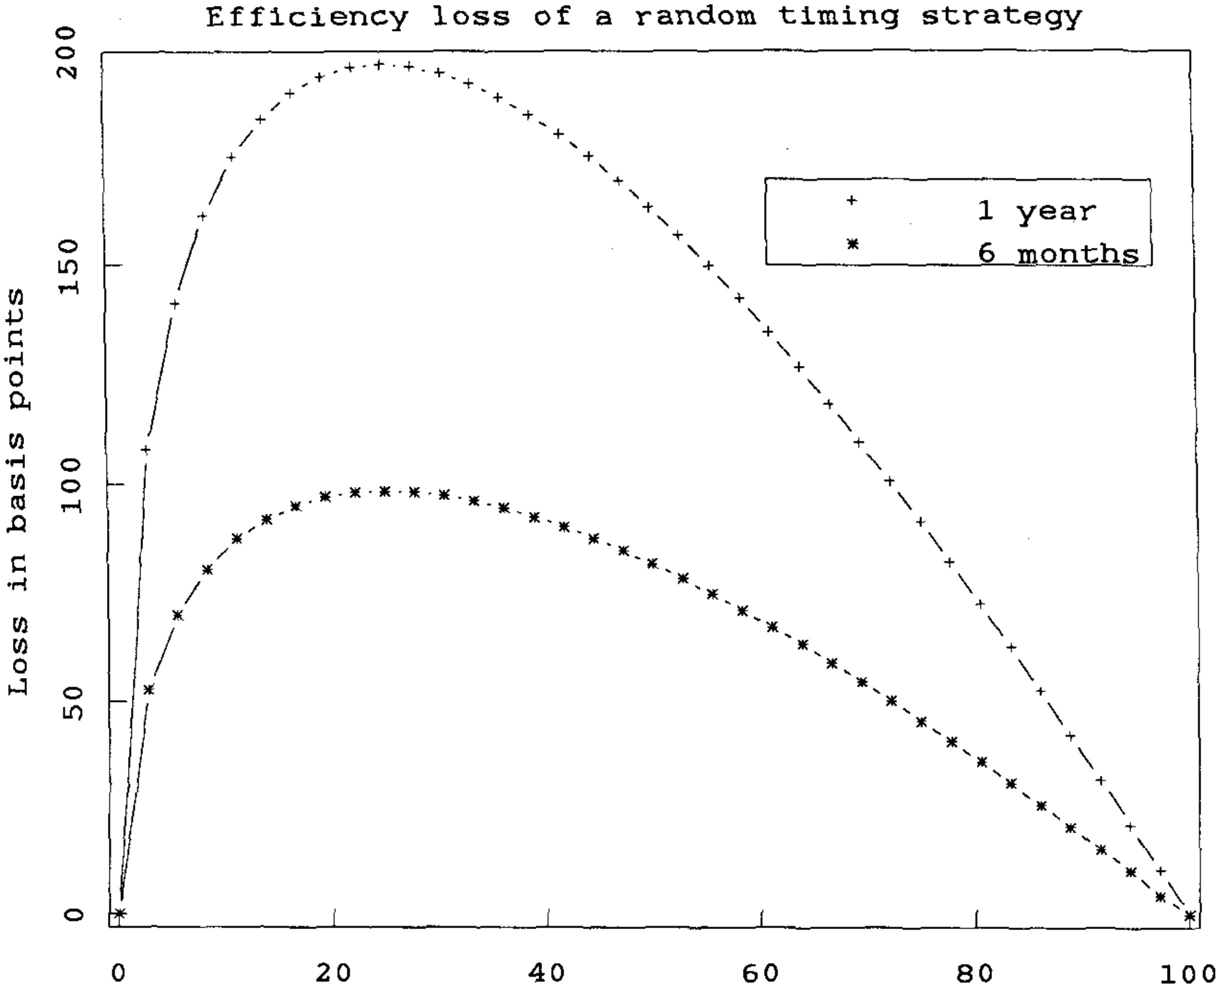

Percentage of time spent invested in stock
r=.08 mu=.16 sigma=.2 dt=1/360

FIGURE 4
Efficiency loss of repeated portfolio insurance

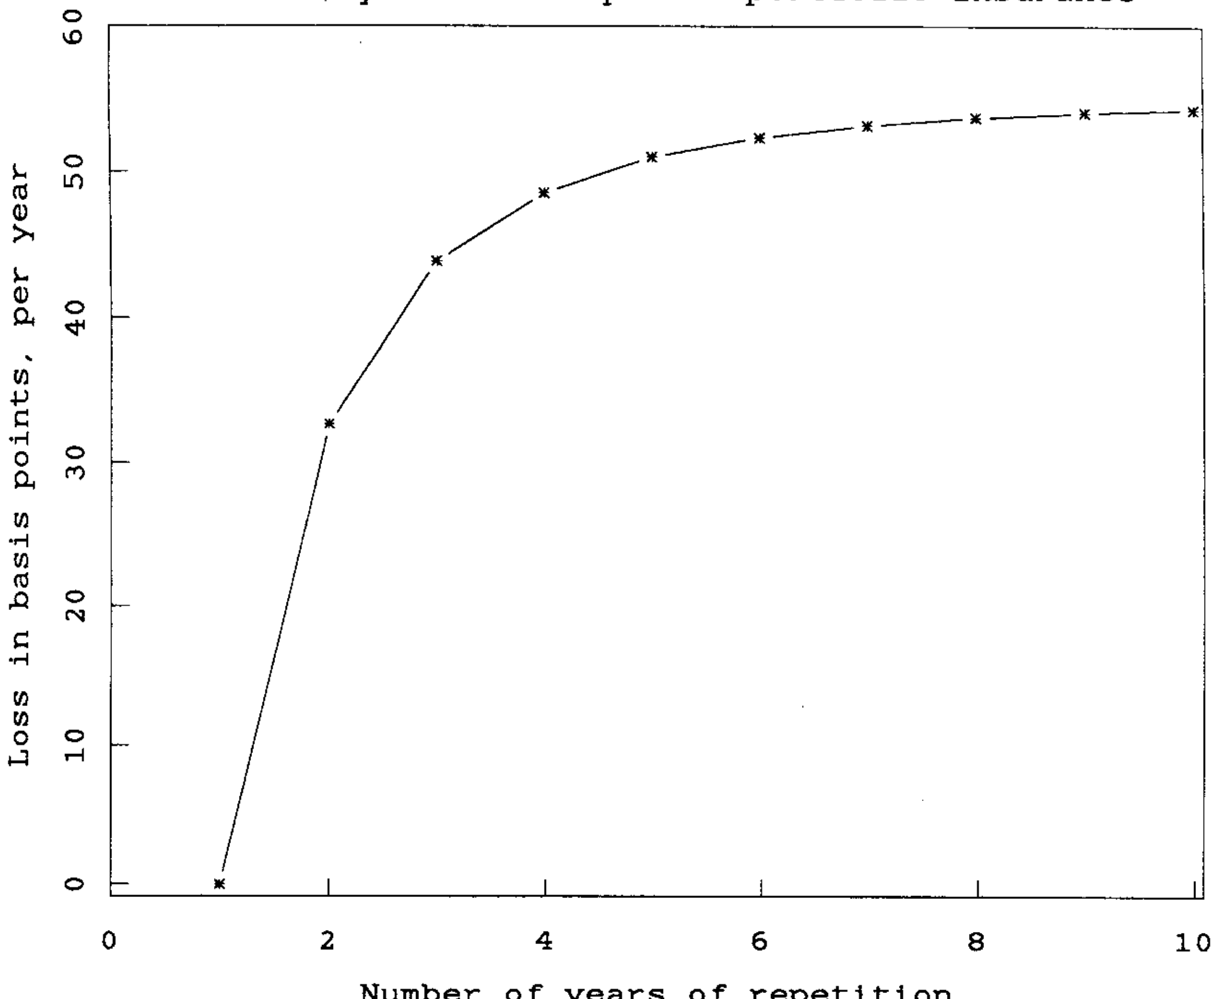

r=.08 mu=.1907 sigma=.2045 dt=1/52