SIMON BENNINGA*

Comparing Portfolio Insurance Strategies

1. Introduction

Portfolio insurance gained both respectability and momentum with the introduction of synthetic put strategies in the early 80's by Leland and Rubinstein. Current estimates of the total size of portfolios managed under portfolio insurance strategies range between 50 and 100 billion dollars (BRENNAN and SCHWARTZ 1987). These estimates cover commercially marketed portfolio insurance strategies, such as synthetic put strategies and the constant proportion portfolio insurance strategies (PEROLD 1986, BLACK and JONES 1986). Stop-loss rules, another strategy which protects portfolio value, are not included in these estimates.

It is difficult to know how the different portfolio insurance strategies are to be evaluated. Neither the stop-loss strategy nor the synthetic put strategies are utility maximizing (BRENNAN and SOLANKI 1981, BENNINGA and BLUME 1986). The unconstrained constant proportion portfolio strategy is utility maximizing for a subset of constant elasticity utility functions (see ZHU and KAVEE 1987), but given the short sale constraints imposed on this strategy in its usual commercial implementation, this strategy, too, is unlikely to be utility maximizing.

Given the almost certain failure - suggested by the previous paragraph - of any attempt at theoretical justification of portfolio insurance strategies, we turn to an empirical evaluation. One empirical approach to evaluation is to test various strategies against actual data. This is the approach taken by JACOBS (1983). A problem with this approach is that it perforce examines only one of the many possible price paths. In a sense every evaluation against actual data, no matter how illuminating, is a special case.

An alternative to testing against actual data is to test the data in a simulated environment. This has the advantage that we can specify the stochastic parameters of asset prices and that we can examine many possible price paths. It is this approach which we take.

Our basic framework is to compare the various portfolio insurance strategies cross-sectionally. We consider an investor who has initial wealth 100 and who can invest in a portfolio composed of a risk-free asset and a risky asset. Suppose that the investor wishes to guarantee that the value of his portfolio at the end of the year be not less than some floor F. In this study the investor can use any one of the three strategies to implement this goal. Using Monte Carlo simulation techniques, we shall relate the strategies for floor F by comparing their expected

* I thank Yakov Bergman, Marshall Blume, Edwin Elton, Günter Franke, Dan Galai, Don Keim, Aris Protopapadakis, Walter Wasserfallen, and an anonymous referee for helpful comments. Final responsibility is, of course, mine. This research has been supported by a grant from the Leonard Davis Institute for International Relations of the Hebrew University.

end-of-period wealths, the standard deviation of the end-of-period wealth, return to variability (Sharpe's measure), skewness, and sensitivity to transactions costs. Replicating these comparisons for a variety of floors generates the cross-sectional comparison.
The stochastic environment which we simulate corresponds to the assumption that risky asset returns are lognormally distributed. In this environment - which we replicate for 500 years of data - we simulate the performance of three portfolio insurance strategies: a stop-loss (SL) strategy, a synthetic put (SP) strategy, and a constant proportion portfolio insurance strategy (CPPI) [1].
The structure of this paper is as follows. In the next section we describe the three portfolio insurance strategies studied in this paper, paying special attention to the implementation of transactions costs and short-sale constraints (for constant proportional portfolio insurance). Section 3 discusses transactions costs and Section 4 discusses the stochastic environment used to simulate the strategies. The results of the simulation tests are presented in Sections 5 and 6.

2. Describing the portfolio strategies

All of our strategies involve an investor who has a one year horizon. The investor starts his year with initial wealth of 100, and can invest this wealth in either a risky asset or in a risk-free asset. The risky asset's log returns are normally distributed, with mean log return μ and standard deviation of log return σ. When we approximate this lognormal distribution discretely, we shall assume that the return distribution is approximately lognormally distributed at intervals of Δt. The risk-free asset earns a continuously compounded risk-free return of r.
The investor invests in his portfolio with the desire that the end-of-year wealth from his investment be not less than some floor F. We investigate three strategies by which this goal may be achieved.

2.1. Stop loss portfolio insurance

A classic portfolio insurance strategy is the stop loss strategy. In this strategy the investor initially invests all of his wealth in the risky asset. If at time the price $P_t \leq Fe^{-r(1-t)}$, then the investor sells all of his holdings of the risky asset and invests in the risk-free asset. In a continuous environment without transactions costs the investor sells the first time that $P_t = Fe^{-r(1-t)}$, and this guarantees that the investor's final wealth will never be lower than F.
In the simulation environment of this study, we assume that the investor sells out of the risky asset at the first t for which $P_t < Fe^{-r(1-t)}$. We further assume that the price at which the investor sells out is P_t [2].

2.2. Synthetic put portfolio insurance

A second portfolio insurance strategy, the synthetic put strategy, was first proposed by LELAND and RUBINSTEIN (1981). In its purest form this strategy uses the Black-Scholes option pricing formula to create a continuously adjusted synthetic European put on the risky asset [3]. Combining the purchase of a risky asset with the purchase of a put on the asset is equivalent to purchasing a continuously-adjusted portfolio which is a combination of the asset and of a risk-free bond. The Black-Scholes option pricing formulas show how to adjust these asset proportions (see details in the Appendix).
An investor following a synthetic put strategy will, at the inception of the period, invest in a portfolio of risk-free bonds and the risky asset. If the price of the risky asset increases, he will increase the proportion of the risky asset in the portfolio; if the price of the risky asset decreases, the proportion of the risky asset in the portfolio will decrease. At the close of the target period, the investor will be wholly invested in either the risky asset or the risk-free bond [4].

2.3. Constant proportion portfolio insurance

An alternative approach to portfolio insurance is the constant proportion portfolio insurance (CPPI). This idea was originally proposed by PEROLD (1986) and BLACK and JONES (1986); see also ZHU and KAVEE (1987) and ROUHANI (1987).

A CPPI strategy is a form of a constant proportion portfolio strategy. In such a strategy the investor invests a proportion γ of his wealth in a risky asset and the remaining proportion $(1-\gamma)$ of his wealth in a risk-free asset. CPPI is a constant proportion portfolio strategy with one added feature. At any time t, the investor who follows CPPI will subtract a floor F from his wealth before computing the amount to be invested in the risky asset [5]. Thus, if W_t is the investor's wealth at time t, then the investor invests $m(W_t-F)$ in the risky asset, and the remainder of his wealth, $(1-m)(W_t-F) + F$, in the risk-free asset. The number m is termed the CPPI multiple. In order to lever the investment in the risky asset, the multiple in CPPI strategies is taken to be greater than 1. In the simulations, m takes alternative values of 2, 3, 4 and 5.

CPPI has some attractive properties as a portfolio strategy. For a variant of the constant proportional risk aversion function, the strategy is optimal (see ZHU and KAVEE 1987). When unconstrained, the strategy is also path-independent (see ROUHANI 1987) [6].

An unconstrained CPPI strategy can easily lead to the investor being short the risk-free asset (when the price of the risky asset is high) or being short the risky asset (when the price of the risky asset is low). In commercial applications it is common to bound the CPPI strategy to avoid short positions in either of the assets. In our simulations we have followed this practice.

Short-sale constraints become particularly important for high multiples and for extremely low or high floors. Thus, for example, an investor having initial wealth of 100, a floor of 70 and a multiple of 5 will wish to invest 5(100-70)=150 in the risky asset and -50 in the risk-free asset at the beginning of each investment period. An investor having a floor of 105 will wish to invest a negative amount in the risky asset at the beginning of the investment period. Short-sale constraints prevent these kinds of strategies (both initially and throughout); in the two examples given in this paragraph, for example, the initial investment in the risky asset would be taken as 100 and 0 respectively.

3. Transactions costs

Both the synthetic put strategy and the constant proportional portfolio insurance strategy typically involve a large number of transactions. Therefore, an important question for the user of any portfolio insurance strategy which involves many changes in the portfolio position is the size of the transactions costs involved. In the simulations we model the transactions costs as a proportion of the change in the portfolio positions. We make the following assumptions:

1. The determination of the initial position (i.e., the time 0 division of wealth between the risky and risk-free asset) does not incur transactions costs. We use this assumption for all three strategies.
2. For SP and CPPI, changes in the allocation of the assets are done so that the ratio of the risk-free to risky asset remains the same after allowing for transactions costs as it would have been without transactions costs.
3. The SL strategy is assumed to incur transaction costs only if the investor sells out his risky asset and purchases the risk-free asset.

Operative details of how these assumptions have been implemented are given in the Appendix.

4. The stochastic environment

Using a technique described in KNUTH (1981) (see Appendix for details), we use a routine which creates normal deviates, which are in turn used to create a pattern of returns which are lognormally distributed with mean μ and standard deviation σ, and time between observations of Δt. A typical example uses $\mu = 10\%$, $\sigma = 20\%$, and $\Delta t = 0.004$ (approximately daily revision of portfolio). Repeating this process N times simulates N years of data.

5. Some simulation results - the case of transactions costs

The graphs present some typical simulation results for the three strategies. The simulations use a reasonable proportional transactions cost of 0.5% (0.005). The simulations employ varying frequencies of portfolio revision: 50 annual revisions, 100 annual revisions, and 250 annual revisions. For a given frequency of portfolio revision, all of the simulations use the same series of prices (discrete approximations to the lognormal distribution with annual parameters $\mu = 10\%$, $\sigma = 20\%$, $r = 7\%$). Given 250 business days per year, these revision frequencies correspond to approximately weekly, bi-weekly, and daily revisions. The results discussed below compare simulations of 500 years of data. Clearly such a wealth of data is available only through the use of simulation techniques. Simulations for other values of μ and σ gave qualitatively similar results.

Comparison of average terminal wealth:
Figure 1 presents comparative data for mean terminal wealth from each of the strategies. Not surprisingly, the highest average terminal wealth results

Figure 1: Comparison of Average Terminal Wealth. (250 annual revisions).

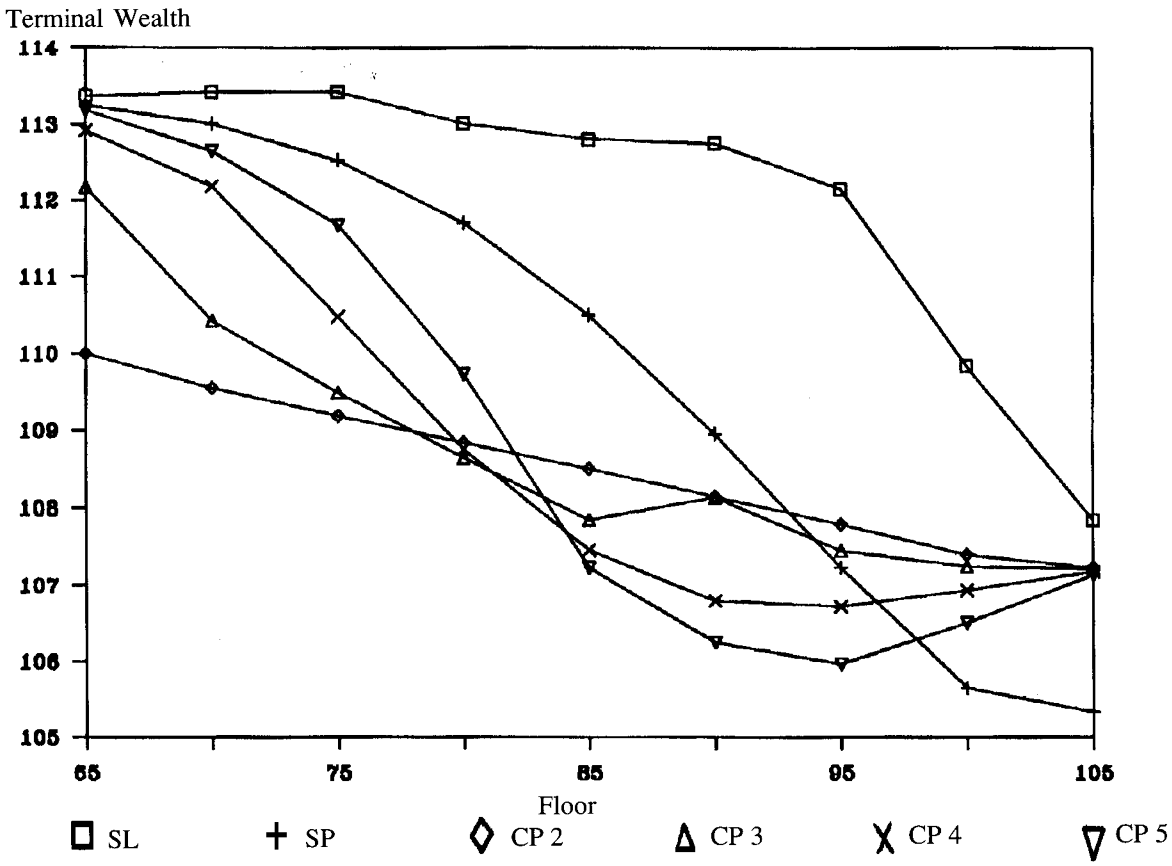

Notes: Average terminal wealth for simulations with $\mu = 10\%$, $\sigma = 20\%$, and transactions costs = 0.5%. Note the dominance of the stop loss strategy for all floors.

Figure 2: Comparison of Sharpe's Index. (250 annual revisions).

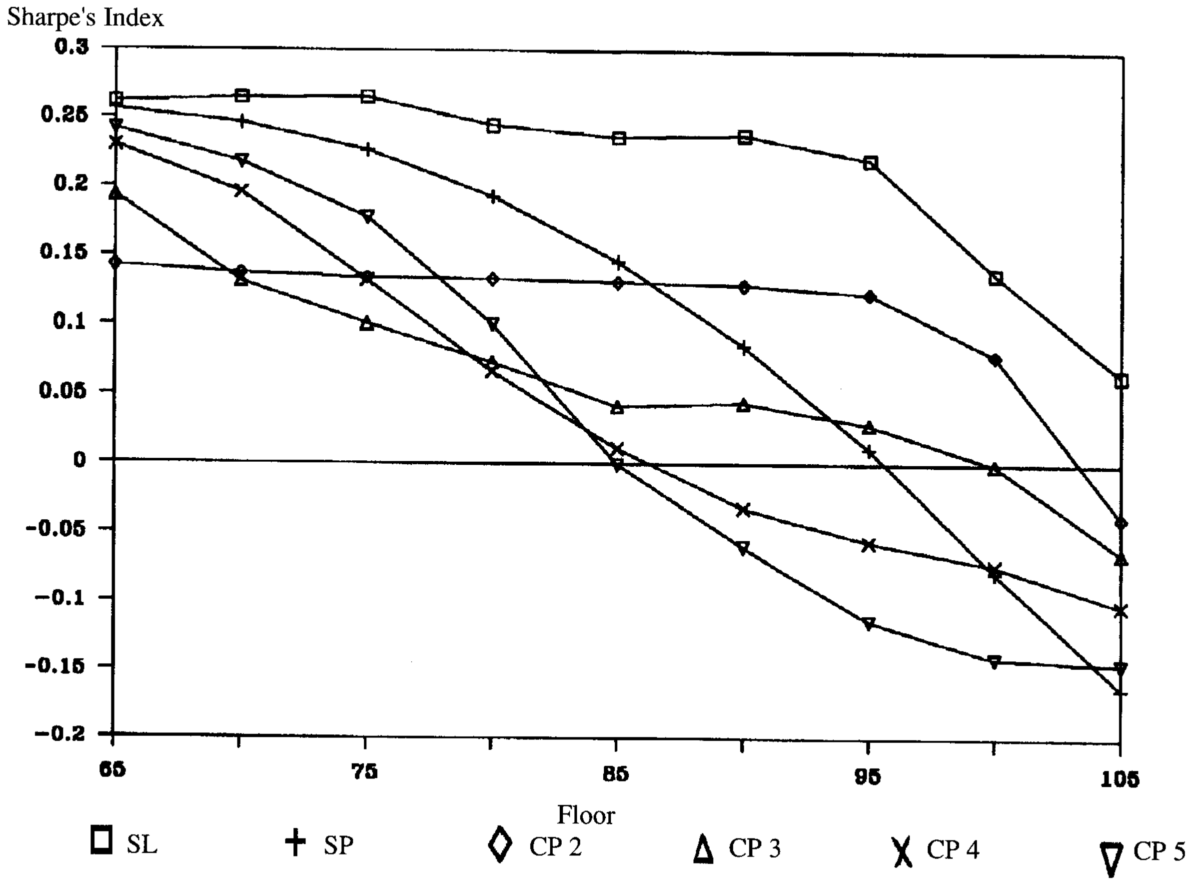

Notes: Sharpe's index (reward to variability) for simulations with $\mu = 10\%$, $\sigma = 20\%$, and transactions costs = 0.5%. For high floors, the average terminal wealth for synthetic put strategies and CPPI strategies with multiples ≥ 3 is less than the average terminal wealth for a strategy invested wholly in the risk-free asset; this explains the negative reward to variability for these strategies when the floor ≥ 100.

from the stop-loss strategy. For floors which approximate the initial wealth of the investor (100), the synthetic put strategy gives the lowest average terminal wealth, with the four CPPI strategies in the middle.

Returns to variability - Sharpe's index:
We calculated Sharpe's index for each of the strategies. This involved calculating:

$$\frac{\text{average terminal wealth-(initial wealth} \bullet e^{\text{risk-free}})}{\text{sigma of terminal wealth}}$$

for each floor and for each strategy. The results are graphed in Figure 2. The stop-loss strategy gives the highest return to variability throughout. For the higher floors, most of the other strategies give a negative reward to variability; this means that, on average, terminal wealth would have been higher had the investor invested only in the risk-free asset.

Skewness:
In addition to Sharpe's index, we calculated the skewness of each strategy (the third moment of the strategy's return, divided by σ^3). Skewness is often claimed to be an advantage of dynamic portfolio insurance strategies. However, except for the highest floor level (105), there do not appear to be any substantial differences between the skewness measures of the stop-loss and the synthetic put strategies.

For a floor level of 105, however, the stop-loss strategy had significantly lower skewness than the synthetic put strategy. The difference between the skewness measures for the stop-loss and the synthetic put strategies for the floor of 105 are explained as follows: At a floor of 105, 86.8% of the yearly runs of the stop-loss strategy end with a stopped out strategy (i.e., conversion of the risky asset into the risk-free). The fact that on average this conversion occurs very early in the year (before the first quarter of the year) means that the stop-loss strategy for this floor level has a low skewness. The synthetic put strategy, on the other hand, leaves the investor always invested (if only partly) in the risky asset, so that he can pick up the advantages of increases in this assets price. This explains the differences in the skewness for the high floor level.

6. The importance of transactions costs

We tested the importance of transactions costs by using three different measures. First, we calculated the level of transactions costs for each strategy. The results for 250 annual revisions are presented in Figure 3 [7].

Several results are of interest:

1. Although the transactions costs for both the stop-loss and the synthetic put strategies rise as

Figure 3: Comparison of Transactions Costs. (250 annual revisions).

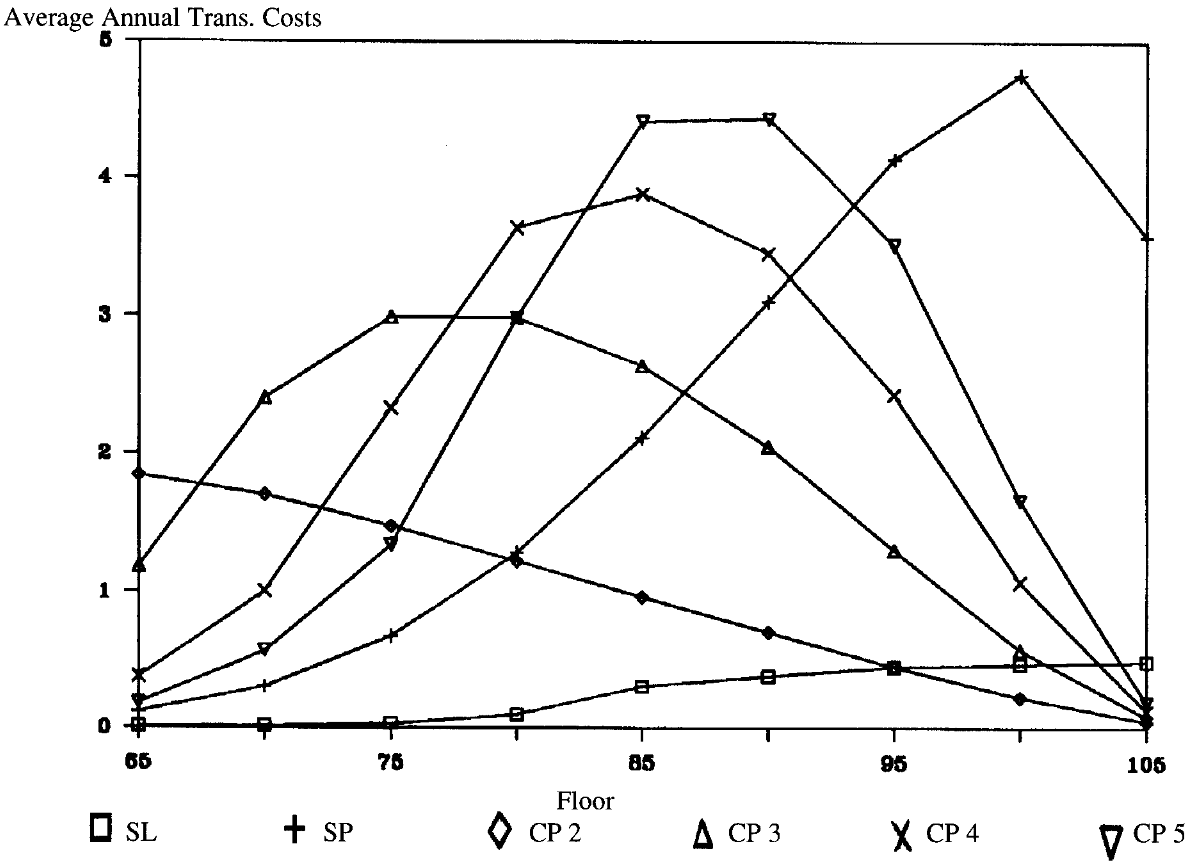

Notes:
Average annual transactions costs for simulations with $\mu = 10\%$, $\sigma = 20\%$, and transactions costs = 0.5%. Transactions costs are totaled over the year. CPPI and SP strategies with low floors tend to be largely in the risky asset and have low transactions volume. With high floors the CPPI strategies tend to be largely in the risk-free asset, which explains their low transactions costs.

the floor rises, the increase in the latter is much more dramatic than that of the former. The explanation for this is that the synthetic put strategy involves, as the floor rises, both a larger proportion of wealth invested in the riskfree asset and substantially greater changes in the trading volume. The greater transactions costs for the SL strategy are explained primarily by the fact that as the floor rises, so does both the probability of selling out of the risky asset and the price of the risky asset when it is sold out.

2. The transactions costs for the CPPI strategies "peak out" at a midrange of floors, and low transactions costs for both low floors and high floors. (This is also true for CPPI with a multiple of 2, despite the fact that it is not evident from Figure 4. The peak of transactions costs for this multiple is for a floor less than 65, and is not shown in the graph).

A second way of judging the effect of transactions costs is to consider the decrease in average terminal wealth for each floor as trading frequency increases. For CPPI strategies, trading frequency makes practically no difference in average wealth when the floor is high. The reason for this is that for a high floor only a proportion of the portfolio is, on average, invested in the risky asset; this follows from the fact that on average W_t-F is low. For low

Figure 4: Average Terminal Wealth, No Transactions Costs.

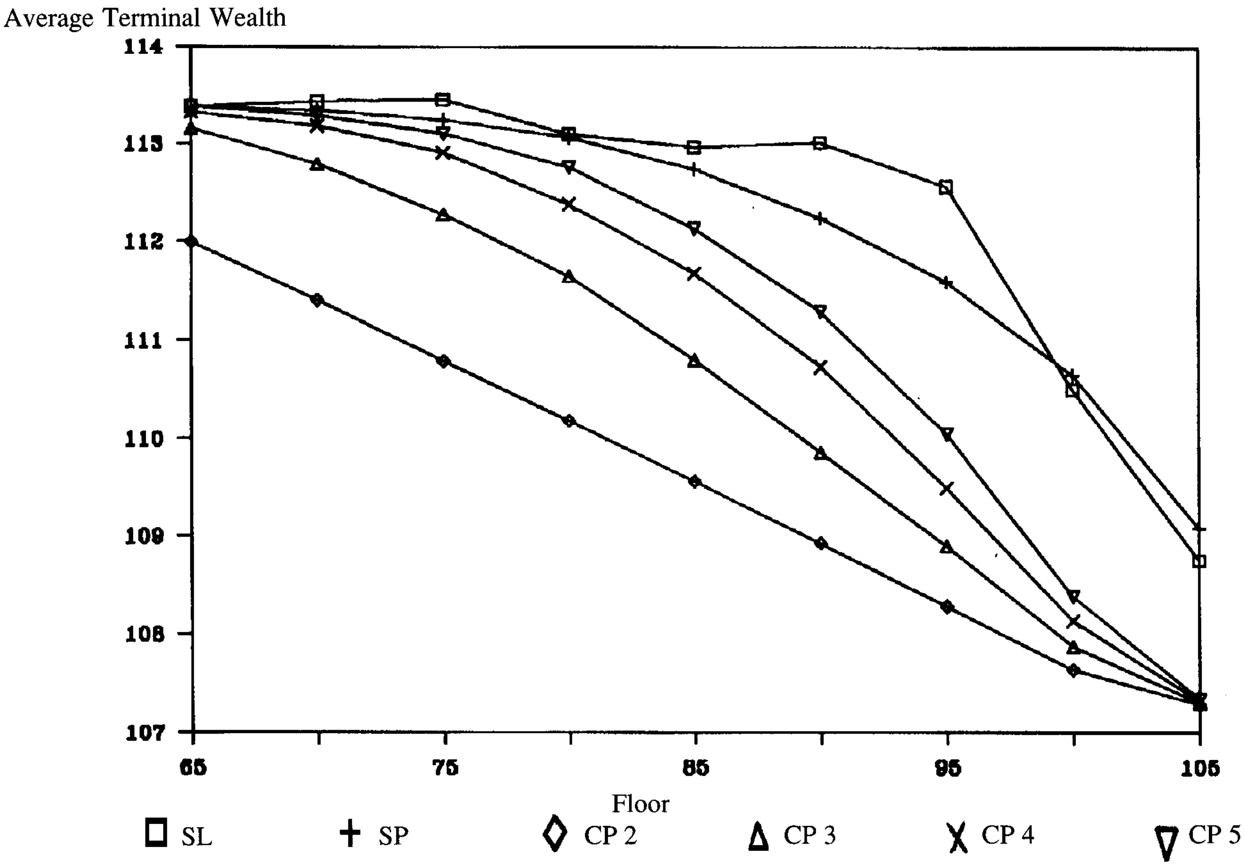

Notes:
Average terminal wealth without transactions costs. The simulation uses $\mu = 10\%$, $\sigma = 20\%$. Note that for all but the highest floors (F=100,105) stop loss strategies produce the highest average terminal wealth.

Figure 5: Sharpe's Index, No Transactions Costs.

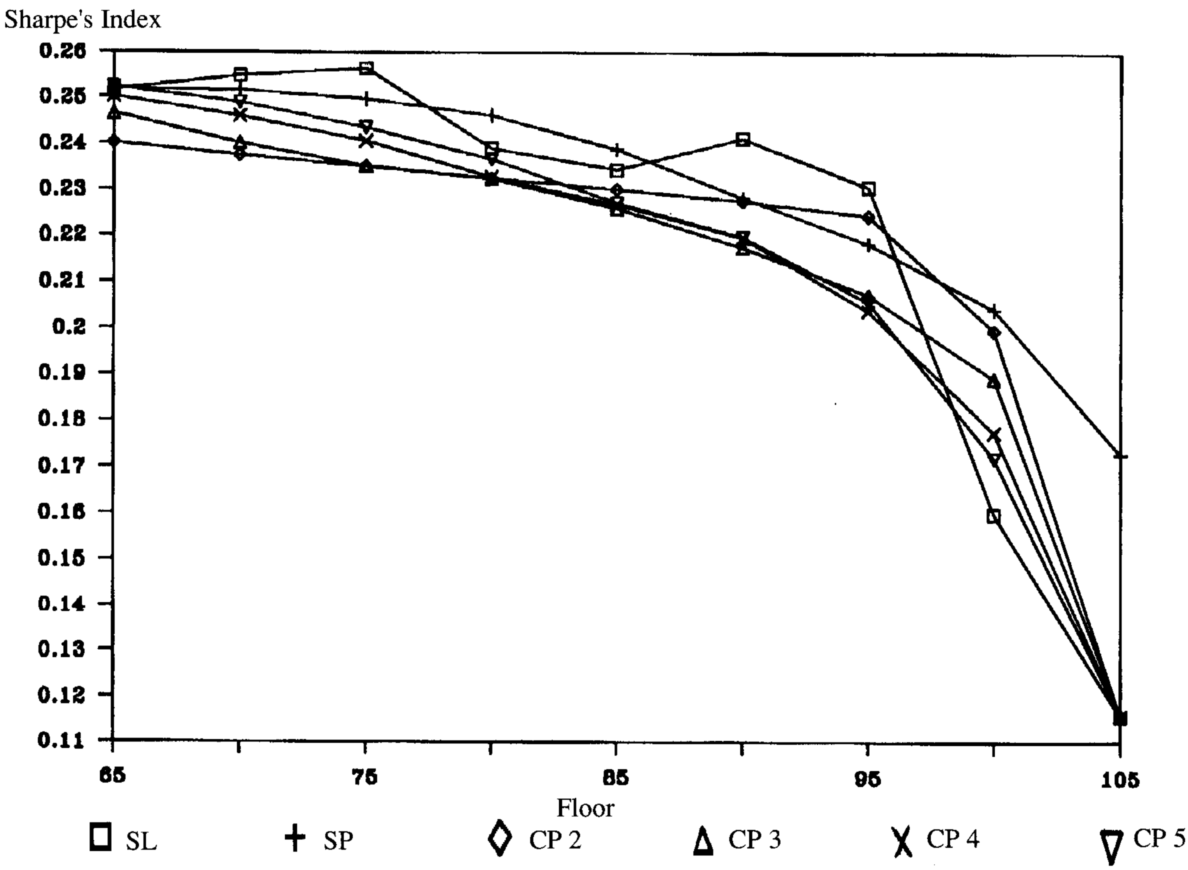

Notes:
Sharpe's index without transactions costs. The simulations use $\mu = 10\%$, $\sigma = 20\%$. A comparison with Figure 2 shows that transactions costs are the primary factor determining reward to variability. Without transactions costs it is difficult to differentiate meaningfully between the strategies for Sharpe's index.

floors, however, increasing the trading frequency has a significant effect (around 1% decrease for an increase from 50 transactions per year to 250 per year) on the average terminal wealth for CPPI strategies.

For synthetic put strategies the picture is the opposite: At low floor levels, these strategies are largely invested in the risky asset, and therefore are less sensitive to revision frequency. At higher floor levels, these strategies become increasingly sensitive to revisions frequency, and we again observe a reduction of up to 1% in the level of average terminal wealth.

A third way of judging the effect of transactions costs is to run the simulation on the same price data with and without the transactions cost. This is done in Figures 4 and 5. Again we note that CPPI strategies are insensitive to transactions costs for high floors.

Note that when we rerun the simulation without transactions costs, the average wealth attained with the stop-loss strategy is still higher than that attainned with the synthetic put strategy for all but the highest floors, although the differences between the two strategies are very much attenuated. Significant differences between Sharpe's index for the two strategies arise only for the highest floors ($F > 95$).

7. Conclusion

In this paper we have simulated three popular portfolio insurance strategies: stop-loss rules, synthetic put strategies, and constant proportional portfolio insurance strategies. All three strategies were compared on a cross-sectional basis; i.e., comparisons were made across strategies which set out to achieve the same minimum level of terminal wealth. Simulations were used because they allow us to test a strategy against a large number of realizations of the same stochastic scenario.

The surprising fact which comes out of the simulations of the three portfolio insurance strategies is that in terms of both expected terminal wealth and Sharpe's index of reward to variability, the stop-loss strategy - a primitive, unsophisticated portfolio insurance rule - dominates both synthetic put and constant proportional portfolio insurance strategies for all floors for a fairly low level of transactions costs.

Even in the absence of transactions costs, it can be seen from Figures 4 and 5 that the stop-loss strategies dominate other portfolio insurance strategies for all but the highest levels of floors.

There are two complicating factors which we did not consider, but which are important in practice: First, when the risky portfolio is composed of many risky assets, the role of transactions costs becomes more important. Second, we have assumed that the risky asset does not pay dividends. When dividend payments are small relative to price movements, they complicate the simulation (and, indeed, the implementation) of the various strategies, but they do not change them materially.

Given the simplicity of the stop-loss rule, perhaps it warrants reconsideration by portfolio insurers.

Footnotes

[1] The most complete simulation study to date is included in ZHU and KAVEE (1987). However their study does not include the stop-loss strategy, nor does it include cross-sectional comparisons of the results of the strategies. Other simulation studies are those of JACOBS (1983), RUBINSTEIN (1985) and ROUHANI (1987).

[2] This assumption is used in the simulations of all the strategies. An alternate assumption would be that the portfolio of risky assets is liquidated at the period t+1 price. Because the price process has a positive trend, this would cause the average terminal wealth of the strategy to rise, though the effect is substantial only when the number of annual portfolio revisions is small (i.e. for Δt relatively large).

[3] Commercial applications often use a position in a futures contract to implement a change in a position in the risky asset. If one assumes perfect correlation between the futures price and that of the underlying asset, this is an innocuous procedure. Given the marking to market on futures and the fact that stock futures are imperfect substitutes for most stock portfolios, alternative assumptions are difficult to postulate (both theoretically and empirically), and we have therefore chosen not to simulate this problem. In actual fact the low correlation between futures and stock prices in October 1987 caused many problems in the execution of portfolio insurance strategies.

[4] For proofs of these assertions, see BENNINGA and BLUME (1985).

[5] Although in principle the floor can be time dependent, we shall simulate only cases in which the floor is fixed over the horizon period. This makes for greater comparability with the other two strategies.

[6] The imposition of short-sale constraints on CPPI makes the strategy path dependent. Note that SP strategies are always path independent (see RUBINSTEIN 1985 for a discussion of this point).

[7] It bears repeating that the actual results presented for terminal wealth, sigma of terminal wealth and skewness are net of transactions costs.

References

BENNINGA, S. (1989): "Numerical Techniques in Finance", Cambridge: MIT Press.

BENNINGA, S. and M. BLUME (1985): "On the Optimality of Portfolio Insurance", Journal of Finance 40, December, pp. 1341-52.

BLACK, F. and R. W. JONES (1986): "Simplifying Portfolio Insurance", Goldman, Sachs, & Co.

BRENNAN, M. J. and E. S. SCHWARTZ (1987): "Portfolio Insurance and Market Volatility", unpublished working paper, John E. Anderson Graduate School of Management, UCLA.

BRENNAN, M. J. and R. SOLANKI (1981): "Optimal Portfolio Insurance", Journal of Financial and Quantitative Analysis 16, September, pp. 279-300.

JACOBS, B. (1983): "The Portfolio Insurance Puzzle", Pensions and Investment Age, August 22.

KNUTH, D. E. (1981): "The Art of Computer Programming, Vol. 2: Seminumerical Algorithms", Reading, MA: Addison-Wesley.
LELAND, H. E. (1985): "Option Pricing and Replication with Transaction Costs", Journal of Finance 40, December, pp. 1283-1301.
LELAND, H.E. and M. Rubinstein (1981): "Replicating Options with Positions in Stock and Cash", Financial Analysts Journal July/August, pp. 3-12.
PEROLD, A. (1986): "Constant Proportion Portfolio Insurance". Harvard Business School.
ROUHANI, R. (1987): "Constant Proportion Portfolio Insurance: Expected Return Estimation and Sensitivity Analysis", Goldman Sachs & Co.
RUBINSTEIN, M. (1985): "Alternative Paths to Portfolio Insurance", Financial Analysts Journal, July/August, pp. 42-52.
ZHU, Y. and R. C. KAVEE (1987): "Performance of Portfolio Insurance Strategies", Journal of Portfolio Management.

Appendix

In this appendix we briefly state some of the technical results used in the simulations.

A1. Deriving the synthetic put strategy

The Black-Scholes put option price at time t, $0 \leq t \leq 1$ is:

$$P_t = - S_t N(-h) + Ke^{-r\tau} N(\sigma\sqrt{\tau} - h),$$

where

h	=	$\ln(S_t e^{r\tau}/K) / \sigma\sqrt{\tau} + 0.5\sigma\sqrt{\tau}$
τ	=	1-t.
r	is	the risk-free rate of interest,
σ	is	the standard deviation of the log return of the risky asset,
K	is	the exercise price of the put.

Buying one share of risky asset and buying a put on the share with exercise price K is therefore equivalent to buying

$S_t[1-N(-h)]$ of the risky asset and
$Ke^{-r\tau}N(\sigma\sqrt{\tau}-h)$ of the risk-free bond.

The total investment required to buy one share of the risky asset plus the put is S_t+P_t. In terms of portfolio proportions, a synthetic put strategy involves investing a proportion

$$\omega_t = \frac{S_t[1-N(-h)]}{S_t+P_t} = \frac{S_t[1-N(-h)]}{S_t[1-N(-h)]+Ke^{-r\tau}N(\sigma\sqrt{\tau}-h)}$$

of the investor's wealth in the risky asset at each time t.

We note that in a portfolio insurance strategy the exercise price of the put K has to be higher than the desired minimal terminal wealth F. This follows from the fact that the put itself costs money. Thus the investor with an initial wealth of 100 faced with the problem of purchasing δ puts with exercise price of K and δ shares whose initial cost per share is 100 will find that he can only afford $\delta < 1$. In order to guarantee that the terminal wealth of the investor is no less than F, it follows that K > F. All of the synthetic put simulated in this paper have had the appropriate adjustment made to the exercise price (see BENNINGA 1989 for details).

A2. Modeling transactions costs for CPPI

Consider a CPPI strategy with multiple m and transactions costs α. Denote by X_t, Y_t the value of the risky and risk-free assets respectively at the beginning of period t and denote by x_t, y_t the value of the allocations to the risky and risk-free assets respectively at the end of period t. Denote by R_t one-plus-the-return on the risky asset (R_t is a stochastic random variable, whose natural logarithm is distributed normally with mean μ and standard deviation σ). If at the beginning of period t the value of the risky assets is X_t and the value of the risk-free assets is Y_t (where $X_t = x_{t-1}R_{t-1}$ and $Y_t = y_{t-1}e^{r\Delta t}$), then we have to find Δx and Δy such that:

$$x_t = X_t - \Delta x + \alpha|\Delta x|,$$
$$Y_t = Y_t - \Delta y + \alpha|\Delta y|.$$

First of all it is clear that $\Delta x + \Delta y = 0$. Next, following the allocation rule for CPPI, we write

$$x_t = m(X_t + Y_t - F - 2\alpha|\Delta x|).$$

Setting the two expressions for x_t equal to one another, we get:

$$X_t + \Delta x - \alpha|\Delta x| = m(X_t + Y_t - F - 2\alpha|\Delta x|),$$

which solves to give

$$\Delta x = \begin{cases} \dfrac{(m-1)X_t + mY_t - mF}{1 + (2m-1)\alpha} & \text{when } \Delta x > 0, \\[2ex] \dfrac{(m-1)X_t + mY_t - mF}{1 - (2m-1)\alpha} & \text{wehn } \Delta x < 0. \end{cases}$$

An exception to these rules is when the strategy hits the no-short-sales constraint. Suppose, for example, that $x_t = 0$ (i.e., the investor sells all of his risky asset at time t - presumably because he would, in an unconstrained CPPI go short these assets). Then

$$y_t = Y_t + X_t(1-\alpha) / (1+\alpha).$$

A similar rule applies for the case where $y_t = 0$.

A3. Modeling transactions costs for synthetic put strategies

We use the same notation as above. At time t, suppose that the desired ratio of risky assets to wealth is ß. Then it follows that

$$\beta = \frac{X_t + \Delta x - \alpha|\Delta x|}{X_t + Y_t - 2\alpha|\Delta x|}.$$

This solves to give:

$$\Delta x = \begin{cases} \dfrac{X_t(1-\beta) - \delta Y_t}{\alpha(1-2\beta) - 1} & \text{when } \Delta x > 0 \\[2ex] \dfrac{-X_t(1-\beta) + \delta Y_t}{1+\alpha(1-2\beta)} & \text{when } \Delta x < 0. \end{cases}$$

A4. Modelling the stochastic environment

The algorithm used to create the normal deviates is described and proven in KNUTH (1981):

1. Generate two random numbers, call them random1 and random2, which are uniformly distributed between -1 and 1.
2. Calculate $S_1 = (\text{random1})^2 + (\text{random2})^2$. If $S_1 \geq 1$, repeat the previous step. Else proceed to the next step.
3. Calculate $S_2 = (-2\ln(S_1) / S_1)^{1/2}$.
4. The two normal deviates are given by $X_1 = S_2$ random1, $X_2 = S_2$ random2.

At each step we generate X_1 and X_2 and two lognormal returns, $\exp(\mu\Delta t+\sigma X_1\sqrt{\Delta t})$ and $\exp(\mu\Delta t + \sigma X_2\sqrt{\Delta t})$. If S_t is the risky asset price at the start of the period, then the next two risky asset prices are $S_{t+1} = S_t\exp(\mu\Delta t+\sigma X_1\sqrt{\Delta\tau})$ and $S_{t+2} = S_{t+1}\exp(\mu\Delta t+\sigma X_2\sqrt{\Delta t})$. Throughout we assume that S_0 (the risky asset price at the beginning of each year) is 100. Repeating the algorithm $1/(2\Delta t)$ times simulates one year of data.