UNIVERSITEIT
GENT

FACULTEIT ECONOMIE
EN BEDRIJFSKUNDE

TWEEKERKENSTRAAT 2
B-9000 GENT
Tel. : 32 - (0)9 – 264.34.61
Fax. : 32 - (0)9 – 264.35.92

WORKING PAPER

Performance evaluation of portfolio insurance strategies using stochastic dominance criteria

Jan Annaert

Sofieke Van Osselaer

Bert Verstraete

July 2007

2007/473

D/2007/7012/44

Performance evaluation of portfolio insurance strategies using stochastic dominance criteria*

Jan Annaert[a,b] , Sofieke Van Osselaer[a,] and Bert Verstraete[c]**

[a] *Ghent University, W. Wilsonplein 5D, Ghent, 9000, Belgium*

[b] *Antwerp University, Prinsstraat 13, Antwerp, 2000, Belgium*

[c] *KBC Asset Management*

Abstract

The continuing creation of portfolio insurance applications as well as the mixed research evidence suggests that so far no consensus has been reached about the effectiveness of portfolio insurance. Therefore, this paper provides a performance evaluation of the stop-loss, synthetic put and constant proportion portfolio insurance techniques based on a block-bootstrap simulation. Apart from more traditional performance measures, we consider the Value-at-risk and Expected Shortfall of the strategies, which are more appropriate in an insurance context. An additional performance evaluation is given by means of the stochastic dominance framework where we account for sampling error. A sensitivity analysis is performed in order to examine the impact on performance of a change in a specific decision variable (ceteris paribus). The results indicate that a buy-and-hold strategy does not dominate the portfolio insurance strategies at any stochastic dominance order. Moreover, both for the stop-loss and synthetic put strategy a 100% floor value outperforms lower floor values. For the CPPI strategy we find that a higher CPPI multiple enhances the upward potential of the CPPI strategies, but harms the protection level in return. As regards the optimal rebalancing frequency, daily rebalancing should be preferred for the synthetic put and CPPI strategy, despite the higher transaction costs.

JEL classification: G11

Keywords: Portfolio insurance; Performance evaluation; Stochastic dominance; Block-bootstrap simulation

---

* We are grateful to Isabelle Bajeux-Besnainou, Oliver Linton, and the participants of the FMA Europe Meeting 2006 in Stockholm.

** Corresponding author. Tel.: +32-9-264-78-95. Fax: +32-9-264-89-95.

E-mail address: sofieke.vanosselaer@ugent.be (S.Van Osselaer).

1. Introduction

In the early eighties[1] Leland and Rubinstein developed the portfolio insurance (PI) technique based on the option pricing formula of Black & Scholes (1973). The idea behind it was that a strategy which would provide protection against market losses while preserving the upward potential should have considerable appeal to a wide range of investors. This payoff pattern can be achieved by synthetically creating a put option on a portfolio. Nowadays, PI techniques, such as stop-loss, synthetic put, and constant proportion portfolio insurance (CPPI) continue to be regularly applied in practical applications, such as capital guaranteed funds, which confirms that the principle still enjoys great popularity (see e.g. the December 2005 issue of *Risk Magazine*).

From a theoretical perspective, Benninga & Blume (1985) show that depending on the investor's utility function, PI strategies may or may not be optimal. Analytical results are also presented in Brooks & Levy (1993). Under the assumption that returns are lognormally distributed they show that (very) risk averse investors may prefer PI to a buy-and-hold strategy. Empirical results are also mixed. For example, Garcia & Gould (1987) argue that PI cannot outperform static mix portfolios in the long run. Conversely, some recent papers, such as Cesari & Cremonini (2003) and Do & Faff (2004), provide empirical evidence on the benefits of PI in bear markets, although they do not perform formal statistical tests. Hence, the continuing use of PI as well as the mixed research evidence suggests that so far no consensus has been reached about its effectiveness.

In general, investment performance is traditionally measured from a mean-variance point of view. However, since PI strategies are developed to provide an upward potential combined with a downward protection, their evaluation should take into account their alleged improvement in the left tail of the return distribution. Therefore, previous research has incorporated measures such as floor protection (Do & Faff (2004)) or downside risk (Cesari & Cremonini (2003)). Surprisingly, the strategies have not yet been analyzed in terms of Value-at-Risk (VaR) and Expected Shortfall (ES), although Acerbi & Tasche (2002) and Jorion (2001), amongst others, have indicated that VaR and ES are more appropriate risk measures in an insurance context. Although the VaR measure – indicating the worst case loss at a certain confidence level – has become one of the standard risk measures in the financial management industry (cfr. Basel II), one of its main drawbacks remains that it is not a coherent risk measure. Therefore, the coherent risk measure ES has become a popular

[1] See Leland & Rubinstein (1988). According to Benninga & Blume (1985) and Bird *et al.* (1990), portfolio insurance contracts were already implemented in the United Kingdom in 1956 and in the United States in 1971.

alternative to VaR. Nevertheless, performance analyses based on these measures are prone to the criticism that the choice of the confidence level remains an arbitrary choice.

Given that the mean-variance framework is only partly useful in an insurance context, and that VaR and ES entail the arbitrary choice of a confidence level, we suggest using stochastic dominance rules for the performance analysis of PI strategies. This implies that we evaluate the performance of the strategies from the viewpoint of an expected utility maximizing investor. To our knowledge, comparing the strategies in a stochastic dominance framework – in which the whole return distribution is considered – is an application that has not been explored in a PI context so far. In fact, applying these rules in an empirical context has typically been marred by sampling error considerations (see e.g. Kroll & Levy (1980), Stein *et al.* (1983), Nelson & Pope (1991), amongst others), which explains why stochastic dominance applications have not seen a widespread proliferation. However, recent theoretical work on sampling error in stochastic dominance tests with iid data (e.g. Davidson & Duclos (2000) and Barrett & Donald (2003)), has triggered a new wave of stochastic dominance literature, e.g. Post (2003). Moreover, thanks to the subsampling method of Politis & Romano (1994) and insights of Linton *et al.* (2005) and Kläver (2005), it is now possible to handle the sampling error in stochastic dominance tests with non iid data. Using these techniques, PI strategies can be evaluated more profoundly by comparing whole return distributions rather than just some selected moments or quantiles.

Moreover, most research is centered on the synthetic put strategy, while little attention has been devoted to a profound comparison of this strategy with a stop-loss and CPPI strategy. Furthermore, simulation exercises have mostly been limited to Monte Carlo simulations or backtesting. However, given the underlying normality assumption of the former and the limited number of possible scenarios of the latter, both approaches fail to correctly assess the performance of the strategies. Finally, only few studies have examined the impact of a different choice of the floor value, rebalancing time, and CPPI multiple so far (e.g. Do & Faff (2004)). This manifests itself in practice, since the choice of the implementation variables to be used in commercial applications of PI strategies still remains quite arbitrary.

The present study reassesses the value of such techniques regardless of the underlying market condition. To this end, we consider stock return data from different markets. We use a block bootstrapping procedure in which we repeatedly select a one-year block randomly and test the performance of stop-loss, synthetic put and CPPI. In this way, the performance of each of the strategies can be examined for different market environments. This block bootstrapping technique is used in order to counter the problems that emerge when using a Monte Carlo simulation or backtesting. Indeed, applying bootstrapping implies that no assumptions have to be made about the distribution of market returns, and due to the resampling principle, we can generate a substantial number of scenarios. Moreover, block bootstrapping preserves both the limited autocorrelation and the substantial heteroscedasticity of real world data in the sample. In addition, we extend the traditional performance analyses of dynamic insurance strategies by considering the VaR, ES, and stochastic dominance efficiency of these strategies.

We follow a three-stage procedure in our effort to examine the desirability of dynamic insurance techniques. First, we examine whether insurance techniques outperform a buy-and-hold strategy in terms of downside protection and risk/return trade-off. Second, we demonstrate how choosing a different floor value and CPPI multiple affects the performance of the strategies and test the importance of setting a specific floor value. Thirdly, the impact of a departure from the continuous rebalancing discipline is analyzed by testing whether an equally good protection can be obtained by using a lower rebalancing frequency.

Our results indicate that portfolio insurance is useful for investment purposes, that is, all three strategies provide significantly better risk-return trade-offs and downside protection than a buy-and-hold strategy, albeit at the cost of a reduction in return. The VaR and ES results confirm that the strategies are indeed able to truncate the downward tail of the return distribution. Moreover, relaxing the rebalancing discipline substantially harms the performance of the strategies. Not surprisingly, lowering the floor value increases the potential upward gain and seems to reduce the frequency of negative excess returns, but unfortunately implies larger (in absolute value) negative excess returns and worse VaR and ES results. Indeed, the stochastic dominance results point out that both for the stop-loss and synthetic put strategy a 100% floor value should be preferred to lower floor values. Furthermore, we find that a higher CPPI multiple enhances the upward potential of the CPPI strategies, but harms the protection level in return. Using stochastic dominance rules, we find that for the CPPI strategy a combination of a 95% floor value and multiple 14 outperforms combinations of this floor value with a higher multiple. In other words, the stochastic dominance tests show that for fixed CPPI floor values, the lowest initial equity allocation should be preferred. Finally, our results point out that both for the synthetic put and CPPI strategy daily rebalancing should be preferred to lower rebalancing frequencies, despite the higher transaction costs.

This paper is organized as follows. In the next section, we discuss the three popular portfolio insurance alternatives stop-loss, synthetic put and CPPI. In section 3 the performance measures are described. Section 4 presents the empirical setup and section 5 contains the simulation results. In section 6 we conclude by giving some guidelines for practical implementations.

2. Portfolio insurance strategies

Using *stop-loss portfolio insurance* (SL)[2] the portfolio is fully invested in the risky asset (which in this paper is referred to as 'equity') as long as its value is above the discounted value of the floor, which is the minimum target level the portfolio has to reach at the end of the investment horizon *T*. Once the portfolio drops below the discounted floor, the portfolio switches entirely to the risk-free asset, ensuring that the target is reached at the end of the period. Hence, this strategy is only subject to a single transaction cost, which can be substantial, given the fact that it is computed on the entire portfolio value.

The *synthetic put strategy* (SP) is a dynamic portfolio strategy that replicates the payoff of a protective put (i.e. long positions in a put option and in the underlying risky asset). The strike price *F* is set equal to the desired floor at *T*. The portfolio invests changing proportions in the risky and risk-free assets. The proportion invested in the risky asset follows from an option pricing model (see e.g. Rubinstein & Leland (1981)). When the traditional Black & Scholes valuation is used Benninga (1990) shows that the proportion invested in the risky asset is given by

$$S \cdot N(d) \Big/ \left(S \cdot N(d) + F \cdot e^{-rT} \cdot N\left(\sigma\sqrt{T} - d\right)\right), \tag{1}$$

where S is the initial stock price, r is the risk-free rate, $N(\cdot)$ is the cumulative standard normal distribution function, and $d = \left(\ln(S/F) + \left(r + \frac{1}{2}\sigma^2\right)T\right) \Big/ \sigma\sqrt{T}$. The standard deviation ('volatility') of stock returns is given by σ. Taking into account transaction costs, we follow the Leland (1985) approach and adjust the option volatility as follows:

$$\sigma^2_{adjusted} = \sigma^2 \left[1 + \left(\sqrt{2/\pi}\right) c \big/ \sigma\sqrt{t}\right], \tag{2}$$

where *c* is the relative transaction cost and *t* is the rebalancing interval.

Portfolio protection can also be achieved by implementing a *Constant Proportion Portfolio Insurance* (CPPI) strategy,[3] which was introduced by Black & Jones (1987). This strategy also implies portfolio positions in both a risky and a risk-free asset. The risky proportion is determined by the cushion, which is the difference between the portfolio value and the floor *F*, and by the multiple *m*. The latter is a constant that represents the desired sensitivity to market changes. CPPI requires investing an amount equal to the product of the multiple and the cushion in the risky asset, while investing the remainder in the risk-free asset. We impose short sale and credit constraints in order to follow the common practice in commercial applications. The risky proportion at time t is then given by:

---

[2] See for example Rubinstein (1985).

[3] See also Black & Perold (1992) or Brennan & Schwartz (1988).

$$\max\left\{\min\left[m\left(W_t - F_t\right), W_t\right], 0\right\}, \tag{3}$$

where W_t represents the portfolio value at time t. Similar to the SL strategy, a time subscript is added to the floor F to indicate the discounted value of the target minimum portfolio.

As is clear from the discussion above, the investor has to choose a floor value, a rebalancing frequency and in the case of CPPI a multiple. Very risk averse investors will prefer relatively high floor values, which imply only limited upward potential for the strategy. Conversely, more risk tolerant investors want to benefit more from upward market movements and will rather implement floor values less than 100%. In this paper we focus on strategies with floors ranging between 95% and 100% of initial portfolio value. Rebalancing can be initiated by a predefined minimal (absolute) market move or it can occur after a given time interval. In the present study, we consider rebalancing based on a time interval.[4] Initially, we study daily rebalancing. Afterwards, we investigate the impact of weekly or monthly rebalancing frequencies on the strategies' performance. Finally, for the CPPI strategy a higher multiple implies a riskier strategy and can serve to achieve higher portfolio values. However, risk averse investors will rather apply low multiples in order to obtain a high protection. To facilitate comparisons, we will set the multiple in such a way that the initial exposure to the risky asset is similar across strategies.

3. Performance measurement

Portfolio theory assumes that investors select portfolios that are optimal to them. Optimality is often understood as maximizing expected utility of wealth. As such, choosing between portfolios amounts to choosing between return distributions. Of course, in the absence of knowing the exact return distributions, such choices are difficult to perform. In practice, investors therefore focus on specific moments or other statistics of these distributions. Examples are average returns in excess of a risk-free rate of return ('excess returns') and the standard deviation of return ('volatility'). The Sharpe ratio combines both by taking the ratio between the average excess return and volatility. Under some conditions, it can be shown that portfolio selection based on comparisons of these statistics is consistent with expected utility (see e.g. Elton *et al.* (2003)).

Unfortunately, for PI techniques, where the focus is on value protection and upward potential, these statistics are not sufficient for adequate selection. High volatility can be due to positive return outliers, which would attract rather than shy away investors. Informal performance measures such as the occurrence of negative excess returns are therefore also

[4] Cesari & Cremonini (2003) find that a market move rebalancing discipline yields the same performance ranking as rebalancing according to a time interval.

contemplated. Recently, VaR, an asymmetric risk measure, has been put forward as an alternative to the symmetric volatility measure. VaR denotes the maximum loss at a certain confidence level. A major drawback of the VaR computation concerns the fact that no indication is given of the magnitude of the losses for the extreme situation in which the limit value is exceeded. The ES provides an answer to this disadvantage, because it expresses the average loss below this limit level. VaR and ES can be computed for different confidence levels. Of course, a higher confidence level will entail more negative VaR and ES values. Hence, choosing a higher confidence level indicates a higher degree of risk aversion. The skewness of the return distribution is another performance measure. Generally, investors prefer right-skewed distributions. Hence, a larger skewness makes a protection strategy more appealing; see e.g. Harvey & Siddique (2000).

The problem with all these performance measures is that they are difficult to associate with expected utility. Moreover, comparing different measures across strategies may lead to contradictory results. Stochastic dominance (SD) rules provide a framework that is explicitly based on expected utility. Moreover, they are based on the entire return distribution, rather than on some arbitrarily chosen statistics. Because of the asymmetric nature of PI techniques the SD framework seems particularly attractive. Yet, up to now it has not been popular because the theoretical distributions are unknown and all inference is plagued by severe sampling error. Short of a statistical framework to deal with this sampling error, researchers have relied on more informal measures like those introduced above.[5] Fortunately, recent work referred to below, has developed tests that allow testing for SD even when portfolio returns are correlated and exhibit time dependency. In the remainder of this section, we first discuss the intuition behind SD rules and then summarize the testing procedure.

3.1. Stochastic Dominance Rules

The goal of SD rules is to partition the investment opportunity set into an efficient and an inefficient set.[6] The latter contains all investment prospects no investor satisfying certain characteristics would ever be interested in. The efficient set contains all investments that are not stochastically dominated by any other investment. This is the set from which investors will select their optimal portfolio. Depending on the number of required characteristics, different SD orders can be formulated. All versions assume investors have von Neumann-Morgenstern utility functions and that they maximize expected utility. In addition, first-order SD (FSD) requires that investors prefer higher returns to lower returns, which implies a utility

[5] In the context of PI, only Trennepohl *et al.* (1988) use the SD approach to compare their strategies, albeit without accounting for sampling error.

[6] This section is based on Levy (2006).

function with a non-negative first derivative. Second-order SD (SSD) also takes risk aversion into account. It posits a negative second derivative (i.e. diminishing marginal utility) of the investor's utility function, which is sufficient for risk aversion. Whitmore (1970) introduced third-order SD (TSD) by adding the condition that utility functions have a non-negative third derivative. This assumes the empirically attractive feature of decreasing absolute risk aversion. It is clear that higher order efficient sets are subsets of the lower order efficient sets.

Testing for SD can be based on comparing (functions of) the cumulative return distributions $F_k\left(x\right)$ of the different investment prospects *k*. A practical characterization of any SD order is due to Davidson & Duclos (2000). Prospect X_1 stochastically dominates prospect X_2 at order *s* if $D_1^{(s)}\left(x\right) \le D_2^{(s)}\left(x\right)$ for all *x* and with strict inequality for some *x*, where

$$D_k^{(s)}\left(x\right) = \int_{-\infty}^{x} D_k^{(s-1)}\left(t\right) dt \text{ for } k=1, 2, \text{ and} \tag{4}$$

$$D_k^{(1)}\left(x\right) = F_k\left(x\right). \tag{5}$$

3.2. Testing for Stochastic Dominance

Of course, the true cumulative distributions are not known in practice and SD tests have to rely on the empirical distribution function (EDF), which is subject to sampling error. Linton *et al.* (2005) show how consistent critical values for testing SD can be estimated from the data with only a limited number of assumptions. Most importantly, their test allows for dependency in the data. Consider *K* prospects $X_1, \ldots, X_K$. Let *N* denote the full sample size of return observations of X_k for $k=1,\ldots,K$, i.e. $\left\{X_{ki} : i = 1,\ldots,N\right\}$. Let *b* denote the size of the subsample, and *s* the SD order. Under the null hypothesis that a particular prospect *k* *s*-th order stochastically dominates all other outcomes, a three-step procedure is suggested. First of all, the test statistic $T_N^{(s)}(k)$ for the full sample is computed:

$$T_N^{(s)}(k) = \max_{l:k \ne l} \sup_{x \in \chi} \sqrt{N}\left[\hat{D}_k^{(s)}(x) - \hat{D}_l^{(s)}(x)\right] \text{ for } s \ge 1, \tag{6}$$

where

$$\hat{D}_k^{(1)}(x) = \hat{F}_{kN}(x) = \frac{1}{N}\sum_{i=1}^{N} 1\left(X_{ki} \le x\right), \text{ and} \tag{7}$$

$$\hat{D}_k^{(s)}(x) = \int_{-\infty}^{x} \hat{D}_k^{(s-1)}(t) dt \text{ for } s \ge 2. \tag{8}$$

Here χ denotes the union of supports of all distributions $\hat{F}_{jN}$ for j=1,…,K and $1(\cdot)$ denotes the indicator function. To find out whether prospect k s-th order stochastically dominates all other prospects, Linton *et al.* (2005) suggest to take the maximum over all l with $k \neq l$ in (6).

Secondly, subsamples of size b are used to recompute test statistic (6) for each of the $N-b+1$ different subsamples $\{W_i,...,W_{i+b-1}\}$, where $W_i = \{X_{ki} : k = 1,...,K\}$ and $i = 1,...,N-b+1$. However, as suggested by Kläver (2005), this procedure implies that both the observations at the beginning and the end of the return series are considered less than the observations in the middle of the return series. Therefore, it is recommended to apply a circular block method by recomputing test statistic (6) for $N-b+1$ different subsamples $\{W_i,...,W_{i+b-1}\}$, where $i = 1,...,N-b+1$ and additionally for the subsamples $\{W_i,...,W_N, W_1,...,W_{i+b-N-1}\}$, where $i = N-b+2,...,N.$ We denote this recomputed statistic by $t_{N,b,i}(k)$. Note that this test statistic is multiplied by the square root of the subsample size b. The underlying rationale is that since each of these subsamples is also a sample of the true sampling distribution, the distribution of the subsample test statistics can yield an approximation of the sampling distribution of the full sample test statistic.[7] Following Kläver (2005), we consider a subsample size $b(N)=10\sqrt{N}$.

Finally, the p value, $\hat{p}_N^{(s)}(k)$, of the test is computed in order to decide whether or not to reject the null hypothesis of dominance:

$$\hat{p}_N^{(s)}(k) = \frac{1}{N}\sum_{i=1}^{N-b+1} 1\left(t_{N,b,i}^{(s)}(k) > T_N^{(s)}(k)\right) \text{ for } s = 1,2,3\,. \tag{9}$$

In section 5 we use the SD tests for several sets of strategies. We proceed in the following way. For each strategy k, we test whether it dominates all other strategies in the set by FSD, SSD, or TSD. This gives us three p values per strategy (one for FSD, SSD, and TSD). We next determine for each strategy the highest SD order for which the null hypothesis of stochastic dominance cannot be rejected.[8] We call a strategy FSD when we cannot reject the null that it first-order dominates all other outcomes. Note that first-order dominance also implies second- and third-order dominance. Likewise, a strategy is called SSD when the null that the strategy dominates all other outcomes by FSD is rejected, but not that it dominates the other outcomes by SSD. Finally, if a strategy is TSD we can reject that the strategy dominates all other prospects by FSD and SSD, but not that it dominates the other strategies in the set by TSD.

---

[7] Under the assumption that $b/N \to 0$ and $b \to \infty$ as $N \to \infty$.

[8] We use the 10% significance level to decide whether or not we can reject the null hypothesis.

4. Simulation setup

In order to analyze the performance of the strategies, we conduct a bootstrap simulation on an elaborate dataset containing equity return data from the US, the UK, Japan, Australia, and Canada for the past 30 years. Although these markets are not independent from each other, we considerably increase the number of possible return scenarios on which to test the portfolio insurance strategies. By doing so it is hoped that we obtain a better approximation of the population distribution. Daily simple returns in local currency are retrieved from Datastream for the time period 1 January 1973 through 31 March 2005, thereby deleting all non-trading days. The respective interest rates from the euro-interbank market are also downloaded. As some risk-free rate series are not available from 1973 onwards, we only use stock return data for which we also have risk-free rates (see table 1). In the bootstrap simulation we assume a year contains 252 days, whereas a week and month contain five and 21 days respectively.

Summary statistics are presented in table 1. The results confirm the empirical regularity that financial return series exhibit fat tails and negative skewness. The Ljung-Box test detects significant serial correlation in almost all countries. Both Engle's ARCH test and the Ljung-box test on squared returns find significant heteroscedasticity in all series.

INSERT TABLE 1 ABOUT HERE

Given that our data are not normally distributed, we do not opt for a traditional Monte Carlo simulation. Alternatively, bootstrapping offers a way of generating simulated return series without making any assumption regarding the return distribution. It retains both the skewness and the fat tails from the original data sets. Even so, a simple bootstrapping procedure would destroy the dependency effects (i.e. autocorrelation and heteroscedasticity) in our data. Therefore, we resort to block bootstrapping to counter this problem (see e.g. Sanfilippo (2003)).

The block bootstrap procedure consists of the following steps. First, we randomly draw a market (UK, US, Japan, Australia, or Canada) with replacement. Because we only have 30 years of historical data, we attempt to better proxy the population of possible stock market scenarios by considering different markets. Next, a random date (i.e. a start date) is drawn with replacement. Starting from this date we analyze the one-year performance of each portfolio insurance strategy for the drawn market, that is, the 252 days following the starting date are used to evaluate the different portfolio insurance strategies. This procedure is repeated for 10,000 'years'.

To implement the SP strategies the volatility of the underlying asset and a risk-free rate are needed. Since we do not know the actual volatility in advance, we are confronted with a forecasting problem. We use the standard deviation of the 252 continuously compounded daily stock returns prior to the randomly drawn date as an estimate of future volatility. We adjust the estimate using equation (2) to take into account the effect of transaction costs. In addition, the continuously compounded one-year risk-free rate on the starting date is applied to compute (1).

For the CPPI strategy, the floor grows daily at the one-year risk-free rate on the starting date. Similarly, we use this one-year rate for discounting the SL floor. During the investment horizon, we use the short-term (S/T) (i.e. 2 days notice) risk-free rate to gross up the risk-free part of the portfolios. Since the one-year risk-free rate is generally somewhat higher than the S/T risk-free rate, this procedure implies that sometimes the floor will be missed by the end of the investment horizon.

In order to address the different issues in a systematic way, we will first compare some basic PI strategies to the buy-and-hold strategy. To do this, we will adopt a daily rebalancing frequency – the best we can do with our data to approximate the theoretically assumed continuous rebalancing. To address the effect of frequent rebalancing, a proportionate transaction cost of 0.1% is taken into account.[9] In a second step, we investigate whether changing the guaranteed floor significantly changes the performance of the PI strategies. For the CPPI variant we also consider different levels of the multiple. Thirdly, we study the effect of lowering the rebalancing frequency.

5. Performance measurement results

5.1. Portfolio insurance versus buy-and-hold

To answer the question whether PI strategies are sensible alternatives to a buy-and-hold strategy for at least some investors, we study the following implementations: a SL strategy with a 100% floor value, a SP strategy with a 100% floor value, and a CPPI strategy with a 95% floor value and multiple 14. For the latter we do not opt for a 100% floor value, since the strategy would only have a very small equity exposure with such a high floor value. As an alternative, we consider a CPPI strategy that has an equal average initial equity allocation as the SP strategy, which was found to be 70%. Since this initial equity allocation can be obtained by different combinations of floor values and multiples, we first examine which combination should be preferred from the following (arbitrary) set: floor 80% and multiple 3.5, floor 85% and multiple 4.67, floor 90% and multiple 7, floor 95% and multiple 14, or

[9] As such dynamic strategies are generally implemented using futures transactions, a 0.1% transaction cost is reasonable. Note that we also analyzed 0% and 0.3% transaction costs with qualitatively similar stochastic dominance results.

floor 97.5% and multiple 28. We find that for the combination floor value 95% and multiple 14 the null hypothesis of third-order dominance relative to the other CPPI combinations cannot be rejected.[10] Therefore, we consider this combination as an alternative to the SL and SP strategies with a 100% floor value.

Table 2 presents the simulation results, whereas Figure 1 plots the pay-off functions of the three PI strategies. The floor values are indicated by dotted horizontal lines and a full 45° line is drawn to facilitate comparisons. The payoff functions confirm that PI strategies provide payoffs similar to a protective put strategy. Note that PI is no free lunch: portfolio insurers pay for the downside protection in terms of a reduction in upside capture. As typical for "buy high – sell low" strategies, the payoff function of the synthetic put and CPPI strategy (with multiple > 1) is convex (see Leland (1980)). We focus on returns in excess of the one-year rate, since we want to asses the excess performance of the strategies relative to a risk-free investment. Tests are performed to check whether there is a significant difference between the performance statistics of the PI strategies and the buy-and-hold strategy. A t-test reveals that none of the strategies is able to outperform a buy-and-hold strategy in terms of average excess return. This foregone excess return illustrates the implicit cost inherent in PI investments. Nonetheless, a Levene test points out that PI strategies deliver a significantly lower risk than a buy-and-hold strategy.

INSERT TABLE 2 ABOUT HERE

INSERT FIGURE 1 ABOUT HERE

We apply the Jobson & Korkie (1981) test to find significant differences in the Sharpe ratios of the PI and buy-and-hold strategies. We find that the Sharpe ratio of the SL and CPPI strategy is significantly lower than the Sharpe ratio of the buy-and-hold strategy. Conversely, the SP strategy reports a (statistically) significantly higher Sharpe ratio than the buy-and-hold strategy. It should be noted that in a PI context the Sharpe ratio is not necessarily an adequate performance measure, since portfolio insurers do not only care about the mean and variance of returns. As emphasized by Leland (1999), PI strategies will typically be undervalued by CAPM performance measures. Likewise, Booth *et al.* (1985) stress that the mean-variance framework is relevant only for investors with quadratic utility.

---

[10] Detailed results are available upon request. Here, as well as in the remaining part of this paper we only report the SD test results for a subsample size of 1,000 (=$10N^{0.5}$). Since other subsample sizes generate similar p values, the robustness of our p value results for different subsample sizes should not be an issue.

Furthermore, table 2 indicates that PI strategies have a rather high frequency of negative excess returns. For the CPPI this is less surprising, as the floor is only 95% of initial portfolio value. But also for the SL and SP the frequencies are rather high, see also figure 1. The results from a t-test on frequencies point out that all PI strategies have a significantly higher frequency of negative excess returns than the buy-and-hold strategy. However, this result should be interpreted with caution, since t-test results confirm that the average negative excess return of PI strategies is significantly better than the one generated by a buy-and-hold strategy. The lower risk of the PI strategies is corroborated by the VaR and ES measures. To test whether PI strategies yield significantly better VaR estimates than the buy-and-hold VaR, we conduct an unconditional coverage test (Christoffersen (2003)). Under the null, the buy-and-hold VaR equals the VaR of a particular PI strategy. A 95% confidence level is used, which implies that under the null, we expect 5% of the observations of a PI strategy to exceed the buy-and-hold VaR.[11] Therefore, if the number of exceedences of the buy-and-hold VaR is smaller than 5% for a given PI strategy, we can reject the null of equal VaRs. For all strategies, the null is rejected at the 1% significance level, implying that the PI strategies have statistically significantly better VaR levels. For the ES, we test whether the PI strategies and the buy-and-hold strategy have a significantly different ES by using a bootstrap approach (again we use a 95% confidence level). First, we randomly draw 5,000 returns pairwise from the full sample of 10,000 empirical PI returns and 10,000 buy-and-hold returns. From these bootstrapped return observations, we calculate the ES for both strategies. This procedure is repeated 10,000 times to obtain two ES distributions. We then use a paired-samples t-test to check whether the two ES distributions are statistically different. We find that all t-statistics are extremely significant. Again, this points to less risky PI strategies. This result is robust to alternative sizes of the sampled return observations. Note that the CPPI strategy shows considerably better VaR and ES levels than the other strategies in the set.

Again we resort to bootstrapping in order to test differences between the skewness of PI strategies and the skewness of the unprotected portfolio. Given the null that the skewness of the buy-and-hold strategy exceeds or equals the skewness of a PI strategy, we first construct symmetric return distributions, using on the one hand the original 10,000 return observations and on the other hand 10,000 adjusted return observations, using the transformation $-(r_i - \bar{r}) + \bar{r}$, where r_i is the *i*th return of a particular strategy for $i = 1...10{,}000$ and $\bar{r}$ is the average return. In this way, all odd sample moments (e.g. mean and skewness) are set to zero, while all even sample moments (e.g. variance and kurtosis) are maintained. This ensures that both series have identical skewness coefficients, i.e. zero. Next,

[11] The results for a 90 and 99% confidence level are available on request.

we simulate the distribution of the skewness difference by randomly drawing 10,000 returns pairwise from the 20,000 returns of a PI strategy and the buy-and-hold portfolio and computing the sample skewness difference. Repeating this 10,000 times, we obtain an approximation of its distribution, which is then used to compute the p value of obtaining an even more extreme skewness difference than the empirical skewness difference under the null of a zero difference. We find that all PI strategies have statistically significantly higher skewness coefficients than the buy-and-hold strategy.

Summarizing, we find that the three PI strategies entail less risk than the buy-and-hold strategy, regardless of how we measure risk – the only exception being the proportion of negative returns. In addition, they also have higher skewness. Both attractive features seem to come at the price of lower average returns. Therefore, it is not clear whether risk averse investors would prefer the PI strategies to the buy-and-hold strategy. As the Sharpe ratio is not necessarily an adequate statistic to compare asymmetric distributions, we investigate whether any stochastic dominance relations exist among the four strategies considered in table 2. We find that the null hypothesis of stochastic dominance is rejected for all strategies and all stochastic dominance orders. That is, none of the strategies in table 2 is able to dominate the other strategies in the set. In economic terms, this implies that some investors prefer one of the PI strategies, whereas others will favor the buy-and-hold strategy. This finding of a lack of dominance between insured and uninsured portfolios corroborates the theoretical results of Benninga & Blume (1985) and Brooks & Levy (1993). In sum, we find that even though PI strategies yield lower returns than a buy-and-hold strategy, the corresponding lower risk (both in terms of symmetric and downside risk) compensates this reduced return, making these strategies valuable alternatives to buy-and-hold investments.

5.2. Changing the floor value

In this section the impact of choosing a floor value less than 100% is examined. The choice regarding the floor value to be protected entails important consequences for each of the protection strategies. In particular, a lower floor value implies that the SL strategy will be relatively longer invested in stocks before switching to a risk-free investment. The SP strategy on the other hand, will initially allocate more funds to the risky asset. Furthermore, the CPPI strategy will behave riskier because of the relatively larger equity exposure in case of lower floor values. We do not attempt to find an optimal floor value, which would be impossible without assuming a particular distribution for the risky asset returns (e.g. Brooks & Levy (1993)). Rather we wish to indicate the impact this particular choice may have on PI performance. The results for the different strategies are presented in table 3, where a daily rebalancing frequency is maintained; results for weekly and monthly frequencies are qualitatively similar and are therefore not reported. The results in each panel refer to the same

PI strategy, where only the floor value is different. The strategies with the lower floors are compared to the benchmark strategies from table 2, whose results are duplicated to facilitate comparisons.

INSERT TABLE 3 ABOUT HERE

In table 3 alternative floor values of 95% and 97.5% are considered for the SL strategy. The lower floor values yield better average excess returns and Sharpe ratios but with a corresponding higher risk, both in terms of standard deviation and downside risk. Moreover, in case of lower floor values the return distributions are characterized by a lower skewness. However, we can reject the null hypothesis that these lower floor strategies stochastically dominate both other SL strategies at all orders. We cannot reject the hypothesis that the 100% floor strategy third order dominates the lower floor strategies. In economic terms, this implies that investors who have a preference for positive skewness would prefer a SL strategy with a 100% floor value to a strategy with a lower floor value. The results for the SP strategy are similar to those for the SL strategy. The lower floor strategies have higher return and risk compared to the 100% floor, but we cannot reject that the latter stochastically dominates the former at the third order. The lower floor strategies on the other hand do not dominate. Again, the effect of the higher positive skewness seems to be an important characteristic.

For the CPPI strategy we change both the floor and the multiple. As in the previous section, we choose combinations that result in the same initial equity allocation as the SP strategy. For each initial equity allocation we first compare several combinations and retain the combination for which we cannot reject the hypothesis that it dominates the other combinations. For instance, to obtain a 70% initial equity allocation the following floor-multiple combinations are considered: (80%,3.5), (85%,4.67), (90%,7), (95%,14), or (97.5%, 28). We find that for the combination (95%,14) the null hypothesis of third-order dominance cannot be rejected. Therefore we consider this combination as an alternative to the SL and SP strategies with a 100% floor value. Similarly, the combinations (95%,15) and (95%,16) are retained for the initial equity allocations of 75% and 80% respectively.[12] The three CPPI strategies that are comparable to the SL and SP strategies with floor values 95%, 97.5% and 100%, are subsequently compared in last panel of table 3. Again the 70% initial equity allocation strategy, combination (95%,14), is preferred. Both other combinations do not stochastically dominate, whereas third order dominance cannot be rejected for the (95%,14) strategy.

[12] Detailed results are available upon request.

5.3. Changing the rebalancing frequency

The results above pointed out which floor value should be preferred for a particular rebalancing frequency. However, this still leaves us with the arbitrary choice of the rebalancing frequency. This choice will be the result of a trade-off between seeking high protection (e.g. daily rebalancing) and reducing transaction costs (e.g. weekly or monthly rebalancing). In order to study this trade-off, the preferred strategies in table 3 are replicated at both the weekly and monthly rebalancing frequency. In table 4 their results are compared to those of the daily rebalancing variant.

We find that both the SP and the CPPI strategy suffer from a significant decrease in average excess return in case of weekly or monthly rebalancing instead of daily rebalancing. In contrast, for the SL strategy we find that a lower rebalancing frequency yields a significantly higher average excess return compared to the daily rebalancing counterpart, albeit with a higher risk. The longer rebalancing interval may prevent too early a switch to a risk-free investment following a temporary stock price fall below the discounted floor, thereby explaining the higher average returns. Since the portfolios are managed less strictly, the average negative excess return as well as the VaR and ES becomes significantly worse for all strategies. Furthermore, the skewness decreases, indicating a less right-skewed return distribution.

INSERT TABLE 4 ABOUT HERE

The SD tests in table 4 point out that for the SL strategy the null of dominance is rejected for each rebalancing frequency. This absence of any dominance relation implies that the choice regarding the optimal rebalancing frequency will depend on the type of investor. Even though daily rebalancing will provide the best protection of the floor value, a monthly rebalancing frequency seems to compensate a lower protection with a higher average excess return for the SL strategy. Even though the SP strategy suffers from high transaction costs in case of daily rebalancing, the SD results in table 4 show that we cannot reject the null that the SP strategy with a daily rebalancing frequency third-order dominates the SP strategy with a weekly and monthly rebalancing frequency. In other words, the benefits from daily rebalancing (i.e. better protection and average excess return) outweigh the advantage of lower transaction costs provided by lower rebalancing frequencies. This implies that investors who have a preference for positive skewness would choose to implement a daily rebalancing discipline. Similarly, we find that daily rebalancing is the preferred rebalancing frequency in case of the CPPI strategy. Based on SD tests we cannot reject the null that a daily rebalancing frequency dominates lower rebalancing frequencies by second-order (and thus also third-order) degree. In economic terms, risk-averse investors with a preference for positive skewness will prefer a daily rebalancing frequency to lower frequencies. Note that since the SP and CPPI strategy require active rebalancing of the portfolio mix, whereas the SL strategy requires maximum one transaction, transaction costs incurred by the former strategies will exceed those incurred by the latter.

6. Concluding remarks

The present study provides an answer to the continuing controversy over PI strategies. Our results are threefold. First, we have demonstrated that these strategies outperform a buy-and-hold strategy in terms of downside protection and risk/return trade-off, but provide lower excess returns in return. The SD results point out that the buy-and-hold strategy does not dominate the PI strategies. This implies that even though PI strategies yield lower returns than a buy-and-hold strategy, the corresponding lower risk compensates this reduced return, making these strategies valuable alternatives to buy-and-hold investments. Second, the highest floor values deliver the best downside protection, i.e. a superior average negative excess return, VaR, ES, and skewness, albeit at the cost of a lower excess return. Based on SD tests, we cannot reject that a floor value of 100% dominates the lower floor values by third-order stochastic dominance. Third, we find that even though daily rebalancing will provide the best protection of the floor value for the SL strategy, a monthly rebalancing frequency seems to compensate a lower protection with a higher average excess return. However, for the SP and CPPI strategy a daily rebalancing frequency should be preferred to lower rebalancing frequencies. Apparently, the benefits from daily rebalancing (i.e. better protection and average excess return) outweigh the advantage of lower transaction costs provided by lower rebalancing frequencies for these strategies. Furthermore, our results point out that based on the more general stochastic dominance analysis, a simple analysis of Sharpe ratios appears to be too short-sighted an approach for selecting the optimal strategy.

It should be noted that the performance of SP strategy strongly depends on the accuracy of the volatility estimate in the Black and Scholes option pricing formula. We considered the volatility of the previous year to be an appropriate proxy for the ex post volatility. The performance of SP strategy may be enhanced by using better volatility predictors. In the present paper, we have not really focused on this volatility issue. Moreover, the strategies can be adjusted in order to take into account a lock-in threshold, so that intermediate gains are locked-in by dynamically adjusting the floor. Another research opportunity is to consider the performance of the strategies in case rebalancing is triggered following a market move instead of a time interval.

TABLES

Table 1 – Summary statistics

Country	*Average return (annualized)*	*Standard deviation (σ) (annualized)*	*Skewness*	*Kurtosis*	*p value autocorrelation (Ljung-Box test)*	*p value autocorrelation (Ljung-Box test squared returns)*	*p value heteroscedasticity (Engle's ARCH test)*
UK	19.69%	16.28%	-0.01***	11.28***	0.00	0.00	0.00
US	15.03%	15.64%	-0.94***	24.45***	0.00	0.00	0.00
Japan	7.25%	17.49%	-0.14***	13.04***	0.00	0.00	0.00
Australia	12.07%	12.74%	-0.31***	9.74***	0.85	0.00	0.00
Canada	13.31%	13.05%	-0.64***	16.89***	0.00	0.00	0.00

Summary statistics for the five markets in our dataset. The variable UK contains 7714 observations, US contains 7655 observations, Japan contains 6647 observations, Australia contains 2050 observations, and Canada contains 7674 observations. The daily average return is reported on an annual basis, using $\overline{r_{annual}} = (1 + \overline{r_{daily}})^{252} - 1$. Daily standard deviations are transformed into annual standard deviations, using $\sigma_{annual} = \sigma_{daily} \cdot \sqrt{252}$. *, ** and *** denote significance at the 10, 5 and 1% level respectively.

Table 2 – Bootstrap simulation results

	Stop-loss [D, F100%]	*Synthetic put* [D, F100%]	*CPPI* [D, F95%, m14]	*Buy-and-hold*
Average excess return	3.52***	4.62***	3.87***	5.62
Standard deviation	15.12***	14.38***	13.14***	17.98
Sharpe	0.23***	0.32**	0.29***	0.31
% < 0	57.55***	43.13***	58.69***	35.95
Average negative excess return	-7.14***	-8.18***	-4.81***	-12.86
VaR 5%	-13.16***	-13.04***	-6.33***	-24.44
ES 5%	-17.12***	-16.48***	-7.50***	-31.65
Skewness	1.20***	0.93***	1.73***	0.14
Stochastic dominance efficiency	No SD°°°	No SD°°°	No SD°°°	No SD°°°

Performance results of three basic portfolio insurance strategies and a buy-and-hold strategy obtained from a block bootstrap simulation. In the simulation we repeatedly draw a one-year block from the return index series (covering the UK, US, Japan, Australia, and Canada). This procedure is repeated for 10,000 years. Returns are in excess of the one-year rate corresponding to the drawn market and date. Transaction costs of 0.1% are taken into account. The VaR is obtained by first sorting the portfolio returns into ascending order and then looking at the return at the 5% level. ES is the average of VaR exceedences. *D*, *F*, and *m* denote daily rebalancing, floor value and multiple respectively. *, **, and *** denote a significant difference between the portfolio insurance and buy-and-hold strategy at the 10, 5, and 1% level respectively. For the stochastic dominance (SD) tests, we consider a set of strategies containing the present four strategies. Under the null, a particular strategy dominates all other strategies in the set. If the null is rejected, no dominance relation is present. The table displays the highest SD order for which the null hypothesis of dominance cannot be rejected. First-order SD (FSD) implies that we cannot reject the null that the strategy first-order dominates all other strategies. Note that FSD also implies second- and third-order dominance. Similarly, second-order SD (SSD) means that we can reject the null that the strategy dominates all other outcomes by FSD, but not that it dominates the other outcomes by SSD (and thus also TSD). Finally, if a strategy is TSD we can reject that the strategy dominates all other prospects by FSD and SSD, but not that it dominates the other strategies in the set by TSD. No SD means that we reject the null of SD by any order. °, °°, and °°° denote the significance level of the hypothesis test for SD identification, i.e. 10, 5, and 1% level.

Table 3 – Impact of changing the floor value (daily rebalancing, 0.1% transaction costs)

Portfolio insurance strategy	*Stop-Loss*			*Synthetic Put*			*CPPI*		
Floor value	*95%*	*97.5%*	*100%*	*95%*	*97.5%*	*100%*	*95%*	*95%*	*95%*
Multiple	-	-	-	-	-	-	*16*	*15*	*14*
Initial equity allocation	*100%*	*100%*	*100%*	*80%*	*75%*	*70%*	*80%*	*75%*	*70%*
Average excess return	4.26***	4.00***	3.52	4.95**	4.81	4.62	3.97	3.92	3.87
Standard deviation	16.98***	16.25***	15.12	15.68***	15.07***	14.38	13.51***	13.35***	13.14
Sharpe ratio	0.25***	0.25***	0.23	0.32***	0.32***	0.32	0.29	0.29*	0.29
% < 0	45.96***	50.49***	57.55	40.50***	41.73***	43.13	58.70	58.87	58.69
Average negative excess return	-10.79***	-9.18***	-7.14	-9.91***	-9.07***	-8.18	-5.09***	-4.95***	-4.81
VaR 5%	-16.67***	-15.04***	-13.16	-16.33***	-14.60***	-13.04	-6.49***	-6.43***	-6.33
ES 5%	-21.13***	-19.35***	-17.12	-19.94***	-18.17***	-16.48	-7.68***	-7.59***	-7.50
Skewness	0.75	0.94	1.20	0.68	0.80	0.93	1.67	1.70	1.73
SD efficiency	No SD°°°	No SD°°°	TSD	No SD°°°	No SD°°°	TSD	No SD°°°	No SD°°°	TSD

Performance results of the strategies obtained from a block bootstrap simulation. Results are shown for a daily rebalancing frequency. In the simulation we repeatedly draw a one-year block from the return index series (covering the UK, US, Japan, Australia, and Canada). This procedure is repeated for 10,000 years. Returns are in excess of the one-year rate corresponding to the drawn market and date. Transaction costs of 0.1% are assumed. The VaR is obtained by first sorting the portfolio returns into ascending order and then looking at the return at the 5% level. ES is the average of VaR exceedences. For the SD tests, we consider the set of strategies within each panel. Under the null, a particular strategy dominates all other strategies in the panel. If the null is rejected, no dominance relation is present. The table displays the highest SD order for which the null hypothesis of dominance cannot be rejected. FSD implies that we cannot reject the null that the strategy first-order dominates all other strategies. Similarly, SSD means that we can reject the null that the strategy dominates all other outcomes by FSD, but not that it dominates the other outcomes by SSD. Finally, if a strategy is TSD we can reject that the strategy dominates all other prospects by FSD and SSD, but not that it dominates the other strategies in the set by TSD. No SD means that we reject the null of SD by any order. °, °°, and °°° denote the significance level of the hypothesis test for SD identification, i.e. 10, 5, and 1% level.

Table 4 – Impact of changing the rebalancing frequency (0.1% transaction costs)

Portfolio insurance strategy	*Stop-Loss*			*Synthetic put*			*CPPI*		
Rebalancing discipline	*Daily*	*Weekly*	*Monthly*	*Daily*	*Weekly*	*Monthly*	*Daily*	*Weekly*	*Monthly*
Floor value	*100%*	*100%*	*100%*	*100%*	*100%*	*100%*	*95%*	*95%*	*95%*
Multiple	-	-	-	-	-	-	*14*	*14*	*14*
Initial equity allocation	*100%*	*100%*	*100%*	*70%*	*70%*	*70%*	*70%*	*70%*	*70%*
Average excess return	3.52	3.76	4.25***	4.62	4.55	4.49	3.87	3.42***	3.33***
Standard deviation	15.12	15.68***	16.43***	14.38	14.32***	14.29***	13.14	12.91***	13.47***
Sharpe	0.23	0.24**	0.26***	0.32	0.32***	0.31***	0.29	0.26***	0.25***
% < 0	57.55	54.58***	49.75***	43.13	42.87	42.51	58.69	60.03***	58.24
Average negative excess return	-7.14	-7.92***	-9.11***	-8.18	-8.25	-8.33*	-4.81	-4.85*	-5.67***
VaR 5%	-13.16	-14.40***	-16.22***	-13.04	-13.28***	-13.86***	-6.33	-6.66***	-8.77***
ES 5%	-17.12	-18.93***	-21.07***	-16.48	-16.68***	-17.81***	-7.50	-8.51***	-12.91***
Skewness	1.20	1.06	0.87	0.93	0.91	0.85	1.73	1.78***	1.54
Stochastic dominance efficiency	No SD°°°	No SD°°°	No SD°°°	TSD	No SD°°°	No SD°°°	SSD	No SD°°°	No SD°°°

Performance results of the strategies obtained from a block bootstrap simulation. Results are shown for a daily rebalancing frequency. In the simulation we repeatedly draw a one-year block from the return index series (covering the UK, US, Japan, Australia, and Canada). This procedure is repeated for 10,000 years. Returns are in excess of the one-year rate corresponding to the drawn market and date. Transaction costs of 0.1% are assumed. The VaR is obtained by first sorting the portfolio returns into ascending order and then looking at the return at the 5% level. ES is the average of VaR exceedences. For the SD tests, we consider the set of strategies within each panel. Under the null, a particular strategy dominates all other strategies in the panel. If the null is rejected, no dominance relation is present. The table displays the highest SD order for which the null hypothesis of dominance cannot be rejected. FSD implies that we cannot reject the null that the strategy first-order dominates all other strategies. Similarly, SSD means that we can reject the null that the strategy dominates all other outcomes by FSD, but not that it dominates the other outcomes by SSD. Finally, if a strategy is TSD we can reject that the strategy dominates all other prospects by FSD and SSD, but not that it dominates the other strategies in the set by TSD. No SD means that we reject the null of SD by any order. °, °°, and °°° denote the significance level of the hypothesis test for SD identification, i.e. 10, 5, and 1% level.

Figure 1 Payoff functions of the three basic portfolio insurance strategies (daily rebalancing, 0.1% transaction costs)

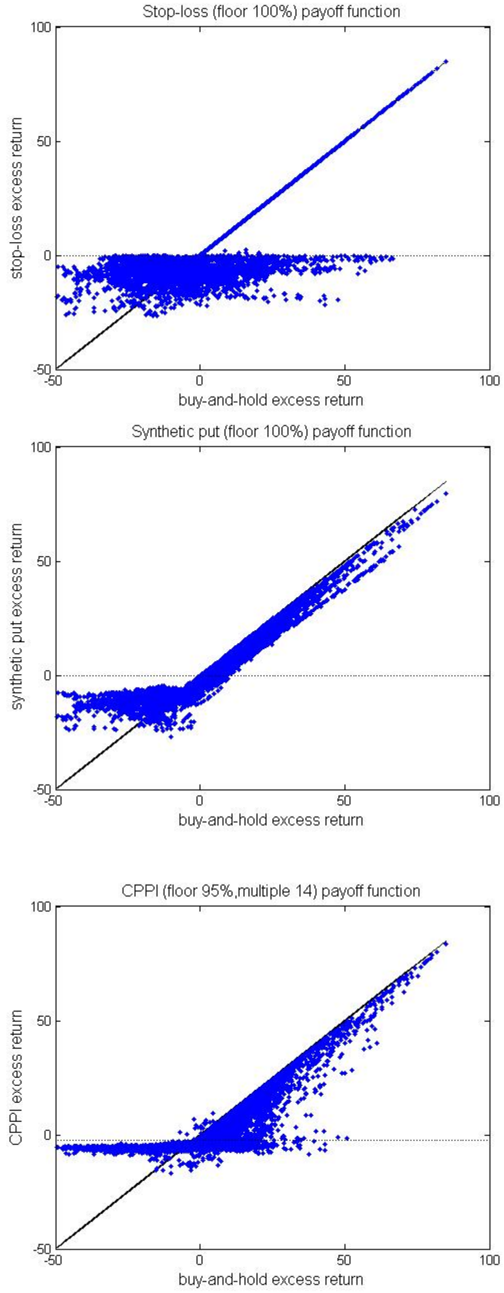

References

Acerbi, C., Tasche, D., 2002. On the coherence of expected shortfall. Journal of Banking and Finance 26, 1487-1503.

Barrett, G.F., Donald, S.G., 2003. Consistent tests for stochastic dominance. Econometrica 71, 71-104.

Benninga, S., 1990. Comparing portfolio insurance strategies. Finanzmarkt und Portfolio Management 4, 20-30.

Benninga, S., Blume, M.E., 1985. On the optimality of portfolio insurance. Journal of Finance 40, 1341-1352.

Bird, R., Cunningham, R., Dennis, D., Tippett, M., 1990. Portfolio insurance: a simulation under different market conditions. Insurance: Mathematics and Economics 9, 1-19.

Black, F., Jones, R., 1987. Simplifying portfolio insurance. Journal of Portfolio Management 14, 48-51.

Black, F., Perold, A.F., 1992. Theory of constant proportion portfolio insurance. Journal of Economic Dynamics and Control 16, 403-426.

Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. Journal of Political Economy 81, 637-654.

Booth, J.R., Tehranian, H., Trennepohl, G.L., 1985. Efficiency analysis and option portfolio selection. Journal of Financial and Quantitative Analysis 20, 435-450.

Brennan, M.J., Schwartz, E.S., 1988. Time-invariant portfolio insurance strategies. Journal of Finance 43, 283-299.

Brooks, R., Levy, H., 1993. Portfolio insurance: does it pay? In: Chance DM & Trippi RR (eds.) Advances in Futures and Options Research. JAI Press, pp. 329-353.

Cesari, R., Cremonini, D., 2003. Benchmarking, portfolio insurance and technical analysis: a Monte Carlo comparison of dynamic strategies of asset allocation. Journal of Economic Dynamics and Control 27, 987-1011.

Christoffersen, P., 2003. Elements of Financial Risk Management. Academic Press.

Davidson, R., Duclos, J.-Y., 2000. Statistical inference for stochastic dominance and for the measurement of poverty and inequality. Econometrica 68, 1435-1464.

Do, B.H., Faff, R.W., 2004. Do futures-based strategies enhance dynamic portfolio insurance? Journal of Futures Markes 24, 591-608.

Elton, E.J., Gruber, M.J., Brown, S.J., Goetzmann, W.N., 2003. Modern Portfolio Theory and Investment Analysis. John Wiley & Sons, New York.

Garcia, C.B., Gould, F.J., 1987. An empirical study of portfolio insurance. Financial Analysts Journal 43, 44-54.

Harvey, C.R., Siddique, A., 2000. Conditional skewness in asset pricing tests. Journal of Finance 55, 1263-1295.

Jobson, J.D., Korkie, B.M., 1981. Performance hypothesis testing with the Sharpe and Treynor measures. Journal of Finance 36, 889-908.

Jorion, P., 2001. Value at Risk: The New Benchmark for Controlling Market Risk. McGraw-Hill, New York.

Kläver, H., 2005. Testing for stochastic dominance using circular block methods. Graduate School of Risk Management. University of Cologne.

Kroll, Y., Levy, H., 1980. Sampling errors and portfolio efficient analysis. Journal of Financial and Quantitative Analysis 15, 655-688.

Leland, H., 1980. Who should buy portfolio insurance. Journal of Finance 35, 581-594.

Leland, H.E., 1985. Option pricing and replication with transaction costs. Journal of Finance 40, 1283-1301.

Leland, H.E., 1999. Beyond mean-variance: performance in a nonsymmetrical world. Financial Analysts Journal 55, 27-36.

Leland, H.E., Rubinstein, M., 1988. The evolution of portfolio insurance. In: Luskin DL (ed.) Portfolio insurance: a guide to dynamic hedging. John Wiley & Sons, New York, pp. 3-10.

Levy, H., 2006. Stochastic dominance. Investment decision under uncertainty. Springer.

Linton, O., Maasoumi, E., Whang, Y.-J., 2005. Consistent testing for stochastic dominance under general sampling schemes. Review of Economic Studies 72, 735-765.

Nelson, R.D., Pope, R.D., 1991. Bootstrapped insights into empirical applications of stochastic dominance. Management Science 37, 1182-1194.

Politis, D.N., Romano, J.P., 1994. Large sample confidence regions based on subsamples under minimal assumptions. Annals of Statistics 22, 2031-2050.

Post, T., 2003. Empirical tests for stochastic dominance. Journal of Finance 58, 1905-1932.

Rubinstein, M., 1985. Alternative paths to portfolio insurance. Financial Analysts Journal 41, 42-52.

Rubinstein, M., Leland, H.E., 1981. Replicating options with positions in stock and cash. Financial Analysts Journal 37, 63-72.

Sanfilippo, G., 2003. Stocks, bonds and the investment horizon: a test of time diversification on the French market. Quantitative Finance 3, 345-351.

Stein, W., Pfaffenberger, R., Kumar, P.C., 1983. On the estimation risk in first-order stochastic dominance: a note. Journal of Financial and Quantitative Analysis 18, 471-476.

Trennepohl, G.L., Booth, J.R., Tehranian, H., 1988. An empirical analysis of insured portfolio strategies using listed options. Journal of Financial Research 11, 1-12.

Whitmore, G.A., 1970. Third-degree stochastic dominance. American Economic Review 60, 457-459.