Evaluating risk mitigation strategies

By Dr. Harald Lohre, David Happersberger and Erhard Radatz

In brief
Risk mitigation strategies seek to create an asymmetric risk-return profile. But benchmarking against the underlying investment is not a valid approach given the potentially stark difference in risk profiles. We discuss how to appropriately calibrate and assess portfolio insurance strategies based on the ensuing return distribution to better fit a given client's risk preferences.

In light of the sustained low yield environment, investors have increasingly taken on more risk to meet their return targets. Yet, their ability to cope with higher risk is limited, which is what makes strict risk management and suitable portfolio insurance techniques so important.

In a previous article[1], we discussed a variety of risk mitigation approaches for a given underlying investment strategy. In particular, we investigated portfolio insurance strategies ranging from static stop-loss techniques to option-based strategies and dynamic portfolio insurance techniques. We concluded that an active portfolio insurance strategy based on a dynamic risk forecast is a cost-effective way to limit a portfolio's maximum loss at a high probability.

In this article we go further and explain how to calibrate such a strategy to individual risk preferences. Since portfolio insurance is meant to accommodate conservative clients' need for an asymmetric return profile, adding a risk overlay ultimately boils down to reshaping the portfolio return distribution. Essentially, the aim is to significantly reduce the probability of suffering from severe tail events while sacrificing some of the underlying strategy's upside potential.

The mechanics of dynamic portfolio insurance
Our preferred dynamic portfolio insurance strategy is rooted in the classic CPPI (constant proportion portfolio insurance[2]) strategy. It typically sets the exposure in a given risky underlying in such a way that a chosen floor level is not breached within a specified investment period. Thus, it is essential to

closely monitor the cushion C_t that represents the difference between the invested wealth W_t and the net present value of the floor $NPV(F_T)$:

$$(1) \quad C_t = W_t - NPV(F_T)$$

To effectively protect the floor,

$$C_t \geq W_t * \text{MaxLoss}(W_t)$$

must hold true. With the investment exposure e_t and the corresponding risky investment $E_t = e_t * W_t$ the above formula can be restated as

$$(2) \quad C_t \geq e_t * W_t * \text{MaxLoss}(\text{risky asset})$$

$$\Leftrightarrow E_t \leq \frac{C_t}{\text{MaxLoss}(\text{risky asset})} = m * C_t$$

This reformulation brings in the notion of the CPPI multiplier m. The multiplier indicates how often the cushion can be invested in the risky underlying without breaching the floor provided the maximum loss assumption holds.

To be on the safe side, one could impose a static multiplier derived from a worst-case risk estimate. But, as we demonstrated in the previous article, such a conservative estimate would severely undermine participation in the underlying. To remedy this issue, we put forward the use of a dynamic forecast of maximum loss. That is, we make use of a dynamic multiplier

$$m_t := \frac{1}{ES_t^{99\%}(\text{risky asset})}$$

labelling this type of risk mitigation DPPI (dynamic proportion portfolio insurance). In this setting, the risk budget and investment exposure dynamically adjust to changes in the estimated expected shortfall (ES) forecast. In particular, participation in the underlying is higher in calmer risk environments, while a pick-up in risk leads to a reduction of investment exposure. Obviously, it is essential to rely on risk estimates that allow for timely modelling of tail risk within the portfolio return distribution.

Panel (a) of figure 1 charts the mechanics and evolution of a DPPI strategy applied to an S&P 500 underlying at an 85% floor level.[3] The dynamic adjustment of the time-varying multiplier m_t follows the expected shortfall forecast derived from a GARCH(1,1)-model. Clearly one can appreciate the role and interaction of floor and multiplier: if the underlying investment is far above the floor, the DPPI tends to have a high investment exposure more or less independent of the risk estimate. With less cushion, the DPPI strategy is more sensitive to risk changes, potentially leading to a complete de-investment.

Over the course of the 32-year backtest, we only observe a few periods of de-investment, of which only four ended in a cash-lock position. While one seeks to avoid cash-lock through the adaptive positioning based on the risk forecast, the success of this approach depends on the specific nature of the corresponding market setbacks. For instance, the minimum daily return of the S&P 500 (-28.6% on 19 October 1987) fully consumed a seemingly comfortable cushion of more than 25%, and induced switching from a 100% investment exposure to cash-lock in just one day. However, in other periods of weak S&P 500 performance, market drawdowns evolved more gradually, allowing the DPPI portfolio time to de-invest and re-invest. The last complete de-investment occurred during the global financial crisis. In the aftermath, interest rates have come down, implicitly elevating the floor level. During high volatility episodes in the equity market, we could observe similar de-risking events within the last decade. Yet these only served to reduce portfolio volatility given quick recoveries in the S&P 500.

Examining the whole sample path, we learn that the DPPI strategy was indeed able to mitigate downside risk. Compared to the underlying investment, the maximum drawdown decreases by approximately 15 percentage points, volatility by 5 percentage points and expected shortfall by 1.5 percentage points under the DPPI strategy (cf. panel (b)). Although these reductions come at the cost of some return potential – the DPPI portfolio earns 141bps less than the underlying – , risk-adjusted measures are in favour of the DPPI strategy.

Figure 1
Performance and allocation of the DPPI strategy

Panel (a)

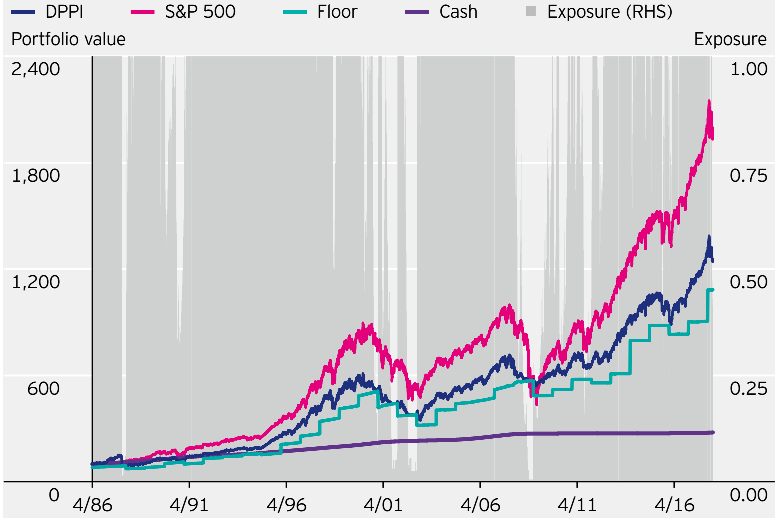

Panel (b)	**S&P 500**	**Money market**	**DPPI**
Return p.a. (%)	9.23	3.20	7.82
Volatility p.a. (%)	19.37	0.22	14.41
Sharpe ratio	0.31	0.00	0.32
Maximum drawdown (%)	-61.17	0.00	-45.80
Expected shortfall 99% (%)	-5.09	0.00	-3.66
Mean exposure (%)	100.00	0.00	86.18

The chart in Panel (a) shows the performance of an equity portfolio (S&P 500) using a DPPI strategy (blue line) in relation to the floor (green line) over time. Exposure is calculated using the cushion (difference between the portfolio value and the floor; here: 85% of the initial annual portfolio value) and the multiplier (based on daily risk forecasting; here: GARCH 99%-ES). For comparison. we have included the performance of the underlying S&P 500 (pink line) and a money market investment (purple line). Panel (b) shows the corresponding performance measures.
Period: 9 April 1986 to 9 April 2018; 9 April 1986 = 100.
Sources: Bloomberg. Invesco. This is simulated past performance and past performance is not a guide to future returns.

Designing DPPI strategies

The preceding example illustrates an important caveat in evaluating a given DPPI strategy, namely, its inherent path dependency. To avoid assessing the strategy based on just one historical path, we rather simulate a large number of alternative price paths and apply the given DPPI-setup. Hence, instead of just one risk and return combination, we obtain a full return distribution.[4] Figure 2 shows portfolio return distributions of yearly returns based on 5,000 simulations, for the portfolio fully invested in the (simulated) underlying S&P 500 as well as for the corresponding DPPI strategy with an 85% floor. The risk estimates required for computation of the dynamic multiplier for the DPPI strategy are based on a simple GARCH(1,1)-model. This model captures the main empirical characteristics of asset returns, such as time-varying volatility, fat tails and volatility clustering.[5]

We observe a left-skewed distribution for the simulated equity underlying. There is tail risk with a non-negligible probability of yearly returns being less than -15%. Applying DPPI results in significantly less tail risk. Yet, one has to note that there is still a small probability of breaching the floor level given that the strategy is adjusted at discrete (daily) intervals.

More importantly, however, figure 2 clearly demonstrates that tail risk reduction, on average, comes at the cost of reduced upside potential. While the historical backtest might suggest an outperformance of the DPPI strategy relative to its underlying, the simulated return distributions more readily articulate that portfolio insurance actually comes at an implicit insurance premium.

Judging by the mean yearly return difference of the two distributions, this premium would amount to some 1.8% (10.5% – 8.7% = 1.8%). At this premium, we can expect to avoid severe tail risk events, 29 of which could be worse than −40% (as simulated in our block-bootstrap analysis).

In the same vein, this framework clarifies the consequences of certain design choices (such as underlying and floor level) for the client's expected portfolio return distribution. For instance, a common theme is that floor levels are set too tight relative to the riskiness of the underlying. Put differently, investors often favour riskier underlyings to achieve certain return targets. Yet, absent a higher risk budget, a riskier strategy will frequently be prevented from breathing freely given that the available cushion is easily consumed. This leads to frequent de-investments or even cash-lock situations triggered by the DPPI mechanism.

To illustrate this issue, figure 3 shifts the floor level from 85% to 95%. As a result, the DPPI return distribution is massively distorted with a lot of return realizations around -5%, i.e. rather close to the floor level. Obviously, this is reminiscent of the fact that, under a too tight floor level, the DPPI strategy frequently de-invests or ends up in cash-lock, disabling it from participating to a meaningful extent in equity markets. The corresponding statistics in table 1 show that the mean exposure reduces to 61%, leading to a significantly lower mean return (6.5% vs. 8.7%) and lower Sharpe ratio (0.24 vs. 0.35) when we shift the floor level from 85% to 95%.[6]

Figure 2
Comparing return distributions

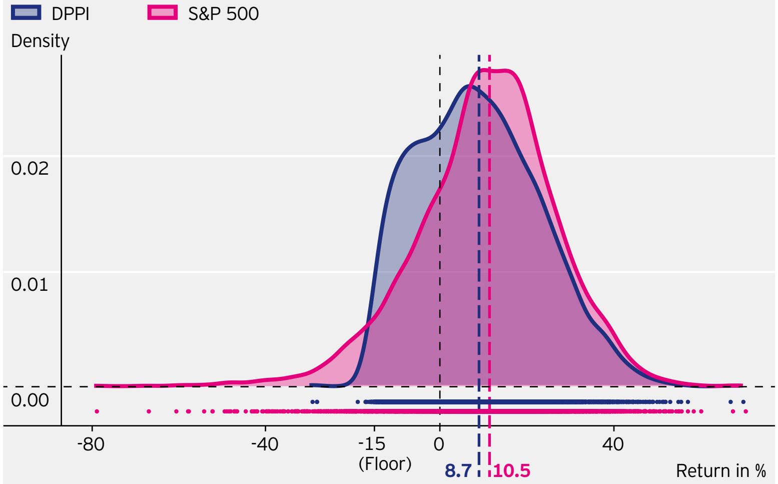

The chart shows the distribution of block-bootstrapped yearly returns (M = 5,000 simulations) of the DPPI portfolio (blue shade) and the one of a pure buy-and-hold portfolio invested in the corresponding simulated S&P 500 (pink shade). The floor level of the DPPI strategy is 85%. Below the two density plots we have added the corresponding support and the mean levels of the return distributions.
Sources: Bloomberg, Invesco.

Figure 3
Comparing return distributions: tight floor levels

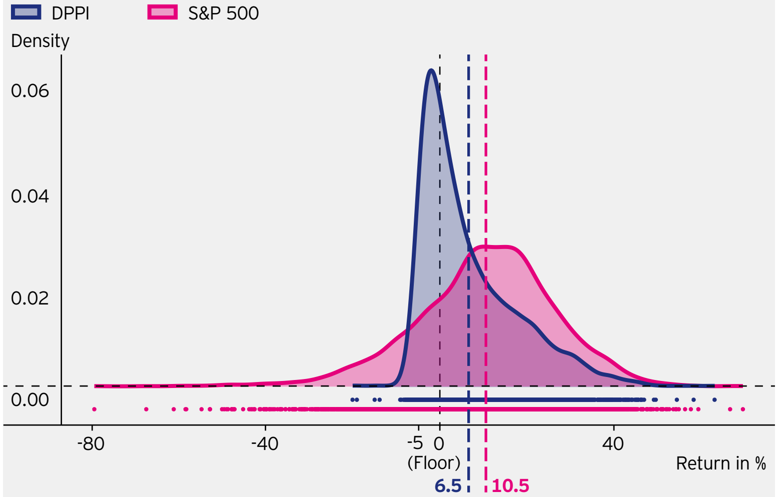

The chart shows the distribution of block-bootstrapped yearly returns of the DPPI portfolio (blue shade) and the one of a pure buy-and-hold portfolio invested in the corresponding simulated S&P 500 (pink shade). The floor level of the DPPI strategy is 95%. Below the two density plots we have added the corresponding support and the mean levels of the return distributions.
Sources: Bloomberg, Invesco.

An alternative benchmark for DPPI strategies

Given the potential for considerable reshaping of the portfolio return distribution through portfolio insurance, it is evident that DPPI should not be benchmarked relative to its underlying. As an alternative, we construct a benchmark with similar risk characteristics. Because we are comparing an asymmetric distribution, a symmetric risk measure like volatility is not viable. Given that risk-averse investors are more concerned about the tails of a distribution, we will base our analysis on the expected shortfall (ES), using a 99% confidence level.

Given the potential for considerable reshaping of the portfolio return distribution through portfolio insurance, it is evident that DPPI should not be benchmarked relative to its underlying.

While there are numerous ways to create a benchmark with a given ES, we opt for an easy and replicable solution. We add cash to the underlying S&P 500 investment to scale down its risk to the pre-defined ES limit of 15%, corresponding to the floor level of the DPPI strategy. We will call this portfolio "ES-target benchmark".[7] As a result, we are comparing two different strategies with similar risk profiles (as defined by their 99%-ES): a portfolio dynamically allocating between cash and the risky underlying (DPPI portfolio) and a static mix of cash and underlying that has an ES similar to the DPPI portfolio (ES-target portfolio).

To achieve an ES of 15% over the sample period, a 39/61 mix of S&P 500 and cash is needed to compute the ES-target benchmark. In figure 4, the ensuing portfolio return distribution is contrasted to that of the underlying S&P 500 and the DPPI strategy with a floor level of 85%. Obviously, the ES-target benchmark return distribution is a compressed version of the underlying S&P 500 return distribution. Most importantly, although its mean return is smaller than the DPPI (6.4% vs. 8.7%), there is still a small probability of significant tail events attached to this strategy (cf. figure 4 and table 1).

Figure 4
Comparing return distributions

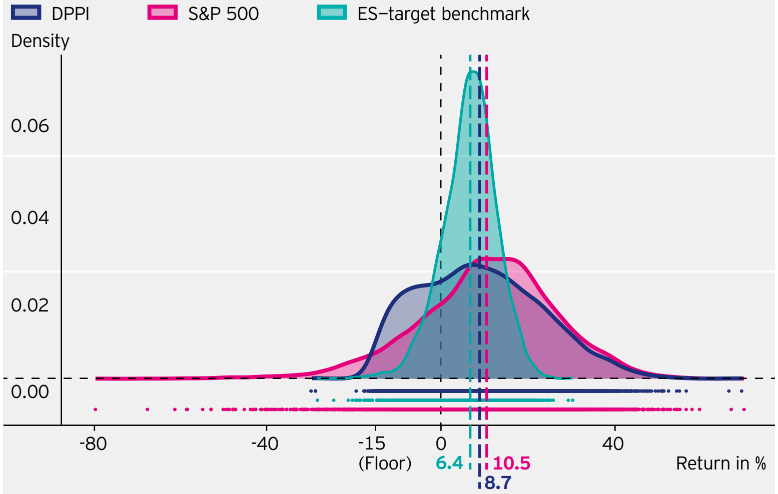

The chart shows the distribution of block-bootstrapped yearly returns of the DPPI portfolio (blue shade) and the one of a pure buy-and-hold portfolio invested in the corresponding simulated S&P 500 (pink shade). The floor level of the DPPI strategy is 85%. The third return distribution applies to a partial investment in the underlying that adds cash such that the average risk level (in terms of the 99%-ES) conforms to the floor level of the DPPI strategy (green shade). Below the density plots we have added the corresponding support and the mean levels of the return distributions.
Sources: Bloomberg, Invesco.

Conclusion

Many investors tend to benchmark the performance of their portfolio insurance strategy vis-à-vis the return of the underlying portfolio. Instead, we suggest the ES-target benchmark strategy. This tail risk-adjusted alternative transforms the underlying's return distribution to better fit the client's risk preferences. Of course, investigating the ensuing portfolio return distributions based on block-bootstrap resampling sheds even more light on the effects of a given portfolio insurance application. We seek to apply this methodology in a future article to investigate the merits of different underlyings in a portfolio insurance framework.

Table 1
Performance of DPPI strategies vis-à-vis the ES-target benchmark

	S&P 500	Money market	DPPI (95% Floor)	DPPI (85% Floor)	ES-Target
Return p.a. (%)	10.49	3.81	6.45	8.71	6.43
Volatility p.a. (%)	15.95	0.96	10.93	14.09	6.30
Sharpe ratio	0.42	0.00	0.24	0.35	0.42
Maximum drawdown (mean, %)	-14.98	0.00	-8.09	-11.77	-3.52
Expected shortfall 99% (%)	-43.83	1.42	-7.85	-16.83	-15.00
Mean exposure (%)	100.00	0.00	61.14	87.28	39.18

The table shows performance measures of a block-bootstrapped DPPI strategy based on an equity portfolio (S&P 500) using different floor levels (85% and 95%). For comparison, we have included the performance measures of an ES-target strategy, targeting the same level of expected shortfall as the DPPI, alongside the underlying S&P 500 and a money market investment. Reported are the mean return, volatility, Sharpe ratio and expected shortfall of the simulated yearly returns, as well as the mean of the maximum drawdowns (which are computed for each simulated path) and mean exposure.
Period: 9 April 1986 to 9 April 2018; 9 April 1986 = 100.
Sources: Bloomberg, Invesco. This is simulated past performance and past performance is not a guide to future returns.

Bibliography

- Andersen, T. G., T. Bollerslev, P. F. Christoffersen, and F. X. Diebold (2013). Financial risk measurement for financial risk management, in Handbook of the Economics of Finance, ed. by G. M. Constantinides, M. Harris, and R. M. Stulz 2(17), 1127–1220.
- Ardia, D., Boudt, K., and Wauters, M. (2016). Smart beta and CPPI performance. Finance, 37(3), 31-65.
- Black, F. and Jones, R. (1987). Simplifying portfolio insurance. Journal of Portfolio Management 14(1), 48–51.
- Black, F. and Jones, R. (1988). Simplifying portfolio insurance for corporate pension plans. Journal of Portfolio Management 14(4), 33–37.
- Happersberger, D., H. Lohre, and I. Nolte (2018). Estimating portfolio risk for tail risk protection strategies. Working paper, Lancaster University Management School.
- Perold, A. F. (1986). Constant proportion portfolio insurance. Working paper, Harvard Business School.
- Perold, A. F. and Sharpe, W. F. (1988). Dynamic strategies for asset allocation. Financial Analysts Journal 44, 16–27.
- Politis, D. N. and Romano, J. P. (1994). The stationary bootstrap. Journal of the American Statistical Association 89(428), 1303–1313.

About the authors

Dr. Harald Lohre
Senior Research Analyst,
Invesco Quantitative Strategies and Visiting Research Fellow EMP/Lancaster University

Dr. Harald Lohre develops quantitative models to forecast risk and return used in the management of multi-asset strategies.

David Happersberger, PhD Candidate
Lancaster University and Invesco Quantitative Strategies

As part of a joint research initiative between Lancaster University and Invesco Quantitative Strategies, David Happersberger is pursuing postgraduate work on practice-oriented issues of financial market econometrics. At the same time, he is actively supporting the transfer of research results into the multi-asset investment process at Invesco Quantitative Strategies.

Erhard Radatz
Portfolio Manager,
Invesco Quantitative Strategies

In his role, Erhard Radatz manages multi-asset portfolios that include the elements factor-based investing, active asset allocation and downside risk management.

Notes

1 See Theory and practice of portfolio insurance, Risk & Reward #2/2017.
2 For more on CPPI strategies, cf. Perold (1986), Black and Jones (1987, 1988), Perold and Sharpe (1988).
3 Throughout the article, and in all figures and tables, we employ the S&P 500 Future as equity investment. For money market investments we use the 3-month US Treasury bill. All asset returns are in local currency. All simulations in this article are provided for illustrative purposes only and are subject to limitations. Unlike actual portfolio outcomes, the model outcomes do not reflect actual trading, liquidity constraints, fees, expenses, taxes or other factors that could impact future returns.
4 In simulating alternative price paths, we use the stationary block-bootstrap of Politis and Romano (1994). We follow Ardia, Boudt and Wauters (2016) in that block lengths are drawn from a geometric distribution with a minimum block length of one day and an average of 15 days.
5 For more on GARCH models, cf. Andersen et al. (2013).
6 As is common in academic literature, the annualized returns, volatilities, and Sharpe ratios shown in Table 1 are based on the 5,000 annual returns from the simulations. So, given the different frequencies, it is not surprising that the historical volatilities shown in Panel (b) of Figure 1 and that are based on historical daily returns, are slightly higher. This effect is exacerbated because, of course, the simulation paths are relatively rare in containing the extreme historical returns realizations, and thus there is a corresponding relativization.
7 See Happersberger, Lohre and Nolte (2018) for an empirical study of ES-target strategies in the context of tail risk protection.

The outputs of the assumptions are provided for illustration purposes only. Unlike actual portfolio outcomes, the model outcomes do not reflect actual trading, liquidity constraints, fees, expenses, taxes and other factors that could impact future return.

Important Information

This document has been prepared only for those persons to whom Invesco has provided it for informational purposes only. This document is not an offering of a financial product and is not intended for and should not be distributed to retail clients who are resident in jurisdiction where its distribution is not authorized or is unlawful. Circulation, disclosure, or dissemination of all or any part of this document to any person without the consent of Invesco is prohibited.

This document may contain statements that are not purely historical in nature but are "forward-looking statements", which are based on certain assumptions of future events. Forward-looking statements are based on information available on the date hereof, and Invesco does not assume any duty to update any forward-looking statement. Actual events may differ from those assumed. There can be no assurance that forward-looking statements, including any projected returns, will materialize or that actual market conditions and/or performance results will not be materially different or worse than those presented.

The information in this document has been prepared without taking into account any investor's investment objectives, financial situation or particular needs. Before acting on the information the investor should consider its appropriateness having regard to their investment objectives, financial situation and needs.

You should note that this information:

- may contain references to amounts which are not in local currencies;
- may contain financial information which is not prepared in accordance with the laws or practices of your country of residence;
- may not address risks associated with investment in foreign currency denominated investments; and
- does not address local tax issues.

All material presented is compiled from sources believed to be reliable and current, but accuracy cannot be guaranteed. Investment involves risk. Please review all financial material carefully before investing. The opinions expressed are based on current market conditions and are subject to change without notice. These opinions may differ from those of other Invesco investment professionals.

The distribution and offering of this document in certain jurisdictions may be restricted by law. Persons into whose possession this marketing material may come are required to inform themselves about and to comply with any relevant restrictions. This does not constitute an offer or solicitation by anyone in any jurisdiction in which such an offer is not authorised or to any person to whom it is unlawful to make such an offer or solicitation.