Adding value to pensions portfolios through effective risk management and hedging

Received: 15th December, 2005

Frances Cowell

has worked in the investment management industry since 1983, starting as an equity analyst for Aetna Investment Management in Australia. Since then she has managed domestic and international equity and fixed interest portfolios, and asset allocation. In each case the focus has been applying quantitative techniques, with an emphasis on combining derivatives and physical investment, to construct and manage investment portfolios with predefined risk targets to meet specified investment objectives. In 1998 she moved to the UK to work for QUANTEC, a major provider of portfolio risk management systems, and in 2002 took up duties as Interim Head of Portfolio Risk at Morley Fund Management. She now works in the Portfolio Risk team, specialising in risk management for derivatives and hedge funds.

Abstract Lurking in a pensions portfolio may be risk about which the investor is unaware. Portfolio construction processes often do not take full consideration of the correlations between their components. Risk management that mimics the portfolio construction process is likely to misestimate risk and confound capital allocation with economic exposure. The result is poor risk budgeting, ineffective hedging and underperformance over time. By contrast, portfolio-level risk measurement and decomposition can distinguish between actual and apparent risk, allowing the investor to distinguish between intentional and incidental risk. This allows the investor to concentrate risk where it is likely to contribute most to return. By eliminating unwanted risk, this process reduces portfolio risk for a given level of return. The practical implication is that the portfolio will be much more closely targeted to its investment objectives, with better returns over time and less unforeseen volatility. Accurate risk measurement and effective risk management result in more cost-effective hedging strategies, which can be critical to funds with low risk tolerances.

Keywords: *risk management, hedging diversification, risk factors, risk measurement risk decomposition, risk targeting*

Introduction

The objective of investment is to take risk: after all, this is where returns come from. This paper is about managing investment risk to achieve the best possible returns. It will discuss the requirements for effective risk management, starting with meaningful risk measurement and decomposition, some observations about actual versus apparent risk in the context of a multi-asset class portfolio, proceeding to

the relative merits of diversification and hedging as risk management techniques, finishing with a discussion of the importance of tailoring risk to the portfolio's investment objectives.

It is imperative to be clear about how much risk the portfolio needs to take. Clearly, it needs to take enough risk to give it a reasonable chance of meeting its return objectives. Clearly, it should also avoid taking more risk than it needs, to avoid excessive volatility and therefore

Frances Cowell
Morley Fund Management
138 Avenue Ledru Rollin
Paris 75011
France

Tel: +33 1 43 67 37 86
Fax: +33 1 46 59 48 29
Email: frances.cowell@morleyfm.com

unacceptable losses. The higher the return objective, the more risk the portfolio needs to take. The greater the appetite for risk, the higher the potential returns.

Risk measurement

In order to manage risk, it first must be measured. Whether the risk is measured as tracking error or value-at-risk (VaR) is less important than the accuracy and validity of the measure selected. The choice of metric should be determined by the type of portfolio and the assets in which it invests. Factors that need to be considered include:

- Is the portfolio aiming for absolute returns or to beat some benchmark or other reference portfolio?
- What is the investment horizon?
- What kind of assets will it invest in?

Both tracking error and VaR are ways of expressing likely portfolio return variability by linking a gain or loss to the probability with which it is likely to occur, and both have advantages and disadvantages. Tracking error is usually, but not always, most appropriate for portfolios that aim to beat some kind of benchmark, have an investment horizon of longer than about a month and invest in assets with symmetric or near-symmetric return distributions, such as equities and currencies. It has the enormous advantage that it provides a closed-form solution, efficiently accommodates covariances between portfolio components and facilitates risk decomposition. It is usually forecast using some variant of Markovitz mean-variance estimation.

VaR is suitable for many types of absolute return portfolios, portfolios with short investment horizons, and those investing in assets with skewed return

distributions, such as corporate and convertible bonds. VaR is usually estimated using some kind of Monte Carlo simulation, which has the advantage that it accommodates non-normal return distributions, such as skew and fat tails. Both techniques have been used extensively in portfolio management and the strengths and limitations of each approach are by now fairly well understood.

Whether tracking error or VaR, risk should be measured as far as possible at the aggregate portfolio level. A pensions portfolio for example, is usually a diversified portfolio investing in a number of different asset classes and across several currencies. In practice, the components are often managed by different investment managers, each with a specific benchmark, investment universe, set of constraints and risk and return targets. Although it is tempting to assume that the risk of the overall portfolio is approximately equal to a weighted sum of the risks of the portfolio components, this would be a mistake, as this assumption takes no account of the correlations between the asset allocation decisions (allocating funds between asset classes and managers) and the stock selection decisions of those managers; and the correlations between the asset classes themselves. The effect of these correlations is that the portfolio-level risk can be significantly higher — or lower — than the weighted average of the components would imply. In short, risk is not additive.

Other things being equal, and subject to gearing constraints, adjusting the total risk of a portfolio is simply a matter of increasing or reducing the size of active positions — or the amount of cash held (because cash is assumed to be riskless). So, for example, if the risk measure indicates that the portfolio has only a small chance of earning the returns

Table 1: Asset allocation contribution to portfolio risk

Asset class	Active allocation (%)	Expected return (%)	Tracking error (%)	AA contribution to portfolio risk (%)	SS contribution to portfolio risk (%)	Interaction (%)
Domestic equities	5.2	9.5	8.0	9.5	2.8	1.3
North American equities	2.1	6.5	8.0	6.3	2.3	1.0
Japanese equities	4.5	8.5	8.0	6.8	2.5	0.8
European equities	1.5	4.5	8.0	6.1	2.7	1.3
Pacific Basin equities	-3.0	3.5	10.0	5.5	4.2	2.1
Emerging markets equities	-4.8	-1.5	10.0	5.6	4.6	3.2
Domestic fixed Interest	-4.3	-0.1	3.0	3.5	1.8	0.9
Global fixed Interest	-3.5	0.0	3.0	4.0	2.2	1.0
Corporates	-1.3	3.8	4.0	5.8	4.5	-1.0
Property	3.6	7.8	8.0	6.5	3.6	1.2
Cash	0.0	0.0	0.0	0.0	0.0	0.0
Active asset class returns		1.3				
Tactical asset allocation		2.5	5.0			
Total	0.0	3.8	8.0	59.6	31.2	11.9

expected of it, then the investor can usually simply increase the size of the risky positions or reduce the amount of cash held. If, on the other hand, the risk estimate indicates that the portfolio is unduly prone to large return variations, then position sizes can be reduced or more cash held.

Economic theory says that the higher the return demanded by the investor, the more risk needed to achieve those returns. But not all risk is rewarded with extra returns: the same theory says that only risk that is not diversifiable is so rewarded. For the most part, non-diversifiable risk means market or common factor risk while diversifiable risk means security-specific risk. So investment risk management is not just about managing the level of risk in the portfolio, but about managing the sources of that risk. The objective is to distinguish between intentional and incidental risk.

For a balanced, multi-asset class portfolio, there are two broad approaches to risk decomposition. The more intuitive, and more commonly used, is to analyse the risk in the same way that the portfolio is constructed; that is, estimate the risk of each component asset class

and then estimate the risk attributable to the asset allocation decision. This is illustrated in Table 1.

The table shows, for each asset class, active asset allocation (the difference between portfolio and benchmark allocation), expected return and asset class tracking error. It also shows how much risk is due to the active asset allocation position, the amount of risk due to security selection within asset classes, and the interaction between the two. It illustrates a common misunderstanding in portfolio risk analysis: that, contrary to popular perception, risk exposure is not the same as capital allocation.

For example, this portfolio is overweight in domestic equities because this asset class is expected to have the highest returns. This decision is estimated to contribute about 9.5 per cent to total portfolio risk, a much higher proportion than that which is implied by the allocation of 5.2 per cent over benchmark. The domestic equities portfolio aims for a tracking error of about 8.0 per cent, and this stock selection risk will contribute a further 2.8 per cent to the total risk of the portfolio. Meanwhile, the domestic equities

portfolio manager believes this asset class will deliver positive returns, and has positioned the portfolio accordingly. This means that his stock selection decisions are positively correlated with the asset allocation decision, and this correlation contributes a further 1.3 per cent to the risk of the overall portfolio.

But the analysis does not quantify the interaction at the stock selection level between asset classes. The portfolio is underweight in corporate bonds, but the corporate bond manager does not believe that bonds will do so badly, and positions his portfolio for a decline in credit spreads and, consequently, positive returns. The asset allocation and stock selection decisions are therefore negatively correlated, and this correlation is subtracting 1.0 per cent from the risk of the overall portfolio.

In addition to the interaction between asset allocation and stock selection decisions, the stock selection decisions of the domestic equities and corporate bond managers are positively correlated with each other, so adding hidden risk to the portfolio.

The effect of the interactions is that portfolio-level tracking error is different from the weighted sum of asset class tracking errors, and the implication of this is that, although risk management for each of the portfolio components may be perfectly adequate, it does not constitute effective portfolio-level risk measurement and analysis, and the portfolio will be either too volatile, as in this example, or not volatile enough; and consequently will fail over time to meet its return objectives. Either way there is waste. If the hidden correlations are positive, then there is unmanaged risk that will contribute to extra returns, and so should be eliminated or redirected to where it will make a positive contribution. If the hidden correlations are negative, then asset selection skill is being wasted by

offsetting exposures, a position that is sometimes called the 'closet index' effect.

In reality the asset class distinction is arbitrary. Asset classes are used by portfolio managers to simplify the job of managing large, diversified portfolios. Most portfolios are selected from a universe of tens of thousands of assets. Conducting detailed research and evaluating the assets of such an extensive universe cannot be carried out by a single individual or even a single team. Instead it tends to be divided by asset type, such as bonds, equities and currencies; and by geographical region. This allows specialisation and therefore superior asset selection, so simplifying the task of portfolio selection. But what works for asset selection does not necessarily work for risk management.

Another, more effective, but less intuitive, way to account for investment risk is to eliminate the asset allocation layer altogether and carry out the risk analysis at the aggregate portfolio level. Although conceptually simple enough, in practice this approach is complex. Apart from the sheer number of assets that comprise a typical portfolio and benchmark, the assets themselves will have important differences so that finding a single risk analysis technique that works tolerably well for all is rarely simple. For example, tracking error works well for equity portfolios but less well for interest rates and corporate credit, for which VaR simulation is usually more effective.

Whatever risk measure is used, the objective is to identify themes, or risk factors that drive returns for the whole portfolio. The relationship between the portfolio and the risk factors, usually thought of as the portfolio's factor exposure, is estimated most effectively using some kind of regression analysis. For this reason, a risk factor is usually defined by a return series, but in other respects can be anything. Examples of

Table 2: Factor contribution to portfolio risk

Factor	Portfolio factor coefficient	Factor contribution to portfolio risk (%)
Global equities	1.08	21.4
Value-growth	0.89	6.1
Market capitalisation	0.84	7.2
Long-term interest rates	1.02	5.0
Short-term interest rates	1.05	1.5
USD	0.91	11.8
EUR	1.07	9.9
JPY	0.82	-2.9
CHF	1.03	10.5
Energy	0.82	8.1
Base metals	0.93	6.8
Residual		14.6
Total		100.0

Table 3: Contribution to portfolio risk

Asset	Active weight (%)	Contribution to residual risk (%)
Ericsson	4.2	19.8
Vivendi Universal	4.3	14.5
Ing Groep	4.1	14.2
Numico	4.5	11.2
Siemens	4.2	10.6
Tietoenator	3.8	10.0
Sodexho Alliance	4.4	9.8
Assa Abloy	4.1	9.1
Akso Nobel	4.4	8.7
Continental	4.5	8.1
Volkswagen	-1.0	-1.8
Next	-2.5	-2.2
KPN	-2.9	-2.5
Reed Elsevier	-2.8	-3.1
Solvay	-2.9	-3.3
Converium	-3.0	-3.4
CIBA Specialty	-2.7	-3.5
Stora Enso	-2.2	-3.7
TF1	-2.4	-4.9
Infineon Technologies	-2.7	-7.1

factors that are commonly applied to diversified, multi-asset class portfolios include the return to global equities, using a global equity index as a proxy; global bonds, with the return to a global bond portfolio as a proxy; currencies, global industries and returns to commodity indices. Other factors, such as value-growth and market capitalisation, which cannot be represented by an existing asset or index, must be defined, usually by means of dummy portfolio for which historical returns can be computed. To enhance

their explanatory power, risk factors should be statistically independent as far as possible, and as relevant as possible to the nature of the portfolio. The more 'normal', or symmetrical the distribution of component asset returns, the more power this measure has. The regression results in a factor coefficient, an R square and some measure of how much the factor exposure contributes to portfolio-level risk. If these factors are well chosen, this common factor analysis will explain a significant part of the overall risk of the portfolio.

A factor-based risk attribution might look like the example in Table 2.

The portfolio in Table 2 has, as expected, a positive bias toward the equity market and shows a strong growth (negative value) bias, which is echoed in its tilts toward small stocks and interest rates, and away from energy and base metals. More surprising is the exposure to Swiss francs, which, if neutralised, would reduce the portfolio's tracking error by 10.5 per cent.

Not all portfolio risk can be explained by common factors. The risk that is left over is called the residual. In the above example, this is contributing nearly 15 per cent of the portfolio-level risk. It is therefore important that it be quantified.

Table 3 shows the top ten security-level contributors to residual risk and the top ten diversifiers. It further illustrates the distinction between capital allocation and risk exposure. This equity portfolio has some interesting concentrations of risk that bear no resemblance to the allocation to those assets, reflecting both the volatility of the individual securities and their correlations with the other assets in the portfolio. For example, the largest contribution to risk is the bought position in Ericsson, contributing 19.8 per cent to residual risk although, at 4.2 per cent, it has by no means the largest allocation of capital.

In this case, the manager is probably unaware of the concentration of risk in Ericsson, and indeed in most of the top ten overweight positions in the portfolio. He is probably also unaware that the relatively modest underweight position in Infineon is reducing risk by as much as it is. This analysis would probably lead him to modify the stock-level risk by reducing the overweight positions so that the risk they contribute corresponds more closely with his views on the stocks concerned.

The power of the aggregate portfolio

approach to risk decomposition is that because it contrasts with most portfolio construction techniques, it can highlight risks that were otherwise not apparent. With a good factor-based risk analysis the investor is in a position to decide which risks are desirable and which should be eliminated.

Favourable or intentional risk by its nature can be linked to expected return, with the magnitude of the risk corresponding to the strength of the investor's view on the likely returns to the factor or asset.

Unintentional risk, whether factor-related or asset-specific on the other hand should be reduced or eliminated. If left, this risk will add further portfolio volatility without contributing anything to expected return.

Risk management

There are two broad approaches to neutralising or otherwise managing factor-related and security-specific risk. It can be either diversified away or hedged.

For example, if the investor wished to modify the growth tilt in the portfolio, it could be diversified by rebalancing the portfolio. In theory, it could also be hedged, using swap or option contracts. Because there is no traded market in growth-value instruments, the hedging strategy would need to be an over-the-counter transaction, such as a swap or an option. This has the advantage that it can be tailored to cover precisely the risk to be modified, and the hedge can be effected and reversed relatively quickly. This is useful if the investor changes his forecast of the returns to the factor. The disadvantage of over-the-counter hedges is that they can be costly. By contrast, diversifying this type of factor risk incurs only the transactions costs of the portfolio rebalance. Re-imposing the growth tilt,

Table 4: Strategies for neutralising portfolio risk

Risk	Instrument	Strategy	Advantages	Disadvantages
Growth-value factor	Physical	Diversify	Cost limited to transactions costs Can be costly	Expensive to reverse Eliminates stock-specific risk
Futures/Swaps	Swap return to growth portfolio for the return to broad equity portfolio or index	Targets factor risk Preserves stock-specific risk		
Put Options	Buy put option on growth portfolio	Targets factor risk Preserves stock-specific risk Maintains exposure to positive returns to growth factor	Can be costly	
Swiss franc factor	Physical	Diversify	Cost limited to transactions costs Can be costly	Expensive to reverse Eliminates stock-specific risk
Futures/Swaps	Swap return to Swiss Franc for the return to broad equity portfolio or index	Targets factor risk Preserves stock-specific risk		
Put Options	Buy put option on Swiss Francs	Targets factor risk Preserves stock-specific risk Maintains exposure to positive returns to Swiss franc factor	Cost of option premium	
Security-specific	Physical	Diversify	Cost limited to transactions costs	
	Futures/Swaps	Swap return to individual security for the return to broad equity portfolio or index	Targets factor risk	Can be costly
	Put Options	Buy put option on individual security	Targets factor risk Maintains exposure to positive returns to stock	Cost of option premium

however, would incur further rebalancing costs, so the rebalancing option works best if the intention is to neutralise the factor risk permanently. The diversification strategy will also diversify away most, if not all, of the returns expected from stock selection associated with the growth tilt.

For the Swiss franc exposure, the investor has greater choice because actively traded markets exist in Swiss franc forwards and options. This means that hedging can be very cost-effective and is easily reversed. In this example, the Swiss franc tilt was incidental, not intentional, resulting from active stock selection. Hedging just the factor risk can

leave much of the benefit of this stock selection in place, preserving the benefit of the stock selection skill.

Unwanted security-level risk can be eliminated either by diversification (rebalancing the portfolio) or hedging using forwards, contracts-for-differences or options, exchange-traded if a suitable market exists for the stocks in question, or over-the-counter if they do not. Single security over-the-counter transactions is very simple to implement, so the cost of this strategy is likely to be minimal, and is easily reversed.

The strategies for neutralising factor-related and security-specific risk are summarised in Table 4.

It is important to note that if factor-related or stock-specific risk is neutralised using forwards, contracts-for-differences or swaps, the result will be a synthetic cash position, perhaps leaving the portfolio underinvested. Options have a similar effect because of the delta effect of the option, which means that, prior to expiry, the changes in the option price offset only imperfectly those in the underlying security.

Given the costs of hedging factor-based or security-level risk, the investor must question whether this cost is well spent. A future, forward, contract-for-difference or swap transaction that completely neutralises the target risk will, in retrospect be seen as very expensive if that factor or security does not deliver the expected negative returns prior to the expiry of the hedge. The hedge will appear to have 'lost' money. Similarly, many investors have underperformed compared with their peers because of the cost of options bought to protect against negative returns that did not occur.

This illustrates the point that, by itself, hedging can be a risky activity. It also suggests that between them, diversification and hedging do not address all portfolio risk management issues.

So far this paper has considered only the assets of the portfolio. For hedging to be effective, it needs to take into account the investment objectives of the portfolio, whether these are explicit or implied. What does this mean in practice?

Matching portfolio risk with investment objectives

This paper started with the observations that the objective of investment is to take risk, and that how much risk should

be taken is a matter of how much return is required. This in turn is a function of what liabilities or promises the fund must meet, and when.

For pension portfolios, the objective is to meet a set of projected liabilities over time. Although the actual liabilities of the fund are rarely known precisely, they can in practice usually be approximated by a series of absolute returns.

The unavoidable link between a pension fund's promises, its investment objectives and the actual portfolio composition is of course the forecasts of returns and covariances, neither of which can be known with certainty. So there is always the possibility that returns will fail to meet the portfolio's objectives. What are the standard approaches to this risk?

One way is to hold reserves that are invested in very low-risk assets, while investing the balance of the fund in a conventional balanced portfolio. The ratio of reserves typically reflects the perceived mismatch between forecast investment returns and projected fund liabilities. Asset-liability matching employs a similar, but more sophisticated approach that more closely models the known liabilities of the fund with bonds of the appropriate duration, while investing the 'surplus' in a conventional multi-asset class portfolio.

A fund that is in surplus or is a defined contribution scheme, does not, strictly speaking, need to hold reserves, and a conventional balanced portfolio can be implemented. The mix of assets, and therefore the risk level, should reflect the risk tolerance of fund members, usually as a function of the time until they start to draw down their investment.

A fund that has very low tolerance for negative returns, perhaps due to very low reserves relative to its known liabilities, may need to consider some kind of protection strategy to avoid

further losses. There are various techniques for implementing such a strategy, all using some kind of option technology. The benefits of this approach are that it allows the fund to participate in positive asset returns while protecting against negative returns. The disadvantage is the cost, which can be substantial depending on market volatility and the level of protection required.

Thus the investment objectives of any fund usually boil down to how much it must return and how much it can afford to risk in achieving that return.

Whatever the circumstances of the fund and its chosen investment strategy, the benefits of sound risk management are enormous. Funds with plenty of risk tolerance can eliminate waste and enhance their returns by taking only carefully targeted risk and eliminating unwanted risk, effectively improving their risk-return ratio. Funds with a reserving or asset-liability matching strategy can further improve the return to the active component of their portfolios by eliminating incidental risk. Improving the risk-return balance of this part of the portfolio will, other things being equal, reduce the need for future reserves, so improving future return prospects.

The benefits of effective risk management are most important however for portfolios with very low risk tolerances. Because effective risk management enhances the returns achievable for a given risk level — or reduces the risk for a given return, the fund has a better chance of achieving a healthy balance of assets and liabilities without needing to deploy expensive protection. If portfolio protection is required, then the cost of that protection is less because:

- The risk of the portfolio to be protected is less for a given return.

Because the cost of protection is a function of, among other things, the risk of the asset to be protected, the protection is cheaper.

- At a given level of risk, the return to the active portfolio are greater, affording the fund a greater participation rate in positive returns to the protected portfolio.

For a protected portfolio, the net effect of good risk measurement and analysis is that the protection itself is targeted so that no portion of it is wasted on risk that is apparent but not real. At the same time, no unwanted risk is overlooked, thus avoiding unwanted portfolio volatility.

Conclusion

Without risk there can be no return, but not all risk is equal, and investors need to be selective about the risk that drives the returns to their portfolios. Obtaining the insight necessary to understand the actual, as opposed to the apparent risks in an investment portfolio depends on posing the right questions.

The benefits of doing so are substantial. A well-considered strategy of diversification and hedging can result in portfolio-level risk that is concentrated where it contributes most to expected returns, while incidental risk is eliminated. The result is reduced portfolio volatility for any given level of expected returns.

For pension fund managers seeking to maximise the value of their fund, while avoiding unacceptable losses, the benefits are significantly better risk-return ratios at all levels of risk. These benefits are particularly important for funds with very low risk tolerances, as good risk management delivers a much higher participation rate in rising

markets and materially cheaper portfolio protection.

Further reading

- Bodie, Z., Kane, A., Marcus, A. J. (1993) 'Investments', McGraw Hill Irwin, New York, NY.
- Chiang, A. C. (1974) 'Fundamental Methods of Mathematical Economics', McGraw Hill, New York, NY.
- Kritzman, M. (1995) 'The Portable Financial Analyst', Chicago, IL, Probus.
- Lakonishok, J., Schleifer, A. and Vishny, R. W. (1992) 'Study of the US Equity Money Manager Performance', Brookings Institute Study, Washington, DC.
- Markowitz, H. M., (1959) 'Portfolio Selection: Efficient

Diversification of Investment', Cowles Foundation Monograph 16. Yale University Press, New Haven, CT.

- Rosenberg, B. (1974) 'Extra market components of covariance in security returns', *Journal of Financial and Quantitative Analysis*, Vol. 9, No. 2, March, pp. 263–274.
- Rudd, A. (1980) 'Optimal selection of passive portfolios', *Financial Management*, Vol. 9, Spring, pp. 57–65.
- Rudd, A. and Clasing, H. K. Jr. (1988) 'Modern Portfolio Theory' (2nd edn). Andrew Rudd, Orinda, CA.
- Sharpe, W. F. (1970) Capital asset prices: a theory of market equilibrium under conditions of risk, *Journal of Finance*, Vol. 19, No. 3, pp. 425–442.
- Sharpe, W. F. (1970) 'Portfolio Theory and Capital Markets'. McGraw Hill, New York, NY.