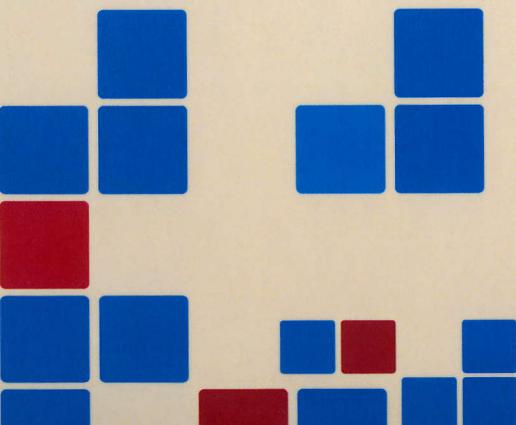


# Foreign Exchange

A Practitioner's Approach  
to the Markets  
Edited by Amy Middleton



# *Risk Assessment in Currency Programmes*

**Michael Shilling and Gavin Francis**

Pareto Investment Management Limited

Currency risk differs from all other risks in a portfolio because it can be added or removed without affecting the leverage of the portfolio. A fully invested portfolio can have anything from zero to 100% foreign currency exposure. The currency allocation, like the asset allocation, spans the entire portfolio. This is why currency risk management is so important.

The role of currency exposure in a portfolio must be based on risk and return. There is little doubt that currencies display large and sometimes rapid moves, both positive and negative. This makes them risky. This chapter explores ways to measure this risk, both when it is held passively and when it is managed actively. The case for a positive expected return from passively holding a basket of foreign currencies is hard to make. However, there is a rationale for holding actively managed currency exposure.

When assessing currency risk in a portfolio of investments, it is important to define the risk-free position. When an international portfolio's currency exposure is fully hedged, the portfolio is, by definition, free of foreign currency risk. This means it must be exposed only to the investor's base currency. So currency risk exposures must be measured against domestic cash. In fact, this is entirely consistent with the way that asset risk exposures are measured, for example, in calculating the Sharpe ratio.

Currency exposure is created by exchanging domestic cash for foreign currencies. This occurs when investors purchase foreign assets. So, unless some hedging takes place, currency exposure is created by hiring a manager of foreign equities or bonds, by investing in foreign property or private equity or by purchasing units in a foreign

commingled fund or hedge fund. An unhedged investment implements two investment decisions at the same time: an asset decision and a currency decision. However, a broad array of modern financial instruments allows institutional investors to make either or both of these decisions independently.

Investors often find themselves exposed to currency as a by-product of holding foreign assets. They do not make a deliberate decision to hold a basket of currency exposures that happens to match, say, the allocation of a capitalisation-weighted equity index. It just comes along, sometimes as unwanted baggage, with foreign equities. With a 20% overseas investment, the resultant currency exposure would have a dramatic negative impact at the total portfolio level if foreign currencies fell by 10%. Such a loss could wipe out a year's worth of alpha generated across the entire portfolio. This makes currency exposure a risk that needs to be addressed by a currency hedging policy.

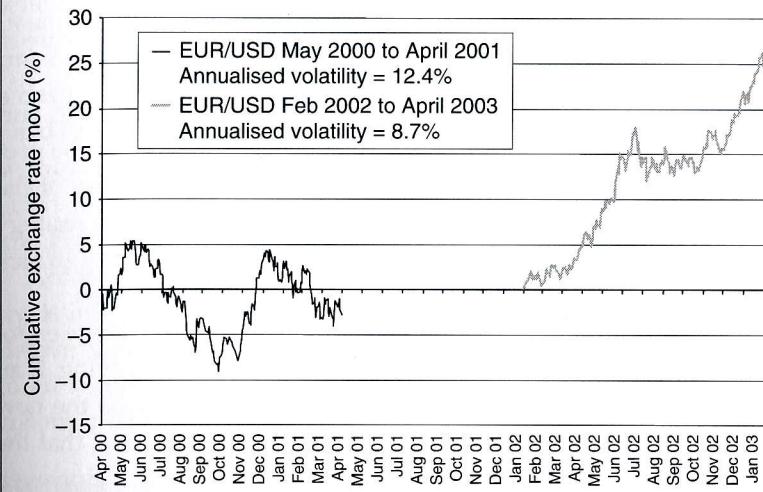
On the other hand, a 10% rise in foreign currencies could add 2% to total assets. The perception that either outcome is equally likely leads to the use of volatility to measure risk. With this symmetric measure, the triumph of a gain is considered to be just as risky as the disaster of a loss. With apologies to Rudyard Kipling, volatility treats those two impostors just the same.

The purpose of currency hedging is not to create gains; its role is to reduce losses. So it is meaningless to consider the risk of the hedges in isolation; an investor must assess the net impact of the hedges in combination with the underlying currency exposure. Hedging controls the absolute level of currency risk at the total portfolio level.

A common method of implementing a hedging policy employs a static or passive approach. But the problem is that in addition to reducing losses, this also reduces gains symmetrically without discretion. Since a passive strategy aims only to replicate a benchmark, not to beat it, any deviation, upward or downward, should be avoided. This makes the volatility of benchmark relative returns (tracking error) an appropriate risk measure for assessing the net impact of passive hedging.

Contrast this with an active hedging strategy. Once again the role is to reduce losses, but now the manager seeks to adjust the hedge ratio over time so that the reduction of losses exceeds the reduction of gains. This results in a positive overall return. The entire purpose

Figure 1.1 Short-term v. long-term volatility



of active management is to create an asymmetry between gains and losses. This does not apply simply to currency. Any form of active management seeks to create asymmetry. Therefore, risk measures that highlight rather than ignore asymmetry should be much more useful than volatility for identifying attractive investment opportunities and, in particular, for defining the investment guidelines of active currency managers.

There is another difficulty with volatility as the measure of risk, particularly in currency markets. It does not scale over time. For example, the annualised volatility of daily currency movements is not the same as the annualised volatility of monthly movements or indeed the volatility of annual returns. This means that the investor's time horizon is critical if any measure of risk is to be meaningful. It is quite possible, for example, to see large cumulative moves over time, even if daily volatility is low and, vice versa, high short-term volatility may not necessarily lead to a significant move over the course of a year.

Figure 1.1 illustrates this point. In the first period volatility is high but there is little cumulative move. In the second period, volatility is low, but there is a substantial cumulative move over time. Which environment should the investor be most concerned about? An investor with a long-term horizon is very much concerned with

prolonged cumulative moves, and much less worried about short-term volatility. Indeed, the main worry is about a long-term adverse move. This concept of risk is therefore not reflected in a measure of volatility that uses high-frequency data and is symmetrical.

This observation has implications not only for assessing currency risk for investors but also in constructing currency alpha programmes, which can aim to exploit different characteristics of the market. It is probably impossible to construct a single alpha strategy that exploits both long-term trends and short-term volatility. In fact, an opportunity on the one hand becomes a risk on the other.

Finally, returning to the unique feature of foreign exchange exposure, currency risk is additive and is therefore not diversifying. If an investor takes on a new asset, risk is diversified, because the new asset replaces some other investment, provided, of course, that the new asset is not fully correlated with the existing portfolio. However, when an investor brings currency risk into a portfolio, it is not replacing any other risk, it is simply adding to it. This means that it will only be diversifying if the currency risk is negatively correlated with the existing portfolio. It is not sufficient merely to be uncorrelated.

So, in assessing currency risk in an investment portfolio and defining the scope for a currency management strategy, we need to be able to address asymmetric returns, we need to determine the appropriate time horizon and we need to fully understand what we mean by diversification.

### DETECTING ASYMMETRY

It is so natural to consider a gain to be more attractive than a loss that it can seem puzzling that investors have been fixated for so long on using the symmetric volatility statistic as a risk measure. Even the proponents of absolute return investing feel an obligation to provide volatility statistics, whereas the outcome they seek to avoid is incurring a loss.

The realm of currency provides some interesting examples of asymmetry, particularly when interest rate differentials are high. Holders of a high-yielding currency benefit as long as the exchange rate remains stable, but this comes at the cost of being exposed to a sharp devaluation. So, from a risk perspective, the asymmetry lies in the large negative tail.

Skilful active managers need to avoid large negative tail events. In fact, their returns would be more attractive if they could create large positive tail events. This is the kind of attractive asymmetry that both preserves and enhances wealth. This means that, in order to achieve a desirable outcome, the investment guidelines of active managers need to balance two conflicting objectives: to limit negative outcomes without limiting positive outcomes. The talent of a genuinely skilful active manager lies in the ability to add value in this way (Okuyama and Francis 2007).

According to this logic, a symmetric risk limit, such as volatility, is inappropriate for an active management strategy. Either the degree of outperformance will be constrained in order to limit losses to an acceptable level or, more likely, losses will be insufficiently constrained so that a high expected return can be justified.

### Asymmetry in active hedging

If hedging is about controlling absolute currency risk, how much active risk should be allocated to a currency hedging programme, and how should it be measured? Consider this scenario. A currency programme is structured to hedge a foreign currency exposure back to the base currency, the euro. Let us take one currency in the exposure, the US dollar. If the US dollar appreciates, we would like the currency manager to keep out of the way and allow the fund to benefit from the move. But let us say that the US dollar falls by 10%. Now we would like to see a significant hedge, preferably up to the maximum permissible. If the currency manager is successful, then a 100% hedge would yield a return within the currency programme of 10% (ignoring interest rate differentials for now). What if the currency move were 20%? Now the same hedge would yield a 20% return, but the manager has done nothing different. Offsetting this, of course, is a corresponding loss in the underlying foreign assets, so the ability to implement a high hedge ratio is very important for the fund.

The volatility of the returns from the hedging programme is dependent on the underlying exposure. When currency translation losses are large, we need the hedging programme to generate large gains. Therefore, it is inappropriate to constrain an active hedging programme by reference to the volatility of its benchmark relative returns. Yet this is what many funds do in practice. Why do they

do this? Because they are really concerned about what can happen when the move is in the other direction, in which case a hedge is expensive and generates a highly visible negative cashflow.

The objective of an active hedging strategy is therefore to create asymmetric and potentially quite volatile cashflows. Ideally, these need to be very large when the base currency is strong and small when it is weak. If you apply a volatility constraint on the manager, the positive will be penalised as much as the negative and the manager will never be able to establish a sufficiently high hedge ratio to protect the fund in an adverse environment (strong base currency). The result is that, by limiting active risk, the portfolio is exposed to a higher absolute risk!

We therefore need to constrain an active hedging programme with a loss budget rather than a traditional volatility budget. Concepts like tracking error go out of the window (after all, we are not actually trying to track anything here), and we simply measure losses relative to a passive benchmark stance. A typical loss budget could be 3%. This would mean that the manager could underperform his benchmark by up to 3% in any one year, but there is no constraint on his upside or on the actual volatility of his returns. Clearly, an option could achieve this objective (3% being the premium), if only it were cheap enough to provide a sufficient hedge when in the money. The ideal active hedge therefore has an option-like payout, but at low cost.

### TIME HORIZON

Now that we are thinking of the probability of loss relative to a neutral benchmark position rather than to the volatility of returns, we can also avoid the problem of volatility not scaling in currency markets (or indeed in many others). In other words, the fact that annualised volatility using high-frequency data bears little relationship to actual movements over 12 months no longer concerns us. If our risk horizon is indeed 12 months (a reasonable assumption for most investors), we can say that the loss budget is 3%, or something similar, over a 12-month period. It is as simple as that.

By implementing this as a set of investment guidelines, it becomes the manager's job to assess the portfolio's position with reference to that budget. Once the portfolio starts to make some money, a greater reserve becomes available to support a larger position. In

this way, the portfolio can build a far greater hedge relative to the benchmark than any measure based on annualised *ex ante* tracking error or volatility would permit.

### DIVERSIFICATION

Investors have two types of risk control available to them: direct and indirect. The direct method is to increase or decrease an exposure. This can be achieved by changing the size of an investment or by adjusting leverage/hedging using derivatives. This rescaling impacts both risk and return. The indirect method is through portfolio diversification which seeks to control risk without diminishing return. But portfolio diversification is a lot harder than it seems. Diversification is a much misunderstood concept, yet it is so fundamental to much within risk management that it is worth exploring further.

It is often argued that the benchmark hedge ratio should be less than 100% or even that currency exposures should be left unhedged, on the basis that currency movements are uncorrelated with underlying asset classes and therefore provide diversification for the fund. We have to disagree, and not because we think there is some correlation with equity markets (although there may be, sometimes). The point is more fundamental even than that. Diversification in investment portfolios is achieved by replacing part of one risky exposure with a new risky exposure having a correlation of less than one. Introducing foreign currency exposure, however, replaces the risk-free domestic currency with risky foreign currencies. This is additive. Taking currency exposure does not reallocate risk; it adds risk.

This means that currency would only be a source of diversification if it were quite significantly negatively correlated with other assets in the portfolio. Most passive baskets of developed currencies show a relatively low level of correlation with asset returns. There is also strong theoretical and empirical evidence that their returns are indistinguishable from a random walk and therefore have no expected return. In this case they fail the diversification test on two counts: no expected return and a correlation that is not reliably negative.

Historic evidence shows that some markets that are heavily influenced by commodities, such as Canada or Australia, have seen negative correlations over quite long periods but this correlation has been unstable over time. Furthermore, these currencies have tended

to have high domestic interest rates, which have made foreign currency exposure more costly. So the benefit of a negative correlation with asset returns is offset by the cost of a negative foreign currency return. Note that a perfectly negatively correlated return is the same as a hedge, and this eliminates the asset return completely!

### The search for independent returns

Investors often state that they are looking for uncorrelated returns. In fact, they are looking for logically independent returns and correlation is just a statistic, which may or may not indicate independence. It is relatively easy to show that independence implies zero correlation, but zero correlation does not imply independence.

Suppose that an analyst concludes that the return on an asset follows a random walk. Imagine a derivative that pays the asset return with a one month lag. The expected return and volatility of the asset and the derivative must be the same, yet the random walk implies that their correlation is zero. It would not be sensible to follow the advice of a mean variance optimiser by investing 50% in the asset and 50% in the derivative.

Another problem is that the correlation coefficient attempts to encapsulate the connection between two sets of observations in a single figure. This approach throws away much of the information contained in a data set as it assumes a linear relationship. Any correlation figure needs to be treated with caution.

Insurance companies can expose themselves to genuinely independent events. For example, a car accident in Miami is unlikely to be connected to a house fire in Seattle. But financial markets are driven to a large extent by liquidity, and this factor is becoming increasingly global. The very act of adding a historically uncorrelated return to your portfolio makes it more highly correlated with your existing holdings. This is not a statistical quirk; it is causation. If you have an urgent call on liquidity, you will have to sell assets wherever they happen to be invested. This also applies to funds of funds and it is particularly likely to occur on the downside, when you need diversification the most.

When we consider currency alpha strategies it is often argued that a key benefit lies in the lack of correlation between currency markets and traditional asset classes. The logic being that currency

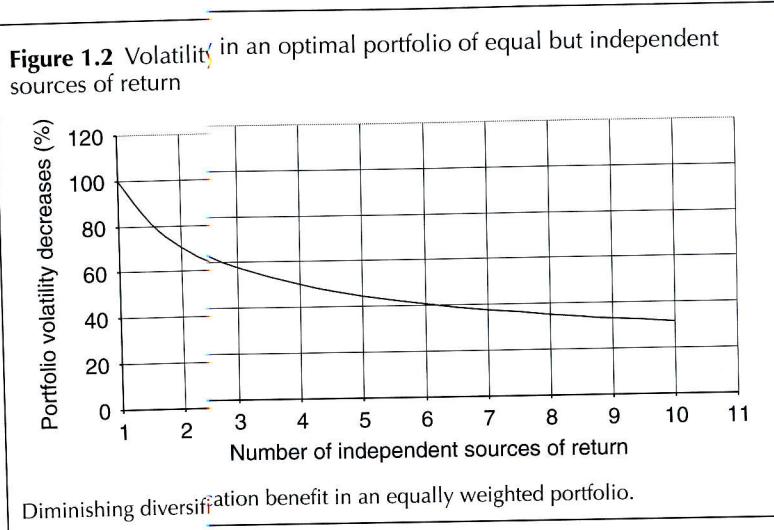
funds (as opposed to additive currency exposure) can therefore provide genuine diversification for a portfolio. A currency fund, however, is really a set of long/short positions, established and managed actively by the currency fund manager. In assessing the diversification benefit, therefore, we are really comparing a market return (the traditional asset class) with a pure alpha return. It is hardly surprising that these are found to be uncorrelated during normal conditions.

However, it is not sufficient to rely simply on this observed lack of correlation. An understanding of the actual currency management strategy is essential in order to judge whether or not this will be sustained in the future, particularly during a time of crisis. Consider what happened in July 2008 as an example. For several years currency funds and equity portfolios had lived happily together. Both had seen positive returns, but short-term movements had displayed low correlation. However, the currency returns were driven very largely by exploiting the interest differential between different currencies (the carry trade). This was a strategy that had become highly levered, and employed by many investors. When the subprime crisis hit, investors pulled back from (risky) equities and the stock market fell. There was then a universal move to reduce risk and this caused a dramatic rise in the Japanese yen (a low-yielding currency) and a corresponding fall in the high-yielding currencies such as Australian and New Zealand dollars. Over a very short period we suddenly saw a massive increase in the correlation between stocks and many currency funds (they both fell). A hidden link was revealed which had not been even hinted at by past historical analysis.

## THE LAW OF DIMINISHING DIVERSIFICATION

### Diversification in an ideal world

As investors tinker with their investment strategies, they inevitably encounter the problem of diminishing marginal returns. In economics, the concept is normally associated with production systems, where each additional unit of production yields less additional output. In investment, the unit of production is the ability to tolerate loss and the objective is to maximise gains per unit of risk. This is achieved by allocating a risk budget over independent sources of return. However, each additional independent return produces a



lower improvement of the Sharpe ratio. This is demonstrated by a simple example.

Suppose an investor finds an asset with the same expected return and volatility as the existing portfolio. Suppose further that the returns of this asset are totally independent of the existing portfolio. A mean variance optimiser would recommend allocating 50% to the existing portfolio and 50% to the new asset.<sup>1</sup> This would preserve the expected return, but scale down volatility by 29%, thereby boosting the Sharpe ratio by 41%. This is clearly a huge and highly attractive benefit of diversification.

However, suppose the investor were lucky enough to find another asset with the same expected return and volatility that happened to be independent of both the original portfolio and the first asset. Now a mean variance optimiser would allocate one third to each. The addition of the second independent asset would bring the overall volatility down to only 58% of the initial portfolio's volatility. However, the marginal reduction in volatility from adding the second asset would be only 18%, compared with the down-scaling of 29% from adding the first independent asset. Since the portfolio becomes more thinly spread, the benefit of diversification decays as you add more independent assets.

Figure 1.2 shows how this effect works in a mean-variance optimal portfolio as more independent sources of return are added.<sup>2</sup>

**Table 1.1** Correlation of monthly asset returns against US equities

	Non-US equities	EM equities	CRB commodities
1988–1997	0.46	0.42	-0.09
1998–2007	0.83	0.73	0.16

Volatility falls because the correlations between genuinely independent returns are zero, but the marginal improvements become smaller and smaller. Even in this idealistic case, it becomes progressively more difficult to benefit from diversification.

### No free lunch

Diversification is often touted as the “free lunch” in the world of finance. The approach has been exploited repeatedly in portfolio construction, insurance, business strategy and the packaging of CDOs. But, as the last example shows, over-reliance of diversification can lead to a significant under-estimate of risk. The magic of diversification relies crucially on the degree of independence between risk exposures.<sup>3</sup> We call this “the law of diminishing diversification” and it suggests that investors actually have to work hard to earn their free lunch. Maintaining diversification through time is tough because the correlation between return sources has a natural tendency to rise.

### Rising correlation

Evidence of the rising correlation between historically uncorrelated returns is plentiful. Tables 1.1 and 1.2 compare the correlation between assets and simple investment strategies over two five-year periods. Note that the availability of liquidity can drive the performance of passive strategies, such as the carry trade, as well as asset prices. The increasing availability and use of commodities is likely to have an important impact on the way prices behave, as liquidity shocks make sharp falls more frequent than observed historically. This rising correlation is strong evidence that the benefits of diversification decay over time.

More recently, particular strategies such as the carry trade as well as hedge funds have become better correlated with equities.

The search for independent sources of return becomes increasingly difficult as portfolios become more diversified. Fortunately, the

**Table 1.2** Correlation of strategies against US equities

	Carry trade proxied by AUD/JPY	Hedge fund index
1998–2002 <sup>2</sup>	0.20	0.39
2003–2007 <sup>7</sup>	0.37	0.64

financial markets constitute an open system, so new opportunities continually become available. In order to maintain diversification, investors must expose themselves to novelty and innovation.

#### Confusions

Investors can sometimes become blinded by correlation to the extent that they forget about return. There is no point in adding an uncorrelated exposure if it does not generate an attractive return. Currency risk is a good example. Few people would choose to be holding a passive equity market weighted basket of foreign currency exposures if it had not come as a by-product of an equity asset allocation decision. Betting on a series of coin tosses is another example of an independent return stream, but no one should add this to a portfolio, because it has no expected return.

Some investors confuse diversification with hedging. Beware of things that are negatively correlated with your portfolio; they are likely to deliver negative returns, assuming that you expect your portfolio to increase in value. The addition of completely negatively correlated exposure is called hedging and that is direct risk control. Hedging is necessary to reduce the impact of unintended or unattractive exposures.

#### THE CURRENCY BENCHMARK

What does all this imply for the benchmark in a currency programme? Let us first be clear about the distinction between hedging programmes and alpha programmes. For an alpha programme we start with nothing and try to generate returns by taking positions and generating net exposures in the currency markets. Very often no funding is required, in which case the obvious benchmark is zero return. If we make money, we are ahead. If funding is required, it will sit as cash collateral, in which case the obvious benchmark is the

passive return on this cash. So far so easy, but what about hedging programmes?

If the strategy is a passive hedge, at what level should this be fixed? If it is an active strategy, what is the position against which loss is measured? In other words, what should the fund's neutral position be relative to its currency exposure? The logical answer is that if currency exposure in itself provides no expected return and offers no diversification benefit, then it should be eliminated with a 100% hedge. An active strategy can seek to improve upon this by reducing the hedge during periods of base currency weakness.

The real world, however, is not as simple as this logic would suggest. There are two main reasons why a fund may opt for a different currency benchmark. We exclude the common misconception of diversification benefit already addressed.

#### Cashflow

Having a 100% hedge as the benchmark, or an active programme that gravitates towards it, will necessarily generate volatile cashflows. These may be highly negative (although there will then be a corresponding unrealised gain in the underlying assets). Negative cashflows can be highly visible and uncomfortable for a fund, especially if a new programme has recently been introduced. They can also be expensive to manage, in terms of both management resources and trading costs (assets will often have to be sold to finance them and then bought back when the cashflows are positive).

A lower hedge ratio reduces these costs and the volatility of cashflows, at the expense of more currency risk at the portfolio level. This trade-off is more extreme for a passive programme. An active programme can start with a lower hedge ratio as its benchmark, but can at least increase the hedge ratio when the base currency is strong, when the cashflows will be positive and therefore more palatable.

#### Opportunity cost and minimising investment error

Reducing risk in a portfolio is all well and good, especially if no return is sacrificed in the process, but where do you then spend that risk in order to actually improve returns? Sometimes an investor's base currency will be weak for a prolonged period of time. If the currency risk has been hedged, this will represent a significant opportunity cost. This could be important if a fund has taken a different

position from its peers or has a view on the direction of the base currency (not recommended but often encountered).

Such arguments can be strong enough to take the benchmark hedge ratio right down to zero. Again, with an active programme there is considerable protection around this decision because the manager is able to bring the hedge back up in the event of a prolonged period of base currency strength.

This concept of opportunity cost is important because, for many investors, it relates to the real risk. For them, the risk includes both cashflow loss and currency translation losses on assets. It is about the failure to capture an opportunity that would normally have resulted in a gain, which we can call "investment error".

Pursuing this line of reasoning, the selection of a currency benchmark is all about minimising investment error. It is an error to be unhedged when foreign currencies are falling and it is an error to be fully hedged when they are rising. If currencies spend half their time going up and the other half falling, both the fully hedged and the unhedged benchmark positions will be wrong half of the time. And a 50% hedged benchmark will be half wrong all the time! We reach the conclusion that, in the long term, all passive benchmarks suffer from 50% investment error.

This is a very different conclusion from the conventional volatility-based approach, because volatility ignores opportunity cost. The mean-variance paradigm is very hard to shift, but the fact is that hedging out a 10% gain has exactly the same impact on your wealth as incurring a 10% loss.

If, from the perspective of investment error, all passive hedging benchmarks are equal, how does this help in defining a currency hedging policy? The only way to reduce investment error is by implementing the hedging policy actively. A skilful active hedging manager can shift the hedge ratio towards unhedged when foreign currencies are rising, and towards fully hedged when they are falling. The benchmark determines the extent to which this can be achieved.

Benchmark selection can be framed in terms of a trade-off between upside capture and downside protection. An unhedged benchmark provides maximum capture of foreign currency gains and an active hedging manager can seek to reduce the magnitude of losses. In contrast, a fully hedged benchmark provides maximum downside

protection and an active hedging manager can seek to increase upside capture. In both cases, the manager's investment guidelines should specify a limit on the degree of benchmark relative underperformance.

## CONCLUSIONS

There is a weak syllogism which is often encountered, particularly in the currency world. It goes like this:

You cannot have an expected return unless you take risk. I am taking risk. Therefore, I have an expected return.

In other words, "You get paid to take risk".

This is a dubious assertion in many areas, but in currency it is simply wrong. You do not get paid, either with a return or with a diversification benefit. However, there is no simple passive solution. Currency risk needs to be constantly monitored and actively managed, and the returns from this activity need to focus on the probability of loss. The traditional use of volatility as a measure of risk is inadequate because it is a symmetrical measure and makes assumptions that are contradicted by the actual nature of the currency markets.

We would go further and say that the traditional approach is to reduce risk passively and spend risk actively; the sophisticated approach is to reduce risk actively and spend risk passively.

- 1 An equally weighted portfolio of independent assets has minimum variance.
- 2 The volatility of an equally weighted portfolio is inversely proportional to the square root of the number of independent assets.
- 3 Investors need to maximise the information entropy of their portfolios.

## REFERENCES

Okuyama, N., and G. Francis, 2007, "Quantifying the Information Content of Investment Decisions in a Multiple Partial Moment Framework", *Journal of Behavioral Finance* 8(3), pp. 121-137.