Optimal hedging of liability risk

Author

Stuart Jarvis
Strategic solutions group
Barclays Global Investors
Murray House
1 Royal Mint Court
London EC3N 4HH
E: **stuart.jarvis@barclaysglobal.com**
T: 020 7668 7029
F: 020 7668 6029

Abstract

This paper discusses investment and risk management strategies for pension funds. Using a simple mean-variance framework for understanding the risks and returns of alternative investment strategies, it is argued that the optimal strategy involves a full hedge of liability risks combined with a mix of optimal alpha and beta portfolios. Although the mix is determined by the plan's risk budget and/or return requirements, the alpha and beta portfolios are, in theory, independent of the plan. We show how the increased use of swaps by pension plans around the globe is enabling these strategies to be practically obtained in a cost-effective manner.

Contents

Introduction

Liability driven investment (LDI) starts from the position that the liabilities should play a central role in the process of setting investment strategy. Increasingly, this approach is being adopted by pension plans in Europe and the US - sources of volatility in the plan balance sheet are now being rigorously identified and managed.

Chief among these sources of volatility are the exposures to real and nominal interest rates that arise from the liabilities. These risks had previously been disguised by actuarial and accounting conventions but are now visible to all and they form the primary driver behind many an LDI discussion.

The natural question arises: how much of this liability risk should be hedged? For example, a typical UK pension scheme is in deficit, and takes significant risks on the asset side of the plan balance sheet, in pursuit of a return in excess of the liabilities. These risks are well understood and often diversified. Should interest rate risk be simply run as one of these many similar sources of volatility? This paper argues that interest rate risk should be treated on the same basis as the other portfolio risks but that the result of an optimal risk/return trade off should involve hedging the full liability.

More concretely, we show that the presence of the liability means that, in the absence of constraints, and with transaction costs ignored, the optimal strategy (in a mean-variance sense) consists of:

- The liability hedge; plus
- An allocation to an optimally diversified portfolio (with the allocation determined by the plan's return target or risk tolerance); plus
- The balance (positive or negative) in cash.

Crucially, the liability hedge is independent of the level of assets or the target return for the portfolio; and the diversified portfolio is independent of the structure of the liability (so, it is the same portfolio as would arise from an asset-only optimisation would produce, for example).

In practice, of course, the transaction costs can be significant and borrowing may not be possible. However, with the increased use of interest rate and inflation swaps by pension schemes, and their availability through pooled funds, these barriers are falling away. It is thus possible to construct portfolios that are significantly more efficient than was possible until recently.

The diversified portfolio can be further broken down into two pieces: a diversified beta portfolio, and a diversified alpha portfolio. Again, these two pieces are independent both of each other and also of the liability.

This 'full hedge' result can be surprising in light of related studies in the context of currency hedging, which show that the 'optimal' portfolio involves only partial hedging of foreign assets. We show that a 'partial hedging' approach could be more appropriate if a fund is more concerned with managing its funding level (assets ÷ liabilities) rather than its deficit (assets – liabilities).

This paper discusses, in turn:

- the optimisation framework required to investigate this problem;
- the impact of the long-only constraint on the optimal solutions;
- the relationship to optimal currency hedging

Risk return trade off

In determining investment strategy, most pension plans and consultants use the mean-variance architecture of Modern Portfolio Theory, due to Markowitz and others. This framework permits risk to be traded off against reward in an intuitive and mathematically tractable way. Further background to the model used here is given in Waring (2004).

In this section, we:

- Describe the framework and in particular set up the optimisation problem
- Solve the problem in the absence of constraints
- Discuss the solution and the way that it splits into three pieces.

In subsequent sections we

- Examine the impact of constraints and newer investment products on the solution
- Relate the result to work on optimal currency hedging

Model and notation

The aim is to model the plan asset value at time t, A_t. The universe of possible portfolios are characterised by their mean and variance. Given a set of basic asset classes, the plan asset allocation can be represented as a holding vector h^A of weights, and the mean and variance of A_t are[1]:

$$\mathbf{E}(A_t) = (h^A)^T \mu$$
$$Var(A_t) = (h^A)^T \Sigma h^A$$

Where μ, Σ are a vector and matrix giving the expected returns of the basic asset classes, and the covariances between these. Together, they encode the assumptions for the model.

It is often useful to further refine these assumptions:

- Split the risks and returns into market returns and active-management returns. Thus, the covariance matrix, for example, splits into market ("beta") risks Σ^β and active ("alpha") risks, Σ^α uncorrelated with the market risks.
- A factor decomposition for the market risks, $\Sigma_\beta = Y_\beta{}^T \Sigma^F Y_\beta$ is often used to simplify the construction of the beta risk model. $\beta^A = Y_\beta h^A$ then gives the factor exposures for the assets.

A pension scheme's liability L_t also has return and risk characteristics: again, these can be split into an exposure to beta factors, β^L and other uncorrelated factors (e.g. related to demographics and mortality).

Thus, the net return and variance become:

$$(Y_\beta h^A - \beta^L)^T \mu^\beta + (h^A)^T \mu_A^\alpha - \mu_L^\alpha$$
$$(Y_\beta h^A - \beta^L)^T \Sigma^F (Y_\beta h^A - \beta^L) + (h^A)^T \Sigma^\alpha h^A + \Sigma_{LL}^\alpha$$

The residual liability risk and return, $\mu_L^\alpha, \Sigma_{LL}^\alpha$ are not impacted by the portfolio holding vector, and so will fall out of the optimisation process, but they are retained here for completeness.

[1] For simplicity, this model is presented in discrete time. It is possible, with more complication, to construct a similar model in continuous time. For further details on the relevant definitions, see the appendix.

Optimisation objective

The form of the objective now follows the usual pattern: the idea is to maximise return for a given unit of risk, which is achieved by optimising

Return – (Aversion to market risk) × (Market risk) – (Aversion to active risk) × (Active risk)

Separate aversions to market and active risks are introduced here in place of the more usual single aversion. Background to this approach can be found in Grinold & Kahn (2000). The intuition is that unlike market returns, positive active returns require 'skill', and empirically investors are (and arguably should be) more averse to taking active risk than they are of taking market risk. An alternative approach could be to apply confidence levels to the active returns, and then take a Bayesian approach to blending these with market returns. This approach is taken by Black & Litterman. The results are similar, but their approach is probably more useful for a portfolio management problem than for a strategic asset allocation problem.

Of course, risk and return need to be measured at the plan level rather than just at the asset level, so that using the model and terminology above, the problem is to find an asset allocation vector h^A that maximises the objective function

$$\left[(Y_\beta h^A - \beta^L)^T \mu^F + (h^A)^T \mu_A^\alpha - \mu_L^\alpha\right] - \frac{\lambda_\beta}{2}\left[(Y_\beta h^A - \beta^L)^T \Sigma^F (Y_\beta h^A - \beta^L)\right] - \frac{\lambda_\alpha}{2}\left[(h^A)^T \Sigma^\alpha h^A + \Sigma_{LL}^\alpha\right]$$

This is less daunting than it seems: collecting terms, the objective function is seen to be a simple quadratic function.

Budget constraint

The budget constraint is of the form $v^T h^A = A_0$, where the vector v encodes the value of a unit allocation to each asset. Usually, v is a vector of ones, but more generality enables swaps to be more readily included in the asset universe.

The liability hedge

To simplify the objective function, and to calculate the solution more easily, it helps to define the liability hedge, as follows. The liability hedge is the portfolio h^L that minimises

$$Risk(h) = \left[(Y_\beta h - \beta^L)^T \Sigma^F (Y_\beta h - \beta^L)\right] + \left(\lambda_\alpha / \lambda_\beta\right) h^T \Sigma^\alpha h,$$

subject to $v^T h = A_0$. We shall later assume that $\Sigma^\alpha h^L = 0$, but for the moment allow for the possibility that the liability beta exposure can only be taking on some active risk. Then, if $v^T h^A = A_0$,

$$Risk(h^A) = \left(h^A - h^L\right)^T \left[(Y_\beta)^T \Sigma^F (Y_\beta) + \lambda_\alpha / \lambda_\beta \Sigma^\alpha\right]\left(h^A - h^L\right) + Risk(h^L)$$

The objective function can then be written

$$(h^A - h^L)^T \left(\mu^\beta + \mu^\alpha\right) + \frac{\lambda_\beta}{2}(h^A - h^L)^T \left[\Sigma^\beta + \lambda_\alpha / \lambda_\beta \Sigma^\alpha\right](h^A - h^L)^T + c,$$

where $c = (Y_\beta h^L - \beta^L)^T \mu^F - \mu_L^\alpha - \frac{\lambda_\alpha}{2}\Sigma_{LL}^\alpha - \frac{\lambda_\alpha}{2}(h^L)^T \Sigma^\alpha h^L - \frac{\lambda_\beta}{2}(Y_\beta h^L - \beta^L)^T \Sigma^F (Y_\beta h^L - \beta^L)$ is a constant term.

An optimal swap

Mathematically, it is tempting to minimise the quadratic objective function immediately by 'completing the square' – but this would not provide any insight into the solution. Rather than this *algebraic* approach, it is preferable to solve a related *financial* problem; a financial condition (no arbitrage) is then invoked rather than an algebraic one (non-singularity).

Set $\Sigma = \Sigma^{\beta} + \lambda_{\alpha} / \lambda_{\beta} \Sigma^{\alpha}$ and $\mu = \mu^{\beta} + \mu^{\alpha}$. Then, as in the asset-only case (see appendix) define an optimal swap s to be a portfolio that maximises - among h satisfying $h^T v = 0$ - the "Sharpe ratio"[2]:

$$h^T \mu \big/ \left(h^T \Sigma h \right)^{1/2} .$$

This maximum must exist, or else the model would admit arbitrage opportunities. The first-order condition for a maximum implies that, for any portfolio h satisfying $h^T v = 0$,

$$h^T \mu = \frac{s^T \mu}{s^T \Sigma s} s^T \Sigma h .$$

The optimal swap is not unique – in particular any positive multiple is also optimal; so we normalise by further demanding that

$$\frac{s^T \mu}{s^T \Sigma s} = 1 .$$

Solution in unconstrained case

In the absence of constraints, other than the budget constraint, the problem can now be solved. The objective function can be rewritten as

$$\begin{aligned} & (h^A - h^L)^T \Sigma s - \frac{\lambda_\beta}{2} (h^A - h^L)^T \Sigma (h^A - h^L) + c \\ & = c + \lambda_\beta^{-1} s^T \Sigma s - \frac{\lambda_\beta}{2} (h^A - h^L - \lambda_\beta^{-1} s)^T \Sigma (h^A - h^L - \lambda_\beta^{-1} s) \end{aligned} .$$

Which is clearly maximised when

$$h^A = h^L + \lambda_\beta^{-1} s .$$

Note that:

- The full liability hedge appears for all risk aversions
- The optimal swap is independent of the liability hedge; it is the same diversified portfolio of risky assets that any investor would add to their minimum-risk portfolio

Alpha beta separation

We now simplify things further, by assuming that the model permits alpha/beta separation. That is, for any portfolio h, there exists a 'beta-only' portfolio with the same beta characteristics. The portfolio h

[2] As normally defined, the Sharpe ratio for an asset is its excess expected return over the risk-free asset divided by its risk. As any nil-value asset can be expressed as the difference between a unique positive-value asset and the risk-free asset; the definition here is equivalent to the Sharpe ratio for that positive-value asset. The definition is slightly more general in that a risk-free asset is not required.

can then be written as a sum of a beta-only portfolio and an alpha-only portfolio. More formally, for every h there exists h^{β} such that

$$h^T \Sigma h = h^{\beta T} \Sigma^{\alpha} h^{\beta},$$
$$(h - h^{\beta})^T \Sigma^{\beta} (h - h^{\beta}) = 0,$$
$$(h - h^{\beta})^T \mu^{\beta} = 0,$$
$$(h - h^{\beta})^T v = 0.$$

Then $h^L = h^{L\beta}$ and the objective function can be split into two pieces:

$$\left[(h^{A\beta} - h^L)^T \mu^{\beta} - \frac{\lambda_{\beta}}{2} (h^{A\beta} - h^L)^T \Sigma^{\beta} (h^{A\beta} - h^L) \right]$$
$$+ \left[\left(h^A - h^{A\beta} \right)^T \mu^{\alpha} - \frac{\lambda_{\alpha}}{2} (h^A - h^{A\beta})^T \Sigma^{\alpha} (h^A - h^{A\beta}) \right] + c$$

The two square-bracketed terms can be separately maximised; and the optimum is therefore the sum of these two separate pieces. Thus, proceeding as before, the optimum is

$$h^A = h^L + \lambda_{\beta}^{-1} s_{\beta} + \lambda_{\alpha}^{-1} s_{\alpha},$$

Where s_{α}, s_{β} are optimal alpha and beta swaps respectively. In other words, the unconstrained mean-variance optimum consists of:

- The liability hedge; plus
- A multiple of an optimal beta swap; plus
- A multiple of an optimal alpha swap.

Where the multiples are determined by the investor's risk aversions.

Constrained optimisation

In practice, constraints are placed upon the set of possible portfolios, so that the optima described above may not be "feasible" investment strategies.

The most common constraint in practice is the long-only constraint. We show that, although this is not a serious constraint in the asset-only or no-liability case, it becomes a serious detractor from portfolio efficiency in the surplus optimisation case. Swaps, or swap-based funds, offer a means to introduce leverage into the liability-piece of the portfolio, and this turns out to be sufficient to dramatically improve the risk/return trade-off.

Leverage and asset-only optimisation

In the usual 'asset-only' optimisation, i.e. with no liability, the usual constraint placed is a 'no leverage' or 'long only' constraint. That is, the the allocation to each asset must be between 0% and 100%.

The graph below shows optimisation results for a UK investor, for a reasonable set of assumptions. The unconstrained efficient frontier is a straight line, starting from cash as the risk-free asset at the origin.

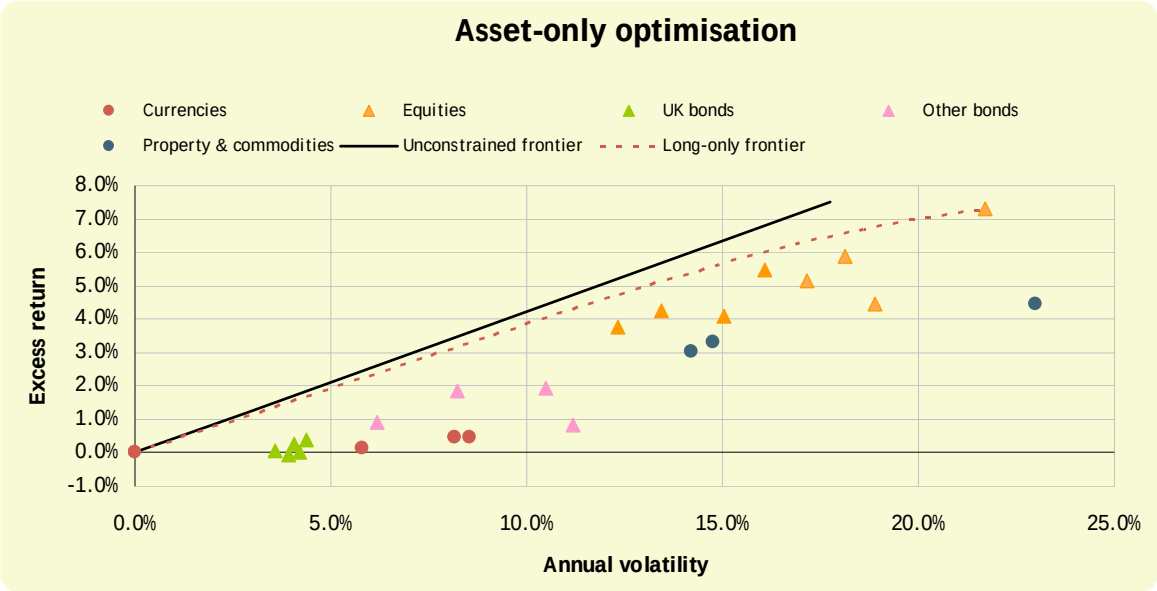

If a long-only constraint is imposed, then the dotted line is the result. It can be seen that the difference between these lines is small for a wide range of return targets. For example, a 4% excess return could be achieved with almost 1% less risk if leverage were permitted.

In fact, in practice, return assumptions are often derived from a long-only presumed-efficient portfolio so that the long-only and unconstrained frontier are actually coincident up to around 4% excess return. See Dopfel (2006) for a full discussion of the impact of leverage on portfolio efficiency in this asset-only case.

Leverage and surplus optimisation

When there is a liability, the situation can become quite different. As described above, even for low-return portfolios the optimal *un*constrained portfolio will tend to involve a liability hedge plus a return portfolio, financed by cash borrowing. This cash borrowing will tend to involve leverage at the portfolio level.

Using a typical UK pension scheme liability profile, the above risk/return plot can be recast in terms of return and risk relative to the liabilities:

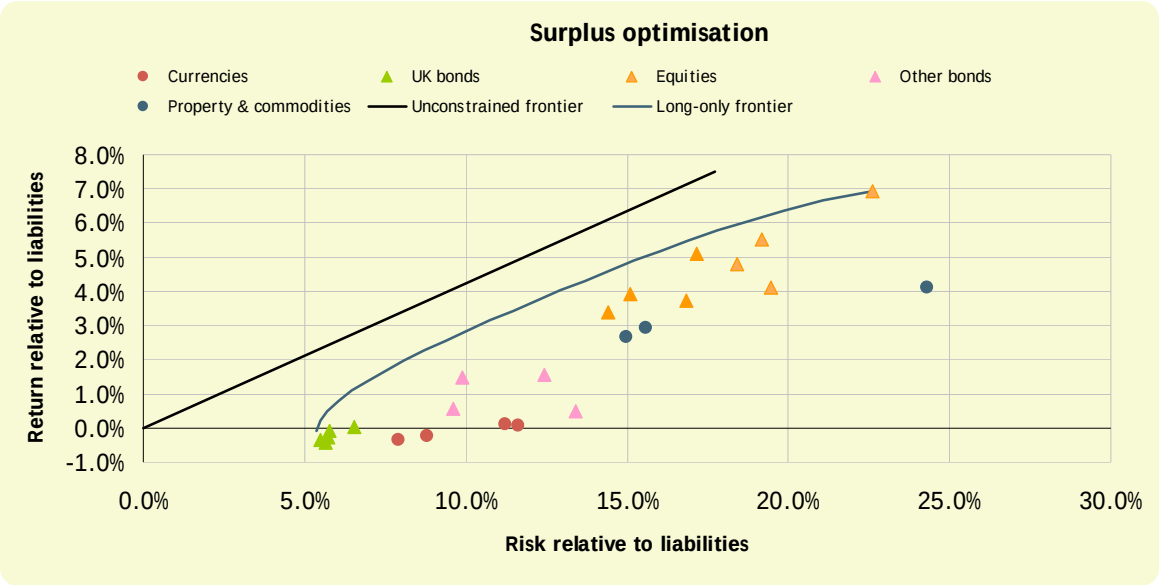

It can be seen that there is immediately a much greater gap between the unconstrained and the long-only frontier – risk is around 3-5% greater across a wide range of target returns. Alternatively, (looking vertically rather than horizontally) for a given level of risk the potential extra return from allowing leverage is over 1%.

From the preceding analysis we know that the problem is an inability to borrow to fund the liability hedge. If extra 'liability driven investment' (LDI) funds are introduced into the mix – these use swaps to access different pieces of the nominal and real term structures and in particular can have implicit leverage to achieve very much longer durations than the traditional bond indices included above – then it becomes possible to achieve a portfolio with characteristics close to the liabilities[3]:

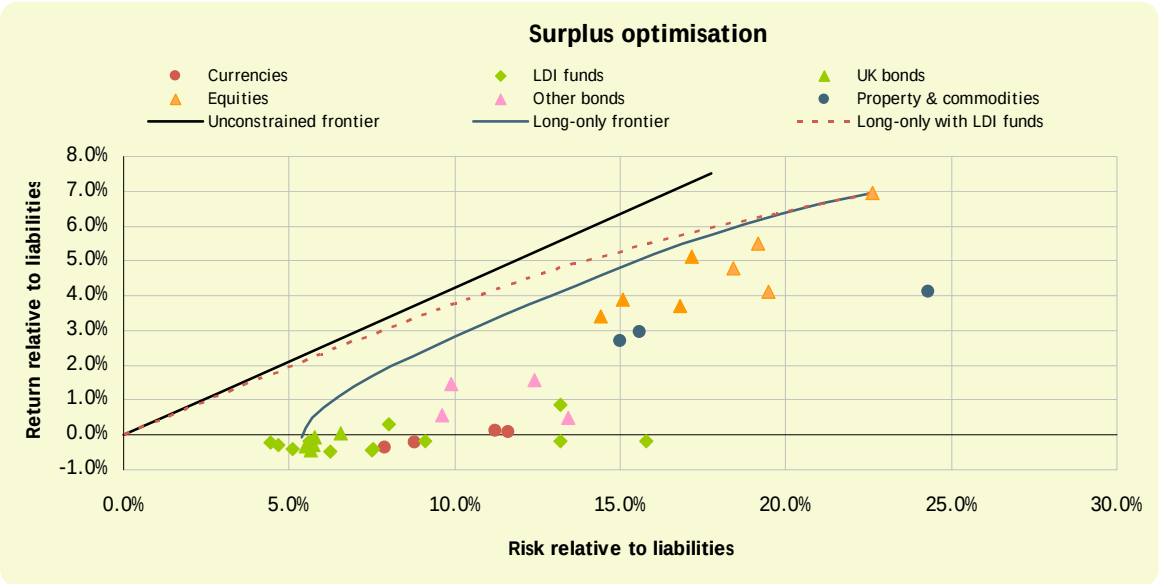

[3] The risks & return associated with non-market features of the liabilities (such as longevity) are explicitly ignored here. In theory the unconstrained frontier would be a hyperbola rather than a straight line, with a minimum risk equal to the volatility of these unhedgeable risks. The entire picture here would effectively switch to the right; but the relations between the different frontiers would be unchanged.

Expanding the investment universe to include swap-based instruments funds is equivalent to allowing leverage in the limited circumstance that this leverage is used to hedge the liability interest rate exposure. The long-only constraint is otherwise retained – this small relaxation is sufficient to bring the constrained frontier close to the unconstrained frontier (as close as in the asset-only case).

If swaps were allowed directly in the portfolio, then the liability hedge would be achievable with no initial capital outlay, and the analysis in the previous sections shows that this would reduce the surplus-optimisation problem to the asset-only optimisation problem. The LDI funds used above are not simply swaps, but they implicitly use swaps and some leverage (limited to 50%) to provide investment vehicles with long duration.

Comparison with currency hedging

The previous sections have looked at what happens when an investor's focus moves from a cash benchmark to a liability benchmark; we have argued that the only difference between the two, for a given risk / return target, is a full hedge of the liability.

Other authors, most notably Fischer Black, have investigated what happens to the efficient frontier when one moves from a cash benchmark in one currency to a cash benchmark in another. In contrast to the 'full hedge' here, efficient portfolios differ by a partial hedge, with the size of the hedge determined by the investor's risk aversion.

In principle, many assets can be used as an accounting currency – or numeraire - for measuring asset prices and returns. It need not be a currency like euros or yen but could be the price of a particular bond. This insight provides the impetus behind much of the modern approach to modelling fixed income markets, often combined with changes of measure linked to these different numeraires.

For an investor with a liability, a possible approach is therefore to use the liability value as numeraire. If the assets were measured instead in liability terms – put another way, if attention were moved from the deficit to the funding ratio – then Black's result would still hold. We briefly derive this result.

Finally we apply these considerations to the design of the optimal alpha and beta portfolios: although optimal beta portfolios in different currencies need to differ by a partial hedge; a full hedge is appropriate for an alpha portfolio.

Liability numeraire

Although the asset and liability values are normally measured in cash terms, and the difference (surplus or deficit) is the important item for plan sponsors, it is common for trustees to focus on the funding ratio (assets ÷ liabilities). This is equivalent to using the liability, rather than cash, as the accounting currency, or *numeraire*. The evolution of the asset value in this numeraire is then simply the evolution of the funding level.

The mean-variance problem above can be restated in this case: again the goal is to maximise

Return – (Aversion to market risk) × (Market risk) – (Aversion to active risk) × (Active risk)

But now the 'return' and 'risk' refer to changes in value of (assets ÷ liabilities) rather than the difference. With no liability – or, with a cash liability – this is visibly the same problem as before. But with a liability, the two-fund separation argument does not hold in quite the same way.

Continuous-time objective function

It is convenient to move to a continuous-time framework, in order to work with ratios of assets more easily. Suppose that an asset X has a price process, measured in sterling, $X_t^{£}$ such that

$$dX_t^{£} = X_t^{£} r_t^{£} dt + \mu_X^{£}(t)dt + \sigma_X^{£}(t).dW ,$$

where W_t is a (vector) Brownian motion process and the first term has been separated out so that $\mu_X^{£}(t)$ denotes the *excess* rate of return at time t. Then, following the same argument as in appendix one, albeit in a continuous time setting, the function

$$\mu_X^{£}(t) - \lambda \sigma_X^{£}(t)^T \sigma_X^{£}(t) ,$$

Is maximised, subject to a budget constraint, at a portfolio Q where

$$h_Q(t) = h_M(t) + \lambda^{-1} h_S(t).$$

Here, the portfolio M minimises the variance of instantaneous returns, $Var(dM_t^{£}) = \sigma_M^{£}(t)^T \sigma_M^{£}(t)dt$ (as measured in the sterling numeraire) at time t. The portfolio S is a portfolio with nil value at time t that maximises the Sharpe ratio $\mu_S^{£}(t) / \left[Var(dS_t^{£})\right]^{1/2}$ and has $\mu_S^{£}(t)dt = 2Var(dS_t^{£})$.

Unconstrained optimum with liability numeraire

Using the liability as numeraire means switching focus from the return & volatility of $X_t^{£}$ to that of

$$X_t^L = \frac{X_t^{£}}{L_t^{£}}.$$

As shown in appendix two, the function $\mu_X^L(t) - \lambda\left[\sigma_X^L(t)\right]^T \sigma_X^L(t)$ is maximised at

$$\begin{aligned} h_Q^L(t) &= h_Q(t) + \left(\lambda^{-1} - 1\right)\left(h_M(t) - h_L(t)\right) \\ &= h_L(t) + \lambda^{-1}\left(h_S(t) + h_M(t) - h_L(t)\right) \end{aligned}.$$

This can be interpreted in a number of ways:

- The 'optimal swap' portfolio, h_S, is not the same for an investor with a cash numeraire as for an investor using a liability numeraire: it has become $h_S + h_M - h_L$.
- In moving from a cash numeraire to a liability numeraire, the optimal portfolios vary by an incomplete hedge - the degree of completeness, $(1 - \lambda^{-1})$, depends on the risk aversion.

This is the same result as appears in theoretical work on currency hedging (see appendix and references).

Diversified alpha and beta portfolios

The fact that the hedge is not complete relates to the fact alternative currencies (i.e. short dated bonds in foreign currencies) are bona fide investments: they have risk and return characteristics. Therefore, in equilibrium, some exposure to these currencies should be included in an optimally-diversified portfolio.

Currencies are 'beta' assets, that is they confer exposure to market risk rather than active risk. If, as above, the optimal portfolio is split into a diversified beta and a diversified alpha component, then the partial-hedging is only relevant for the beta portfolio.

A diversified alpha portfolio, with no beta exposure, provides only 'active' returns in excess of the risk-free rate. A full hedge is needed to create a diversified alpha portfolio in a different currency – the hedged portfolio will then have no beta-risk from the perspective of the new currency.

Appendix 1: Asset-only unconstrained efficient frontier

This appendix summarises a proof of Tobin's two fund separation theorem. The form chosen here is that the unconstrained efficient frontier can be described in terms of:

- The minimum-variance asset amongst those that satisfy the budget constraint; *plus*
- Multiples of the zero-value asset with the highest Sharpe ratio (the optimal swap)

(Traditionally, the separation theorem is expressed in terms of a risk-free asset rather than a minimum-variance asset; this version is given to emphasise the symmetry with the liability case, where the liability hedge is not risk-free versus the liability.)

Terminology and objective function

The framework described in the main part of the paper is used here, but for this appendix the liability and the separation of risk into alpha and beta components, are not used. Thus:

- Portfolios are described by a vector h of allocations to basic asset classes;
- The expected return and variance of the portfolio are $h^T\mu$ and $h^T\Sigma h$; consistent in the sense the model is assumed to be free of arbitrage opportunities;
- Impose a budget constraint $h^T v = A_0$. Denote the set of portfolios h satisfying this budget constraint as $F(A_0)$.

The efficient frontier is the set of portfolios in $F(A_0)$ that maximise $h^T\mu - \frac{\lambda}{2}h^T\Sigma h$ for different values of the risk aversion parameter λ.

Building blocks: the minimum variance asset and the optimal swap

The minimum variance asset, m is the asset that minimises $h^T\Sigma h$ over $F(A_0)$. It follows that, for any portfolio $h \in F(A_0)$,

$$h^T\Sigma h = m^T\Sigma m + (h-m)^T\Sigma(h-m).$$

The optimal swap s (we use the term "swap" loosely here to denote a nil-value portfolio, ie elements of $F(0)$) is a portfolio that maximises, for $h \in F(0)$, the "Sharpe ratio"

$$h^T\mu \Big/ \left(h^T\Sigma h\right)^{1/2}.$$

Note that this maximum must be finite, since otherwise we have a nil-risk, nil-value portfolio with positive return, which would be an arbitrage. It follows that, for any portfolio $h \in F(0)$,

$$h^T\mu = \frac{s^T\mu}{s^T\Sigma s}h^T\Sigma s.$$

Note that the swap s is not unique: any positive multiple of s has the same Sharpe ratio, and if there is a nil-value combination of assets that has nil variance (and therefore nil-expected return by the arbitrage-free condition) can be added to s.

It is useful below to normalise s by focussing on the positive multiple s_1 that has:

$$\frac{s_1{}^T\mu}{s_1{}^T\Sigma s_1} = 1.$$

Efficient frontier

Consider the objective function $h^T\mu - \frac{\lambda}{2}h^T\Sigma h$ for some fixed λ. Write $h = m + k$; then:

$$\begin{aligned} h^T\mu - \frac{\lambda}{2}h^T\Sigma h &= m^T\mu - \frac{\lambda}{2}m^T\Sigma m + k^T\mu - \frac{\lambda}{2}k^T\Sigma k \\ &= m^T\mu - \frac{\lambda}{2}m^T\Sigma m + k^T\mu - \frac{\lambda}{2}k^T\Sigma k \\ &= \left(m^T\mu - \frac{\lambda}{2}m^T\Sigma m\right) + k^T\Sigma s_1 - \frac{\lambda}{2}k^T\Sigma k \\ &= \left(m^T\mu - \frac{\lambda}{2}m^T\Sigma m + \frac{\lambda^{-1}}{2}s_1{}^T\Sigma s_1\right) - \frac{\lambda}{2}\left(k - \lambda^{-1}s_1\right)^T\Sigma\left(k - \lambda^{-1}s_1\right) \end{aligned}$$

The terms in the brackets are constant, i.e. independent of k, and the final term is ≤ 0, with equality when $k = \lambda^{-1}s_1$. Thus, the objective function is maximised when $h = h(\lambda) = m + \lambda^{-1}s_1$.

Sure enough, then, the efficient frontier is the set of portfolios $m + \lambda^{-1}s_1$ for varying λ. An efficient portfolio consists of the minimum-variance portfolio (which varies linearly with the budget constraint A_0) plus multiples of an optimal swap s.

Note that if there is a risk-free asset then the ratio of risk to return of the optimal portfolio is simply the risk aversion parameter:

$$\frac{h(\lambda)^T\mu}{h(\lambda)^T\Sigma h(\lambda)} = \lambda .$$

Comments

Note that the two-fund separation result in this model:

- Does not depend upon the existence of a risk-free asset;
- Does not require the covariance matrix to be non-singular; although it
- Does require absence of arbitrage, which is a strong though sensible constraint, since without it there is no solution.

It is also worth noting that the mathematics benefits from not specifying returns and risks in percentage terms: in particular, looking at assets with nil value sheds light on the overall solution.

Appendix 2: Multi-currency modelling

This appendix summarises a continuous-time framework in which a multiple-currency economy can be explored.

Terminology and notation

Denote by $X_t^{£}$ the value of an asset X at time t in currency £.

Define an asset to be a *numeraire* if it has positive price and pays no dividends. With this definition, a currency £ is not a numeraire, but the money-market account for that currency, is: define

$$B(£)_t^{£} = \exp \int_0^t r_s^{£} ds$$

Where $r_s^{£}$ is the short rate for the currency.

Define the excess return $R_X^N(s,t)$ for asset X relative to numeraire N from s to t to be

$$R_X^N(s,t) = \frac{X_t^{£} N_s^{£}}{X_s^{£} N_t^{£}}.$$

Note that this is independent of the unit of account (here, £) in which the individual assets are priced. The rate of excess return, $dr_X^N(t) = dR_X^N(0,t)/R_X^N(0,t)$, will typically be of the form $\mu dt + \sigma . dW_t$ for a (vector) Brownian motion process W. The stochastic term can be quantified by calculating covariances: define

$$\sigma_t^N(X,Y)dt = Cov_t^N\left(dr_X^N(t), dr_Y^N(t)\right)$$

The "dt" term can be quantified by taking expectations:

$$\mu_t^N(X)dt = \mathbf{E}\left(dr_X^N(t)\right).$$

By a slight overuse of notation, for a currency \$, denote excess returns $R_X^{B(\$)}, dr_X^{B(\$)}$ as $R_X^{\$}, dr_X^{\$}$.

Switching currency: impact on risk & return

Suppose that M, N are two numeraires. Then $R_X^M(s,t) = R_X^N(s,t)/R_M^N(s,t)$ and so by Ito's lemma,

(A.1) $$dr_X^M(t) = dr_X^N(t) - dr_M^N(t) + \sigma_t^N(M,M)dt - \sigma_t^N(X,M)dt.$$

Taking covariances of this:

(A.2) $$\sigma_t^M(X,Y) = \sigma_t^N(X,Y) - \sigma_t^N(X,M) - \sigma_t^N(M,Y) + \sigma_t^N(M,M)$$

And, as a special case,

(A.3) $$\sigma_t^M(X,N) = \sigma_t^N(M,M) - \sigma_t^N(X,M).$$

Equations (A.1) and (A.3) can be combined to form the simpler relation:

(A.4) $$dr_X^M(t) = dr_X^N(t) - dr_M^N(t) + \sigma_t^M(X,N)dt.$$

(A.2) and (A.3) provide the means for moving from excess return rates and covariance matrix in one currency, to the covariance matrix and excess return rates in another.

Combining portfolios

Consider an investment Z consists of proportions x, y in two underlying portfolios X and Y, with the residual $(1-x-y)$ in a numeraire asset N. Assume that the portfolio is rebalanced frequently; then the excess return versus the numeraire N will be of the form

$$dr_Z^N(t) = x.dr_X^N(t) + y.dr_Y^N(t).$$

The process for Z is not quite $N.(X/N)^x.(Y/N)^y$ due to covariance effects, although it can be helpful to proceed informally as if it were. For example, if X is an asset whose base currency is sterling, then the same asset fully hedged into dollars is informally $X.(B(\$)/B(£))$ but more accurately it is the asset whose excess return process relative to the dollar numeraire $B(\$)$ is

$$dr_X^{\$}(t) - dr_{B(£)}^{\$}(t).$$

Optimal portfolios and portfolio betas

A portfolio Q is risk/return optimal for a numeraire N if it maximises the ratio

$$\frac{\text{return}^N(X)}{\text{risk}^N(X)} = \frac{\mu_t^N(X)}{\left[\sigma_t^N(X,X)\right]^{1/2}}$$

The optimality condition can be used to impose conditions on any other portfolio X: the ratio above for Q must exceed the ratio for a portfolio consisting of Q plus a proportion α in X (and $-\alpha$ in the numeraire):

$$\frac{\mu_t^N(Q) + \alpha\mu_t^N(X)}{\left[\sigma_t^N(Q,Q) + 2\alpha\sigma_t^N(Q,X) + \alpha^2\sigma_t^N(X,X)\right]^{1/2}} \le \frac{\mu_t^N(Q)}{\left[\sigma_t^N(Q,Q)\right]^{1/2}}.$$

In particular, this is true for positive and negative α and so the derivative of the left hand side at $\alpha = 0$ must be zero. That is:

$$\frac{\alpha\mu_t^N(X)}{\left[\sigma_t^N(Q,Q)\right]^{1/2}} - \frac{\mu_t^N(Q)\sigma_t^N(Q,X)}{\left[\sigma_t^N(Q,Q)\right]^{3/2}} = 0$$

So

$$\begin{aligned}\mu_t^N(X) &= \frac{\mu_t^N(Q)\sigma_t^N(Q,X)}{\sigma_t^N(Q,Q)} \\ &= \phi\sigma_t^N(Q,X) \\ &= \frac{\sigma_t^N(Q,X)}{\sigma_t^N(Q,Q)}\mu_t^N(Q) = \beta(X)\mu_t^N(Q)\end{aligned} \tag{A.5}$$

Thus, the expected return on a portfolio can be expressed either as a constant multiple of its covariance with the optimal portfolio Q or as a multiple of the expected return on Q, with the multiple being the portfolio's beta with respect to Q. In a CAPM context, the portfolio Q would of course normally be identified with the market portfolio. As described in the previous appendix, the parameter ϕ, the ratio of return to risk for the optimal portfolio, can be identified with the risk-aversion parameter λ.

Switching currencies: impact on optimal portfolio

The terms in (A.5) can be replaced with versions relative to a new numeraire M using the relations (A.2) to (A.4):

$$\begin{aligned}\mu_X^M(t) &= \mu_X^N(t) - \mu_M^N(t) + \sigma_t^M(X,N) \\ &= \phi\left(\sigma_t^N(X,Q) - \sigma_t^N(M,Q)\right) + \sigma_t^M(X,N) \\ &= \phi\left(\sigma_t^M(X,Q) - \sigma_t^M(X,N)\right) + \sigma_t^M(X,N) \\ &= \phi\left(\sigma_t^M(X,Q) - \left(1-1/\phi\right)\sigma_t^M(X,N)\right) \\ &= \phi\left(\sigma_t^M(X,\tilde{Q})\right)\end{aligned}$$

Where $\tilde{Q}$ is the portfolio whose excess return satisfies

$$dr_{\tilde{Q}}^M(t) = dr_Q^M(t) - \left(1-1/\phi\right).dr_M^N(t).$$

The portfolio $\tilde{Q}$ is therefore risk/return optimal portfolio for the numeraire M.

Comparing this with the discussion above on combining portfolios, we see that informally $\tilde{Q}$ is the portfolio $Q.(M/N)^{1-1/\phi}$ rather than $Q.(M/N)$ which would be the portfolio fully hedged from numeraire N to numeraire M. This "partial hedging" result, where the degree of hedging is independent of the currency (numeraire) is due originally to Fischer Black; see also the discussion in Grinold (1996).

Including active management

As described in the main paper, it is common to split risk & return into market-exposure and active-management components. For simplicity, we retain the notation (μ,σ) above for the market risks and add active return & risk components denoted by α,ω. Thus:

$$\mathbf{E}\left(dr_X^N(t)\right) = \mu_t^N(X)dt + \alpha_t^N(X)dt$$

$$Cov_t^N\left(dr_X^N(t), dr_Y^N(t)\right) = \sigma_t^N(X,Y)dt + \omega_t^N(X,Y)dt.$$

We work with numeraires that are 'beta only'. With this assumption, equations (A.2) and (A.4) remain unchanged, and are augmented by

$$\omega_t^M(X,Y) = \omega_t^N(X,Y).$$

As well as a portfolio Q that maximises the (beta) Sharpe ratio, we also assume an alpha-only portfolio Q^α that maximises $\alpha_X^N(t)/\left[\omega_t^N(X,X)\right]^{1/2}$. Let $\theta = \alpha_Q^N(t)/\left[\omega_t^N(Q,Q)\right]$. Then $\alpha_X^N(t) = \theta\omega_t^N(X,Q^\alpha)$ and

$$\mu_X^M(t) + \alpha_X^M(t) = \phi\left(\sigma_t^M(X,Q) - \left(1-1/\phi\right)\sigma_t^M(X,N)\right) + \theta\omega_t^M(X,Q^\alpha).$$

It follows that, although the 'optimal beta' portfolio transforms by a partial hedge as described above, an optimal alpha swap requires no hedge (if in fund rather than swap format, i.e. in the form of a fund benchmarked against LIBOR in the base numeraire, then a full hedge will be required).

Bibliography

Fischer Black, "Equilibrium exchange rate hedging", *Journal of Finance* July 1990

Frederick Dopfel, "Leverage and the limits of the possible", *Journal of Portfolio Management*, Spring 2006.

Bjorn Flesaker and Lane Hughston, "International Models for Interest Rates and Foreign Exchange", *NetExposure* 1997

Richard Grinold, "Domestic Grapes from Imported Wine", *Journal of Portfolio Management*, Special Issue 1996

Richard Grinold and Ronald Kahn, *Active Portfolio Management*, Irwin 1995 (2nd ed., 2000)

Richard Grinold and Richard Meese, "Strategic asset allocation and international investing", *Journal of Portfolio Management*, 2000

Kevin Kneafsey, "Wealth Destruction and Portfolio Construction", Internal BGI paper 2006

Harry Markowitz, *Efficient diversification of investments*, Blackwell 1959 (2nd ed., 1991)

James Tobin, "Liquidity Preferences as Behavior Towards Risk", *Review of Economic Studies*, 1958

Barton Waring, Duane Whitney, John Pirone and Charles Castille, "Optimizing Manager Structure and Budgeting Manager Risk", *Journal of Portfolio Management*, Spring 2000

Barton Waring, "Liability-Relative Investing II", *Journal of Portfolio Management*, Fall 2004.