Controlling the growth of your hedge

Draft Report of a Working Party on Dynamic Investment Strategies

Abstract

Whereas Modern Portfolio Theory seeks to determine optimal investment strategies over a single time interval, and is well known to actuarial practitioners, the use of its multi-period equivalents seems to have been restricted to academic actuaries. This paper sets out the basic utility-based theory used to study multi-period problems, and applies this theory in some simple pension and insurance contexts. Results from the finance literature that deserve to be better known are recovered, and some time is spent showing how investor preferences can be reflected in the utility specification and what this means for the optimal strategy.

Investors with a regulatory surplus (e.g., insurers, Dutch pension plans) have to manage their investment risk budget through time in order to ensure continuing solvency. This problem is akin to that of an investor seeking capital protection, a problem to which known solutions already exist.

The equivalent problem for investors with a deficit is of course more challenging. UK pension plans are the obvious examples for us here. They are under regulatory pressure to set out recovery plans, and the natural context for this question is a dynamic one, where rebalancing is possible during the recovery period. We seek to establish for simple cases the structure of the optimal investment strategies. For pension plans, we find that risk budgets should not be spent evenly over the recovery period, but should diminish over time; and that the use of simple option strategies can result in an improved risk/reward trade off.

Working Party members' details

The members of this working party are:

Matthew Barnes, Pension Corporation
Alvar Chambers, Insight
Anthony Earnshaw, Schroders
Stuart Jarvis, Barclays Global Investors
Richard McMahon, Friends First
Peter Ridges, Pension Capital Strategies

Controlling the growth of your hedge

1: Introduction

The working party has investigated dynamic investment strategies. These are strategies where the investment policy changes through time in some defined manner, in order to better meet an overall investment goal. Whereas Modern Portfolio Theory, which forms a suitable framework for single-period investment decisions, is by now well known to actuaries, the multi-period theory seems not to be. Yet, it is clear that the typical problems faced by our employers or clients are often multi-period in nature, and it is often the case that breaking the problem down into a series of single-period problems means that some opportunities for a better solution are missed.

The usual output from Modern Portfolio Theory is an efficient frontier of investment strategies: there is a hedge portfolio, a growth portfolio, and the choice of a particular point on this frontier (the amount invested in the growth portfolio) is a matter of investor judgement. This choice might typically be expressed in terms of either a return target, risk target or risk aversion. When it comes to updating this view at the next review (which might be between a week and three years later), should this judgement change? Describing a framework in which it is sensible for the point on the efficient frontier to move, and the manner of its movement over time, is one of our objectives here.

Dynamic strategies offer the opportunity to manage risk across time, but this also leads to a sudden explosion in the number of possible strategies. Fortunately, this is not a new problem in financial theory: quite the reverse. The dynamic investment problem for individuals has been the subject of a huge volume of financial theory, starting, in its modern form, with Merton (1969) and Merton (1971). While the initial model economies had constant interest rates, more recent work has moved this into more complex (stochastic interest rate) models, where the risks of assets and liabilities can be more fairly treated together.

A further theme of which we have seen more in recent years has been a desire to time entry into markets, particularly bond swap markets, as part of a de-risking strategy. We have reviewed some of the common approaches taken here, and sought to suggest efficient ways in which this can be achieved.

A final motivation behind our work was to provide a framework within which option strategies can be considered. There has been an increased interest among (the advisers to) pension schemes in the use of options to tailor the distribution of future funding levels. Although these can be fairly easily presented with plausible arguments in their favour, it is usually very difficult to see why a particular strategy has been chosen above another, unless perhaps the profit to the writer of the options is taken into account. We wanted to support the general thrust behind these option strategies, while providing a means whereby less expensive and potentially more liquid options can be used to meet the same objectives.

This paper therefore seeks to meet several goals:

- To provide a brief description of the utility-based approach often used to analyse multi-period problems, and to give an overview of some of the results in this significant vein of the finance literature;
- To provide examples from the finance literature, from life insurance and from pensions where a dynamic policy is often considered;

– To show how some of the insights from the utility approach can be applied to these problems; and

– To analyse the improvement that could result from the use of more dynamic investment strategies, and in particular the use of options within pension plans.

A word is perhaps in order on our title. As noted above, there is often a (low-risk) *hedge* portfolio and a (higher-risk) optimal *growth* portfolio. Determining the optimal dynamics for how the investments in these two portfolios should shrink or grow is ultimately our main purpose. In mathematical terms, the theory of dynamic optimisation is usually known as *stochastic control.*

2: Dynamic investment strategy

2.1. Single-period investment decisions

What is now known as Modern Portfolio Theory, as initially developed by Markowitz in the 1950s, provides a way to determine optimal strategies over a single investment horizon. The notion of an *efficient frontier* of investment policies was introduced – for different levels of risk aversion λ, a portfolio is identified that offers the optimal trade off between risk and return of the final wealth w_T.

$$\textit{Maximise} \quad \mathbb{E}\left(w_T\right) - \lambda \mathbb{E}\left(w_T{}^2\right) \tag{1}$$

When working over a single period, this problem can (completely equivalently and more commonly) be stated in terms of returns $R = w_T / w_0$. However when using multiple periods, with the option to rebalance, looking at returns can be misleading, as the risk aversion depends on the wealth level.

In a well-known 'separation theorem', the efficient frontier can be expressed as linear combinations of two portfolios: the minimum-risk portfolio and an optimal risky asset portfolio.

2.2. Expected utility

A more general version of (1) is to maximise the expected value of a utility function of wealth

$$\textit{Maximise} \quad \mathbb{E}\left(U\left(w_T\right)\right). \tag{2}$$

When each period is independent and identically distributed (i.i.d.) it is natural to hope that maximising this expected utility over each sub-period will result in portfolios that are optimal over the entire period. It turns out that this is only the case when the utility is in the CRRA (constant relative risk aversion) class:

$$U(w) = \frac{w^{1-\gamma} - 1}{1-\gamma}. \tag{3}$$

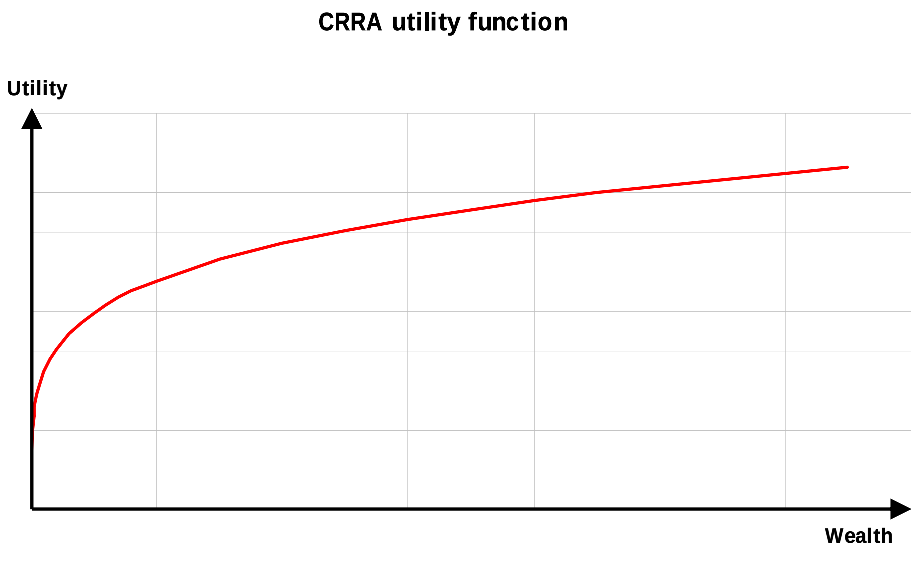

Surprisingly, although (1) is amazingly successful in the single period case, this does not generalise neatly to multi-period investment policies, and maximising a function of the form (3) becomes the more natural choice. The results of using (3) in simple models are also much more intuitive, of which more below.

Appendix A provides some brief background on utility functions often found in the literature.

2.3. Optimising over multiple periods

In a discrete setting, the dynamic problem can be solved sequentially starting with the final period: for given starting levels of wealth, the expected utility can be maximised given optimal expected utilities already calculated for levels of wealth at the end of the period. That is, the function:

$$V(w,t) = Max\ \mathbb{E}_t\left(U(w_T)\middle| w_t = w\right)$$

can be calculated iteratively starting at T, since $V(w,T) = U(w)$:

$$V(w,t) = Max\ \mathbb{E}_t\left(V(w_{t+1,}t+1)\middle| w_t = w\right).$$

This recursive equation is an example of a Bellman principle. The multi-period model is broken down, in a straightforward way, into single-period problems. In simple models, this 'dynamic programming' problem can often be solved explicitly; and more generally a numerical approach can be used.

At each stage (for each time interval), the problem is to find the asset strategy that will maximise expected utility. Mathematically, this strategy is known as a (stochastic) 'control' and the full problem as a stochastic control problem.

2.4. Continuous-time modelling

The continuous-time setting is used very commonly in the academic literature. This is of course less realistic than the discrete-time setting, and requires much stronger model assumptions and mathematical sophistication, but allows some impressive results to be derived. For example, a separation result very similar to the single-period result mentioned above can be proved.

One particular gain that the continuous-time framework can provide is due to Cox & Huang (1987). The dynamic programming problem is reinterpreted as two steps, each of which is considerably easier than the original problem, as follows. The utility-maximisation problem (2) can be interpreted as a problem to find a final wealth w_T for each possible asset price path ('state'). Convex optimisation methods (notably, the formation of a Lagrangian) can be used to determine this wealth function. Then as a second step, a search can be undertaken for an investment strategy that has this wealth function as payoff. The second step is familiar from option pricing theory, and in a complete model the existence of such a strategy is often assured by the martingale representation theorem.

2.5. A brief word on the use of simple models

We have restricted much of our analysis to models that are analytically tractable. This enables clean results to be derived, and intuition to be gained, but ultimately more realistic models are needed to make decisions in practical situations. We seek to identify general principles (e.g. on the risk budgeting process over time) that are expected to hold when more realistic asset pricing, discrete rebalancing, a reduced investment set and transaction costs are introduced. A final addition is to include non-investment strategy (e.g. contribution rate strategy; benefit improvement strategy) into the optimisation problem. Although consumption was an integral part of Merton's original work in this area (see below), this appears to raise the level of complexity considerably.

A particular problem that we want to avoid is the naïve use of complicated models. It would be easy to create a very complicated model for a series of assets and write some software to solve the corresponding dynamic program. But the result may be very dependent on the assumptions and constraints used in the program. Inconsistent assumptions may lead to arbitrage or near-arbitrage ('good deal') opportunities in the model. As an example of the impact of constraints, in recent years, we have seen examples of investment advice given to pension funds that is derived from models that only allow a small set of possible bond portfolios. Yet when long-dated bonds or swaps have been included, there has been a step change in the advice, as the capital efficiency of using swaps to manage interest rate risk (without disturbing the 'growth' portfolio) is suddenly appreciated. Here, therefore, we want to identify optimal strategies with as few constraints as possible, and only then proceed to see whether additional constraints are needed and, if so, what impact they have.

2.6. A special case of Merton's portfolio problem

Merton proposed the following problem. A CRRA utility function (2) and an initial level of wealth w_0 are given. The problem is to find investment strategy and consumption (or drawdown) c_t to be taken from the portfolio at each time t in order to maximise

$$\int_0^T U\left(c_t\right) dt + U(w_T) .$$

It turns out that in the case of a Black-Scholes model set up (constant interest rates, lognormal Wiener process for the single risky asset) the optimal policy is to invest a constant fraction of wealth in the risky asset and to consume wealth at a constant rate.

Similar results have been derived in more complicated models. For example, suppose we work in a model driven by Brownian motion where asset volatilities are deterministic functions of time. There is a portfolio process h_t that has the highest expected rate of growth $\ln\left(h_T / h_t\right)$ - up to a multiple, this is just the inverse of the pricing kernel (state price deflator) for the model. (In the Black-Scholes model this is a leveraged holding in the risky asset.)

If we ignore consumption, expected CRRA utility is maximised by a portfolio $X(\gamma, T)_t$ that invests $1/\gamma$ of wealth in the growth-optimal portfolio h_t and $(1-1/\gamma)$ in a bond maturing at time T .

2.7. Further reading

For a fuller introduction to the use of utility functions and dynamic investment strategy problems, we recommend Nielsen (2006).

2.8. Approach taken

In the remainder of the paper we focus on the choice of utility function. This function is supposed to express the investor's preferences and is therefore worthy of serious consideration. We show what plausible functions might look like, and identify the resulting investment strategies. As far as we know, this focus on the utility function is essentially new, and appears to us to offer interesting insights.

3: Example 1: Minimum solvency

As an initial twist to the usual CRRA utility function, we show how a minimum guarantee or minimum solvency can be imposed on the utility function. This leads immediately to two well-known methods of portfolio insurance: put options and Constant Proportion Portfolio Insurance (CPPI).

3.1. Introduction

Unit linked products with guarantees have proved popular with retail investors. They are somewhat similar to with profit business but crucially tend not to have smoothing and thus are more likely to be judged transparent. Policyholders gain exposure to equity markets whilst also benefiting from a guarantee on their money. The guarantee may vary depending on the product involved – it may simply be a return of premium or the guarantee may ratchet up over time depending on how the fund performs over the term of the investment. For this paper the nature of the guarantee is not important. The CPPI method of dynamic hedging is the most common way by which the guarantee is managed.

3.2. Modifying the utility function

Starting from a CRRA utility function, we examine two potential ways of introducing a minimum wealth level M . The first is simply to shift the utility function to the right:

$$U(w_T) = \frac{(w_T - M)^{1-\gamma} - 1}{1-\gamma}.$$

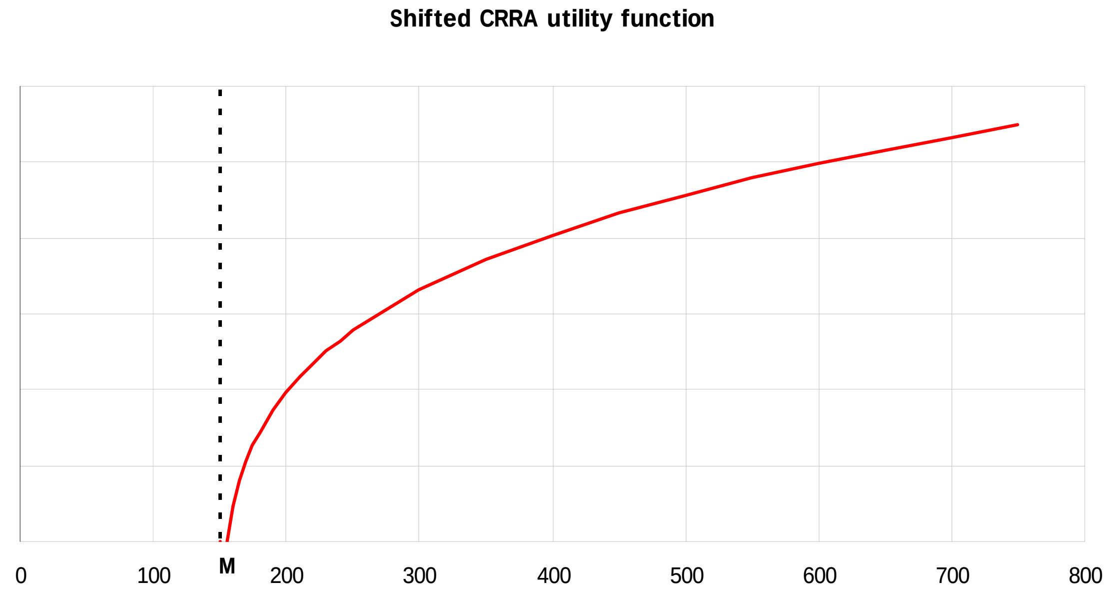

The second way is to truncate the utility function at the desired level (and have utility of $-\infty$ below this level):

$$U(w_T) = \begin{cases} \dfrac{w_T^{1-\gamma} - 1}{1-\gamma} & w_T \geq M \\ -\infty & w_T < M \end{cases}.$$

Truncated CRRA utility function

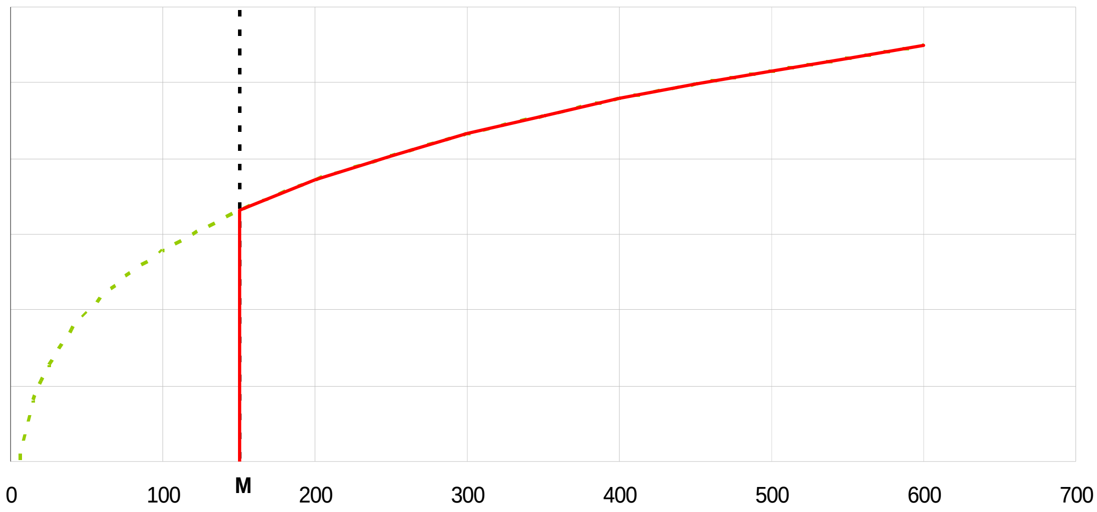

Maximising expected utility in either of these examples will lead to a final wealth that is always above the guarantee level M . However, the distribution of outcomes is quite different. In the first case, we see significant additional benefit for small increases in wealth above the guarantee level (the derivative becomes infinite as we approach the minimum level M). In the second case, there is only a small difference in utility from small increases above the minimum M (the derivative is finite).

3.3. Optimal wealth function

Denote by X_t the value at time t of an investment strategy that maximises the (unmodified) expected utility given by:

$$U(w_T) = \frac{w_T^{\,1-\gamma} - 1}{1-\gamma},$$

As noted above, in typical cases, X consists of a continuously-rebalanced mixture of the growth-optimal portfolio (with weight $1/\gamma$) and a bond maturing at time T (with weight $1-1/\gamma$). It is easy to calculate the result of maximising the modified expected utility functions in terms of this portfolio.

Case 1: Modifying the utility function (ie a shift right), the optimal strategy is of the form:

$$M.P(t,T) + \xi X\left(\gamma,T\right)_t \qquad (4)$$

for some ξ determined by the initial wealth. $P(t,T)$ has the usual meaning of the value at time t of a payment of 1 at time T. Note that if the initial wealth is insufficient to purchase the value of the guarantee $M.P(t,T)$, then the problem is insoluble, so $\xi > 0$). Splitting X into its components, this then amounts to investing ξ/γ of $\left(w_t - M.P(t,T)\right)$ in the growth-optimal portfolio, and the rest in the bond expiring at time T . In other words,

$$\text{Proportion in risky asset} \propto w_t - M.P(t,T).$$

The constant proportion is called the *multiplier*, and the simple investment strategy that we have derived is known as CPPI.

So the investment strategy that is optimum for a shifted utility function is CPPI.

Case 2: modifying the utility function (ie a truncated utility function), the optimal portfolio consists of an investment ξX_t in the portfolio X and ξ put options on X, with strike price M/ξ. Again, ξ is determined by the initial level of wealth.

So the investment strategy that is optimum for a truncated utility function includes put options.

3.4. Comparison

In the CPPI case, the final portfolio value is always at least *M*, with the distribution of possible values being close to *M*. In the put option case, the final portfolio typically has a high chance of being exactly *M*, with the residual distribution being more widely spread than the CPPI case.

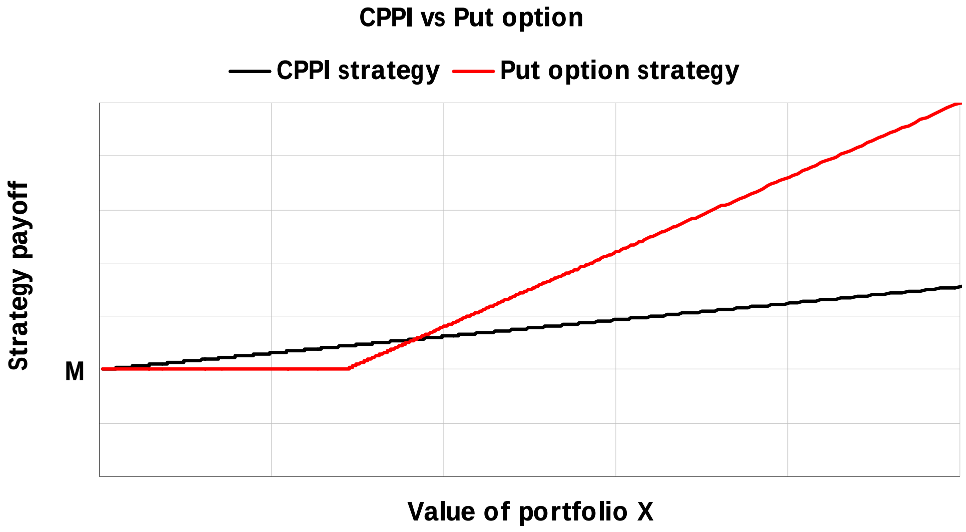

At the risk of stating the obvious, each of these strategies is optimal for the particular utility function that was used to derive it, and sub-optimal for the other. An example of the expected utility calculated under these two functions is shown below for a variety of funding levels (initial asset value ÷ value of guarantee):

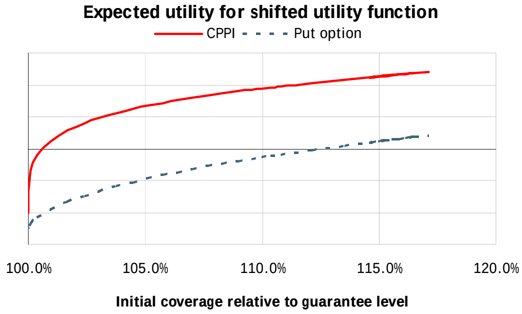

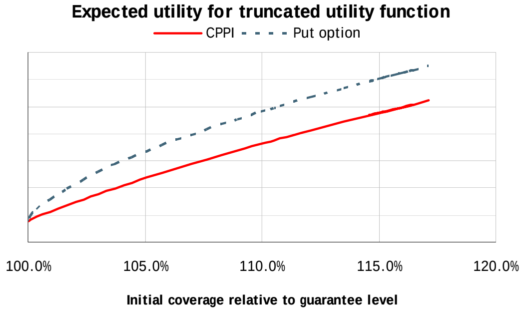

At each funding level, the same two strategies are being used, but the two graphs show how they measure up against the two expected utility measures. In each case the alternative strategy is a long way from being optimal.

As a further comparison, we note that CPPI is much more common in practice. This seems to be largely due to the following two reasons:

– *Transparency*. The CPPI strategy is very simple to explain, and therefore for investors to understand. The CPPI strategy also offers a level of transparency for the writer of the product e.g. a shareholder selling unit linked policies with premium return guarantees. In volatile markets the shareholder who manages these guarantees in-house can see what is happening with the asset mix in the fund. This is easy to reconcile and explain to policyholders. In contrast put option pricing is much less transparent (although the final payoff is of course much easier to define and explain!). This problem is compounded by:

– *Option availability*. The portfolio X is a typically a leveraged investment in a number of asset classes. Liquid option markets tend to be available on individual assets or asset indices, but not on portfolios, and are typically available over short periods rather than offering protection over the 3+ years that an investor might be seeking.

For long-term protection, a CPPI style strategy therefore tends to be preferred in practice. The next section collects together some outline considerations in connection with CPPI strategies.

4: Constant proportion portfolio insurance

CPPI was derived as an optimal strategy under a particular utility function in the previous section. We set out here to put a practical operational background to the theory outlined so far.

4.1. Sample Product Structure

- Usually single premium savings/pension business but can also be regular premium savings/pension business.
- There exists usually a selling period after which the product is closed to new money.
- Guarantee can be initial premium only or;
- Guarantee may ratchet up over time according to some predefined formula e.g. 50% of the maximum growth to date is locked in on pre specified lockin dates.
- Asset mix usually comprises of 'risky' assets such as equities and 'non risky assets' such as cash and government bonds. If the guarantees are managed in-house liquidity of the assets used is key.
- Fixed term – say 5 years. Again for simplicity of management the guarantee would only apply to policyholders who stay in the fund until maturity.
- Surrenders may be allowed.

4.2. How dynamic hedging works

A cost is incurred to the party writing the business if at maturity the fund is insufficient to meet guarantees. At any point in time the fund's asset mix is set so that if an immediate market fall of **up to** a pre specified level occurs there will be sufficient assets remaining to meet the amount promised at maturity. This assumes all assets are moved to bonds/cash once the pre specified market fall happens. A very simple example may help illustrate:

Single premium = 100

Guarantee in 5 years time = 100

Current unit price = 1

Current number of units = 100

Current interest rates/5 year bond yields = 4.5%

Value now of 100 guaranteed in 5 years = 100 * 1.045^(-5) = 80.245

Maximum pre specified market fall = 25% (inverse of this the multiplier, ie the multiplier is 1/25%). This is subjective and depends on the level of risk the writer of the policy is willing to take.

Equity content = (Fund value – value now of guarantee) × (1/25%)

= (100 – 80.245) × 4 = 79.02

Bond/cash = 100 – 79.02 = 20.98

Lets assume the maximum fall in markets we think will happen actually does take place:

If equity markets fell by 25 then the fund value would be:

= 79.02 × (1 - 0.25) + 20.98 = 80.245

The asset mix is then rebalanced. The equity content here will be 0 (fund value of 80.245 – guarantee value of 80.245) × (1/25%): the 80.245 will be invested in bonds which will be sufficient to meet the guarantee in 5 years' time at current rates of interest.

Obviously this rebalancing does not happen only when only a fall of 25% in equity markets occurs. Any market fall less than 25% will result in equities being sold to bring the maximum potential loss that can be incurred back to 25%. Any market rises will result in equities being purchased to also bring the maximum potential loss that can be incurred back to 25%. How often this rebalancing is done depends on the volatility of markets and the company's risk appetite and risk control structures.

We can illustrate the strategy by graphing the return expected from the portfolio as the funding level ($w_t/(M.P(t,T))$) varies:

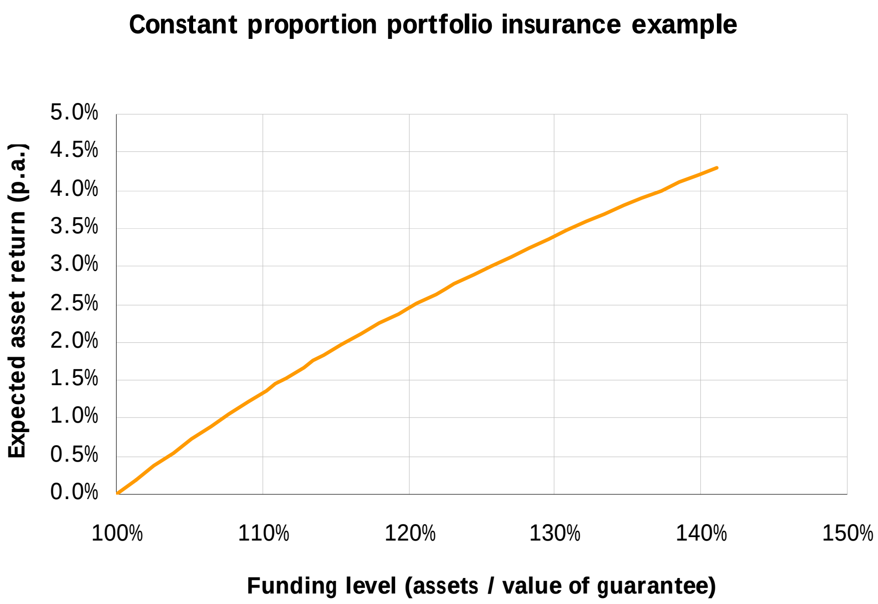

This graph does not depend upon time to maturity: the amount of risk taken at any point depends purely upon the funding level. As the funding level increases further, the risk and return approaches a maximum level determined by the risk aversion parameter in the utility function (more concretely, it is the expected return on the risky asset multiplied by the multiplier; around 15% in the case graphed above). This would of course require leverage, and a 'vanilla' CPPI product would allow at most 100% allocation to the risky asset. (In fact many CPPI products do now have a maximum allocation of more than 100%, but only up to 150% or so.)

4.3. The shareholder guarantee

We have seen that the objectives of the policyholder can for example be described in terms of a shifted CRRA utility curve. This expresses the policyholder's desire for the guarantee implicit in the CPPI structure. If the curve were not shifted, then it would be more optimal for an investor to buy a policy without a guarantee – a unit-linked fund, probably with no dynamic asset allocation.

Although the fund will very likely be managed in accordance with the CPPI method, the shareholder is free to fund the guaranteed payment at maturity at other ways that may make sense.

The detail of this is unimportant to the policyholder - the terms of the CPPI policy will define the payout at maturity.

Parallels can be drawn between this and a traditional with profits approach. The latter, however, is complicated by subjective smoothing rather than being subject to formulaic definition. Both CPPI and with profits have common variants that allow the guarantee to be increased annually to partially lock-in positive performance. This may lead to a different dynamic asset strategy as the increasing guarantee will rarely permit the strategy to gain a very high proportion of risky assets. We consider only the simple case here.

A further parallel may be drawn with insurers providing non profit endowments or annuities where a guaranteed benefit is to be paid and a solvency margin held above the level of assets required to pay the guarantee in a best estimate scenario. The insurer will wish to maintain a level of free assets above the minimum solvency margin and run an investment strategy unlikely to prejudice that margin, as it will be forced to top-up the margin by raising new capital, or be forced to cease writing new business.

The insurance company must decide to either lock-out risk as markets fall, to protect the guarantee, or to use new capital to permit continuation of a risky strategy. The CPPI policyholder has a similar option, as he could surrender the policy that may have become almost risk- (and return-) free and purchase a new policy which would, of course, include a lower guarantee (or require a greater premium for the same guarantee) but would reintroduce some upside potential. It is interesting also to observe regulators sometimes being willing to relax such margins in stressed market conditions to avoid the insurers being forced sellers of risky assets (and forced buyers of risk-free assets) in the way that the CPPI manager is prescribed to do.

Rather than taking the policyholder's perspective, we can take the perspective of a company, or its shareholders, where the company has written a CPPI policy. What is the optimal investment strategy from their perspective? Clearly they have promised the payout from a particular dynamic strategy, but subject to a guarantee, which could be called upon (albeit extremely unlikely in a simple model).

A fund value below the cost of guarantee is a cost to the shareholder. We see this as giving rise to negative utility. The shareholder would prefer not to incur this additional cost and use this money for something else instead of paying for guarantees. The cost is clearly just the difference between the fund value and the value of the guarantee.

Above the guarantee level, the bulk of the benefit of increases in the fund value accrues to the policyholder. There may be a small increase to the shareholder due to fees being payable on the fund value. If the fund is managed in line with the CPPI strategy, then the residual utility curve for the shareholder will therefore look something like the graph below.

4.4. Shareholder Utility Curve – guarantee managed using CPPI

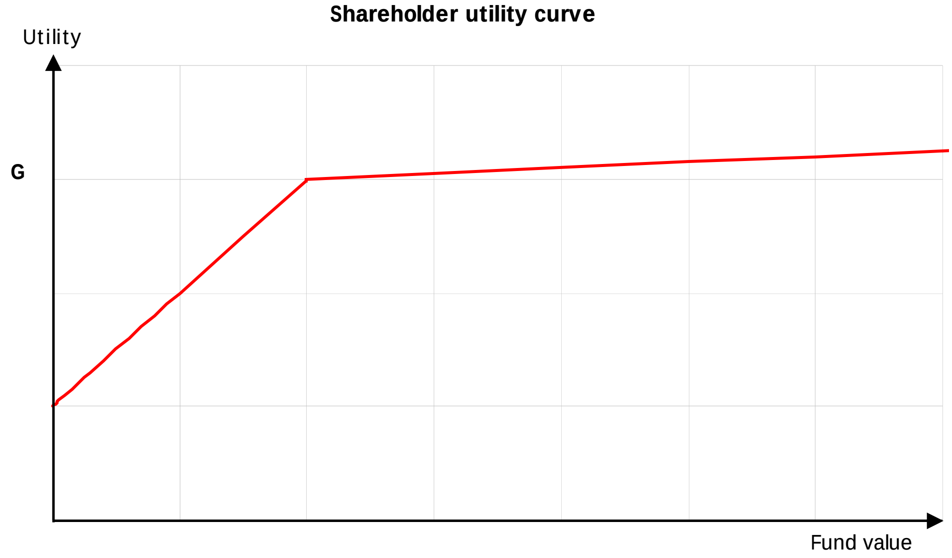

4.5. Further points

CPPI strategies suffer a notorious feature: when markets are rising, the buffer will rise and more of the risky asset will be purchased. Conversely, when markets are falling, the risky asset will be sold. If the trades are sufficiently large, this “buy high / sell low” strategy can magnify the market movements. For example, this effect was widely blamed for exacerbating the stock market fall on ‘Black Monday’ in October 1987. An investor may also dislike this buy high / sell low feature if they believe that markets exhibit short-term mean reversion.

5: Practical issues of providing CPPI products

Asset liquidity.

The primary aim of any shareholder writing this business is to protect itself from having less in the fund at maturity than is promised to the policyholder. The ability to sell out of equities when required
and buy bonds to match the term of the guarantee is key to managing the process. Restrictions on asset sales and a lack of availability of suitable bond can make the process unmanageable.

Pipeline business post sales period.

It is easier to manage these guarantees using CPPI if there is a fixed selling period. Sales staff may push for policies to be issued post the end of the selling period. Their priorities are not protecting the shareholder or managing the guarantee. Policies issued at a higher price than the guaranteed unit price used to monitor guarantee exposure are essentially getting a higher guarantee than is being managed. This is an exposure for the company. Strict communications to sales areas is needed to ensure this is not permitted.

Fund based commission and management fees.

It needs to be clear if these are guaranteed and will continue to be taken from the unit fund if the fund goes all into cash/bonds. If so then the initial equity content of the fund will be reduced considerably. Fund based commission and other guaranteed charges can be allowed for in the discount rate used to value the guarantee. Again clear communication to sales areas and pricing teams is required.

Operational risk

The management of guarantees using CPPI can be a manual process depending on systems that are in place, the volume of business written and the nature of the guarantees offered. Regular rebalancing involves manual changes to asset mixes. Therefore necessary internal audit checks and consistency checks need to be built into the process. These need to be monitored and regular breaches reported.

Reinvestment risk.

It is preferable to use gilt strips to hedge the guarantee. They remove the reinvestment risk on future
coupons and also meet the liquidity requirements. Fixed deposits with investment banks do not meet the asset liquidity requirements outlined above and also carry a risk premium and so are not risk free in nature.

Hedge in house or hedge externally.

If hedged externally the credit rating and level of service from the counterparty need to be considered. This is especially relevant in volatile market conditions where frequent valuations are needed. Also the cost of external hedging needs to be considered. This cost needs to be factored into the initial price.

Concentration risk

Multiple tranches of the product investing in the same underlying open equity fund could mean the pricing basis of this fund may move from a buying basis to a selling basis very often especially in volatile markets or downward markets. The buy sell spread in the fund will have a bearing on whether this is significant issue or not. Frequent movement of the basis is not ideal from an operational viewpoint.

Risk factor used

The risk factor used will depend on the writers risk tolerence levels and whether they are willing or not to set up maturity guarantee reserves to offset scenarios when guarantees bite.

Another factor to consider is the existing solvency regulation which may lead to a further cost of capital which potentially needs to be factored into product pricing. Current products could be set up with a risk factor of 25%. However the proposed new Solvency II regulations will require an equity shock of 32%.

Amount hedged

Invested premium or office premium? Invested premium is net of allocation and commission. Usually, but not always, office premium will be the minimum guarantee offered at point of sale to potential policyholders. This is easier to sell and market to policyholders. One needs to ensure that the guarantee offered policyholders is what is used in the regular rebalancing calculations.

Frequency of rebalancing.

Rebalancing in the extreme may have to be done daily in volatile markets. Weekly rebalancing should usually be fine otherwise but again this depends on the risk appetite of the shareholder and also on what was communicated to policyholders at point of sale along with resources available.

Bomb out scenario

This is when 100% of policyholder money is invested in bonds and cash. Do we offer the customers back their money? Do we offer them an alternative unit linked product without guarantees? Have policyholders reasonable expectations been met? Again adequate communication at point of sale is key to addressing this risk.

6: Example 2: Aiming for a liability target

We show how a similar approach can be used to analyse the situation where an investor is seeking to generate return to achieve a liability target. In a sense this is the opposite problem to the minimum solvency example. This section is preparatory to the next section where, although the economics become more complex, we examine the challenge faced by typical UK or US pension plans.

6.1. Introduction

Techniques of portfolio insurance have been well studied since initial work was done in the early 1980s, and as we have seen good solutions exist to help meet guarantees. In contrast, when under-funded relative to a known payment or stream of payments, there is of course no (investment) strategy that will bring the investor up to a full level of funding with certainty. The question of how best to take risk in order to achieve this target is therefore both a more urgent and a more interesting problem.

We set out a simple utility function that can be used to formulate the problem. The optimal solution can again be framed as an option strategy. However, as noted above, the option is unlikely to be available in the market at a reasonable price, and we therefore consider in some detail the performance of alternative strategies that could be used to approximate this optimal one. This emphasises the benefits to be had from regular rebalancing and the use of simpler (more liquid) options.

6.2. Modifying the utility function: no additional utility from a surplus over a target level

Rather than focussing on a minimum wealth constraint, a more relevant example involves satiation: more wealth is good up until a target is reached, at which point additional wealth generates no additional utility.

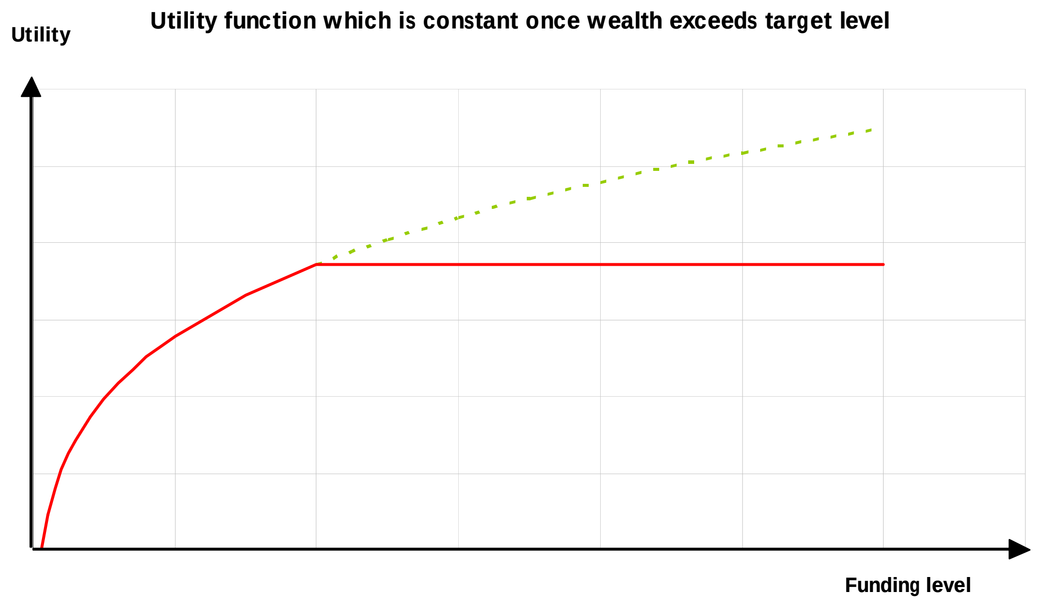

The above picture shows a simple function with these properties, of the form:

$$u(w_T)=\frac{\min(w_T,L)^{1-\gamma}-1}{1-\gamma}.$$

Here L need not be certain: for example, it could be the uncertain value (at time T) of payments due after time T. Indeed, mathematically, this makes any analysis slightly clearer, as the separation of the optimal strategies into three pieces (a liability hedge, an optimal risky portfolio and a bond maturing at the target horizon) becomes more evident.

This utility function gives no credit for having a funding level of more than 100% relative to the target wealth. Low funding levels are penalised without rewarding high funding levels: an asymmetric view of risk is the effective result. In a way, this is similar to the minimum-solvency examples, but crucially differs in that the target level is now above the current asset level, whereas in the minimum solvency cases the guarantee level was below the current asset level. A more flexible boundary was therefore required in this case than in the minimum case, where a hard constraint was imposed.

6.3. Optimal strategy: sell out of the money call options

It is straightforward to show that the final result from an optimal strategy can again be expressed in terms of the solution X to the unmodified utility function (see discussion above in connection with minimum solvency). The portfolio consists of an investment ξX_t in the portfolio X and selling (writing) ξ call options on X, with strike price L/ξ. As before, ξ is determined by the initial level of wealth.

The intuition here is simple: below the target level, returns should be generated in an optimal way, but any returns that bring the funding level over 100% are not required, so these can be sold (by writing call options) – and the proceeds used to generate slightly higher returns.

Of course, as noted previously in connection with downside protection using put options, the options required here are unlikely to be readily available at a reasonable price in the market. But purchasing an option is only one optimal strategy. A dynamic strategy (delta hedge) should in theory be able to replicate the same payoff as the option. Of course, in practice, transaction costs and market jumps mean that perfect replication is neither possible nor desirable, but an approximate delta hedge based on infrequent rebalancing should come close to optimality. Studying this approximation is the task to which we next turn.

6.4. Annual stock/bond rebalancing strategy

The theoretical continuous delta-hedging strategy involves taking less risk as the funding level relative to the target rises, and taking more risk (up to a limit that would obtain if there were no target) as the funding level falls.

The same structure persists if we optimise based on annual rebalancing rather than continuous rebalancing. The following graph illustrates this in the case where only two non-cash assets are available: a risky stock and a liability matching bond, and where only annual rebalancing is possible.

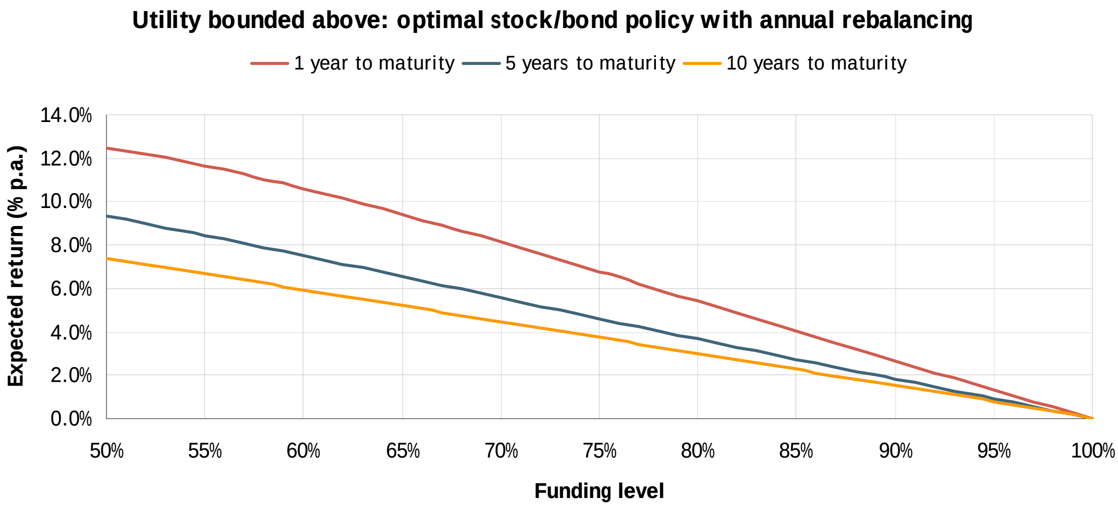

In contrast to the minimum-solvency example, the amount of risk taken does now depend upon the time to the target horizon. This makes sense of course: the longer the time left, the less risk is required to have a good chance of meeting the target.

Interestingly, the risk budget is not expected to be evenly spread over time. This can be read off the previous graph: for example, with a 85% funding level, and 10 years until the target date, a return of less than 2% per annum over 10 years is sufficient to raise the funding level to 100%. But the optimal strategy suggests a higher return in the first year. Conversely, for a one year horizon, the return is insufficient to bring the funding level up to 100%. This is illustrated below – where each year we calculate the expected funding level for the following year:

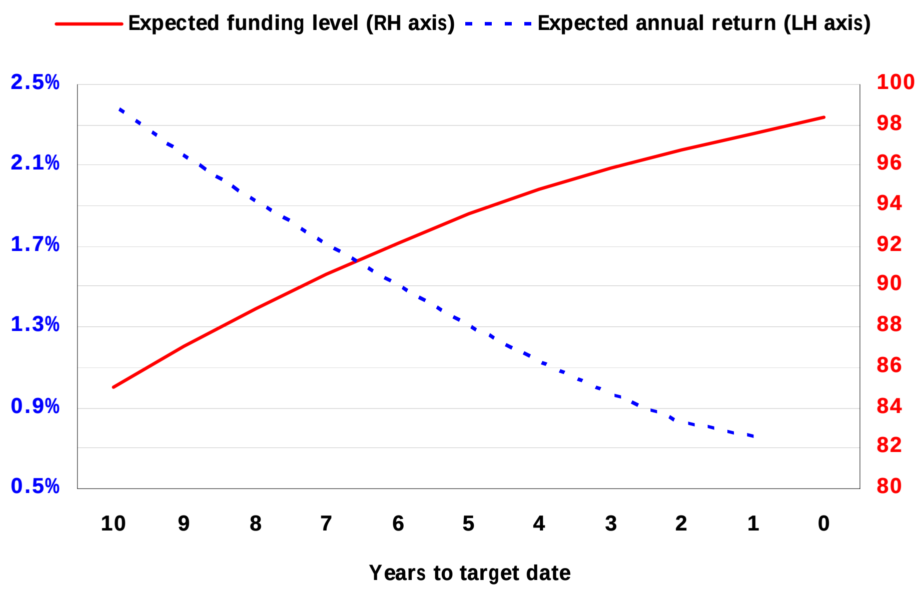

The pattern here is due to the utility function: when the investment horizon is long, then taking more risk than required to achieve the target is likely to be rewarded with a higher utility, and there is scope to reduce risk later. Yet when the horizon shortens, the risk of overshooting the target (with no increase in utility) leads to a reduction in risk-taking.

We can illustrate the efficiency loss from discrete rebalancing by calculating the expected utility level from the optimal annual-rebalance strategy compared to the theoretical (continuous-rebalancing) case. The following graph compares the expected utility with 5 years until the target date:

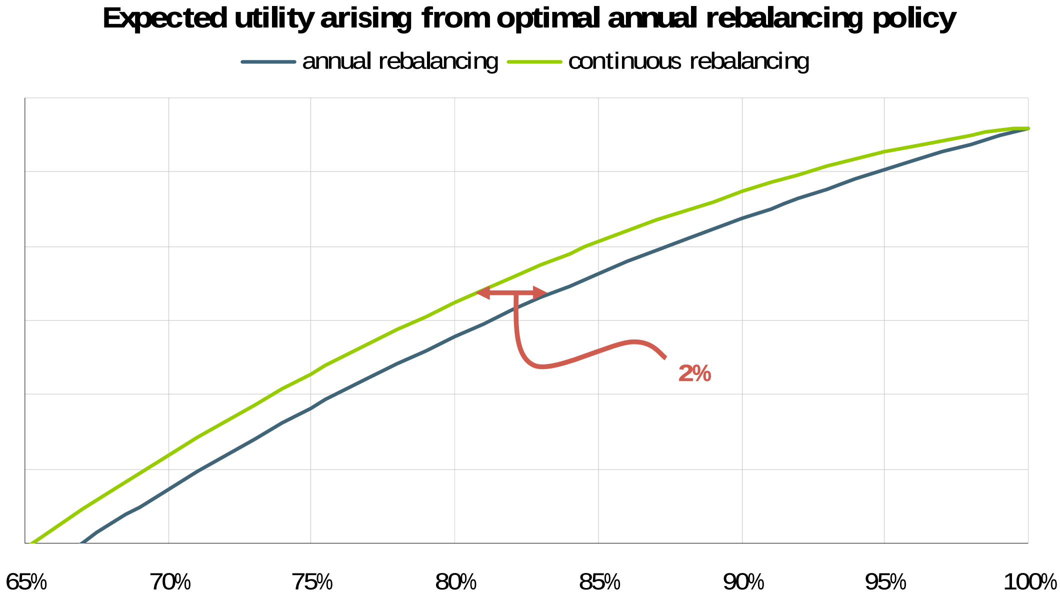

The utility scale is deliberately not shown here: changing the scale or the origin of utility has no economic effect. But we can nonetheless, as illustrated, calculate the wealth-equivalent of the utility loss. At many funding levels (and reducing as we approach 100%), we need a 2% higher funding level in order to achieve the same expected utility as we would have been able to achieve with continuous rebalancing.

One simple way to close this gap faced by the annual rebalancer is to use call options.

6.5. Annual rebalancing strategy with bonds, stocks and call options

As noted above, an investor doing rebalancing at annual intervals faces a dilemma:

– too much risk, and the funding level will often exceed the target level, wasting resources that could have been directed to improving the funding level in bad scenarios;

– too little risk, and the expected improvement in the funding level will be too modest.

The continuous rebalancing strategy would take risk only when the funding level is below target, and then switch into bonds the moment that the funding level reaches 100%. Although this optimal strategy is equivalent to buying and holding a suitable option, this option will have a term of several years and is unlikely to be attractive in practice. However, the dilemma above suggests the use of a strategy using shorter-term options that has a better chance of being close to the continuous optimum.

The idea is simple: sell out-of-the-money call options, in order to sell the unwanted upside. For example it may be that any equity return over 10% would be sufficient to reach 100% funding: in this case, write an option agreeing to pay away any equity return over 10%. (The actual calculations are a little more complex, as a premium is received for this option, and after taking this into account, the threshold may be a little lower than 10%, but the basic idea is the same.) The portfolio value at the end of the year will thus lie between the value of the bonds (*B* in the following picture) and the value of the liability target (*L*):

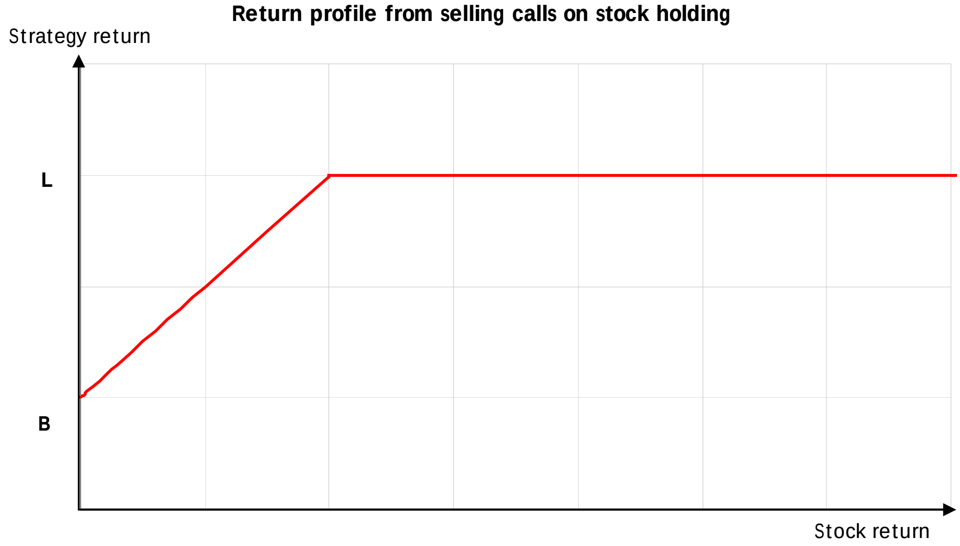

The resulting expected utility is shown below for this option strategy. Note that each year, options with a one-year maturity are used, i.e. expiring on the next rebalancing date. This of course means that we are trading in a liquid part of the market, where transparent pricing is more likely to be available.

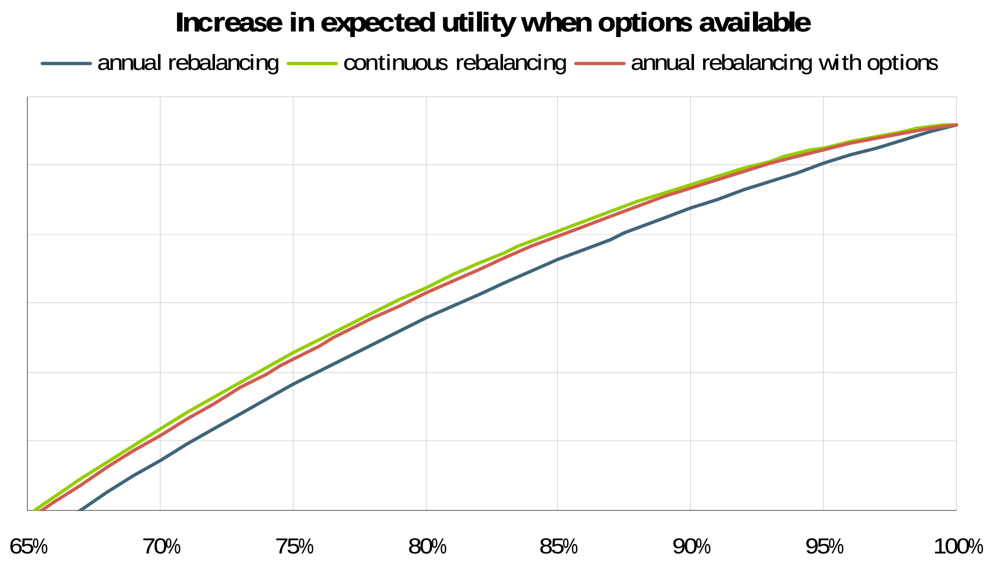

As can be seen, the use of options has allowed us to come very close to the theoretical optimum of continuous rebalancing. The rebalancing requirement has essentially been outsourced via the option.

The resolution of the previous dilemma can also be illustrated: as we now are no longer running the risk of overshooting the target funding level, we are able to take more risk (generate more return).

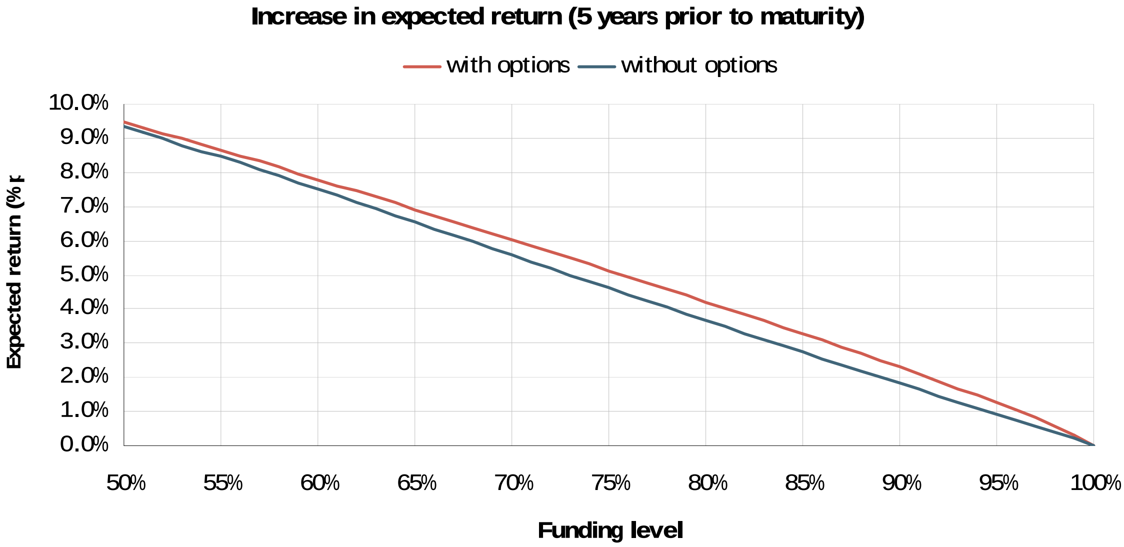

The excess return taken on is more significant as we approach maturity: equivalently, the benefit of the availability of options becomes more useful as the likelihood of the option being taken up increases.

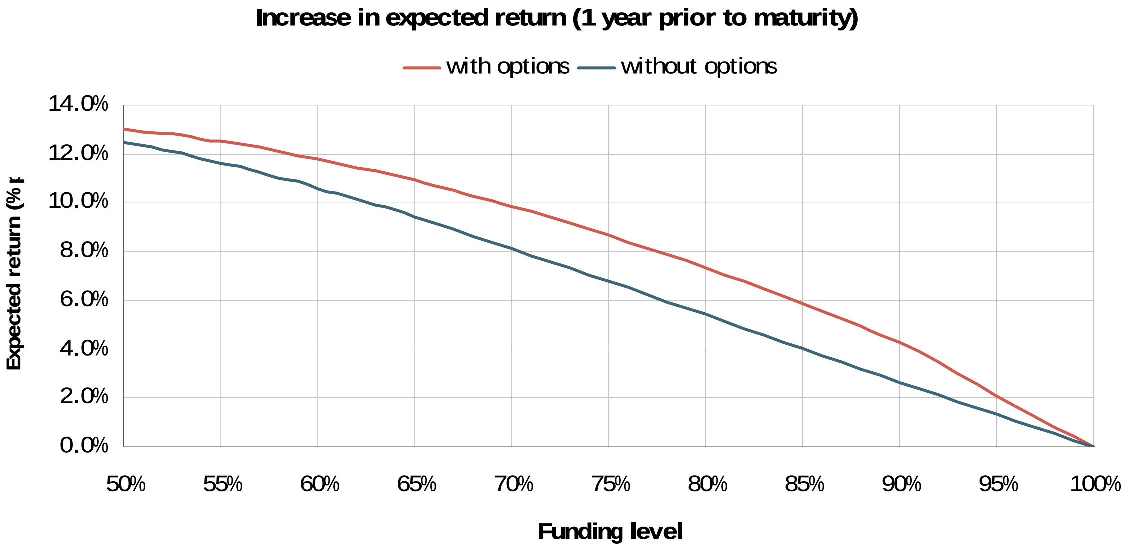

6.6. Comparison with alternative rebalancing rules

The objective of maximising expected utility thus gives rise to a rebalancing rule: for any funding level, target horizon and aversion to risk, an asset allocation strategy can be read off. It is of course interesting to compare this with alternative rebalancing (or ‘trigger’) mechanisms. For example:

– *Use internal rate of return.* A straightforward approach would be to calculate the rate of return required to exactly bring the asset level up to the target liability level at the target horizon. The asset allocation can then be set to have expected annual return equal to this rate.

– *Reverse CPPI.* A more complicated approach could be to invest a constant multiple of the deficit in the growth portfolio, with the rest in a hedge portfolio.

It is more common in practice to have ‘one way only’ variants: that is, if the mechanism above suggests taking risk off the table, then do so, but if it suggests taking more risk then simply retain the current level of risk. The rationale for this is not entirely clear, but presumably reflects a view that the existing risk budget needs to be reduced, but some excuse in the form of positive performance is required.

In the table below, we show some sample statistics – all for an example that starts with an 80% funding level with 12 years to go.

Rebalancing rule	Internal rate of return	Optimal utility	Reverse CPPI	One-way IRR	One-way opt utility	One-way rev CPPI
Average utility (wealth equivalent)	90.1%	91.3%	87.1%	90.4%	90.6%	87.8%
Average funding level	97.2%	94.5%	100.5%	92.7%	92.8%	99.7%
95% funding level tVaR	53%	60%	46%	66%	65%	48%
Standard deviation funding level	14%	11%	17%	10%	10%	16%
Prob funding level >100%	72%	34%	80%	31%	20%	75%

95% tVaR is the average of the worst 5% funding levels.

The general picture is similar:

– As the optimal-utility approach tends to reduce risk near 100% funding (reduced upside to taking risk), the average final funding level tends to be lower than other rebalancing methods. This is consistent with the objective of this approach: utility is optimised but the utility function is used to determine a trade off with lower funding levels. This trade-off can be seen in better standard deviation and value at risk statistics.

– The reverse CPPI method tends to take more risk than required and therefore to overshoot the 100% funding level on average (a lower multiplier than the 4 employer here could be used);

– Introducing a ‘one way’ gate to the CPPI and IRR methods tends to improve their performance, and the overall performance of the one-way IRR strategy is fairly similar to the optimal-utility strategy. In contrast, the performance of the optimal-utility approach is impaired by the one-way gate.

Overall, this is encouraging: a sensible rebalancing rule, consistent with those developed empirically by consultants, has emerged. An ad-hoc one-way gate was not required to derive a sensible strategy. And the trade-off between achieving the funding target and taking risk en route seems an attractive one.

One question that remains is the extent to which increasing risk is desirable as the funding level falls – although the prevalence of ‘one way’ trigger rules suggests that it is not, pension plan economics suggests otherwise. This may be particularly the case in the presence of government-sponsored insurance. This discussion is picked up in the next section, where the joint presence of minimum and maximum constraints is considered.

6.7. Further comments

The calculations have for simplicity been done using bullet payments at the target maturity date. However, we have expressed the results in terms of the ‘funding level’ relative to the target payment as the same general picture applies when we use a stream of liability payments. The mathematics is essentially identical, and now the funding level, or deficit, in the sense we would normally understand it is the crucial determinant of the optimal investment policy.

We now turn to the task of constructing a more realistic model of a utility function for a pension scheme with a deficit. The example in this section may be applicable in simple cases, but more generally the impact of government-sponsored insurance adds an additional layer of interest and complexity.

7: Example 3: DB pension fund in deficit

The trustees of a UK DB pension fund with a deficit often seek to make up the bulk of this deficit using investment returns. In addition to the liability target (discussed in the previous section), a fund also benefits from protection from the support of its corporate sponsor and, in extremis, the Pension Protection Fund (PPF). These complicate the problem considerably.

7.1. Background

Around the turn of the millennium, many pension plans moved very quickly from surplus to sizeable deficits. The path back to surplus has proved to be a much longer one in most cases, but the journey has so far been characterised by many welcome improvements in plans' behaviours. In particular, there has been a substantial change in understanding of risk and an increased willingness to use the full armoury of risk management techniques offered by the financial markets.

Risk has often been discussed in terms of short-term objectives, such as managing the volatility of the position shown on the sponsor's balance sheet in a year's time. Techniques based on modern portfolio theory – hedging and diversification – have been used to effect welcome reductions in plans' risk budgets.

The purpose here is to take a longer term perspective by placing a stronger focus on the time dimension. This is consistent with the regulatory structure, which now encourages the formulation of plans to move to full recovery of a deficit within a defined period (e.g. 10 years). The other main change in the regulatory environment in recent years has been the creation of the Pension Protection Fund (PPF). The PPF provides plan members with a minimum level of benefits if the corporate sponsor of the plan becomes insolvent.

7.2. Identifying the time horizon

The approach developed so far in the paper does not apply entirely naturally to the pensions example. Despite the regulatory encouragement, there is not really a fixed horizon. A pension plan's only true horizon is the death of its last pensioner (and that carries demographic uncertainty), or perhaps the date upon which the benefits are bought out with an insurer. One must therefore create a synthetic horizon for the pension plan. There are some obvious choices: the "recovery period" used under UK scheme-specific funding regulations; an alternative road-map to buyout that many closed funds are considering; or simply the period until the next valuation review.

However, in each of these cases, time may remain a variable. Reaching only 98% of target at the end of the horizon is worse than reaching 100% but it is not a disaster if one can opt to continue for a couple more years. An interesting question arises if a plan is ahead of schedule part way through the period: do the Trustees maintain risk level in the hope of reaching the target sooner, or do the Trustees remain focussed on the original objective and therefore progressively de-risk as the above discussion suggested? The utility approach allows this difficult decision to be abstracted, but is probably just as useful in identifying a range of internally-coherent dynamic strategies.

7.3. Identifying the investor

A further problem is to identify who is taking risks & benefiting from returns in a pension plan? Plan members, acting through the trustees as their agents, have a defined benefit, so that, provided the sponsoring employer does not default, they have no exposure to the funding level of the plan. Of course, this is a strong provision: it is common for companies to go insolvent. On the other hand, the sponsoring employer - or rather its shareholders - is exposed to the funding level of the plan: as the funding level goes up or down, less or more future contributions will be required. But shareholders will measure utility at the level of their entire portfolio; the question for them is one of maximising the value of the company rather than trading off risk and return.

This theory is of course at variance with common practice: trustees are very concerned to improve the funding level of their pension plan, so that if insolvency occurs members will receive a higher level of benefits than those promised by the PPF. And finance directors are typically also keen to improve funding / reduce risk – they will not be willing to place any weight on the possibility of company insolvency. We therefore take the point of view of these agents, albeit with an awareness that this perspective may not be appropriate.

If the funding level of a plan improves to be over 100%, then further complexities arise: is the surplus fully owned by the sponsor (e.g. released in the form of reduced future contributions or a return of surplus after a penal tax rate has been applied) or shared with the plan beneficiaries? In a closed scheme, sponsors will therefore tend to avoid surpluses, and in an open scheme we should expect their utility will increase less fast above 100% funding than below. The trustee / member attitude to a surplus will tend to be determined by the extent to which beneficiaries share in future surplus.

Clearly, having a higher funding level always adds utility to the employer, other things being equal. However, achieving a higher funding level by paying extra contributions creates a corresponding negative utility in the employer's bank account. We do not intend to explore the question of whether cash is better in the scheme or out of it. We focus only on changes of utility derived from investment returns.

In order to navigate these waters, we outline some simplified situations, where the roles of the parties are more clear-cut, and allow similar themes to emerge. The utility function charts are intended to help inform an investment decision. We have suggested utility as a function of the funding level of the scheme at some target horizon. We have assumed that the rate of employer contributions is independent of the investment strategy – in practice, of course, future sponsor contributions are an important asset of the fund, and can have a strong influence on behaviour.

7.4. Scenario 1: Scheme very large relative to employer; any surplus fully owned by employer

The first simplification we can undertake is to minimise the impact of employer insolvency by assuming that there will be negligible contributions whether or not the employer remains solvent, and to have any surplus returned to the employer.

At wind-up, utility for the beneficiary should therefore:

– be static if the funding level is below the PPF insurance level;

– increase above that level, as higher benefits become payable; and

– become static again above the promised-liability level, as surplus is repaid to the sponsor.

This is illustrated below.

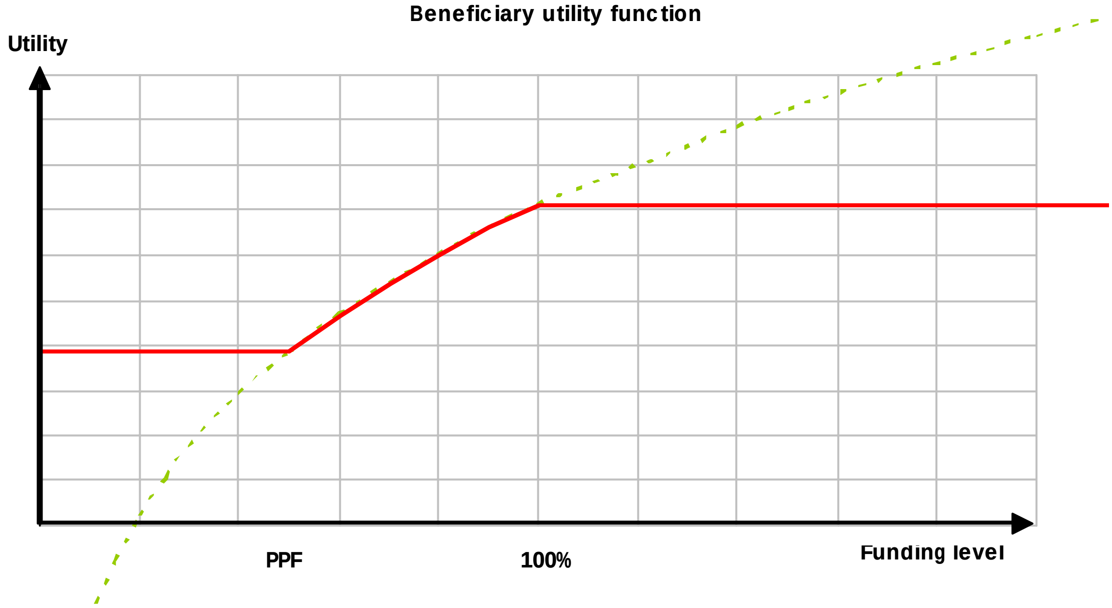

Note that all “funding levels” in these slides are expressed as a percentage of insured buyout. We recognise that this is not a totally objective measure and can only really be identified from time to time by testing the market. However, we have assumed that a suitable proxy can be developed to aid decision-making. On the funding level axis, the “PPF” value corresponds to the level of benefits that would be provided by the Pension Protection Fund (PPF).

We have emphasised, in the example above, the circumstances of wind-up. This provides for a neatly discrete outcome. For example, the scheme falls into the PPF or it doesn’t; the scheme has a surplus to distribute or doesn’t. However, it is likely that the sponsor has survived and, if so, time will go on. In this situation, at low funding levels, the members carry a positive option value (against the PPF) that reflects the chance of recovery by the ongoing scheme. Similarly, at funding levels above full buyout, if the scheme does not wind-up and guarantee the benefits, then the members retain a small disutility due to the chance of slipping back below 100%. In this scenario we had assumed that the members could not benefit from any surplus. This is illustrated by the black line in the chart below.

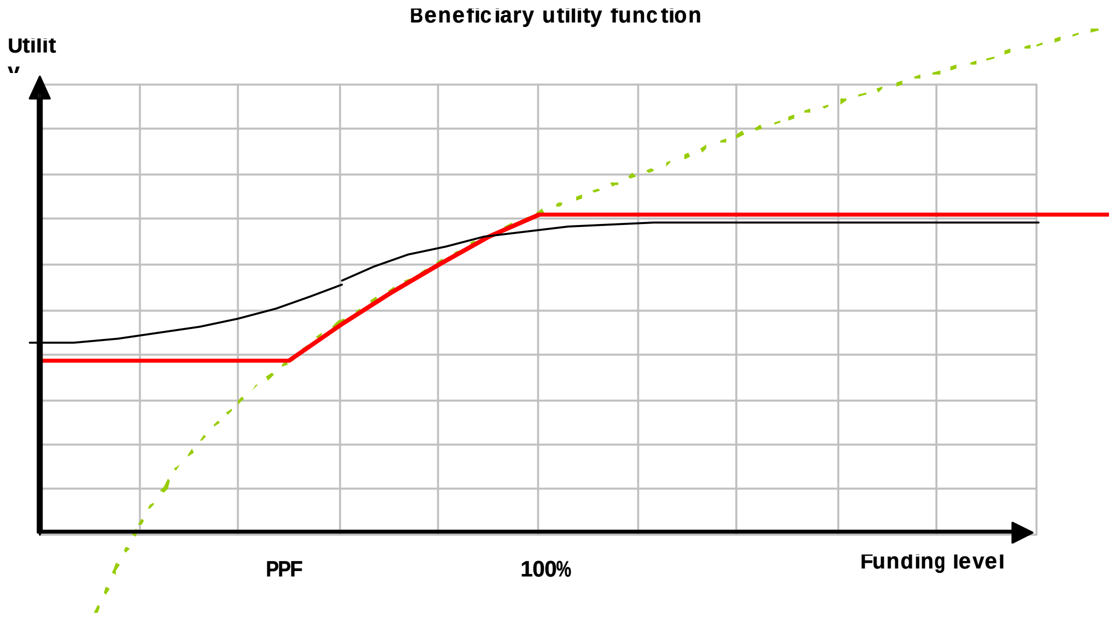

For the sponsor, on the other hand, if the deficit in the fund is greater than the value of the company ("C"), any further losses are irrelevant to the company. Yet if the funding level increases beyond 100%, this has value to the company, even if subject to a penal tax rate. Therefore (at wind-up):

- Utility should be static if the assets are less than Value (Liabilities) – Value (Company);
- Utility should increase above this level; but at a slower rate above 100% funding due to the differential tax treatment of surplus and contributions.

This is illustrated in the chart below.

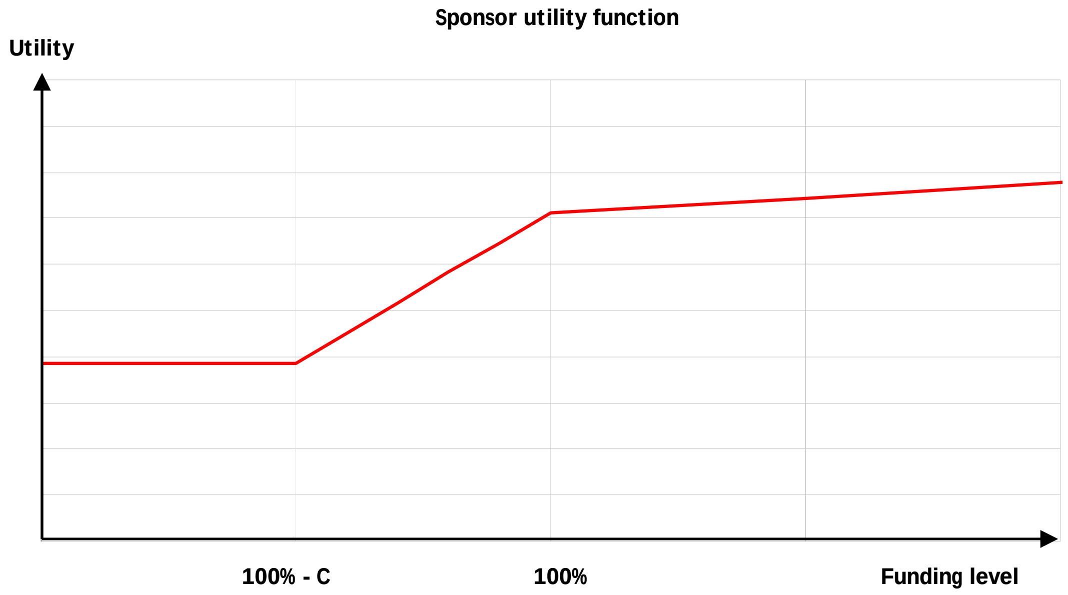

Again, a smoother line could be drawn through this function to reflect the situation of the scheme still being ongoing at the chosen time horizon.

7.5. Scenario 2: Scheme small relative to strong employer; surplus shared with employer

An alternative scenario might be where the sponsor is very strong, and much larger than the pension fund. Suppose now that any surplus that arises is shared between the sponsor and beneficiaries.

Then, the trustee utility function should:

– be static if the funding level is below the promised liability level, as there is a high degree of confidence that these benefits will be paid come what may;

– increase above that level, as improvements are made to benefits.

We are assuming here that members are a long way from breaching their Lifetime Allowance.

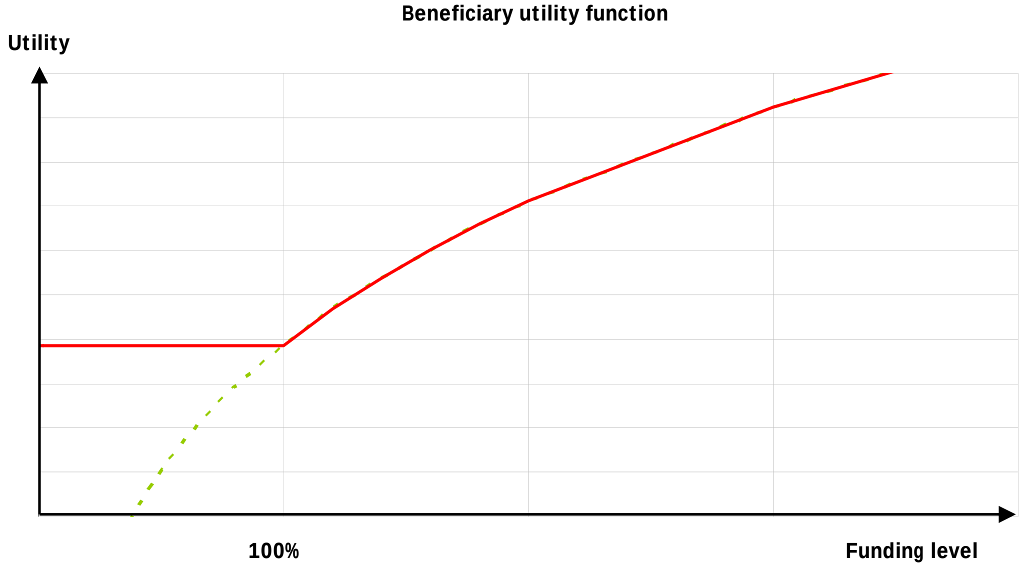

The sponsor utility function should:

– Increase uniformly as the assets increase up the promised liability level;

– Increase more slowly thereafter, as some of the value of any surplus is shared with the beneficiaries.

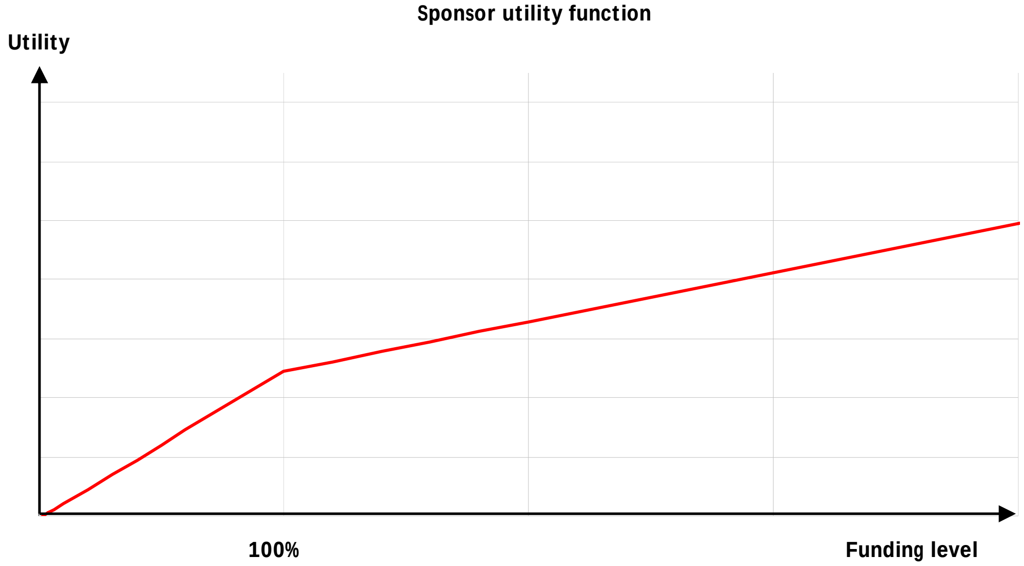

7.6. Scenarios 1 and 2: the PPF's perspective

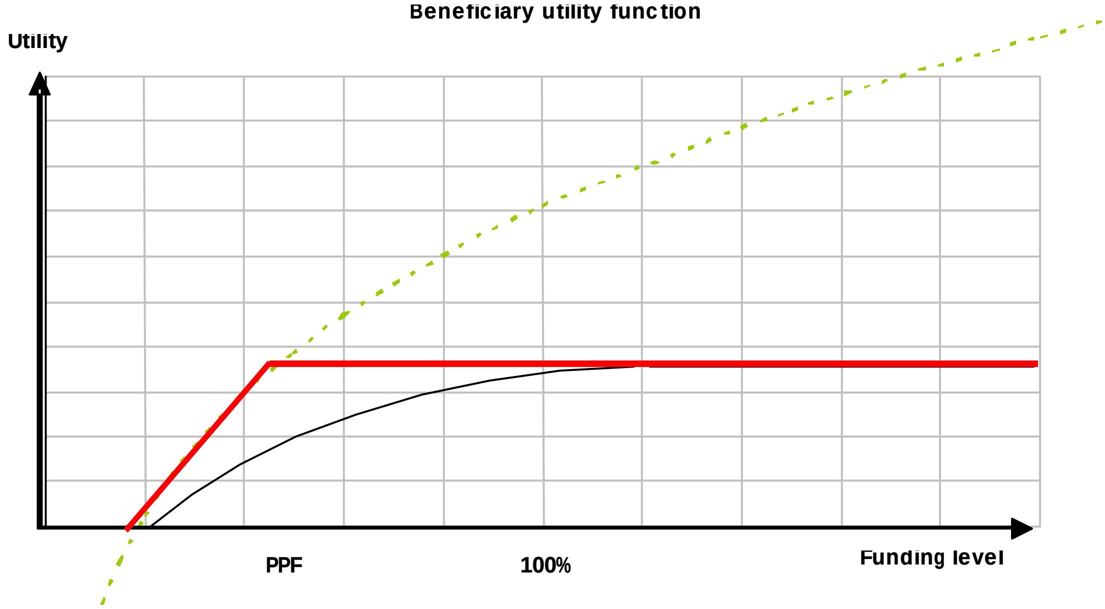

It is interesting to observe that the charts drawn from the Trustees' and members' perspectives assume an indifference to further falls in funding level once the PPF level has been reached. We may anticipate that the associated optimal investment strategy for this utility curve includes a high risk approach when the funding level lies in this area.

Conversely, however, the PPF, if in the same position, might wish the scheme to behave more like an insurer wishing to protect a guarantee. The CPPI approach, for example, would be more appropriate from its point of view. The PPF is currently open to a significant moral hazard from schemes attempting to gamble their way up from low funding levels. Trustees will be motivated to run virtually the opposite strategy to the PPF, increasing both the risk of the plan falling into the PPF and the likelihood of the assets passed across in such a situation being very different to what the PPF would want to hold.

Part of this conflict arises from the Trustees' focus on delivering the benefit promised under the Plan Rules and the PPF providing a lesser benefit. The Trustees' utility curve flattens at or after

100% (of the scheme's benefits) funding level but the PPF's curve flattens at a lower funding level (100% of the PPF benefits). Better alignment could be achieved if the Trustee were permitted to define the PPF benefit level as the "guarantee" above which it aspires to deliver a greater benefit (perhaps that promised by the Plan Rules but preferably without being too anchored on that). Having done this the Trustees would be led to a more CPPI-like strategy. The PPF would perhaps, then, not even need to exist.

Currently the PPF sets levies that are broadly proportional to funding level (as well as being linked to the sponsor's default risk). The rising levy at lower funding levels is not captured in the above scenarios and might go some way to increased alignment of interests in cases where the levy is large enough to have an impact on strategy.

The PPF does not currently allow for a scheme's investment strategy when setting the levy. However, the PPF has begun to consult on adding this refinement. Let us assume for a moment that regulatory intervention of some sort will be applied more aggressively in future (by the PPF or by The Pensions Regulator to schemes with low funding levels that take high levels of investment risk.

In this case, the sponsor, trustees and beneficiaries may be keen to ensure that a minimum level of benefits, over the PPF minimum, is payable, while still targeting a level of benefits that is above those currently covered by the assets. A picture similar to that below will result: this is effectively a combination of two of our previous examples: a solvency-minimum style floor and a liability-target style cap.

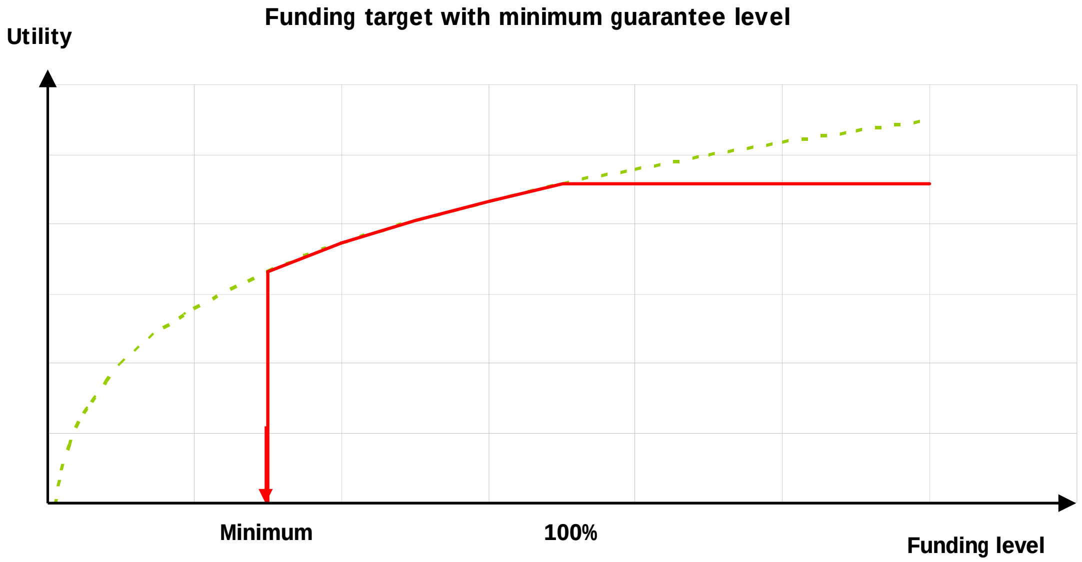

7.7. Common features

A recurring picture here is a minimum level (M), below which the government-sponsored insurance kicks in, followed by a rise to a further level (L), beyond which increases are nil or small relative to rises below this level.

There may also be an absolute minimum level of assets allowable (G)– any approach to this level of assets will lead to intervention by a regulatory authority. Finally, risk may be measured relative to a floor (F) different from zero.

Generic class of modified utility functions

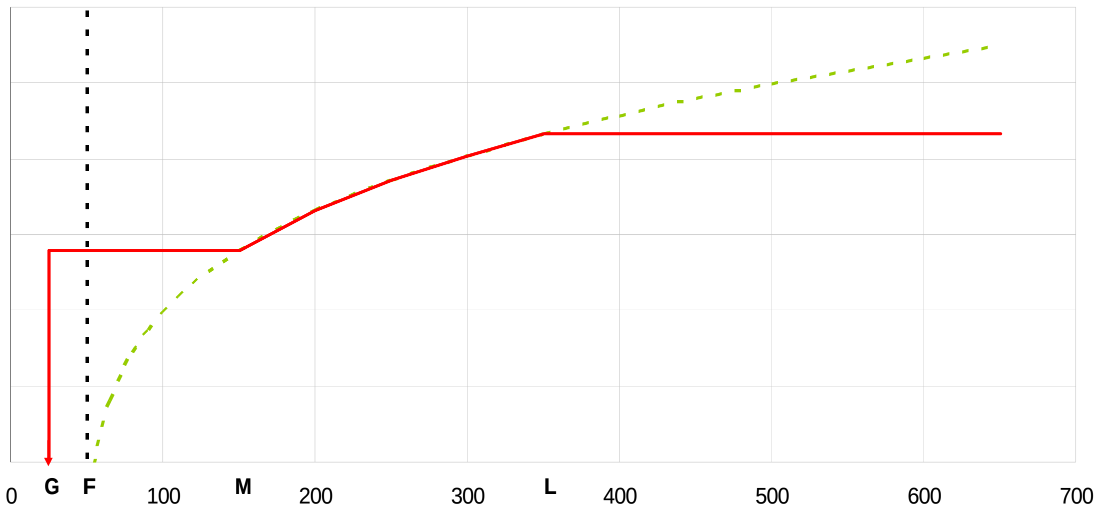

Different relative levels of these quantities will lead to most of the shapes of utility function discussed above.

7.8. Optimal strategies

The optimal strategies for these utility functions are similar to those encountered before. The solution can again be expressed in terms of the solution to the unmodified (CRRA) utility function – now, in addition to the put and call options of the previous sections there is also of a binary option:

- a call corresponding to the maximum level *L*;
- a put corresponding to the minimum level *G;*
- a binary option corresponding to the flat utility function between *G* and *M*.

The binary option captures the fact that it will always be suboptimal to have a final level of wealth between *G* and *M* (indeed, between *G* and *M'* for some *M'* between *M* and *L*).

As before, the optimal strategy for a given level of initial wealth can be constructed by buying a suitable portfolio and set of options on that portfolio at the outset. In this case, no rebalancing through time is required: if the optimisation is done after an interval using the realised portfolio performance, then the optimal strategy will be to purchase exactly the portfolio and options already held. Alternatively, the optimal strategy can be created using a continuously-rebalanced portfolio, effectively delta hedging the payoff from the option portfolio.

7.9. Discrete rebalancing strategies

As in the liability-target example above, it is interesting to explore how well discretely-rebalanced strategies perform relative to this optimum. Whereas in that case, the inclusion of short-dated call options enabled the development of a strategy that was very near-optimal, a much larger gap can now remain. This is because the optimal strategy involves taking a lot more risk when the funding level falls below M, but the risk of breaching the minimum guarantee level (zero in our examples) is then too high.

The simplest solution to this is to allow the use of put options, that constrain the range of future outcomes to always be greater than the guarantee level. Risk can then be increased as the funding level falls, without fear of the total assets falling below zero.

The utility gain is illustrated in one example in the picture below, which shows the expected utility from following an annual rebalancing program with bonds and stocks, or with bonds, stocks and stock options. All are over a five year time horizon.

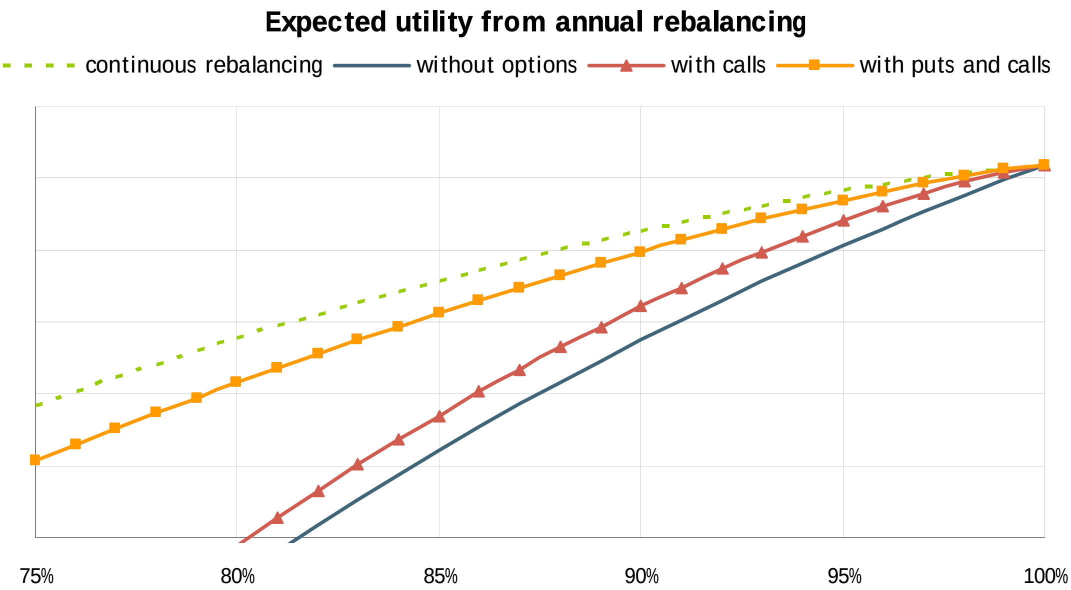

As in the earlier case, adding annual call options increases the (wealth-equivalent) expected utility by around 2% (e.g. adding call options means that the same expected utility can be reached from a funding level of 89% as was only attainable from a 91% initial funding level without options). But the continuous optimum is now much further away (the same expected utility could be achieved with a starting level of around 82% of liabilities). Most of this gap can be closed if put options are brought into the mix.

As before, a very simple strategy is used: the puts and calls have annual maturities, i.e. expiring on the next rebalancing date, and their exercise prices are such that the funding level in a year's time is constrained exactly between zero and 100%.

A binary option (or, approximately, a pair of call options) could be used to further improve this solution, but the point is already clear: in this utility framework there is a clear benefit from using options: risk can be taken in a much more targeted way than was available before. Note that there is no 'view' on the future direction of equities or interest rates implicit here: the option strategy follows more or less immediately from the objectives of the fund, as expressed in the utility function.

The increase in return and risk once options are available is similar to the discussion above for the liability target. The picture below shows the expected return for the various stock/bond/option portfolios discussed above, with 2 years until maturity.

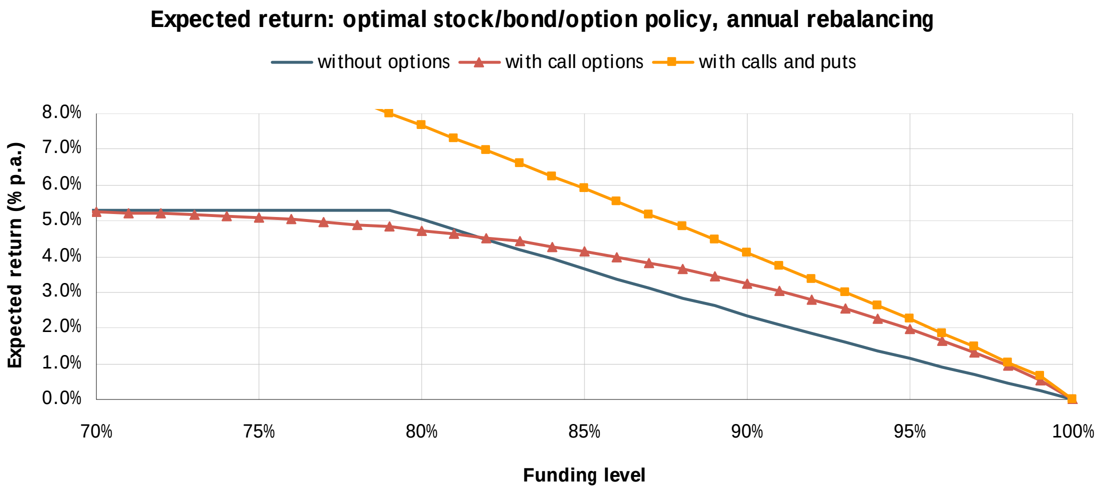

The general picture is similar at other time horizons: the maximum amount in the risky asset portfolio ('stock') is limited to 100% if there is no downside protection, but additional risk-taking is both possible and desirable once put options are available to provide protection. As the funding level increases, protection is not required but the benefit of the upside cap becomes more relevant. In both cases, options allow the profile of future returns to be structured so as meet the investor's requirements fairly precisely.

8: Practical issues of hedging pensions

In this section we consider some of the practical issues that need to be considered in the process of designing an LDI hedging strategy. We do not consider in detail the practical issues of the actual implementation, including the costs, choice of delivery vehicle, the mechanics of collateral management, managing best execution etc, as these are beyond the scope of this paper, and have been covered elsewhere.

8.1. Choice of discount curve

The choice of discount curve for valuing the liabilities is a key issue in designing a liability hedging portfolio for a pension plan. The funding position of the plan may be monitored on a range of different measures, which might incorporate a range of discount curves (for example: gilt or AA rated bond yield or a market consistent swaps based approach). Whilst the gilt or AA rated yield has been the traditional valuation basis, swaps hold clear advantages in terms of:

- *Investibility*. A benchmark is produced which could, in principle, be matched precisely.
- *Flexibility*. A specific liability profile can be hedged; the hedge can be tailored to account for caps and floors, scheme increase inflation reference dates, etc
- *Capital efficiency*. Underlying assets are not tied up, allowing freedom to invest in diversified return seeking assets alongside the hedge.

However, a swap discount rate cannot be said to be a risk free basis. Counterparty risk can be largely mitigated by having a robust collateralisation process, but generating LIBOR to fund the floating legs of the swaps requires taking investment risk, as recent market conditions have shown. The alternative "risk free" discount curve would be gilts, however using gilts have insufficient flexibility to allow precise tailoring to a specific liability profile, and of course require significant capital outlay.

AA corporate bonds have been historically popular as a hedge of the accounting basis, but there are numerous practical difficulties in using this as a basis for hedging, and in recent market conditions this basis has also produced a very high volatility in the liability value because of the movement in credit spreads within the AA corporate bond market.

In some cases, the basis for calculating the inflationary increase within the liabilities is to take a single point from the break even inflation curve and apply this rate across all years. Hedging this inflation rate with inflation swaps can lead to relatively poor results in the short term (ie hedge is not that effective). However, as the liability cashflows are directly linked to actual inflation (eg RPI or LPI) rather than a single point on the break even curve, actual benefit is greater than the volatility of the hedged funding level might suggest.

Hedging liability risks with swaps whilst valuating the liabilities on a different basis can result in "basis risk", ie the risk that the swaps and liabilities do not move exactly in line.

8.2. How much of the liabilities to hedge?

Having chosen a discount curve, the next question is how much of the total liability to hedge. A starting point for trustees might be to consider hedging the whole of the interest rate and inflation risk in the plan's liabilities. The trustees then might consider the following issues in deciding whether to hedge less than the total liabilities:

1. If the plan is underfunded, the trustees may wish to restrict the hedge size to the value of the assets available, to avoid gearing at a plan level. If the full value of the assets were hedged, this would protect the funding level as a proportion of assets against moves in interest rates

and inflation. If the full value of the liabilities were hedged, this would protect the size of the deficit against moves in interest rates and inflation (however the funding level might change).

2. The plan may hold other assets that can help to hedge interest rate and inflation risk. Clearly, the interest rate and inflation exposure in any bond assets should be included when designing a liability hedge. Some trustees or advisers also choose to reduce the size of the hedge to allow for an assumed correlation between other asset classes (e.g. equities) and interest rates/ inflation. However, there is an alternative view that while such correlations may exist in historic time series data over longer periods, the correlation cannot be relied upon because it often breaks down during periods of market stress, when the hedge is most important.

3. There may be considerable uncertainty in the projected liabilities, and large exposure to other risks that cannot be effectively hedged. It may be felt undesirable to hedge the full interest rate and inflation risk in the liabilities, as over-hedging could result.

8.3. Market timing risk

The trustees may be concerned about the market timing risk of making a switch to an LDI hedging strategy. They may decide only partially to hedge the liabilities initially, perhaps with a view to increasing the size of the hedge later. Market timing risk can be managed in a number of ways, some of which can be used in combination:

- Phased implementation of hedge (mechanistic or dynamic)
- Phased switch out of existing assets (e.g. equity de-risking)
- Use of instruments incorporating optionality to achieve desired balance of de-risking and retaining risk exposure during the process
- Active or passive management of the transition to the hedge strategy

8.4. Phased implementation of hedge and dynamic strategies

Dynamic strategies take account of market conditions in the pace of hedge implementation, and typically incorporate maximum or minimum trigger levels, above or below which a phased hedging strategy might be started or suspended, or accelerated or decelerated.

The strategy can be applied in a mechanistic or judgemental way. For example, trigger points could be set on the level of real yields seen in the market, thus not hedging when (for example) inflation is high.

If a dynamic hedging strategy is adopted, there might be a case for dynamically adjusting the asset portfolio. Increasing a liability hedge will increase the LIBOR requirement on the client, and therefore it might be prudent to switch assets into a LIBOR or LIBOR ‘plus’ seeking investment strategy. Of course, this depends on the asset mix of the portfolio and the extent to which the client is prepared to take risk vs LIBOR.

8.5. Active or passive management of the hedge

If the hedge is passively managed, this can be implemented with bonds, or through a combination of cash and swaps. In the later case, the cash return backing the swap hedge will not be completely risk free, as it is not possible to invest in a cash fund that guarantees LIBOR (there are products that pay a LIBOR based return, but not without a degree of credit risk and liquidity risk).

Alternatively the hedge can be actively managed. If it is actively managed by investing in bonds, the range of active views is limited, as the universe of possible active views is limited by the range of possible hedging instruments. If a swap overlay approach is used, there are a number of ways active views can be incorporated:

- Replace some of the underlying assets backing the overlay with diversified return seeking assets, rather than cash, managed against a LIBOR benchmark.

The choice of active or passive management of the hedge, and how it is implemented, needs to be considered in the context of the overall risk budget of the pension plan.

8.6. Pooled vs segregated solutions

Given the uncertainty surrounding the actual future liability cashflows, there is a trade off between precision in hedging and simple, practical solutions. Possible drivers for a particular solution could include: the degree of optionality within the cashflows (eg LPI or non standard increases); desire to closely match pensions in payment in order to move towards buy out; the funding level; and the level of growth asset risk. Pooled solutions can offer a practical way to implement and manage a hedging solution whilst providing a material risk reduction benefit.

8.7. Demographic risks

The hedging solution here refers to interest and inflation rate risks. There is of course a demographic risk due to changes in longevity expectations. Most pension schemes at present are factoring in an underpin with the medium cohort expectations but in future these mortality tables might need to be strengthened further. This could result in the liability profile changing significantly and will affect the accuracy of any hedging strategy that doe s not hedge this risk.

The longevity market is still in its infancy but advances are being made:

- Longevity swaps dealt with investment banks.
- Development of longevity indices (e.g. JP Morgan LifeMetrics and Credit Suisse).
- Reinsurance/ stand alone longevity insurance.
- Partial insurance buyouts

These advances are not part of this Working Group’s remit. We understand that these are being explored further by other Working Parties.

9: Conclusion – using utility theory in practice – what is my utility curve?

If you, the reader, are involved in making investment strategy decisions, or an advisor to decision-makers, you may be wondering how you could apply in practice the approaches we describe. The illustrative utility curves are intended to take into account some of the considerations of real situations but construction of a utility curve is clearly a subjective, personal matter. The illustrative curves are not supposed to be 'cut and pasted' into live situations. So the reader must consider how to construct their own curve from first principles.

Rather than offer a detailed framework for doing this, we have provided commentary with each chart explaining our thought processes in constructing the curves. We hope that doing so has made the charts more plausible for the reader as well as providing a guide as to how others may construct their own curves. Broadly speaking, we found it necessary to identify the points to which the curve tended as the funding level approached zero or infinity, as well as considering any interesting points in between, such as the level of any guarantees. Having done this, it remained only to join the dots in a realistic manner. An advisor to a decision-maker could perhaps construct a questionnaire to help the decision-maker through a similar process.

Our group considered that the creation of the utility curve need not be a one-off preliminary exercise. Given the subjectivity of the process, and the unfamiliarity of the user with the concepts, the exercise may need to be iterative. Having pursued one utility curve through to an optimised investment strategy, the implications should be assessed. This assessment would include the usual range of metrics including expected return, volatility and value at risk. For example, the optimal strategy, when presented, may feel as though it has uncomfortably excessive down-side risk. The decision-maker could, in this case, be criticised for having described inadequately their risk aversion but this need not be a failing. Having perhaps never tried to quantify their risk aversion, errors are very likely but quite useful. The decision-maker may learn a lot about themselves by going through this process, even if it does not lead them to refine their investment strategy.

It is not always the case, particularly in the management of pension plans, that the decision-makers will consider "what will I do when my funding level rises/falls to x". The utility-based approach described in this paper provides a good discipline to such cases, in generating some forethought. It may lead the decision-maker to better consider the really crucial scenarios to hedge against, or to determine when risk-taking has paid off and it is time to stop.

10: Bibliography

Cox, J and Huang, C-F (1989). Optimal consumption and portfolio policies when asset prices follow a diffusion process. Journal of economic theory 49, 33-83.

Merton, R (1969). Portfolio selection under uncertainty: the continuous-time case. The review of economics and statistics, 51(3) 247-257.

Merton, R (1971). Optimum consumption and portfolio rules in a continuous-time model. Journal of economic theory 3, 373-413.

Black, F and Perold, A (1992). Theory of constant proportion portfolio insurance. Journal of economic dynamics and control 16, 403-426.

Leland (1980). Who should buy portfolio insurance? Journal of finance 35(2).

Nielsen (2006). Optimal investment strategies: a short survey of classical and recent results. Presented to the International congress of actuaries.

11: Appendix A: Common utility functions

We provide some more detail on the utility functions that are regularly used in the finance literature.

11.1. Definition

The basic idea of a utility function is to express mathematically the personal value that will arise from different levels of wealth. Generally, utility functions have two reasonable properties:

– *Increasing*. More wealth is always better.

– *Convexity*. The increase in utility for each additional dollar of wealth diminishes as the initial level of wealth increase.

11.2. Measures of risk aversion

The Arrow-Pratt measure of *absolute risk aversion* (ARA) at a particular level of wealth *w* is:

$$ARA(w) = -\frac{U''(w)}{U'(w)}. \qquad (5)$$

This effectively expresses the trade off required between additional return and variance when starting from a given level of wealth.

The *relative* risk aversion is:

$$RRA(w) = w.ARA(w) = -\frac{w.U''(w)}{U'(w)}. \qquad (6)$$

This measures the aversion to risks that are proportional to the initial level of wealth.

Arrow argued that investors' aversion should decrease with wealth (i.e. *ARA* is a decreasing function of *w*), i.e. that willingness to face a risk of a given size increases with wealth. But if the risk rises in the same proportion as the wealth then the willingness falls (i.e. *RRA* is an *increasing* function of *w*).

11.3. Constant relative risk aversion (CRRA) utility functions

As noted in the main text, the CRRA family of utility functions can be defined via:

$$U(w) = \frac{w^{1-\gamma} - 1}{1-\gamma} \qquad (7)$$

It is easy to check that this family does indeed have constant relative risk aversion γ .

Since it is the expected value of a utility function that is maximised, changing a utility function by the addition of, or multiplication with, a constant has no effect on the investment strategy that is followed. A utility function $U(w) = w^{1-\gamma}$ would therefore be equivalent. The form of (6) is used so that the $\gamma = 1$ limit exists: it is simply $U(w) = \ln w$.

Note that:

– The $\gamma = 0$ case is linear

– The cases $0 < \gamma < 1$ are bounded below

– The $\gamma = 1$ case is logarithmic

– The cases $\gamma > 1$ are bounded above.

This family is sketched below for different levels of the relative risk aversion γ.

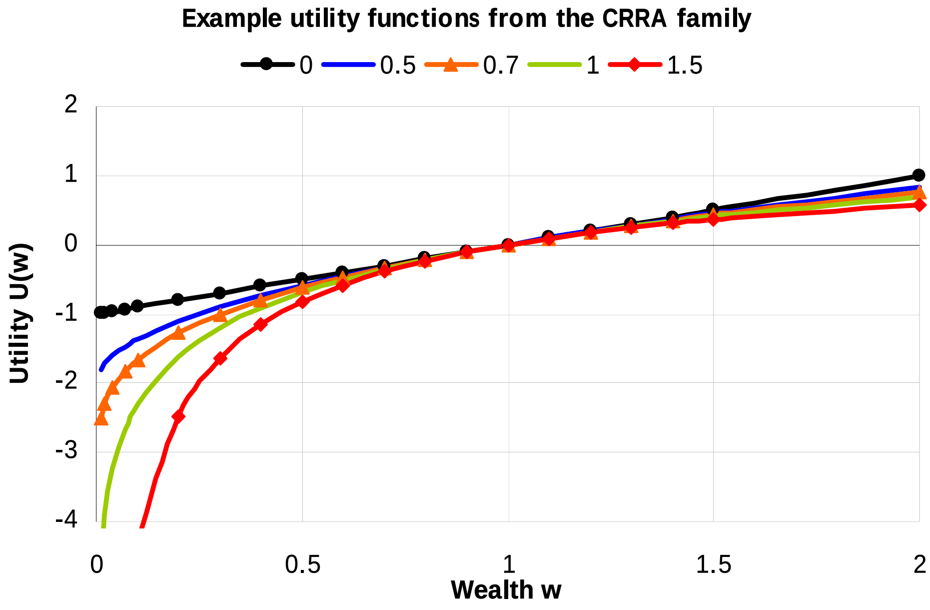

11.4. Hyperbolic absolute risk aversion (HARA) utility functions

The absolute risk aversion of the CRRA family is γ / w, a function whose graph is a hyperbola. Other hyperbolic absolute risk aversion functions, of the form $ARA(w) = a + (bw + c)^{-1}$, define the HARA family of utility functions. These utility functions are also commonly used in academic papers.