INVESTMENT STRATEGY AND VALUATION OF WITH-PROFITS PRODUCTS

Prepared for 2000 Investment Conference

Andrew Smith
of the
Applications of Financial Theory Working Party
andrew.d.smith@bacon-woodrow.com

Abstract:

This brief note outlines an option pricing approach to valuing stakes in a with-profits fund. We depart from previous work in this area by not making direct use of Black and Scholes formula. Instead of trying to approximate the benefit structure with the kind of options that trade on exchanges, we seek to model the liabilities accurately, and then to price the resulting flows in a manner consistent with option pricing theory and market prices.

This is a preliminary report of work done to date. We hope to update this at the Life Convention, and, in the Spring of 2001, to release the underlying models into the public domain.

PART I: ORIENTATION

Idea Behind the Project

The idea of this project is to develop new ways of looking at a with-profits fund. We have concentrated on traditional UK with profits business.

The work is based on a cash flow projection model, described in the appendix. This enables us to build up balance sheets and revenue statements, in a stochastic manner. This model will be demonstrated at the conference.

The model also produces state price deflators. This is the tool required to value cash flows in a stochastic setting, in a way which is consistent with market practice in pricing options and also with the market prices produced by the stochastic asset model.

Traditional ALM

Traditional ALM focuses on distributions of outcomes, and on probabilities of various adverse events. For example, users may be interested in the distribution over time of:

- Free assets
- Free asset ratios
- Policy payouts
- Split between reversionary and terminal bonus

Insurers may seek to optimise strategy, for example by maximising mean policy payouts subject to an upper limit on the probability of ruin.

What is Missing?

The missing ingredient from traditional ALM is a coherent concept of value. Suppose I want to compare a risky cash flow to a risk free flow, to establish which is more valuable. By how much should I reduce the risky flow to allow for the risk? Are there any situations where the risky flow could be more valuable? Is the value of a flow a function only of its distribution, or is more information needed?

Too often, ALM gets bogged down in these imponderables. We now outline a solution based on finance theory.

PART II: VALUING CASH FLOWS USING DEFLATORS

The Algorithm

As well as producing equity returns, bond yields and so on, our stochastic asset model (the reduced Smith Model, or RTSM) produces deflators. These take the form of an additional time series, which is generated along with the more usual outputs.

Suppose we wish to estimate a market price for a cash flow X receivable at time t. We denote by $X(sim)$ the value of X in simulation *sim*. The estimated value of X is then

$$\text{value} = \text{average}\{\ \mathit{deflator}(\mathit{sim},t) * X(\mathit{sim})\ \}$$

The deflator is completely unconnected with the cash flow being valued. If we wanted to value many different cash flows, all payable at time t, we would use the same deflator variable in each case. We can think of $\mathit{deflator}(\mathit{sim},t)$ as the weight that the market places on simulation *sim* when setting market prices.

For pricing streams of cash flows, we value the flow at each future date and then sum the present values.

Validating the Algorithm

To validate the algorithm, we need to ensure consistency with market prices. In particular, suppose that we have a portfolio whose value at time t is $V(\mathit{sim},t)$. We assume that this is an index of total return, and that the starting value $V(\mathit{sim},0)$ is constant for all simulations.

To get consistency of our prices to market values, we need to be sure that for any such portfolio,

$$V(\mathit{sim},0) = \text{average}\{\ \mathit{deflator}(\mathit{sim},t) * V(\mathit{sim},\ t)\ \}$$

In particular, this has to work for both equity and bond portfolios. This provides a straightforward series of tests that can be used to validate the output of RTSM. Of course, as we are using Monte Carlo methods, we cannot expect these means to tally exactly, but the tow sides should be within sampling error bounds given the number of simulations run.

As a special case, we can consider a portfolio consisting of a single zero coupon bond for some future date t. Then for that value of t we would have $V(\mathit{sim},t) = 1$. Then

$$\text{zero coupon bond price} = V(\mathit{sim},\ 0) = \text{average}\{\ \mathit{deflator}(\mathit{sim},t)\ \}$$

Relation between Risk and Return

We can write the deflator test as

$$
\begin{aligned}
V(sim,0) &= \text{average}\{\ deflator(sim,t) * V(sim,\ t)\ \} \\
&= \text{average}\{\ deflator(sim,t)\ \} * \text{average}\{\ V(sim,\ t)\ \} \\
&\qquad + \text{covariance}\ \{\ deflator(sim,t)\ ,\ V(sim,\ t)\ \} \\
&= \text{PV of mean flow at risk free rate + risk adjustment}
\end{aligned}
$$

If the portfolio V earns a risk premium relative to the risk free rate we must have a negative risk adjustment in this equation. In other words, the portfolio value must be negatively correlated with the deflator. The deeply discounted initial price can then be explained by the fact that the market gives below average weight to the scenarios where the asset performs best.

Notice that this is not the same as a risk adjustment you would get from utility theory. It makes no sense to ascribe a utility function to a market. Utilities belong to individual investors. If a market has many investors with different utilities, the aggregate market does not have some sort of average utility. It has a deflator. This deflator reflects the combined preferences of all the investors.

Reconciliation to Option Pricing Theory

In this section, we briefly explain how deflators can be reconciled to option pricing theory. Suppose we consider derivatives paying some function $h(x)$ of an underlying asset x. Under RTSM, maybe x has probability density function $f(x)$. But the option prices are computed under some different risk neutral density $g(x)$. Denoting the zero coupon bond price by z, we then find that

$$
\begin{aligned}
\text{option price} &= z\mathbf{E}^g[h(x)] = z\int_0^\infty h(x)g(x)dx \\
&= \int_0^\infty h(x)\left[z\frac{g(x)}{f(x)}\right]f(x)dx = \mathbf{E}^f\left(h(x)\left[z\frac{g(x)}{f(x)}\right]\right)
\end{aligned}
$$

We can interpret the term in square brackets as the deflator. The full power of deflators uses multivariate distributions, but the principle is the same.

Other Measures of Value

The literature bristles with definitions of value, most of which are inconsistent with market prices. There may be situations where off market values are useful, particularly in the market for excuses. However, in this paper we focus wholly on the insights to be gained from market-consistent value metrics.

PART III: VALUING STAKES IN A WITH PROFITS FUND

The Parties Involved

Traditionally, asset-liability models have focused on the benefits to a particular party, such as policyholders or shareholders. There is rarely any attempt to reconcile the values to all stakeholders into some grand total, as we do here.

In our model, we do precisely that. In addition to the obvious stakeholders, we include the Inland Revenue and also a residual item for whatever of the estate remains at the end of the projection period. If we were to model other items of cost, such as expenses, we would need to include the recipient of these expenses (for example, employees) as a further stakeholder class.

The Role of Accounting Numbers

The management of any office requires some accounting of who owns what. At the most basic level, we can break up the office assets on a statutory or simple accounting basis, for example allocating asset shares to generations of policyholders, a tax provision (if any) to the tax man, and all the free assets to the estate. It is convenient to refer to all of these accounting numbers as generalised asset shares. Thus, we use asset share in its conventional sense when applied to policyholders, but we also refer to the estate as the "estate's asset share" and any provision for unpaid tax as the "tax man's asset share". The shareholders' asset share is set to be zero.

Such accounting allocations are necessarily arbitrary. They are really there to help keep track of what assets are where, and do not imply any distribution of economic worth. They gloss over many important issues, such as guarantees and future charges. But what is important is that the balance sheet balances. In our example the total of all asset shares at any point in time, and on any simulation path, adds up to the total market value of the company assets.

Valuing Cash Flow Streams

We can improve the measure of economic wealth looking instead at cash flows. In other words, the stake of a party would be valued with respect to the cash flows they derive from the insurer, rather than with relying on an initial balance sheet.

The particular cash flows we need to consider, together with our sign conventions, are as follows:

- For policyholders: premiums (positive) and benefits (negative)
- For shareholders: share of cost of bonus (negative)
- For the tax man: tax payments on income and realised gains (negative)
- For the estate: nil

Picking a Time Horizon

We cannot project cash flows for ever, and the infinite sums defining some stakes (for example, those of shareholders) may converge only slowly. So in practice, we project to some future, fixed, time horizon, and value any residual stakes using the accounting method at the terminal date. If our time horizon is short, the cash flows have a modest effect and we effectively recover our initial accounting balance sheet. But as the projection horizon is increased, the value of cash flows progressively supplants the value deriving from the terminal balance sheet, and the values converge to an economic measure of wealth.

A more sophisticated idea, which is commonly used in deterministic models, is to compute valuations at several time horizons and then to extrapolate to infinity, to that point where the terminal balance sheet has nil effect and the valuations reflect cash flows alone. We have not yet investigated whether this approach is useful in a stochastic setting, or whether the extrapolation would simply serve to amplify sampling noise in our value estimates.

Pricing Using Transfers

An alternative method for pricing stakes involves monitoring the flow of wealth between different parties, and valuing the wealth transfers. For each party, the transfer is defined by the equation:

asset share(t) = asset share(t-1) * [1 + gross fund return(t)] + cash flow(t) + transfer(t)

There is zero transfer if the change in the asset share is explained entirely by investment return and cash flows. Any other non-cash charge to the fund qualifies as a transfer. An alternative way of measuring a stake holder's stake is therefore to take the initial asset share and then add in the effect of any future transfers.

- Particular transfers we need to consider, with our sign conventions, are:
- For policyholders: guarantees when they bite (positive), guarantee charge (negative), share of tax transfer (negative), shareholder cost of bonus allocation (negative)
- For shareholders: shareholder cost of bonus allocation (positive)
- For the estate: guarantees when they bite (negative), guarantee charge (positive), share of tax transfer (negative)
- For the tax man: increase in tax provision (positive), return on tax provision (negative – this is the effect of tax deferral), actual tax payments (positive)

We can immediately see that these transfers add up to zero in total.

Equivalence of Transfer and Cash Flow Basis

We now have two ways of measuring stakeholder value - one based on cash flows plus a terminal balance sheet, and one which adjusts an initial balance sheet for future transfers. It is plainly important to establish whether these approaches do in fact give the same answer.

Numerical investigation indicates that these methods do indeed reconcile. The cash flow method gives the same stakes (to within simulation sampling error) as the transfer method. The mathematical proof uses the fact that the deflator correctly prices the gross fund return, and is in effect a stochastic analogue of the classical proof equating prospective and retrospective life insurance reserve calculations.

Given that these wealth measures coincide to within simulation tolerance, the criteria for choosing one method rather than another hinge on issues of communication and convenience rather than on financial theory. One practical issue is that transfer streams are often more regular than cash flows. This means that if any form of extrapolation for long horizons is to be attempted, then it is usually more convenient to use the transfer method.

We could take this one step further using hedge portfolios. Let us suppose that our accounting basis for asset shares was already close to a market basis for the cash flow streams for a particular stakeholder. Let us suppose further that we re-define the transfers to be relative to a hedge portfolio return instead of the overall fund return. This alternative definition would still preserve our identity between cash flow valuations and transfer valuations. But by the definition of a hedge, we would have

asset share(t) ≈ asset share(t-1) * [1 + hedge return(t)] + cash flow(t)

which implies that

transfer(t) ≈ 0

To the extent that transfers are non-zero, this must be either because the asset share basis does not fit the fair value of the cash flows, or because the investment risk is spread unevenly between parties, driving a wedge between fund return and hedge portfolios. So we can see transfers as being a balancing item to compensate for imperfect accounting or risk sharing. To the extent that accounting or risk sharing work well, we would expect transfers to small numbers relative to cash flows. This explains why it tends to be more convenient to value transfers than cash flows.

At the other extreme, if the accounting basis is meaningless then valuing transfers may be no more convenient than valuing the profit streams. A case in point is the shareholders' stake. Under our chosen accounting basis, this is valued at zero - no part of the estate is attributed to shareholders. As a result, the transfers are equal and opposite to the shareholder dividends, and it makes no difference which we choose to value.

PART IV: TENTATIVE MODEL CONCLUSIONS

Prototype Model

Although our model has significant complexities, it is still a gross simplification of reality. As a result, the model is more useful for educational purposes and for explaining concepts than for direct application to a particular insurer.

Most of the additional enhancements that would be required to make our model realistic are points of detail from a theoretical perspective. Allowing properly for expenses, more detailed tax calculations, mortality and other decrements, perhaps monthly projections, multiple policy terms in the same fund, and so on, increase complexity but raise few additional points of principle. The natural question is whether the basic numerical building blocks also scale up to purpose built, detailed, model office systems.

This is a reasonable question, because many of us will have seen models based, for example, on spotting hedges by inspecting analytical expressions or by the use of binomial lattice constructions. The chief difficulty with these methods is that they work fine in simplified examples, but as the products get more complicated, an increasing degree of ingenuity is required to spot exactly the algebraic transformation required to build the hedge or construct the lattice. For all practical purposes, products and financial market models of the complexity usually employed in life ALM, are not easily subjected to these brain-intensive pricing algorithms.

Deflators, however, are completely scaleable. There is no limit on the number of complexity of cash flows that can be valued. While our current model is intended for eventual distribution via the internet, and hence is written in the widely available Excel/VBA software, we who developed it became increasingly frustrated with the contorted process we had to pursue to get output in an easily analysed form. Fortunately, at least one of the proprietary model office systems available on the market offers a stochastic investment model supporting deflators.

Winners and Losers

One fundamental conclusion which our analysis has revealed is the way in which the total of all stakes must come back to the initial value of the company assets. There is simply no way around this - we could not get free lunches out of this model. If we change strategy for example with regard to equity backing ratio, bonus policy or guarantee charges, we may transfer wealth from one stakeholder to another. But we do not create wealth over all.

As a result, there is no way of making everybody better off at once when a market value method is used. This is the principle of no free lunch. This requires us to rethink some common myths.

The Value of Investment Freedom

We regard as a myth the idea that greater free assets allow more investment freedom and hence higher returns, which benefits policyholders in the long term. In contrast, under our model, moving policyholders to a higher risk / higher return position neither creates nor destroys value, any more than selling bonds and buying equities creates value. £100 of equities with worth the same as £100 of bonds, even if the payoffs are measured over a long time horizon.

This has some unpalatable consequences. Many demutualisations have been permitted, on the grounds that greater policyholder returns are possible by the adduction of shareholder capital. The long term value to policyholders arising from greater investment freedom is a benefit which outweighs the immediate cost of sharing of policyholder bonuses with a new class of shareholder. Our model, however, reveals only a transfer of wealth from policyholder to shareholder on demutualisation. The balancing "value" supposedly derived from greater investment freedom is a risk-return illusion - this value is actually zero on a market basis.

Inherited Estate

We also investigated the consequences of allocating some the inherited estate to shareholders. Accepted wisdom dictates that much of this inherited estate is not needed, and so the allocation has no detrimental effect on policyholders. Our experiments show a contrasting result, that if £1m is taken from a with-profits fund then the other stakeholders are, in aggregate, exactly £1m worse off.

One of these other stakeholders is, of course, the value of the estate carried forward at the horizon. But the further ahead we look, the smaller this item becomes, if we assume the estate grows less quickly than investment returns. This means that ultimately the benefit of that estate would have accrued to other parties. In our projections, it is surprisingly common that policyholders need to call on the estate to meet guarantees. In this sense, the last twenty years of astoundingly high equity returns have been exception, and we cannot count on a repeat performance. The maintenance of a healthy estate as important as ever.

This does not imply that all representations made to regulators on the subject of inherited estates, are necessarily misplaced. For example, an Appointed Actuary has no duty to future potential policyholders until they actually become policyholders. So if it can be demonstrated that a large inherited estate would otherwise have been used to subsidise future unprofitable business, a shareholder allocation may be justifiable without compromising the reasonable expectations of today's policyholders.

Measuring Fair Costs of Guarantees

The issue of identifying, and correctly charging for, the costs of with-profits guarantees has been recognised as a problem for at least twenty years. One common way of approaching the solution is to examine a long run average cost of applying guarantees over many generations of policies. This would be equated to a percentage charge on asset shares, applied either on maturity or every year.

Unfortunately this long term average logic is dangerously flawed. The problem is that to value guarantees properly we need to know not the mean cash flows but their mean deflated values. And here we run into the biggest problem - that while we collect charges in good scenarios (low deflator) we pay the guarantees when the deflators are highest. So on a market basis, even if we think we break even in some magical long term, we could be dramatically under-charging for the guarantees.

In contrast, our deflator method of pricing transfers gives us a market-consistent method of pricing both the guarantees and the guarantee charges. This puts a higher value on the guarantees than do traditional methods. Given that current product designs have often not been developed with market-consistent pricing in mind, our results have shown that in many cases there are net wealth transfers out of the estate. The rate at which the value of the estate is depleted varies according to the value of the guarantees. The guarantee cost is increased by higher equity backing ratios and by the declaration of high reversionary bonuses. Indeed, the main effect of changing investment or bonus strategy seems to be to affect the pace at which the estate is eroded for policyholders' benefit; the shareholders are affected to a smaller degree.

The high value that we put on the guarantees perhaps explains why with profits guarantees are popular with customers. The trouble is that past experiments by the life industry in giving away under-priced guarantees have not always seemed so clever with hindsight. Will the whole risk sharing pyramid unravel if the next generation of capital-providing policyholders fails to show up on time?

Tax and Regulatory Effects

Our model has also enabled us to pick up some more marginal effects relating to the effect of investment policy. A shift to a higher equity backing ratio can in fact create value overall at the expense of the tax man. This is because of the way in which dividend income is already franked, while bond income falls directly into the income tax net. We can only detect this effect because we have correctly adjusted for the market risk. Our reason for preferring equities lies in a tax arbitrage, which is risk-free, rather than in the notion of long term risk premiums, which should be discounted on the grounds of the extra risks that come with them.

A further subtle reason why shareholders may benefit from equity investment lies in the calculation of profit sharing. The shareholder share of bonus is computed with reference to the FSA's valuation basis. If more equities are held in the fund, this FSA basis becomes stronger, putting a greater value on the a given policyholder bonus, and therefore allocating a greater economic proportion to shareholders. This effect is to the detriment of policyholders, but fortunately the other benefits of equity investment, namely the increase in guarantee value and greater tax efficiency, outweigh by many times the distorting effect of FSA valuation bases.

Fair Value is Achievable

The final, and perhaps most important, result we have obtained is the fact that we could get results at all! In particular, all our deflator valuations are consistent with the fair value criteria laid down by the International Accounting Standards Committee. This is significant, because as far as we are aware, and apart from the current project, nobody has actually succeeded in constructing a valuation which satisfies all the criteria. Indeed, some have gone so far as to claim the IASC criteria are too tough, and that some less robust and more flexible approach such as embedded value or US GAAP should be preferred instead. We believe the current work refutes such claims; the IASC criteria can be applied in practice, even to complicated assets and liabilities; furthermore, our model does it.

PART V: POSSIBLE FURTHER INVESTIGATIONS

Preliminary report

This paper is a preliminary report of work in progress. There are many more issues to investigate, and we hope to report more progress at the Life Convention. Some issues of current interest are listed below.

Dynamic Investment Strategies

We have so far only investigated the effect of strategies with a constant equity backing ratio. In practice, many insurers try to replicate option characteristics by matching more closely in the face of falling solvency levels. We suspect that this gives rise to subtle wealth transfer effects, and look forward to using the model to investigate this.

Cohort Analysis

We have up till now investigate output where all liabilities are reported in aggregate. However, we believe much could be learned by examining the wealth of cohorts separately. This is likely particularly to be especially illuminating with regard to issues of equity between generations. As we see it, there are currently two popular definitions of fairness:

- The social pooling concept, that prices are equitable if everyone sees the same price. An example of equitable pricing would be to offer (non-profit) annuities at the same rate every year
- The market concept, that prices are equitable if everyone bears the same charges relative to a market return. In this case, unit linked policies are equitably priced if they all bear the same charges.

One man's equity is another man's cross subsidy. For example, with the social pooling concept of equity, unit linked policies are inequitable because those who get poor returns in some sense are subsidising the good returns of the next generation who buy in at a knock-down price. Equally well, the market equity concept would rule out constant annuity pricing bases, because at times of high interest rates customers could get a better deal elsewhere.

You can't have both kinds of equity at once. This is the toughest fact in with profits business. Furthermore, social pooling equity is vulnerable to adverse selection. It can only be sustained if you have an imperfect market, for example with compulsory entry or asymmetric information. To the extent that WP contains an element of social pooling, the advent of greater transparency is a threat to WP as traditionally practised.

It is likely that the advent of stakeholder, and new international accounting standards, will increase transparency, thus creating more informed customers (or IFAs) and increasing the threat of adverse selection against the fund. For example, if the smoothing mechanism were better known, policyholders might time their entry to follow a market rise, or target offices who might be distributing estates.

The rational response to this is to remove the social pooling element of with-profits. It is still possible to provide the cross-generational guarantees, but they must be equitably charged for, in a market sense. This can be accomplished by dynamic re-pricing, changing premium rates year by year. Our model shows how this can be done. This produces a number of advantages; as well as a defence against adverse selection, the demonstration of market-sense equity could be the path to a stakeholder badge. It is also likely to provide a clearer focus on fair values and shareholder impacts.

Sensitivities and Hedges

We plan to extend our model to compute not only fair values of different stakes, but also the hedge portfolio. Where we now have an allocation of the company asset value between different stakeholders, in future we will be able to allocate the assets themselves. In other words, if the company has £100 of gilts and £200 of equities invested, we will be able to investigate who bears the risks (and gets the returns) from each asset class. This might reveal, for example, that a young policy generation has an effective exposure of minus £10 in gilts and +£20 in equities, the gearing effect arising from the provision of guarantees to maturing policies. This will help us to measure whether the policyholders are actually getting the economic exposure they have been led to believe – a fundamental question when considering PRE issues.

Other Types of Profit Sharing

Finally, we look forward to applying our methods to other forms of profit sharing contract, in particular, unitised with profits and continental-style profit sharing funds.

Acknowledgements

I am deeply indebt to Paul Coulthard who built most of the liability model underlying this note. I am also grateful to Shyam Mehta, Jeroen van Bezooyen and Jon Exley for many comments and bright ideas on earlier versions of this paper. Any errors remain my own.

APPENDIX: DETAIL OF MODEL CONSTRUCTION

Setting up a With-Profits Model.

The following describes the general approach, tailored specifically to traditional with-profits business. The model consists of a life insurance fund with several parties, including:

Each cohort of policyholders
Shareholders
The Inland Revenue ("Tax man")
The estate

Conventions

We are interested in identifying cash flows to each of these parties. Our convention is that a cash flow into the fund and from one of the parties is a positive flow. So, for example, benefits are a negative flow to policyholders and premiums are a positive flow. Cash flows, as distinct from the transfers we discuss below, to and from the estate are always zero. For ease of illustration, al cash flows are assumed to arise at the end of each year.

Total assets means the total of all assets inside the fund, calculated at market value ("MV"). These are measured immediately after all cash flows.

The equation for updating assets is that

Total assets(t) = Total assets(t-1) * [1 + gross return(t)] + total cash flows(t)

Asset Shares

In addition to cash flows, we maintain a quasi balance sheet, in which the total assets are broken down into notional stakes ("asset share") for each party. In addition to the asset shares for each generation (or "cohort") of policyholders, we include the asset shares of the estate, the tax man and the shareholders. The estate is treated as another, anonymous, party who receives no outward cash flows. We do not try to apportion ownership of the estate, so under our accounting convention, the shareholders' asset share is always zero. The tax man asset share is the provision for CGT.

We define a series of transfer variables to the pot of each party *i* by

Asset share(*i,t*) = Asset share(*i,t*-1) * [1+gross return(*t*)] + cashflow(*i,t*) + transfer(*i,t*)

Our arithmetic guarantees that the total of all transfers across stakeholders is always zero. Examples of transfers include asset share deductions for expenses, tax or for costs of guarantees. Our sign convention is that all of these would be negative transfers for the policyholder, or positive transfers to the tax man. As the shareholder always has zero asset share, we must assume that the (negative) dividend cash flow is always offset by an equal and opposite positive accounting transfer.

Transfers also include any difference between final payouts and the asset share, due to guarantees for example. In this case, the amount of benefits (in excess of asset share) would be a positive transfer to the policyholder and negative transfer to the estate.

Economic Values

We can value ("PV") all flows in the asset share equation to obtain:

$$PV_0[\text{asset share}(i,t)] = \text{asset share}(i,0) + \text{sum}\{\ PV_0[\text{cashflow}(i,s)]: 1 \le s \le t\}$$
$$+ \text{sum}\{\ PV_0[\text{transfer}(i,s)]: 1 \le s \le t\}$$

This has the natural interpretation that everyone gets out what they put in, plus or minus anything they obtain from other parties.

Model Time Horizon

We are interested in policies of a fixed term, *polterm*. Initially, at time zero, we will still have on the books policies that incepted on –*polterm* + 1.

We project through to *horizon*. At that point, we will have to take account of some business that has just incepted. Consequently, we need to model asset shares etc for cohorts incepting between –*polterm* + 1 to *horizon* inclusive.

Model Variables

The idea of the model is to have a total office loop which progresses from year to year. From the start of a year, immediately after all cash flows, we must then apply the following steps.

- apply the gross investment return to all asset shares
- allowing for purchases and sales of shares ("turnover"), update asset book value
- compute tax payable and the CGT reserve
- deduct tax and the increase in CGT reserve from asset shares
- compute the raw bonus rate such that, under the Actuary's own valuation basis ("BRV"), PV of future RB's is equal to a proportion of (asset shares + PV future premiums)
- Compute actual bonus rate by applying maximum and minimum changes relative to last year's rate.
- Award RB, and compute its cost on the statutory ("FSA") basis
- compute the terminal bonus ("TB") cost, then settle benefits on maturing cohort
- Deduct 10% of the cost of bonus from the asset shares and pay it to shareholders
- receive new premiums
- rebalance to the desired asset allocation (proportion invested in equities, or equity backing ratio ("EBR").

Then the whole process starts over.

Initialising the Variables

To start with, we need to know the asset share and accrued bonuses for existing policies. This should reflect the actual experience of the past, and so requires a table of historic bonus rates and of historic returns (net of deductions such as tax, shareholder profits and so on).

The only complicated issue is setting up the model initial book value. The asset share and estate will be computed net of any CGT provision. We also take as given an equity backing ratio, and the following equations apply:

policyholder asset shares + estate = bond MV + equity MV – CGT provision
equity backing ratio = equity MV / (bond MV + equity MV)
book value ratio = equity BV / equity MV

We also have the equation:
CGT provision = tax rate * max(equity MV – equity BV, 0)

This gives us four equations, with four unknowns (bond MV, equity MV, equity BV, CGT provision)

Substituting for the book value ratio into the CGT definition, we have

CGT provision = tax rate * equity MV * max{ 1 – book value ratio, 0}

Then our equations become:

policyholder asset shares + estate
= bond MV + equity MV * (1 - tax rate * max{ 1 – book value ratio, 0})

= (bond MV + equity MV) * (1 – equity backing ratio
+ equity backing ratio * (1 - tax rate * max{ 1 – book value ratio, 0}))

This is sufficient to determine bond and equity MV's, from which other quantities can readily be determined.

Rolling up Returns gross and Net of Tax

Rolling up returns is more complex than it looks. To make matters simple, we ignore indexation relief on capital gains tax. At the start of the year, on the balance sheet date, we are given the following:

Equity market value: EMV_0
Equity book value: EBV_0
Bond market value: BMV_0
There is a CGT provision of taxrate * max(EMV_0 – EBV_0,0). The market value of assets less this provision, are attributed to asset shares and the estate.

Over the year, we observe the following percentage returns from the capital market generator (all income occurring at the year end):

Equity capital return: ECR
Equity income return: EIR
Bond total return: BTR

Thus, just before the cash flow payments, we have

Equity market value: EMV_0 * (1 + ECR)
Dividends received: EMV_0 * EIR
Equity book value: EBV_0
Bond market value: (including reinvested coupons) BMV_0 * (1+BTR)

Next, we allow a proportion *churn* of the equity fund to be sold and bought back, also within this split second between year end and striking of the accounts. This results in:

Equity market value: EMV_0 * (1+ECR)
Dividends received: EMV_0 * EIR
Equity book value: (1- churn) * EBV_0 + churn * EMV_0 * (1+ECR)
Bond market value: (including reinvested coupons) BMV_0 * (1+BTR)

This results in the following tax flow out of the fund:

$$\text{Tax flow} = -\,\text{taxrate} * \max\{0,\ EMV_0 * EIR + \text{churn} * [\,EMV_0 * (1+ECR) - EBV_0\,] + BMV_0 * BTR\,\}$$

We assume that all items of taxable income are offsettable against other items. We ignore any tax credit which might be taken forward.

We can then pay out any tax, and set up the new CGT provision. This tells us the total assets.

Apportioning Investment Returns and Transfers

Rolling up investment returns and allowing for tax implies a set of transfers between various parties, in order to make the equations balance. This requires us to attribute a gross return to each party. We can now compute revised asset shares, apportioning between generations and the estate according to the starting position and the gross return.

Theoretically, it should not matter what gross return we use, provided that in aggregate we get the gross fund return. The simplest approach is to attribute the same return series to all stakeholders.

Reversionary Bonus Calculation

The next step is to compute the value of future reversionary bonuses, according to the actuary's "realistic" basis.

The idea is to project future reversionary bonuses to maturity (allowing for the policyholders share to be credited). For a declared bonus rate of *b*, our bonus reserve valuation does not explicitly project future shareholders' transfers out from the asset share, but uses an assumed bonus rate of *b*/(1-shxfer).

This needs to start from last year's cumulated guarantees. These projected values are then discounted at a gilt yield (net of tax), and plus or minus specified prudence margin) , and compared to a target proportion of (asset shares + PV future premiums). In practice, different life offices will have different formulae and approaches for setting the target.

We calculate the reversionary bonus rate that hits the target, or gets as close as possible subject to our bounds on bonus rate movements from year to year.

Terminal Bonus Calculation

The terminal bonus is calculated as

max{ asset share – guarantee deduction – sum assured – reversionary bonuses, 0}

The guarantee deduction is a proportion of the asset share.

The extra guaranteed benefits, that is, benefits in excess of asset shares, have to be financed from somewhere. The first port of call is the estate; when this is used up we deduct from other policyholders' asset shares (in proportion). If these are used up, then the office distributes asset shares to all policyholders and then closes to new business.

As a tedious book-keeping point, we must remember to set the asset shares for the maturing policies to zero.

Shareholders' Distributions

The cost of bonus is computed as:

- the amount of terminal bonus distributed
- plus the cost on the FSA basis of reversionary bonus declared

The second item is the present value of the increase in sum assured, discounted at the FSA discount rate. The FSA discount rate is the weighted average of gilt yields and equity yields, according to the specified mix, and netted down for tax. (we don't allow for games with matching rectangles).

A proportion shxfer of the cost of bonus is deducted from the asset share. This is applied on a per cohort basis.

Resolving the Circularity

Unfortunately, there are some circularities here. These are as follows:

For calculating the BRV, should we really be equating to the asset share after the deduction for the shareholder transfers relevant to the year just past?
Also, in calculating the terminal bonus, we should deduct a proportion of TB too.

The first concern is in fact misplaced. If, for example, shxfer = 10% and we declare a 4.5% bonus, the BRV will have already projected on a 5% bonus – starting from one year ago. So the 0.5% margin for transfers is already allowed for.

Addressing the second concern results in the equation:

TB = max{ asset share – gtee deduction - shxfer * TB – sum assured – RB, 0 }

The solution to this equation is:

(1+shxfer) * TB = max{ asset share – gtee deduction - sum assured – RB, 0 }

Rebalance to target mix

Finally, we have to rebalance the assets to the required EBR for next year.

We have been careful to make sure all along that the total value of asset shares (including CGT provision and estate) is equal to the total assets by market value. We also need to maintain and record book values.

Our assumption is that capital gains are triggered only by the operation of churning. Other than this, the value of unrealised gains is unchanged by cash flows. This is equivalent to assuming that benefit payments etc are made out of assets just churned (or out of premiums received and not yet invested). This assumes we have churned enough assets to allow this – ie ignores the possibility that further CGT may be triggered by additional asset sales. Implicitly we are assuming that in this case, the CGT can be deferred by transferring it to any retained assets. This is in effect what companies do when they realise under-performing assets first in order to minimise CGT crystallisation. Clearly, the code could be rewritten to allow for alternative capital gains tax strategies.

Receive new premiums

This is fairly simple. We need to make sure that asset shares are credited, and that guaranteed sums assured are set up.