Presented to the Institute of Actuaries Students' Society

on 17th March 1981

FINANCIAL OPTIONS IN LIFE ASSURANCE POLICIES

by

R. G. Gillespie

Presented to the
Bristol Actuarial Society, 11th March 1982
Faculty of Actuaries Students' Society, 6th December 1982

INTRODUCTION

The purpose of this paper is to describe very broadly the various methods of dealing with options in life policies. In particular the traditional methods have been contrasted where appropriate with the new methods developed to deal with maturity guarantees. In the descriptions the emphasis has been on the broad concepts rather than the detailed results and methodology. The reason for this is that a lot of the detailed work for the new methods has been described in the Appendices - including the bibliography - of the Maturity Guarantees Working Party Report JIA Vol 107 - Part II.

Basically, options occur when there is a choice of benefits or benefit formulae available. The choice may be the policyholder's or he may automatically get the higher of the values produced by the benefit formulae. It can be seen that this will include guarantees. The choice may also involve a change in premium.

The various types of option have been categorised for the purposes of this paper. As can be seen later some of the categories have attributes of some of the others. However, it is hoped that the distinctions drawn between the categories will show some of the ideas and problems that should be thought about when dealing with options in practice.

The views expressed in this paper are my own and not necessarily those of the office I work for.

A CLASSIFICATION OF OPTIONS

I. MORTALITY OPTIONS

In these, the assured has the right to take out a policy of a specified class or classes at ordinary rates at a specified time or period as a result of taking out a policy. The main advantage to the policyholder is that, whereas the original policy is underwritten in the normal way, in the second policy this procedure is by-passed. This option is paid for by a special extra premium or it may be included at no extra charge in the ordinary rates.

An example of such an option is that included in the Convertible Term Assurance, where the original policy is a standard term assurance with the right to effect another policy (say a Whole Life) at any time during the currency of the original policy. The premium rates for the Whole Life policy would be those current at the time of conversion.

II. ANNUAL PREMIUM RATE GUARANTEES

Here, the policyholder has the right to change the terms of his policy at a specified time or during a specified period. The rates of the altered policy are guaranteed at the time the original policy is taken out. This includes the taking out of a new policy at the termination of the original. Examples of this type are a 10 Year Option Policy, which is a Whole Life Policy with guaranteed conversion terms to an Endowment Assurance in the first years, and a Child's Deferred Assurance with its guaranteed terms at vesting date.

III. MATURITY GUARANTEES

Although easy to describe, this guarantee has caused the profession great argument and heart-searching. It is the guarantee of a sum of money as a lower limit of the proceeds of units from a unit-linked policy at maturity. Despite its name this guarantee is usually extended to payment of proceeds on earlier death.

IV. SINGLE PREMIUM RATE GUARANTEES

Here the policyholder has the right to effect a single premium policy at a specified time, using rates specified at the time the original policy was effected. Usually the specified time is at maturity and the single premium is limited to the proceeds of the original policy. Examples of this class are the Guaranteed Annuity Option of an Endowment Assurance and the Guaranteed Annuity of a Deferred Annuity which has a Guaranteed Cash Option.

V. BENEFIT GUARANTEES FOR UNIT-LINKED POLICIES AVAILABLE OVER A PERIOD

This is an enhanced version of the Maturity Guarantee above whereby the guarantee is available not only at maturity but also at surrender before maturity and, perhaps, also if the proceeds can be left invested in the units after maturity.

VI. BENEFIT GUARANTEES FOR NON UNIT-LINKED POLICIES AVAILABLE OVER A PERIOD

In this, the benefits can be taken when the policyholder chooses and he has a guarantee as to the amount. Examples of these classes are Endowments with guaranteed surrender values, Flexible Endowments and Self-employed Deferred Annuities.

COMMENTS ON THE OPTION CATEGORIES

I. MORTALITY OPTIONS

Valuation

Essentially, in a Mortality Option, the policyholder is given the choice of taking out a policy on standard rates with standard underwriting procedures and having the same policy on the option rates with no further underwriting. The cost of the option to the office can be analysed into the following factors:-

1. Because the normal underwriting procedures are being by-passed, the mortality of those who exercise the option is likely to be worse than that assumed in the standard rates. If the option rates are equal to the standard rates then the mortality difference is going to be the difference between ultimate and select assuming that the policyholder does not select against the office. However it is likely that there is some selection against the office so the mortality difference will be larger still.

2. Because the normal underwriting procedures are being by-passed, the initial expenses will be less than normal. So if the option and standard rates are the same there will be a saving available to offset the mortality cost. This factor is of increasing importance because of the cost of standard underwriting procedures.

3. If exercising the option involves the discontinuance of the existing policy, (as in a Convertible Term Assurance), then the release of reserves can also reduce the cost.

In calculating the cost of this option all factors must be taken into account but usually the most significant of these is the first factor, mortality. Assuming there is enough data, the traditional way of estimating the cost of the mortality part is to carry out, firstly, an investigation to find what proportion of the policyholders exercise the option and, secondly, a mortality investigation in which the mortality of the option exerciser is compared with that assumed in the premium rates.

So the average proportion of policyholders exercising the option and the average mortality difference has been calculated. This is really a comparison of averages between the average mortality used in the option rates and the average mortality observed in the investigation.

The cost for those who have exercised the option can be found by assuming that all members of the age group actually experience the observed average rather than that assumed in the rates. The cost for those who have not yet exercised the option is calculated by assuming costs consistent with the cost for those who have already exercised the option and multiplying by the proportion who exercise. This will have to be discounted from the time of the option being available to the present.

Now in practice, it is unlikely that the experience would be large enough for such an investigation so a reasonable approach might be to take a conservative set of average mortality rates consistent with such data that is available. This would provide a margin to cover:-

(a) errors in the estimates because the investigation covered only a small number of lives;

(b) errors in the estimates because the population of option exercisers is small and hence subject to random fluctuations.

Why does such a method work?

It is reasonable for us to assume that the mortality of the lives are independent of each other. (Luckily most people do not die of airliner crashes or multiple car accidents).

This has two consequences:-

1. Assuming that the averages are correct and that the experience is large enough, those that have below average mortality will tend to balance those that are above average (i.e. the random variations will tend to cancel each other out).

2. The larger the number of lives involved the smaller the dispersion about the average. This means that large fluctuations about the average are only likely when the number of lives involved is small. It is at that time that the total risk is smaller and hence more approximate methods are appropriate.

These result from the application of the Central Limit Theorem which states:-

> Given a series of "n" independent random variables whose probability distribution has a mean (a) and a finite variance (v), then the arithmetic average of the "n" random variables will tend to be distributed Normally about the mean (a) with variance (v/n) as "n" gets larger.

Consequence 1 comes from the fact that the Normal distribution is symmetrical. Consequence 2 comes from the fact that (v/n) gets smaller as "n" gets larger.

Premiums

The extra premium required is calculated in the traditional way so that, at the start of the policy, the present value of the expected extra liability calculated as above is equal to the present value of the expected extra income assuming that the assumptions about interest, expenses etc are true.

II. ANNUAL PREMIUM RATE GUARANTEES

A classic example of this type is the guaranteed option to convert a Whole Life policy into an Endowment Assurance.

As all the premium rates available in this option are set at the start of the original policy, it can be looked at as a predictable alteration. Hence the guaranteed premium rates available can be calculated in the same way as premiums on alteration, (say, by equating the policy values before and after the alteration.)

The rate of interest to be used is the same rate as used in the pre-option rates unless there is a significant difference in the length of the two possible policies which may result from the option. Normally, the same rate is used.

The mortality basis should also remain the same because selection against the office should be avoided by not giving the option of continuing the same cover at a lower premium rate (e.g. Whole Life to Endowment is acceptable but Endowment to Whole Life is not).

There should be a small charge on the policy value to cover the cost of the alteration if the option is taken.

It can be seen that, if the policy with the option is valued on the premium rate basis before the option is exercised, then the liability ignoring the charge above will be the same whether or not the option is assumed to be exercised. Hence the liability is covered. However, it is likely that the valuation basis will not be the premium basis and therefore this equality will not be there. The traditional approach is to value on the basis of the alternative that gives the higher reserve.

This is fine, but there is a problem. Earlier the possibility of using a different rate of interest if the option had a significantly different length of liability was mentioned. Such options will have a matching problem and this can be serious because a mismatching reserve will be required. (As will be seen later there are problems in calculating a mismatching reserve). It is best if the problem is avoided by not offering an option which will cause a serious mismatching. Luckily endowments of term shorter than ten years are not generally offered.

III. MATURITY GUARANTEES

A common example of this type of guarantee is when the assured under a unit-linked policy receives the higher of the value of the units and the premiums paid at maturity or earlier death.

Valuation

This type is different from the mortality option because:-

1. No real choice is made by the policyholder so the fact that their circumstances may differ does not affect the issue.

2. The option depends on the unit price and the premiums paid only, hence all policies taken out at the same time and maturing at the same time will be affected in the same way - they will not be independent of each other.

3. The random variable in this option is the unit price, not mortality as before, and its variance does not change as the number of policies increases.

Reasons 2 and 3 mean that the reserve required to cover random fluctuations does not diminish in importance as the portfolio gets larger.

By what criteria should the size of the fluctuation reserve be chosen?

In both solvency and profitability valuations, this reserve should be large enough to give a good measure of security to the portfolio (i.e. it should only be exhausted under very rare conditions). Traditional actuarial methods make allowances for fluctuations by taking margins in the normal assumptions. This is not appropriate here, so what method should be used?

This is very similar to other problems where there is a structure which should only fail under rare conditions; examples are risk factors in civil engineering and general insurance. Risk theory has been developed to solve these problems.

The risk theory approach is to calculate a reserve on the basis that the probability that the reserve will be exhausted is acceptably low. Because the risk associated with each policy is not independent of other policies, a portfolio reserve is more appropriate than valuing each policy independently.

<u>Why is a portfolio reserve appropriate?</u>

This can be seen by comparing two simple Portfolios A and B which consist of 5 identical unit-linked policies with maturity guarantees of 5 years' duration and whose policies have a different relationship to each other.

<u>PORTFOLIO A</u>

```
POLICY NO

5                   -------------------------------
4                -------------------------------
3             -------------------------------
2          -------------------------------
1-------------------------------

0     1     2     3     4     5     6     7     8     9     10   YEARS
```

<u>PORTFOLIO B</u>

```
POLICY NO

5             -------------------------------
4             -------------------------------
3             -------------------------------
2             -------------------------------
1             -------------------------------

0     1     2     3     4     5     6     7     8     9     10   YEARS
```

In A the maturities of the policies have been spread over a number of years; in B they all share the same experience of the unit's performance.

Consider a simplified unit price "performance" which is likely to cause a claim on the maturity guarantee reserve; this is a stable unit price for six years followed by a steep fall in the seventh year. Under this experience:-

In portfolio B: All policies cause a strain on the reserve.

In portfolio A: Policies 1 and 2 are unaffected because they mature before the event.

Policy 3 will cause a strain on the reserve in the same way as the policies in portfolio B.

Policy 4 will cause a strain unless there is a substantial enough recovery in the eighth year.

Policy 5 will cause a strain unless there is a substantial enough recovery in the eighth and ninth years.

So portfolio A is affected by this experience less than portfolio B. Against that, portfolio A is exposed to such an event for a longer period. This means that a guarantee reserve which covers fluctuations depends on the relationship of one policy to another and therefore a portfolio reserving method is appropriate.

<u>What is an acceptably low probability for this reserve?</u>

It could be calculated on the basis that it is impossible for it to run out. This would mean taking the largest possible maturity guarantee claim for each policy in the portfolio. Now the largest possible claim for a policy is to assume that the unit value is zero and hence the maturity guarantee amount is the guaranteed sum assured. These can be valued in the same way as a pure endowment. This method has problems.

1. A consistent continuing value of zero for a unit price is not practical; the fund would either be wound-up and the policies made paid-up or the continuing flow of premiums should make the price recover.

2. The resulting value of the liability is going to be very large, being of the order of size of that for an Endowment Assurance. Remember that this is only a minor part of the valuation of the policy, the major part being the value of the units already bought. This liability cannot replace the value of unit reserve, as the zero price assumption might imply, since a guarantee reserve must be available before claims are payable and should be set up whatever the value of units.

A reserve of this size is only justified if a substantial proportion stands a chance of being called upon in practice.

<u>How can it be tested whether a large proportion of such a reserve is even going to be needed?</u>

The probability distribution of the reserve required for the portfolio must be estimated. It can be seen that this depends on estimates of the probability distribution of the unit price at various times in the future. Because of the lack of data this last probability distribution can only really be estimated by "modelling" Stock Exchange prices. Both evidence and the fact that the unit fund's portfolio is a sample of Stock Exchange securities suggest that this is a reasonable approximation.

There are problems in modelling Stock Exchange prices.

1. There are the usual philosophical problems associated with any time series about the extent to which the future is related to the past.

2. If I or anyone else were to find an accurate Stock Exchange model, they would be most unlikely to publish it but rather to retire to a paradise of their choosing!

3. Were an accurate Stock Exchange model to be published, then the predictions of the model would be anticipated by the market and this would change the prices in the market. So to continue to be accurate, the model has to be able to predict the result of its own prediction being known to the market. This would be difficult.

These problems are formidable but I believe they can be avoided by not being too ambitious. Models which have problems 2 and 3 are models which predict the time of price moves as well as the direction. What is required is a model which will help us predict the probability distribution of the maturity guarantee claims. This does not require an accurate prediction of prices but an estimate of the probability distribution of price with sequences of prices that are reasonable.

Choosing the best model for this purpose is a very large task and is a subject in itself. (See the work and bibliography of the Maturity Guarantees Working Party JIA Vol 107 - Part II mentioned earlier).

<u>Having estimated a probability distribution for the size of the claim for the portfolio, what relation does the claim bear to the maximum possible?</u>

All models of unit funds which automatically have their dividends and investment interest reinvested (accumulation units) tend to have a built-in upward trend because:-

1. it is indicated by the historical data;
2. a trend in any other direction implies negative interest and would be altered by the operation of market forces or by governments.

Hence the expected value of the claim is small and the results of virtually all the models I know show claims to sum assured ratios no greater than the order of $\frac{1}{3}$ to $\frac{1}{2}$, even at an extreme probability of 1/5000. Hence it is extremely unlikely for a large proportion of the maximum reserve to be used so it can be released without weakening the reserve significantly.

<u>How is the probability distribution of the amount of the reserve used calculated, having chosen the Stock Exchange model?</u>

Let

	$R(t)$	=	the price of the accumulation unit on policy anniversary "t",
	n	=	the term of the policy,
	G	=	the guaranteed sum assured,
and	P	=	the premium invested in the units each year.

Then the proceeds of the policy (excluding guarantee) at time "n",

$$= P\left(\frac{R(n)}{R(0)} + \frac{R(n)}{R(1)} + \frac{R(n)}{R(2)} + - - - + \frac{R(n)}{R(n-1)} \right)$$

and the amount of the guarantee claim

= the maximum of 0 and (G - Proceeds)

To calculate the probability distribution of the amount of the guarantee claim using algebra, the distribution of the policy proceeds must be found first. Looking at the formula above it can be seen that this is not easy because in most models the R's are not independent. (The easiest way to test that is to change R(t) to R(t)+e. If any of R's for values greater than "t" change, then the series of R's are not independent).

Actuaries tackling this problem have found an easier approach to this problem than that, namely, simulation. This is a process whereby a string of independent random variables is used to generate a string of "future" unit prices using the formulae of the Stock Exchange model. From these not only the policy proceeds but also the amounts required for the guarantee claims of the portfolio can be calculated. This process is repeated with another string of independent random variables to produce another "future". The process is repeated again and again to produce a whole series of futures based on the assumptions and parameters of the models. These futures are independent of each other and so can be classified by the size of their guarantee claims. Hence a frequency table can be produced and this can be used to estimate the probability distribution of the claim amounts. Obviously, the number of futures used should be large enough so that adding a further set of futures does not alter the estimate significantly, bearing in mind the accuracy which is being used.

These probability distributions show that even for extremely strong models at extremely low probabilities, a large proportion of the full Guaranteed Sum Assured is not required. Hence reserving on that basis is a waste of money to all concerned (to the office and to the policyholders whose premiums would have to service the reserve with interest).

There are other factors that affect the size of the guarantee claims reserves. These are:-

1. <u>The Interest Rate</u> This is used to discount the simulated claim to the present time. The assets backing a reserve of this nature should have security of capital value because it will be required when security values are low. This means that the interest earned will be related to the short end of market; therefore the rate of interest used should be the long term rate that is expected to be available for short term money and should not fluctuate with the short term rate available in the market. The method of application is to produce a present value of guarantee claims for each "future" before it is incorporated into the frequency table mentioned above. This means that the probability distribution shows the probability of the present reserve being exhausted at sometime in the future, for varying sizes of reserves, allowing for future growth through interest and excluding any other input to the reserve or new policies to the portfolio.

2. <u>The Mortality Rate of the Lives Assured</u> The effect of mortality is to reduce the number of maturities and hence the amount of guarantee claims. However, it will also introduce a new series of death guarantee claims because if there is a maturity guarantee there is usually a similar guarantee on earlier death. The probability distribution for the reserves can be altered by adjusting the portfolio's pattern of maturities. This is done on the traditional

"expected" method using the same mortality as used for valuing other parts of the unit-linked policies.

i.e. Let G(x,t,d) = The Guaranteed Sum Assured maturing at time "t" for those lives at present aged "x" with duration "d"

then the value at maturity of

$$G(x,t,d) \quad \text{is} \quad \frac{lx+t}{lx} \cdot G(x,t,d)$$

The same process is applied to all the projected elements, such as premium income and units bought.

In addition, there is a new element

Let D(x,t,d) = The Guaranteed Sum Assured on Death at time "t" at present aged "x" with duration "d"

$$\text{then the new element} = \frac{dx+t}{lx} \cdot D(x,t,d)$$

This "expected" method can be used for the same reason it worked for the Mortality Option.

3. The Withdrawal Rate If it is appropriate to take explicit account of withdrawals in the valuation, (here the important thing is to be consistent with the valuation of the other parts of the same portfolio of policies!) then the rates can be used to reduce the guaranteed sum assured at maturity and the units. Both these are treated in the same way as mortality. However, I consider that only the early duration withdrawals should be allowed for because at later durations the withdrawal rate may not be independent of the unit price level. If there are depressed prices, those policyholders who know that, barring an economic miracle, they are going to get a

guaranteed maturity value greatly in excess of the value of the units are less likely to withdraw. Withdrawals at these durations may not necessarily reduce the guarantee claims - so I think it prudent to assume none.

Having obtained the probability distribution of the guarantee claim amount for the portfolio, the valuation reserve for the guarantee is calculated by choosing an amount large enough so that the probability distribution indicates that the chance of it being exhausted is less than a chosen acceptable low value (say 1/1000).

If the probability distribution has been calculated using simulation methods then it should be realised that the frequency table is only an estimate of the underlying distribution. So care must be taken to make sure that the estimate is accurate enough; this is done by making a large enough number of simulations to reduce the confidence interval for the estimate to an acceptably small range. (Table E1.3 on P196 of JIA Vol 107 - Part II shows the sizes of these confidence intervals for various numbers of simulations).

How is the acceptable low probability of the guarantee reserve being exhausted chosen?

This probability is basically an indication of the strength of the valuation; the lower the probability of being exhausted the stronger the reserve. So the choice is purely a professional judgement without there being much objective information to help. What problems should be thought about in choosing the probability?

1. It should produce a valuation of appropriate strength to fit in with the rest of the valuation. This is difficult at present because, say, what probability is equivalent to a 4% valuation rate of interest? I don't know the answer.

2. The consequences of the reserve being exhausted should be considered. I think that a stronger standard would be appropriate when the consequence would be, say, that of being forced to raise money in the Capital Market than if the reserve could be restocked from general reserves and margins. The reason for this is that, in the latter case, the explicit reserve is only part of the reserve backing the option.

3. The probability should be reasonable bearing in mind the assumptions used to calculate the probability distribution (the price model assumptions and whether a withdrawal rate has been used) For example, I think that 1 in 10 is too weak a probability because, if the life offices were to have statistically independent experiences, (They do not; they share the same economy,) this basis would produce inadequate reserves for 10% of the offices. This does not sound nice. Similarly I think that 1 in 1,000,000 is too strong a probability.

4. Maybe we do not know enough about this subject to make effective use of my next point, or maybe, when we know more about the shapes of the probability distributions, it will prove to be irrelevant. Anyway,

here it is. I think that the policyholders and the life offices should get "good value for money" for their reserves. I can best describe this by an example. Imagine an office with a large maturity guarantee portfolio. A reserve, r, is calculated by the method above and the probability of the reserve being exhausted is 1/1000. Using the same methods a reserve of 2r is equivalent to a 1/1100 probability. I think it would be "bad value for money" to reserve at the 2r level because the extra r (which is a lot of money) does not buy a significant amount of extra security.

How does this sort of valuation work from year to year?

Because the Working Party wished to stabilise the reserve from year to year yet still adapt to new business, maturities and changes in the unit price level, they recommended the following procedure for linking the valuation in one year with that of the next.

The valuation is calculated by accumulating the reserve used in the previous year adding in the premiums' contribution covering the guarantee and removing any guarantee claim amounts paid out in the year.

To this is added a reserve in respect of the new business taken on in the year. This reserve is calculated using the method described above with a probability of 1/100. This reserve is calculated allowing for any interaction between the new business and the existing portfolio. This is done by:-

Valuation (New Business)

= Valuation (Existing & New Business) - Valuation (Existing)

All valuations are on the same basis.

The resulting reserve is tested against the probability distribution to find with what probability it will be exhausted.

If that probability is stronger than 1/1000 then reserves can be released; if the probability is weaker than 1/50 then the reserves must be strengthened.

The probabilities quoted above are the no "withdrawal" rates and if withdrawals are included explicitly then the values would be different.

Learning more about these new methods

The process for valuing these guarantees has been described above and can be seen to be relatively complicated compared with the traditional methods, bearing in mind that it is only part of a policy.

The reason for this complication is that it is early days in valuing these options and the short cuts and approximations that simplify most valuations cannot be made yet because not enough is known about what the reserving basis ought to be or what approximations can be made to that basis. If all this long procedure has got to be gone through and, even then, there are a lot of imponderables, why not simplify the position down to one big imponderable and use a simple arbitrary basis?

I think such a simplistic approach is reasonable if:-

(a) It is possible to estimate the approximate magnitude of the reserve if the methods above were used.

(b) The reserve is not significant in the overall valuation of the office.

Where these are not the case then for a reserving basis to command respect it must:-

1. have a logical basis

2. be reconcilable with any relevant experience

3. not predict implausible futures (except as an approximation to a plausible one)

4. not present any unnecessary financing problems.

I think that the Working Party's recommendation should be used in the same way as a standard mortality table. At this stage of development the recommended model should be used cautiously; as an example those portfolios with a large number of maturities in the next few years should have those maturities valued differently from those further in the future. If necessary other bases and methods should be used as well so that the profession can learn:-

(a) how the method works in practice and in different circumstances

(b) how the method compares with other methods

I believe this approach will result in improved, and hopefully simpler, methods in the future.

Premiums

A portfolio reserve of the sort described above is different from the traditional approach because the premium covering the individual policy risk will not finance the reserve. This is a consequence of the reserving method. An example will illustrate. There is a set of new business taken out in a year and the guarantee reserve is calculated on the basis of a 1% chance of being exhausted. Let us assume that all the assumptions will be true; then there is a just under 99% chance that some of the reserve will remain unused. In fact, it can be seen that the average amount that will be required will be considerably less. The maturities of this set of policies will release reserves inside the Life Fund; these reserves are not payable to the policyholders.

The new business reserve must be financed by shareholders or other policyholders or the maturities of existing business. The guarantee premium must cover:-

1. the individual policy's risk of using the guarantee;

2. the cost of interest charges for the guarantee reserve. As the size of the reserve depends not only on the amount of new business but also its relationship to the existing figures, assumptions about the pattern of new business to be sold must be made when setting the premium rates;

3. a contribution towards the reserve.

The last element is difficult because the shareholders and other policyholders will not want their money locked up in this reserve permanently. Unless a contribution is made by the new policyholders, the only way in which the money can be recovered is to stop writing the business and wait for the maturities because the reserve will be in short term securities. The chances of investment windfalls would be very slim and are inappropriate to seek because of the risk to capital values.

IV. SINGLE PREMIUM RATE GUARANTEES

Let us consider the classic case of the Guaranteed Annuity Rate Option of an Endowment Assurance. The position is very similar to that of the Maturity Guarantee because

1. The cost of the option is (should the guaranteed rate be better than the market rate at maturity) the capital sum required to augment the maturity proceeds so that the guaranteed annuity can be brought on the books at the office's then current rates. This amount will depend on the annuity rates at maturity which in turn depend mainly on the interest rates available (but also on mortality and expenses).

2. Hence all policies taken out at the same time and maturing at the same time will get the same advantage from the option. There will be no cancelling out in the same way there is with mortality.

3. Just as the reserve of the unit-linked policy invested in units cannot be used to cover the maturity guarantee, neither can the traditional policy reserve of the Endowment Assurance be used to cover the additional cost of the option.

It is different from the maturity guarantee in that, even when there is an obvious financial advantage in exercising the option, not all policyholders will actually do so because

1. they may want the money as a lump sum and not an annuity

2. they may be in bad health so an annuity may not be good value

3. the policy may have been used to secure a loan so a large part of the proceeds would be used to repay the loan.

Those policyholders most likely to exercise the option would be:-

1. Those whose policies are part of a pension plan (here the vehicle may well be a Deferred Annuity).

2. Those whose policies are only one part of an investment or tax planning exercise. Here the advisors are likely to take advantage of any option if it can be fitted into the plan.

Reserves and premiums for these options can be calculated in the same way as those for Maturity Guarantees with:-

1. The Stock Exchange model adjusted to forecast the appropriate interest rate rather than the price of an accumulative equity unit fund. The model would have to be fitted to appropriate data such as gilt-edged prices.

2. The simulated costs calculated using the forecast interest, mortality and expenses with margins taken to allow for improvement and inflation.

3. The reserve so calculated adjusted by a pessimistic estimate of the proportion exercising the options. If there is a large portfolio of policies with such options it might be a good thing to use different proportions for each class of policy, or better, for each class of purpose for which the policy was bought (the latter information is probably not on the valuation file!).

The office does have some control over the cost of this guarantee in so far as it sets the guaranteed rates at the time the policy is sold.

There is an alternative approach to this and that is to treat these guarantees in the same way as the annual premium financial guarantee described in Section III. I think that this is less appropriate because

1. The annuity market (and other single premium contract market) is very closely associated with the investment market with very frequent rate changes to reflect the interest rate available in the investment market, hence a method that allows for fluctuations is needed.

2. One of the alternatives in the option is a guaranteed cash sum which means that both alternatives cannot be matched in the traditional way.

V. BENEFIT GUARANTEES FOR UNIT-LINKED POLICIES AVAILABLE OVER A PERIOD

As an example there are some policies which allow the policyholder to continue the policy in paid-up form past maturity with the maturity guarantee continuing; the proceeds thereafter can be taken at any time within a period. Here there are all the problems of the maturity guarantee at maturity date only and in addition the policyholder has the option of when to take the money.

What is the approach to this?

The same methods of valuation as suggested for the Maturity Guarantee, with the pattern of maturity of the portfolio adjusted to produce a pattern of when the policy proceeds are cashed, could be used.

This means predicting when the policyholder will take his option. This is very difficult; at least when estimating the unit proceeds there are records of unit prices and Stock Exchange data available. Here there is nothing. As far as I know, the offices concerned with this sort of business have not compared "cashing" rates.

Some assumptions must be made. It could be assumed that:-

1. these rates are correlated to the price level of a non-accumulative unit fund as shown in the Ford and Masters' paper (JIA Vol 106 - Part II). They did in fact use a mixture of correlated and uncorrelated withdrawals;

2. these rates are correlated to the rate and direction of the change of price level.

There are many plausible assumptions that can be made, the problem is to determine what assumptions can command general confidence. There is one traditional assumption which would give the upper limit of possible answers and that is to assume that the policyholders have foreknowledge and choose their "cashing time" as the time which gives them the largest guarantee claim.

Reserves can be calculated on this basis using the methods described above for maturity guarantees assuming all assumptions other than the "cashing rate" are correct; this must be an over-estimate because:-

1. Maximising the guarantee claim is not the same as maximising the proceeds of the policy which ought to be the aim of the policyholder. (A case where it would be is if the policyholder could re-invest the proceeds in a similar investment).

2. As described above, the reserve is calculated on a portfolio basis to take account of the interaction between policies. Widening the spread of maturities in a portfolio tends to reduce the reserve per policy, yet the effect of the cashing assumption would be to concentrate the maturities with the "bad years" of each simulation. Unless there is a special circumstance such as a lapse and re-entry option about, I believe that giving the policyholders the option of when to take their money should in practice widen the spread of maturities still further compared with the single option at maturity. So maybe the actual cost of a spread guarantee is about the same or even lower than for the single one. However, I do not think it prudent to reserve on the base of the last remark.

I think that the reserve should be calculated firstly on the assumption of cashing at the first opportunity (= to the single maturity guarantee) and secondly on the assumption of the "maximised" guarantee claim (which, as far as I know, has not yet been published). If these two values are of the same order, then the reserve should be made on the higher basis. If they are not, then the answer probably lies in between the two values; where, I do not know.

The premium for this guarantee should be calculated using similar principles to those for the Maturity Guarantee.

VI. BENEFIT GUARANTEES FOR NON UNIT-LINKED POLICIES AVAILABLE OVER A PERIOD

Valuation

The position is more flexible than that for a unit-linked policy because the "investment premium" is not locked-up in the assets of the relevant unit fund. There is still a problem because the premiums cannot be invested in an appropriate dated gilt-edged security because the date on which the proceeds are required can vary so widely and therefore no gilt could match the policy. Hence the portfolio of policies is vulnerable to interest rate changes.

Guaranteed surrender values have caused problems in the UK and other countries.

<u>What methods are available?</u>

1. This option could be incorporated into the traditional reserve calculation by choosing a basis which assumes that the benefit is taken at the last possible time, and so that the reserve per policy is always above the guaranteed benefit available to that policy at valuation time.

This might be considered suitable for an Endowment with guaranteed surrender values because the surrender value is likely to build up from year to year in the same way as the reserve. The problem with this method is that it sets a limit on the valuation assumptions that can be used, i.e. if asset values fall and interest rates rise then it is likely to be appropriate to use higher rates in the valuation. If the higher rate is used then there is a possibility that the guaranteed surrender value is uncovered by the valuation reserve. This inflexibility of valuation reserve is really an indication of a deeper problem, namely the problem of matching, mentioned earlier. Continuing the example, assume that the lower valuation rate is maintained in order to keep the surrender value covered; this produces a strain on the surplus (or deficit). It might be argued that this is fine, since the difference between the two valuation bases represents the mismatching reserve.

Unfortunately there is still a problem. What happens if the valuation interest rate should go still higher? Under this method the problem is that this produces is an unknown chance of future strains having to be financed in order to cover this business.

This method is reasonable if:-

(a) conditions generally are very stable and the valuation rate is well below the rate causing the strain. This means that the chance that extra capital will be required is very low indeed.

(b) the office has a large amount of free reserves and the business is a small proportion of the total business. Those conditions that produce extra strain can be covered without difficulty.

2. A variation of the above method is to build withdrawal rates into the valuation basis in the same way as mortality when the guarantee is above the reserve. This is acceptable provided the withdrawal rates will be stable in the future. If there is no confidence about that, then there are still problems especially from a matching point of view.

3. Another approach is to build up the reserve in the traditional way to the point when the guaranteed benefit is first available and then treat the liability in a similar way to a bank; that is to say, the benefits are available to be taken any time, like a bank account. Just as the bank knows that not everyone will call for their deposits at once and hence only a proportion of the liabilities need be backed by cash or "near cash", so the life office could operate in the same way.

The problem is that all such investments must have capital security in the same way as banks. Hence there is the problem of matching the guarantees based on interest rates set long ago, when the policy was taken out, with investments which, to have capital security, will have rates related to the short term market in each year.

The banks control their position through such weapons as their overdraft rates. (A rise in this produces repayments of overdrafts. This converts an investment to cash available to cover the depositors). Can a life office use short term controls in the same way to cover its position? What would the weapon be? Terminal Bonus? Would it work?

Probably not, since I do not think that a reduction in Terminal Bonus would reduce withdrawals; a loss of confidence could be caused with the result being an increase in withdrawals rather than a reduction. An increase in Terminal Bonus would not delay withdrawals - it may well increase them as the policyholders take advantage of the higher proceeds!

So the life offices do not have the appropriate weapon for control of this sort.

4. Another method is to use the risk theory approach. This has been applied to Flexible Endowments by Messrs Ford and Masters (JIA Vol 106 - Part II). Broadly they used the same methods as described above for the Maturity Guarantees. Because the policy is not unit-linked the purpose of the Stock Exchange model is not the forecasting of the unit prices but an estimate of the general price levels of the assets backing the policy. If all these assets were invested in equities then the position would be very similar to the unit-linked position.

If the assets are invested in a mixture of gilts (in an attempt to match) and equities then the position is more complicated since the model would be forecasting both gilt and equity price levels. However if the gilts concerned were all longdated then a model forecasting the market long term rates of interest could be used to predict gilt prices (straight forward) and equity prices (an assumption about long term dividend growth would be needed).

The withdrawals are forecast as a mixture of rates based on duration and rates correlated to the price level produced by the simulation.

This basis does produce a valuation which includes a mismatching reserve calculated on a logical basis. There are a lot of assumptions and the more assumptions there are; the more difficult it is to have confidence in the assumptions; hence the answers may be suspect. However, research will show which assumptions are the important ones and should also be able to set likely limits on the values to be used.

As can be seen, all the methods discussed have problems yet, if the liabilities have been taken on, they must be covered in the valuation. So what should be done?

If the guarantee is low compared with the policy reserve, then Method 1 is the best, provided tests show that there is a wide enough margin between the valuation rate and the rate that uncovers the guarantee; i.e. that the chance of needing extra reserves is acceptably low. The danger here is that the test ought to be repeated periodically and a system of this sort can lead to the option being "forgotten".

If the guarantee cannot be so covered then Method 4 should be used. Experiments should be done to find withdrawal rates conservative enough to give confidence. The price model used should be based on the historical data with, possibly, a small margin in the width of the funnel of doubt.

CONCLUSION

As can be seen, options present many difficulties; most of these are caused by lack of data from which the future of the policies can be predicted. The methods described produce results which indicate that the sum of these options can be very expensive indeed. However, the market and the public are increasingly demanding more options to give flexibility under increasingly unstable conditions. I believe it to be our job to design products which give the public the flexibility it wants without burdening the Life Office with expensive options which must in the end be charged to the policyholder.

I would like to thank my friends and colleagues who commented on the many drafts of this paper and the girls in my department who had to type them.