

MEASURING AND MANAGING THE ECONOMIC RISKS AND COSTS OF WITH-PROFITS BUSINESS

BY A.J. HIBBERT AND C.J. TURNBULL

[Presented to the Institute of Actuaries, 2 June 2003]

ABSTRACT

The approaches to liability valuation, assessment of prudential capital and measurement of profit for life offices are undergoing radical change. A common thread runs through all of these proposed changes — each change represents a move away from traditional actuarial approaches towards a more economically coherent, market-consistent approach. These changes should encourage a general improvement in the life industry's risk management processes. However, they will come at a cost. The measurement of the economic risks generated by the complex guarantees written by life offices is far more difficult than applying the latest resilience test equity fall. This will require a step change in the sophistication of life offices' risk and capital measurement and management know-how. The measurement of value, risk and capital will soon demand the application of 'stochastic' modelling tools. In this paper, we explore some of the issues raised by the application of these approaches to the valuation and risk management of with-profits business.

KEYWORDS

Market-Consistent Valuation; Stochastic Asset Modelling; With-Profits; Risk-Based Capital; Option Pricing; Risk-Neutral Pricing; Dynamic Hedging; Deltas

CONTACT ADDRESSES

A.J. Hibbert, C.J. Turnbull, Barrie & Hibbert Limited, 48 Melville Street, Edinburgh, EH3 7HF, U.K.

Tel: +44(0)131-625-0203; Fax: +44(0)131-625-0215; E-mail: info@barrhibb.com

1. INTRODUCTION

1.1 *Introduction and Background*

1.1.1 The proper measurement and management of risk and cost and the allocation of prudential capital are fundamental to the management of a life office. However, over the past decade life office managers have struggled with a number of risk management challenges. They have failed to meet the challenges posed by annuity options, with-profits business and pensions business. This failure has led to a fundamental review — led by accountants

and regulators — of the valuation and risk management techniques used by life offices. This review has set in motion huge change. In future, the new ‘realistic’ thinking means that actuaries must ‘mark liabilities to market’, i.e. use market prices as a reference point for valuation. This is the principle underlying the fair value accounting requirements proposed by the International Accounting Standards Board (IASB) and the new six-monthly realistic accounting requirements of the Financial Services Authority (FSA). In the coming years, prudential capital will be set within the same framework. As a result, actuaries will need to start thinking (and acting) like the risk managers of a bank — quantifying and then eliminating risk exposures from the balance sheet. Of course, life companies are not banks. We will show that implementing this new thinking poses some major challenges for actuaries.

1.1.2 Fair value or ‘realistic’ accounting methods are concerned with establishing the *economic value* of an insurer’s liabilities. It is important to understand that this definition of value is radically different to the conventional actuarial approach to valuation. Actuarial philosophy towards the valuation of liabilities has traditionally been based on the notion of funding. In other words, given a schedule of liabilities and a specified investment policy, the funding approach will tell us what reserves and/or contributions are required to meet the liabilities with a given level of confidence. By contrast, the thinking behind the economic valuation of a liability is very different. The economic value is defined as the sum of money required to establish a portfolio of assets that — provided that they are invested in a particular way — will replicate the liability as closely as possible. This special portfolio is called the hedge portfolio.

1.1.3 In this paper we consider how the new ideas will be applied in practice. We consider some of the questions and issues that will arise in estimating the economic value of long-term insurance liabilities. A valuation basis for this valuation is developed in Section 2. Section 3 applies this basis to value a sample conventional with-profits contract. This will highlight how fair valuation (we use the terms fair value, realistic value, economic value and market-consistent value interchangeably throughout this paper) will require detailed assumptions, not only on the asset side of the balance sheet, but on the liability side too (e.g. regarding smoothing, discretion, guarantees, and policyholder behaviour). However, the overall approach to the valuation of guarantees is generic, and can be applied to the investment guarantees of any product.

1.1.4 In discussing the issues surrounding the fair valuation of insurance liabilities, we will address the following broad questions:

- How can asset models be calibrated to calculate fair values for cash guarantees?
- What assumptions are required in terms of liability/policyholder behaviour?

- What size of guarantee cost might we expect for a ‘typical’ conventional with-profits contract?
- How might the fair value compare with statutory reserves?
- How sensitive is this cost to the assumptions that we have made regarding future asset and liability behaviour?

1.1.5 After discussing the determination of fair values for long-term insurance liabilities, the remainder of the paper moves on to discuss the risk management ramifications that this may have. We will show how the hedge portfolio for different types of cash guarantees can be discovered and managed (similar approaches could be applied to interest rate guarantees). The Black-Scholes-Merton analysis and its application to the simple case of a unit-linked fund guarantee are discussed in Appendix A. Through this example, the basic logic and techniques of dynamic hedging are developed and discussed.

1.1.6 In Section 4 the dynamic hedging ideas are extended to the more complex arena of with-profits business. We revisit the cash guarantees discussed in Section 3’s valuation case study. The complexity of with-profits guarantees makes finding the appropriate hedge portfolio far more difficult — the convenient formulae developed by Black-Scholes-Merton no longer apply. However, we show how the market-consistent stochastic model developed for valuation can also be used to find the guarantees’ hedge portfolio.

1.1.7 Section 5 takes the analysis of Section 4 and applies it to the management of fair value profits. We start with the dynamic hedge portfolio identified in Section 4, and then extend the analysis by identifying some alternative risk management solutions. The relative merits of the different solutions are compared.

1.1.8 As well as considering the impact of the risk management strategies on profit volatility, we also analyse briefly the possible implications for capital requirements. In the past, the link between liability valuation (especially statutory valuation) and risk management has been weak. Statutory valuation rules have often been arbitrary and *ad hoc*, resulting in a difference between the valuation and the underlying economic reality. This is unfortunate, and it has provided life offices with incentives to pursue risk management solutions that provide the greatest short-term statutory relief, without necessarily being appropriate for the genuine long-term economic risk exposures faced by offices. However, future changes in regulations (as discussed in the FSA’s Consultation Papers 136 and 143) should see future statutory capital calculations moving towards a basis that reflects the economic exposures. As a result, the realistic liability valuation and the statutory capital requirement will become intrinsically linked. In effect, the statutory capital requirement can be thought of as the realistic valuation plus a mis-matching reserve. The size of the mis-match reserve may also be calculated using a stochastic model. We consider the implications of the risk

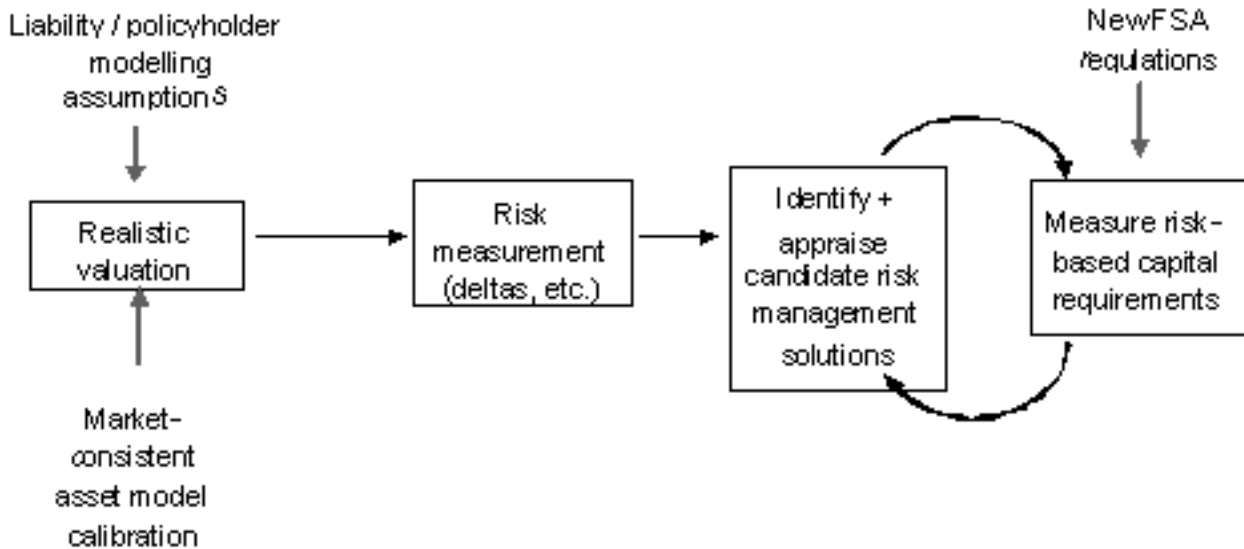


Figure 1.1. Applications of the market-consistent valuation model in the valuation, risk management and capital measurement processes

management strategies that we develop in Section 5 for the risk-based capital requirements of the sample with-profits policy.

1.1.9 All of the above suggest that the market-consistent valuation model may be used for much more than ‘just’ the market-consistent valuation. It can also be an invaluable tool in measuring and managing risk and prudential capital. The most significant areas of application for the market-consistent valuation tool in the valuation, risk management and capital measurement processes are summarised in Figure 1.1.

2. A POSSIBLE FAIR VALUE BASIS FOR VALUING CONVENTIONAL WITH-PROFITS GUARANTEES

2.1 Introduction

2.1.1 This section develops the valuation basis for the market-consistent valuation of a conventional with-profits contract. Before we enter into the detail and its accompanying issues, it is worth pausing to consider what this valuation approach entails. It can be summarised as follows:

- (1) Choose a stochastic asset model that is capable of market-consistent valuation. This calls for a model that is arbitrage-free (Technically, this requires that a probability measure exists under which all discounted model security prices are martingales. This probability measure is known as the equivalent martingale measure. See Harrison & Pliska, 1981). Calibrate the model to relevant market prices. That is, when the model is used to value relevant market-traded instruments, the value calculated by the model should be close to the market price.
- (2) Make assumptions about the behaviour of the liability (e.g. bonus rates

and investment policy), and policyholders (e.g. lapses) in each possible asset model scenario.

- (3) Project the liability cash flows and discount them at the market-consistent discount rates. (The appropriate discount factors required to recover the market-consistent value of future uncertain cash flows is closely related to (1). The market-consistent discount factors, together with the probability measure used in the valuation, will imply that discounted security prices are martingales. Note that this means that the market-consistent discount factors are functions of the risk premia assumed in the stochastic process assumed in the projection of assets. Changing the risk premium, therefore, has two effects on the valuation: it changes the distribution of cash flows; and it changes how these cash flows are discounted. These two effects offset each other exactly: market-consistent valuation is not affected by the choice of risk premium. If the asset model is set up in ‘risk-neutral’ mode — i.e. where it is assumed that all assets have an expected drift equal to the short-term risk-free interest rate — the discount factors will be the simulated cash roll-ups. Alternatively, state price deflators will be applied.)
- (4) The mean discounted value of the liability cash flows is the estimate of the market-consistent liability value. Under suitable assumptions, it can be viewed as the cost of the replicating portfolio for the liability.

2.1.2 This method of valuation is radically different to traditional actuarial approaches. However, the application of option pricing ideas to the analysis of with-profits policies is not new. (See, for example, Wilkie, 1987; and Kemp, 1997). Indeed, when reviewing these papers, it is difficult for the reader not to feel rather bewildered by the lack of application of these ideas in managing with-profits business in recent years.

2.1.3 The application of this valuation approach is not limited to with-profits business. It can be applied to the valuation of any financial asset with guarantees (e.g. guaranteed annuity options, guaranteed minimum pensions, etc.). The remainder of this section discusses each of the above stages in turn.

2.2 *Assets*

2.2.1 As for all fair valuations of insurance liabilities, we need to develop a set of assumptions for the asset behaviour on which the liability valuation depends. Further, these assumptions need to be consistent with the market prices of assets that exhibit similar characteristics to the liabilities. There is likely to be some debate over which assets are the ‘right’ target for calibration of the asset model. After all, typical life office guarantees are likely to crystallise at significantly longer maturities than the options traded on exchange markets. To obtain an estimate of market-implied levels of long-term asset volatility, we will therefore need to consider the use of over-the-

counter (OTC) derivative prices. These will be needed for each of the asset classes in which asset shares are invested. We will confine ourselves here to equities and government bonds. (This makes life a little easier. This exercise is likely to be considerably more difficult for asset classes such as property and corporate bonds.)

2.2.2 *Equities.* The asset model calibration process will involve trade-offs between model simplicity and pricing accuracy. We briefly explore these by means of an example. At the end of June 2002, ten-year at-the-money OTC FTSE 100 options were quoted with an implied volatility of 23% p.a. Short-term (exchange-traded) implied volatilities were slightly lower — one-year options (trading on LIFFE) had an implied volatility of around 21% p.a. Now, what model of equity behaviour should be used to recover these option prices? We should aim to use a model that is no more complicated than it needs to be. Our starting point might therefore be the standard lognormal (random walk) model. The lognormal model can be used to replicate the option value for any given option strike price by ensuring that the assumed lognormal distribution has a volatility consistent with the implied volatility of the option. (Note that ‘implied volatility’ numbers are based on the assumption that the underlying asset is lognormally distributed). So, for a given term and strike of guarantee, we can find the fair value by choosing a volatility parameter that is consistent with the Black-Scholes implied volatility of a similar, market-traded option.

2.2.3 However, for a given option term, the market prices of options with different strike prices typically exhibit different implied volatilities. In this case the standard lognormal model will not be capable of simultaneously recovering these market prices. In determining the fair value of a book of guarantees, we are therefore left with some choices:

- (a) A more sophisticated equity model could be employed that could price all these different options with a single calibration.
- (b) We could ignore this problem and simply use a volatility that was roughly consistent with all the observed market-implied volatilities.
- (c) The lognormal model could be used many times, each time with a different calibration to value options with a particular strike price.

2.2.4 The choice made here will depend on the range of term and ‘moneyness’ of the guarantees being valued, the range of implied volatilities observed in market prices, and the degree of accuracy required in the valuation. Note that this implies that different offices/funds may use different calibrations if the ‘shape’ (i.e. term and moneyness) of guarantees are significantly different. As we will only consider a single example policy in Section 3, we can conveniently side step this issue and use the standard lognormal model.

2.2.5 The behaviour of implied volatilities also needs to be considered in a second dimension. As well as considering how implied volatilities vary across strike prices for a given term, we also need to consider how implied

volatilities vary by term. This is easier to address. To capture the term structure of equity volatility, we could use a deterministic (but not constant) volatility function in a lognormal model. That is, the volatility used in a stochastic projection would be a function of time, but would not be random. This would allow us to retain the simple lognormal model structure, and replicate option prices of all terms with a given strike price. (A similar extension to the volatility function could be made to capture variation in implied volatility by strike price as well as by term. This would require modelling volatility as a deterministic function of both underlying asset price and time).

2.2.6 Alternatively, we could use the standard lognormal model with constant volatility, and choose a parameter that is consistent with the term of the guarantees. The example policy that we shall use in the example set out below has a term to maturity of ten years. We therefore use a volatility parameter of 23% p.a. to be consistent with ten-year equity option prices. For simplicity, we assume all equity holdings are United Kingdom equities. Where a global portfolio of equities is held, an assumption regarding the correlations between (the sterling returns of) different equity markets will be necessary. It may be difficult to establish a market-implied value for this correlation, and some judgement is likely to be necessary.

2.2.7 From the above brief discussion of equity modelling assumptions for the valuation of cash guarantees, it is clear that there may be a number of awkward issues for standard setters to consider. These issues lie beyond the question of whether risk-neutral pricing or state price deflators are the preferred implementation method. For standard setters, this is really a non-issue, as both methods recover the same answer under the same modelling assumptions. The far bigger problem will be the determination of what does and does not constitute a *market-consistent model calibration*.

2.2.8 We should not depart from the topic of the equity model choice without making a final important point. A model that performs well for the purposes of market-consistent valuation is not necessarily a model that is most suited to the purposes of developing and evaluating hedging strategies. For market-consistent valuation, it is possible to take a fairly basic model (e.g. the lognormal model), and find model parameters (e.g. asset volatility) that recover market prices. For the purposes of risk management, where an understanding of the nuances of the future ‘real world’ distribution of asset returns is crucial, a more demanding modelling requirement exists.

2.2.9 *Risk-free interest rates and bonds.* The analysis in this paper employs a two-factor Black-Karasinski model. This can be considered as a lognormal specification of the two-factor Hull-White model. See Hull & White (1994). A brief description of the model specification and the parameters used in this paper can be found in Appendix B. In calibrating the interest rate model, attention needs to be paid to both the level of the risk-free yield curve, and a market-implied level of volatility (as the volatility of

bonds will impact on the asset share volatility and hence the value of guarantees). In calibrating to the risk-free yield curve, the following analysis assumes swap rates can be used as the reference traded risk-free asset. In measuring market-implied interest rate volatility, swaption (i.e. options to enter swap contracts) prices are used. In our example valuation, we will assume that the with-profits fund is invested in ten-year risk-free 5% coupon bonds (which are rebalanced every year). As a swaption can be regarded as an option on a ten-year par bond, the term structure of implied volatilities of swaptions on ten-year swaps will be the relevant measure for calibration purposes. Figures 2.1 and 2.2 show the market-implied swap rates and implied volatilities at the end of June 2002, and the corresponding values generated by our model calibration.

2.2.10 From Figure 2.1 you can see that swap rates were around 5% at the end of June 2002. The model and our calibration does a reasonable job of replicating these swap rates. Figure 2.2 shows that market swaption prices were implying that ten-year swap rates have an implied volatility of around 13%-14% p.a. at virtually all terms to maturity. Note this is a proportional volatility. So for a ten-year swap rate of, say, 5.5%, the option prices imply an annual standard deviation of around 0.75% for the ten-year swap rate.

2.2.11 Figure 2.2 suggests that the interest rate model calibration provides only a rough approximation to the term structure of volatility implied by interest rate option prices. An exact fit would require a more




Figure 2.1. Sterling swap rates (end June 2002)

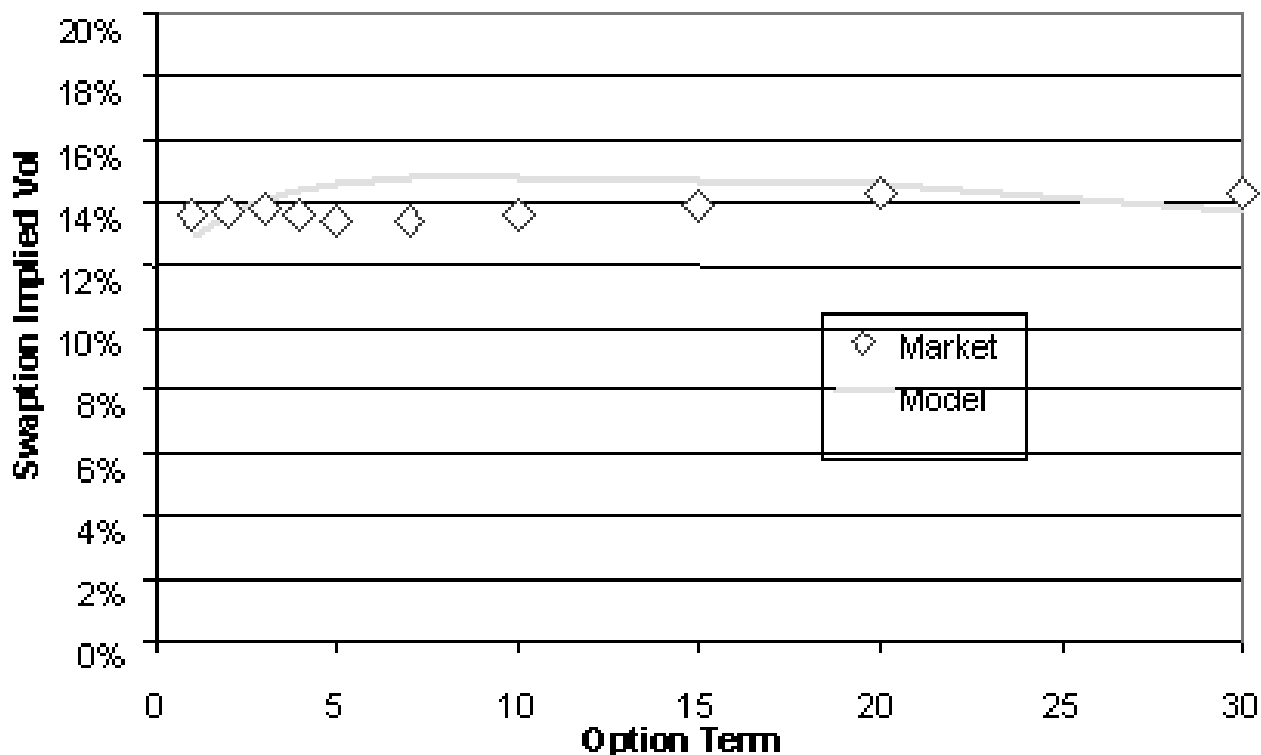


Figure 2.2. Implied volatilities of ten-year sterling swap swaptions (end June 2002)

sophisticated interest rate model than that used in this analysis (for example, the Libor market model, see Hull & White, 2000). It is also worth noting that this calibration has been a best fit to the entire 30-year volatility term structure. In reality, offices may be more concerned with replicating prices in a narrower range of guarantee maturity.

2.2.12 Finally, an assumption is required for the correlation between the equity and bond asset classes. It is very difficult to infer a market-consistent value for this assumption. We therefore prefer to use a value that is reasonable and consistent with long-term historical observation. We assume equity returns in excess of cash, and that movements in the short-term interest rate have a correlation of -0.33 . With the interest rate model employed here, this implies a correlation between equity returns and long bond returns of around 0.25.

2.3 Liabilities and Policyholder Behaviour

2.3.1 Having considered the choice of asset model and its calibration, some assumptions regarding the behaviour of with-profits liabilities are also required. There is, of course, significant discretion available to some offices regarding how with-profits liabilities are managed. In pricing a conventional with-profits (CWP) guarantee, this discretion manifests itself in two material ways: bonus policy and investment policy. The estimated guarantee cost can

be quite sensitive to assumptions regarding how both of these areas of discretion are used, and, in particular, how the use of discretion interacts with asset behaviour. As it is very likely that decisions about both of these areas will depend on experienced investment returns, dynamic rules will be required to determine how bonus rates and investment policy relate to movements in asset share. (We will consider how sensitive the valuation is to the application of such rules in Section 3). Naturally, any such rules will have to respect policyholders' reasonable expectations and the concept of fairness to policyholders. Determination of these rules will therefore be office-specific. The dynamic rules assumed in the example developed in Section 3 now follow.

2.3.2 *A dynamic reversionary bonus policy rule*

Firstly, let us consider bonus policy. Suppose that the reversionary bonus rate is revised once a year, according to the following rules (and the current reversionary bonus rate is 4% p.a. compound):

- Set the reversionary bonus rate equal to the long gilt yield, less deductions for expenses (of 1% p.a.) and tax (at an assumed rate of 20%).
- Changes in the reversionary bonus rate are restricted to a maximum of 1% in any one-year period.
- Further, if the projected value of asset share at maturity (projected at long gilt yield less an allowance for expenses) is less than the projected guaranteed sum assured (based on the current bonus rate), the reversionary bonus rate is reduced by 1% (subject to a minimum of 0%).

2.3.3 *A dynamic investment policy rule*

Suppose that the investment of asset share is in equities and in long gilts, and that the equity backing ratio (EBR) is revised once a year according to the following rules (and the current EBR is 80%):

- If the projected asset share is less than the projected guaranteed sum assured at maturity, reduce EBR by 10% (subject to a minimum of 0%).
- If the projected asset share is greater than the projected guaranteed sum assured at maturity, increase EBR by 10% (subject to a maximum of 80%).
- The bond portion of asset share is invested in a 15-year risk-free bond that pays an annual coupon of 5%. This bond holding is assumed to be rebalanced back to its starting maturity at the end of every year.

2.3.4 Clearly these are simplistic rules. However, these rules do a reasonable job of capturing the essence of the discretion available to life offices in managing with-profits business (and the limits placed on this discretion by factors such as policyholders' reasonable expectations). 'Real-life' implementation of fair value calculations is likely to use much more complex codifications of discretionary behaviour. Indeed, this area is likely to prove a significant challenge for fair valuation of with-profits. We shall see later that results can be quite sensitive to these potentially subjective assumptions.

2.3.5 Smoothing

We also need to make some assumptions about how terminal bonuses are determined (unlike current statutory valuations, future terminal bonuses will form part of the fair valuation of liabilities). If we suppose that a smoothed asset share is calculated according to some specified rule, and the terminal bonus is paid to bring the policy payout to the greater of $x\%$ of that smoothed asset share and the accrued guarantees, then we can write the value of the policy payout as follows (note this does not preclude smoothed asset share being calculated differently for different decrements, e.g. surrender values may be smoothed to a lesser degree than maturity or death payouts):

$$\begin{aligned}\text{Fair value} &= \text{Value} [\max(x\% \cdot \text{Smoothed asset share}, \text{Guarantee})] \\ &= \text{Value} (x\% \cdot \text{Smoothed asset share}) + \text{Value} [\max (0, \\ &\quad \text{Guarantee} - x\% \cdot \text{Smoothed asset share})].\end{aligned}$$

2.3.6 From the above it can be seen that the value of the policy can be considered as the market-consistent value of $x\%$ of a policy's smoothed asset share plus a put option on $x\%$ of the smoothed asset share at a strike equal to the guaranteed sum assured at the time of termination (whether due to death, maturity, surrender, etc.). Taking this a step further:

$$\begin{aligned}\text{Fair value} &= \text{Asset share} \\ &\quad + \text{Value} [x\% \cdot \text{Smoothed asset share} - \text{Asset share}] \\ &\quad + \text{Value} [\max (0, \text{Guarantee} - x\% \cdot \text{Smoothed asset share})].\end{aligned}$$

2.3.7 Now we can see that the value of each policy is the asset share of the policy, plus the market-consistent value of the difference between $x\%$ of smoothed asset share and asset share at termination, plus the guarantee cost (the put option on $x\%$ of smoothed asset share).

2.3.8 The policyholders' benefits can be considered to have three components:

- (a) the asset share;
- (b) the economic value of smoothing; and
- (c) the economic value of guarantees.

2.3.9 Furthermore, the size of these three components may not be independent of each other. For example, suppose that the office applies a smoothing policy such that the expected smoothed asset share is equal to the expected asset share of the policy, but is less volatile (and $x = 100$). In this case the guarantee cost will have been reduced by smoothing — as the volatility of the guarantee's underlying asset has been reduced, but a smoothing cost will have been created. This may seem odd: why has a smoothing cost been created if the average smoothed asset share is equal to

the average 'raw' asset share? Because the cost depends on a risk-adjusted (i.e. market-consistent) present value, and the discount rate applied to the (less volatile) smoothed asset share will be lower than for the 'raw' asset share. Put another way, the office's smoothing policy is transferring some risk from the policyholder to the office, and that risk has an economic cost. (Note that this analysis assumes that smoothing really does involve risk reduction for the policyholder. If policyholders bear the smoothing profit/loss of earlier generations of policyholders, it is not clear that smoothing reduces risk for policyholders (where by risk we mean total variability of payouts). This has the potential to be a rather messy area for valuation, and it may also create some interesting conclusions regarding the effectiveness of smoothing rules in delivering smoothed returns.)

2.3.10 Let us consider an extremely simple example to illustrate this idea. Suppose we have a one-year policy, where asset share is 1, and the sum assured is 1. No reversionary bonus will be paid before maturity, and no one will die or lapse during the next year. The fund is invested in assets with an expected return of 8% p.a., with a volatility of 15% p.a. Further, suppose that the risk-free interest rate is 5% p.a. Now, in the absence of smoothing, what is the fair value of the policyholder's benefits? Using the Black-Scholes option pricing formula, the put option has a value of 0.085, giving a total fair value of policyholders' policy benefits (FVPPB) of 1.085.

2.3.11 Now suppose that smoothing is applied such that the smoothed asset share at maturity is 1.08 (i.e. the expected asset share) with certainty. In this case the guarantee term is worth nothing, as the guarantee will never exceed the smoothed asset share. What about the smoothing term? Well, as the smoothed asset share is a risk-free cash flow, it should be discounted at the risk-free interest rate to find its present value. Thus, the present value of the smoothed asset share is $1.08/1.05 = 1.029$. The smoothing cost is therefore 0.029, and the FVPPB is 1.029. So, in this example, the reduction in guarantee cost is greater than the increase in smoothing costs and the total market-consistent policy value has been reduced. Why has this happened? Well, in this case smoothing has removed the upside potential for the policyholder, whilst the additional downside protection provided by smoothing is not so valuable, given that the guarantee was already providing a significant amount of such protection. (Of course, this will not always be the case. The extent to which this is generally true will depend on the 'moneyness' of the guarantee, the volatility of the underlying asset share and the nature of the smoothing rule.)

2.3.12 Clearly, smoothing is a significant complication in the appraisal of the value of policyholders' benefits (this will be especially so where emerging smoothing profits and losses are borne by adjustments to remaining policyholders' asset shares rather than being funded by 'external' assets). In principle, however, incorporating the impact of smoothing into the valuation is only as difficult as specifying an algorithm for the smoothing process.

This algorithm would then be used in the modelling projections for the valuation. For the sake of simplicity, the remainder of this paper will assume that no smoothing is applied, and we will concentrate on the valuation of guarantees in the absence of smoothing. That is, we assume $x = 100$ and smoothed asset share is equal to raw asset share for all decrements.

2.3.13 *Mortality and lapses*

Assumptions are also required with respect to mortality outcomes. As the same guarantee is available on death and maturity, this assumption turns out to be second order. In the following example, we will assume 90% of PMA92 with CMI 17 improvements.

2.3.14 Policyholder lapse rates must also be specified. High lapse rates will tend to reduce the guarantee cost, since surrender values are not guaranteed (and we are also supposing that the office does not apply any smoothing). This brings us on to another interesting point — should lapse rates be modelled dynamically? i.e. if asset shares fall significantly, should we assume that policyholders are less likely to lapse their now-very-valuable guarantees? For long-duration policies, the guarantee cost could be quite sensitive to the answer to this question. In our fair value basis, we will assume that lapse rates are dynamic. It will be assumed that the ‘normal’ lapse rate is 4% p.a. However, whenever asset shares fall below the guaranteed sum assured, the lapse rate is assumed to be reduced by the proportions shown in Table 2.1.

2.3.15 *Cross-subsidies and participation in profits*

Our assumptions are now sufficiently developed to allow us to project the cash flows and guarantee strains for an example policy. However, a fundamental question that remains to be addressed is: “Who actually bears the cost of the projected guarantee strains?” For example, can it be assumed that the asset shares of other with-profits policyholders are reduced in the event of a guarantee shortfall emerging? If so, can credit be taken for this in the calculation of the fair value of the fund?

2.3.16 The different answers to these questions could make a very significant impact on the fair value of the fund. It will also have implications for the complexity of the projection and the assumptions that it requires — e.g. on how guarantee profits and losses are spread across policyholders. In the interests of brevity and simplicity, our example adopts the convenient assumption that all guarantee losses are borne by the office. In other words,

Table 2.1. Dynamic lapse rate adjustments

Ratio of asset share to sum assured	Proportional lapse rate reduction
$0.9 < AS/SA < 1$	10%
$0.75 < AS/SA < 0.9$	25%
$0.5 < AS/SA < 0.75$	50%
$AS/SA < 0.5$	75%

other policyholders' benefits are not reduced when the policy's guarantees 'bite'.

2.3.17 We now have sufficient information to tackle the market-consistent valuation of with-profits guarantees. Despite the objective nature of a market-consistent valuation, it was necessary to make a number of assumptions and judgements to determine the valuation basis. Whilst fair valuations will necessitate a very major change in actuarial valuation techniques, the demand for actuarial judgement may turn out to be as great as ever. Section 3 takes a specimen conventional with-profits policy and examines the fair value of the guarantees generated by the model developed in this section.

3. A VALUATION CASE STUDY

3.1 This section applies the modelling assumptions developed in Section 2 to an example policy. In the example, we consider a 50-year-old male policyholder. The policy is assumed to have ten years remaining to maturity, a current asset share of £13,000 and a current sum assured of £16,000. The policy is assumed to be a regular premium policy with premiums of £800 payable annually. We assume that expenses of 1% of asset share p.a. are generated by the policy, and that these expenses are charged to the asset share of the policy. The sum assured is payable on death or maturity of the policy (we assume that asset share is paid on surrender).

3.2 *The Simulated Distributions of Guarantee Strains*

3.2.1 Before we calculate the fair value of the death and maturity guarantees, let us first consider the projected distribution of cash flows generated by the guarantees (i.e. the additional cash flows that the policyholder receives as a result of the guarantees), as implied by the model set-up developed in Section 2. These are shown for the death and maturity guarantees in Figures 3.1 and 3.2 respectively. (For these projections, an assumption regarding the size of the equity risk premium is required. We assume that it is 4% p.a. However, as we mention below, the fair valuation does not depend on the assumed size of the equity risk premium.)

3.2.2 Figure 3.1 plots the distribution of cash flow strains resulting from the death guarantee each year. (The segments in the centre of the distribution represent the middle two quartiles of the simulated probability distribution, the bordering segments illustrate the 25th to 5th percentiles and 75th to 95th percentiles, and, finally, the outer segments on the extremity of the figure show the 5th to 1st percentiles and 95th to 99th percentiles. The mean outcome is plotted with a solid box.) For example, the simulation model suggests that the frequency of a death guarantee cash flow strain being more than £5 at the end of the first year of the projection is around 1%. (The

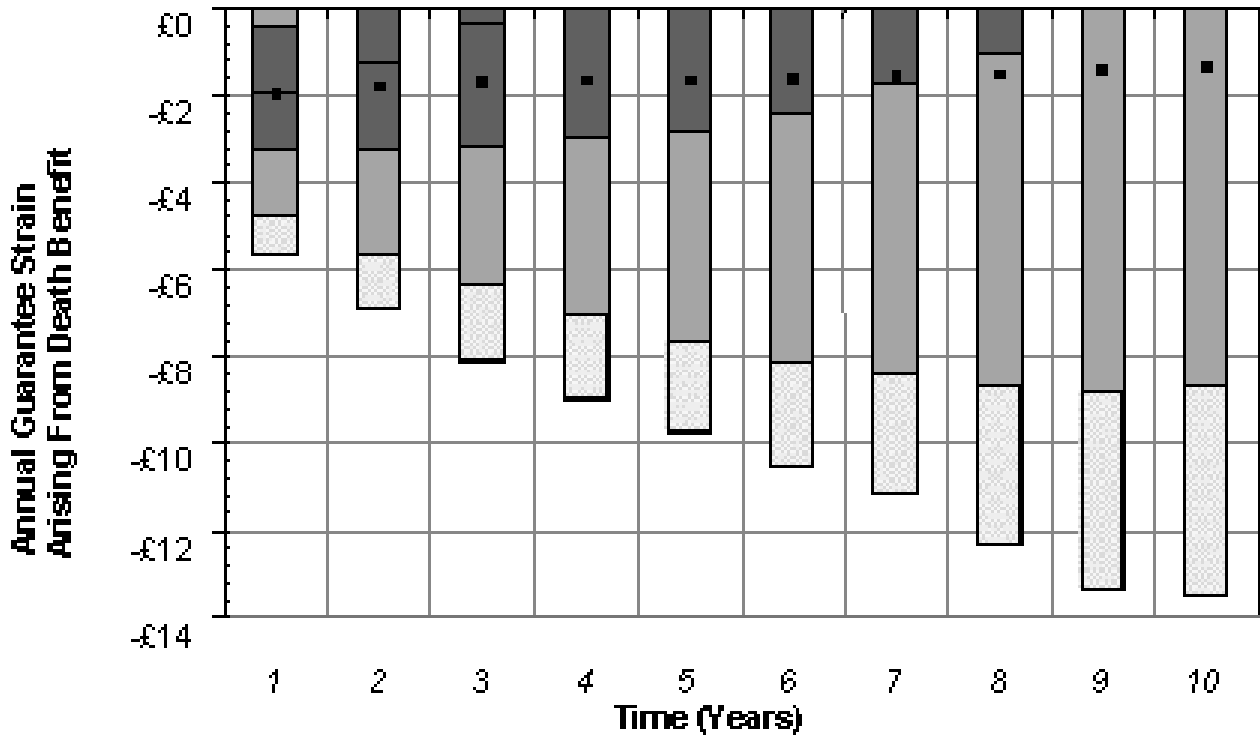


Figure 3.1. Simulated distribution of annual cash flow strains arising from death guarantee

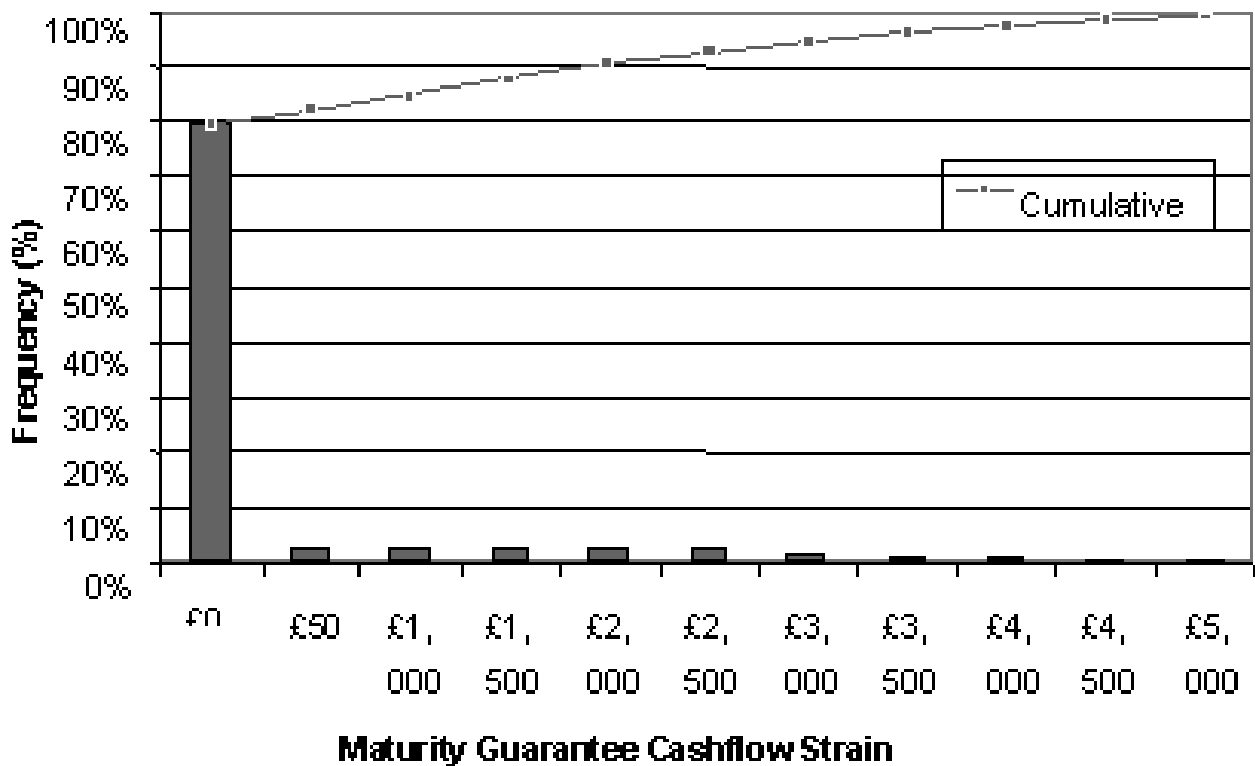


Figure 3.2. Simulated distribution of cash flow strains arising from maturity guarantee at year ten

mortality basis implies that the year one mortality rate is approximately 0.001. At the end of the first year, the 99th percentile asset share is around £11,000. This implies a sum at risk on death of £5,000, and a guarantee strain of $£5,000 \times 0.001 = £5$.) The median cash flow strain arising from the guarantees tends quite quickly to zero, reflecting the reductions in the expected death strain at risk that result from the receipt of future regular premiums. However, despite this reduction in expected death strain, the downside tail of the risk continues to expand with time, reflecting the greater potential for significant asset under-performance as the time horizon extends. The analysis suggests there is a 5% chance of a strain of £7 (or more) arising from the death guarantee in the 5th year of the projection. Figure 3.2 shows the distribution of guarantee strains occurring at maturity of the contract.

3.2.3 Figure 3.2 shows that, for policies that reach maturity, our asset and liability assumptions imply that this specimen policy has a 20% chance of maturing with an asset share less than the final sum assured (including reversionary bonuses). We now turn our attention to the fair values implied by these cash flow distributions.

3.3 *Calculating the Fair Value of Guarantees*

3.3.1 In calculating the fair value of the guarantees described above, market-consistent discount rates are required to calculate the market-consistent value of the cash flow distributions shown above. Two approaches have been developed to make these valuation calculations. One possibility is to develop scenario-specific discount factors known as ‘state price deflators’. Alternatively, and more simply, we can employ risk-neutral valuation methods. Risk-neutral techniques will give the same answer and require a little less work.

3.3.2 Our application of risk-neutral pricing to the specimen policy described above estimates the mortality guarantee value at £19 and the maturity guarantee at £629. (These results were derived from 10,000 simulations. Antithetic variables were used as a variance reduction technique. The standard error of the total guarantee cost estimate of £648 was approximately £6.) Recall that this is the additional impact of the guarantees on the fair value of policyholders’ benefits. To put this into context, let us consider the total fair value of policyholders’ benefits of this policy and its components, as given in Table 3.1.

3.3.3 In this example it was assumed that the policyholder bears all the expense risk, so this item of the liability is zero.

3.3.4 Our hypothetical life office may have a number of options open to it to reduce this fair value deficit. It could revise its assumptions for how it intends to manage the discretionary features of the liability (this is discussed in Section 3.5). Or, more simply, it could reduce current bonus rates or the equity backing ratio (EBR). Alternatively, it might decide to pay less than

Table 3.1. Fair Value of Policyholders' Policy Benefits and Expenses

Component	Fair Value
Asset share	13,000
Expenses net of expense charges to asset share	0
Guarantees	648
Total	13,648

full asset share when the guarantee is not biting. This would reduce the asset share component of the fair value liability. A very crude calculation suggests that the payout would need to be reduced to around 96% (i.e. $1 - 648/17,165$. 17,165 is the asset share + PV of future premiums less PV of future expenses) of asset share to remove the fair value deficit. In actual fact, the reduction would need to be slightly greater than this, as such an adjustment will increase the guarantee value by effectively reducing the underlying asset value used in the option valuation.

3.4 A Market-Consistent Guarantee Charge

3.4.1 Another means of removing the fair value deficit would be through the application of an annual guarantee charge, which is assumed to flow out of the with-profits fund. Calculating the required size of this number will be an iterative process, as the introduction of a guarantee charge to the asset share will increase the guarantee value, as calculated above. Figure 3.3 plots the relationship between guarantee costs and guarantee charges.

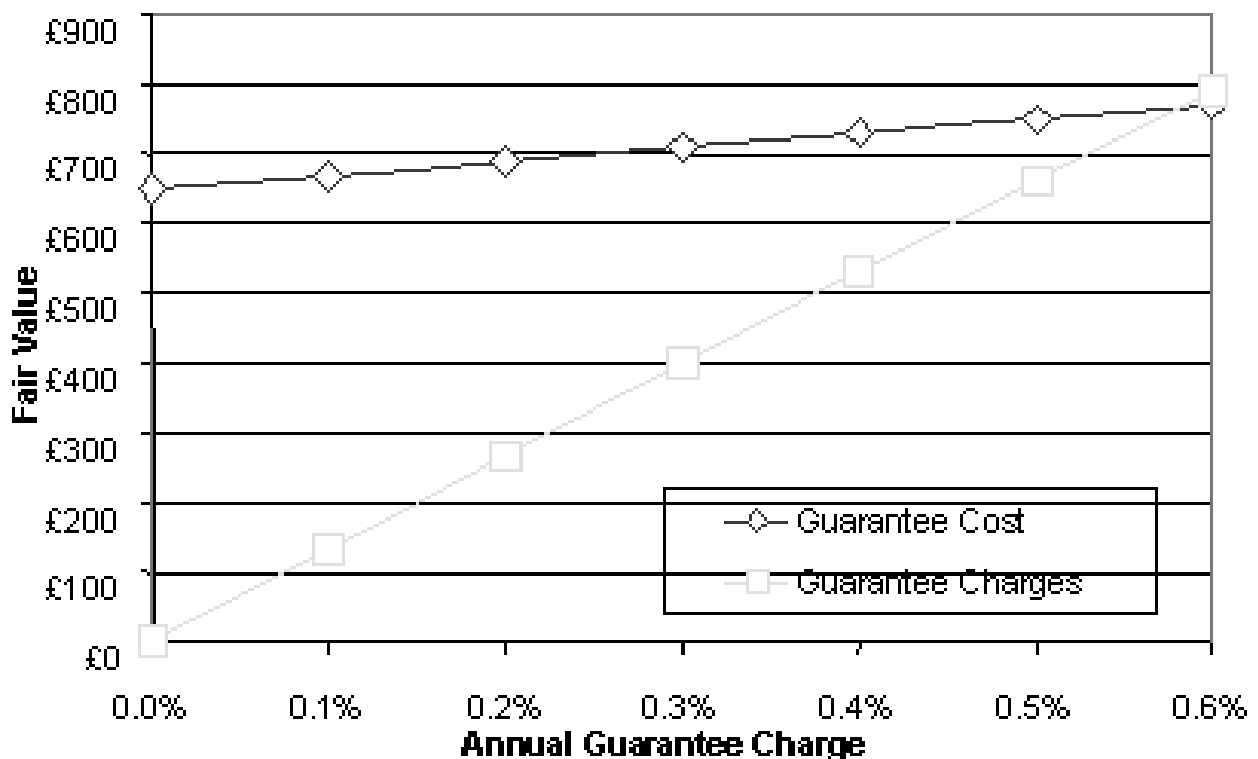


Figure 3.3. Fair value of guarantee and guarantee charges

3.4.2 You can see that an annual guarantee charge of almost 0.60% p.a. of asset share is required to equate the fair value of the guarantees to the fair value of these charges. This can be interpreted as a market-consistent charge for the guarantees (assuming that no previous charges have been made).

3.4.3 But what does this ‘market-consistent’ tag actually mean in this context? It represents the amount that needs to be invested in order to fund the guarantee costs with certainty (give or take some second-order effects) i.e. it is the cost of matching out the guarantees. This is slightly different from conventional matching in two key ways: the asset allocation of the asset share is not affected, the matching is done with the guarantee charges; and the asset allocation of the guarantee charges is managed dynamically and is dictated by market movements. We discuss this idea more fully in Sections 4 and 5.

3.5 *Statutory Reserves*

3.5.1 We saw above that the fair value of policyholders’ benefits was £13,648 and the current asset share of the policy was £13,000. How is this likely to compare with statutory valuations of this policy? Naturally, such a valuation will depend on the level of prudence assumed in the valuation. However, given current approaches to statutory valuation, it would be somewhat surprising if the guarantees of the specimen policy were to result in a valuation in excess of the asset share of the policy. This may seem a rather odd state of affairs, given that the statutory reserve is supposed to have margins for prudence whilst the fair value is not. The lower size of the statutory reserve arises primarily because a statutory valuation ignores prospective terminal bonus, whilst the fair value will include it (as the fair value takes account of the *total* policyholder payout rather than only the guaranteed portion). There is likely to be very little correspondence between changes in the two valuations. A rise in the fair value need not imply an increased statutory valuation, or vice versa. A move towards a risk-based capital approach to statutory reserving is likely to result in a stronger relationship between the fair value and the statutory reserve (a topic that we return to in Section 5).

3.6 *Examining Sensitivities*

3.6.1 It is worth considering the sensitivity of the above results to some of the modelling assumptions made in the valuation. In configuring our model, we had relatively little lee-way in the asset calibration — our asset model needed to be consistent with market prices. However, a number of assumptions were made regarding the behaviour of the with-profits liability, especially regarding the future behaviour of bonus rates and investment policy. As there could be a degree of flexibility in the determination of these assumptions, it may be of interest to analyse the sensitivity of the above

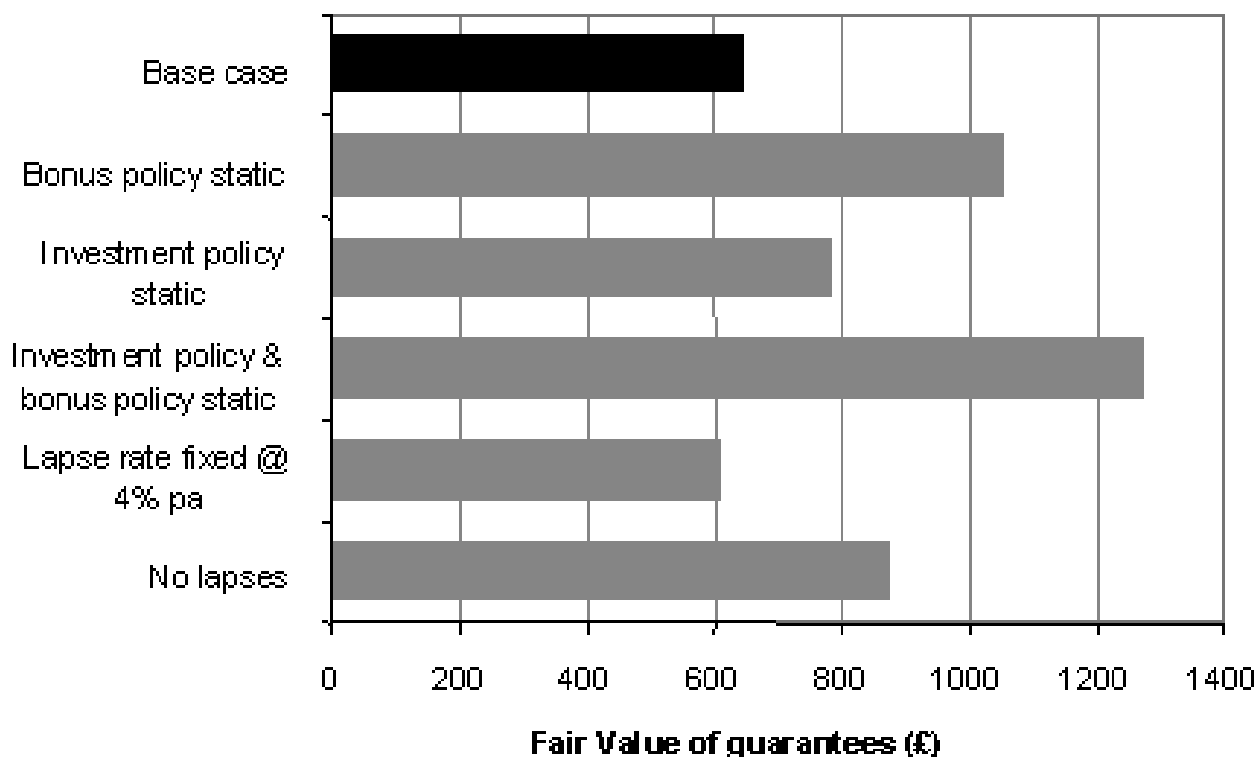


Figure 3.4. Sensitivity of fair value of guarantee to bonus policy and investment policy

results to these assumptions. You will recall that, in our example, bonus rates and equity backing ratios were assumed to be reduced in adverse investment scenarios (where the largest guarantee strains would otherwise emerge). Figure 3.4 shows how the total guarantee fair value is affected when these dynamic assumptions are ‘switched off’ — i.e. the bonus or investment policy is fixed irrespective of the investment returns experienced.

3.6.2 Interestingly, Figure 3.4 shows that the guarantee value of the specimen policy is very sensitive to these assumptions. Indeed, assuming bonus and investment policies are fixed almost doubles the value of the guarantees. Clearly then, the assumptions for the behaviour of with-profits liabilities can have a very significant impact on the value attached to policyholders’ benefits. The valuation tool can provide useful management information on the cost implications of the intended approach to the management of the with-profits business. Managers can explore the impact on the guarantee value when different dynamic management policies are put in place. For example, the guarantee value of £648 was calculated under the assumption that EBRs would not be changed by more than 10% in any single year. Figure 3.5 illustrates how the calculated guarantee value depends on this variable.

3.6.3 A similar analysis can also be made of the sensitivity of the guarantee value to the rate at which bonus rates can be changed. The

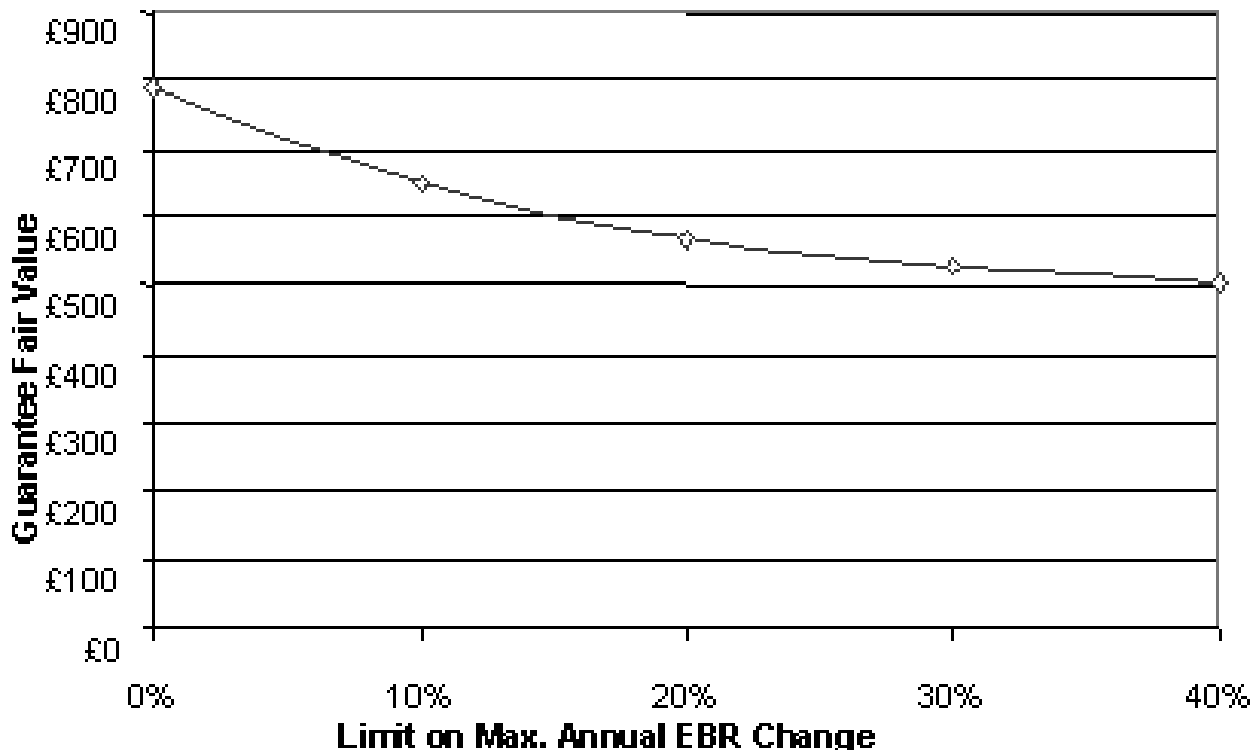


Figure 3.5. Fair value of guarantee and the assumed amount of investment policy flexibility

guarantee value of £648 assumed that reversionary bonus rates could not be changed by more than 1% in any year. Figure 3.6 shows how the guarantee value varies with this parameter (assuming a maximum EBR change of 10% in any single year).

3.6.4 The type of investigations presented above can help actuaries and managers understand the sensitivities of the costs of guarantees to the management policies applied to with-profits business.

3.6.5 The fair value of guarantees will also depend on movements in asset values and changes in assumed asset behaviour (i.e. market-implied levels of volatility). Table 3.2 shows how the fair value of the guarantees of the specimen policy change for different assumed levels of asset share movement and implied equity volatility.

3.6.6 As you might expect, the worst outcome for the guarantee valuation occurs when a significant asset fall is accompanied by rises in levels of option-implied volatility. This ‘double-whammy’ results in an increase in guarantee value of over 50%.

3.6.7 This brief analysis suggests that the selection of liability assumptions could be just as important to the pricing of with-profits guarantees as the asset model calibration (if not more so). However, in terms of managing the on-going variability of the fair value profit, the liability assumptions will only impact on fair value if and when they are changed (and

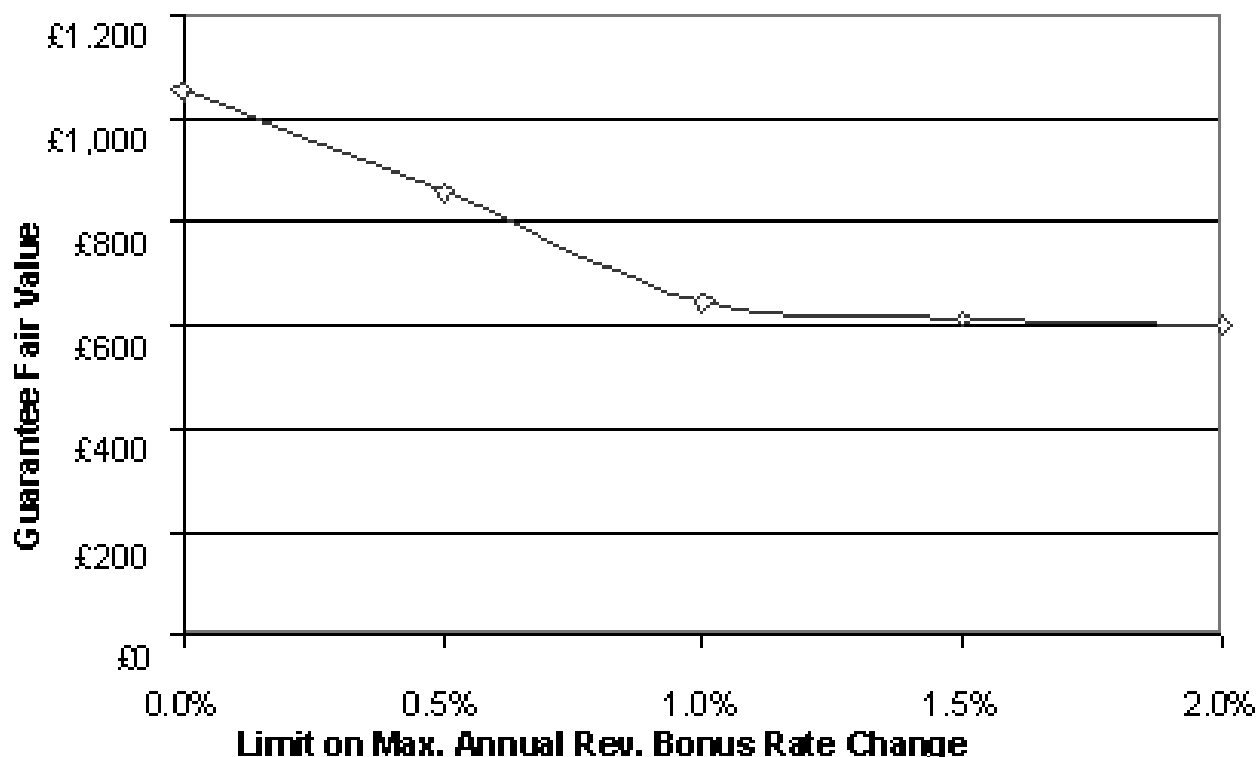


Figure 3.6. Fair value of guarantee and the assumed amount of reversionary bonus policy flexibility

Table 3.2. Sensitivity of fair value of guarantee to asset movements and changes in market-implied asset volatility

		Asset Share Movement		
		-20%	0%	+20%
Equity	20%	712	511	346
Volatility	23%	848	648	478
	26%	985	790	612

also to the extent that actual experience diverges from that assumed in the basis). No such control can be exercised over the asset values or the market-consistent model calibration. Unless hedged, movements in asset shares and market option costs will be a significant source of fair value profit and loss. We examine this risk, and how it can be managed, in the following section.

4. HEDGING WITH-PROFITS GUARANTEES

4.1 Introduction

4.1.1 So far, this paper has been concerned with market-consistent valuation of insurance liabilities, with particular attention being paid to with-

profits. This section moves beyond the valuation exercise to consider its implications for risk management. Those readers unfamiliar with the Black-Scholes option pricing and hedging arguments should now read Appendix A, which considers how we might price and hedge ‘plain vanilla’ guarantees. This demonstrates how fairly simple formulae can be applied to gain insight into how to hedge such guarantees. Unfortunately, the guarantees written by life offices often resemble no more than distant cousins of these standard guarantees. They are generally much more complex. In particular, the contingent, path-dependent nature of both the ultimate level of the guarantee and the conduct of investment policy complicate matters considerably. The convenient formulae of Appendix A turn out not to be very useful. Nonetheless, progress has still been made. There are some fundamental insights at the heart of the Black-Scholes-Merton analysis discussed in Appendix A that can be applied to virtually any form of investment guarantee.

4.1.2 One of the most useful of these insights is the notion of risk-neutral valuation. You may note that the Black-Scholes option pricing formula discussed in Appendix A does not include the expected return on the underlying asset as a parameter. This is because option values (by which we mean the implicit cost of replicating the option payoff) do not depend on the risk preferences of the option buyer. Option values are determined by the current market price of the underlying asset, the asset’s volatility and the risk-free interest rate yield curve. This turns out to be a very useful insight when we set out to estimate option values using Monte-Carlo simulation methods. It means that — for the purpose of option valuation — we can assume that investors are risk-neutral, and therefore that all assets have an expected return equal to the risk-free interest rate. This allows us to discount all cash flows at the risk-free interest rate in calculating option values. There is, of course, a great deal of mathematical rigour behind this result. (We suggest Chapters 1-3 of Baxter & Rennie (1996) as an excellent introduction to interested readers. Nor is risk-neutral valuation the only approach to market-consistent valuation of contingent claims.)

4.1.3 There is also a second important insight that we can gain from the initial analysis. At any given point in time, the hedge portfolio for a put option/guarantee will have two components:

- $\Delta(t) \times \text{Underlying asset } (t)$ invested in the underlying asset; and
- $\text{Guarantee } (t) - \Delta(t) \times \text{Underlying asset } (t)$ invested in a risk-free asset of duration equal to the term of the guarantee.

where $\text{Guarantee } (t)$ is the market-consistent guarantee cost at time t and $\text{Underlying Asset } (t)$ is the value at time t of the asset on which the option has been written.

4.1.4 Note that this hedge portfolio does not affect the investment policy of the asset share — it applies to the component of the realistic value which results from the guarantee only. The dependence works in the other

direction — the investment policy of the asset share will be an important determinant of the guarantee value and the composition of the appropriate hedge portfolio.

4.1.5 Since the guarantee cost can be calculated using risk-neutral valuation, as discussed above, the only unknown in the composition of the hedge portfolio is the ‘delta’ ($\Delta(t)$). In order for the above portfolio to offset changes in the guarantee value, the delta must be equal to the rate of change of the guarantee value with respect to the underlying asset value. So, we can use our guarantee valuation model to estimate the delta as well as the current guarantee value by calculating the ratio:

$$[G(S + \delta S) - G(S)]/\delta S$$

where S is the current value of the underlying asset, $G(S)$ is the value of the guarantee when the underlying asset is worth S , and δS is a small change in the underlying asset value. (Note that the same random number stream should be used in the two guarantee valuations required for this calculation. See Boyle *et al.* (1997) for a wider discussion of the application of Monte-Carlo simulation for valuation.)

4.1.6 So, although the complexities of the with-profits world have robbed us of the convenient formulae of the Black-Scholes-Merton analysis, it is still possible to apply their insights to find the guarantee value and the composition of the hedge portfolio for a with-profits contract using Monte-Carlo simulation. We are back on track.

4.2 The Guarantee Costs

4.2.1 Suppose that the ideas set out above are now applied to the example with-profits policy discussed in Section 3. You may recall that this policy is backed by a current asset share of £13,000, and has a guarantee valued (on an economic basis) at £648. The guarantee value takes account of the office’s discretion to reduce its equity content and reversionary bonus rates (to some extent) in times of poor asset performance. This discretion is a very significant factor in the guarantee valuation. Figure 4.1 highlights this further. It shows how the current guarantee value varies with asset share, both with and without the office’s discretion to adjust investment and bonus policies. Notice that the non-discretion case is always more expensive than where the office retains discretion. Although note how the differences become smaller as the guarantee becomes less valuable. You can see that at current asset share (£13,000) the removal of discretion would nearly double the value of the guarantee. A fall in asset share to £9,000 would increase the value of the guarantee to just less than £1,000.

4.3 Finding the Hedge Portfolio

4.3.1 Using the simulation approach described above, the current delta

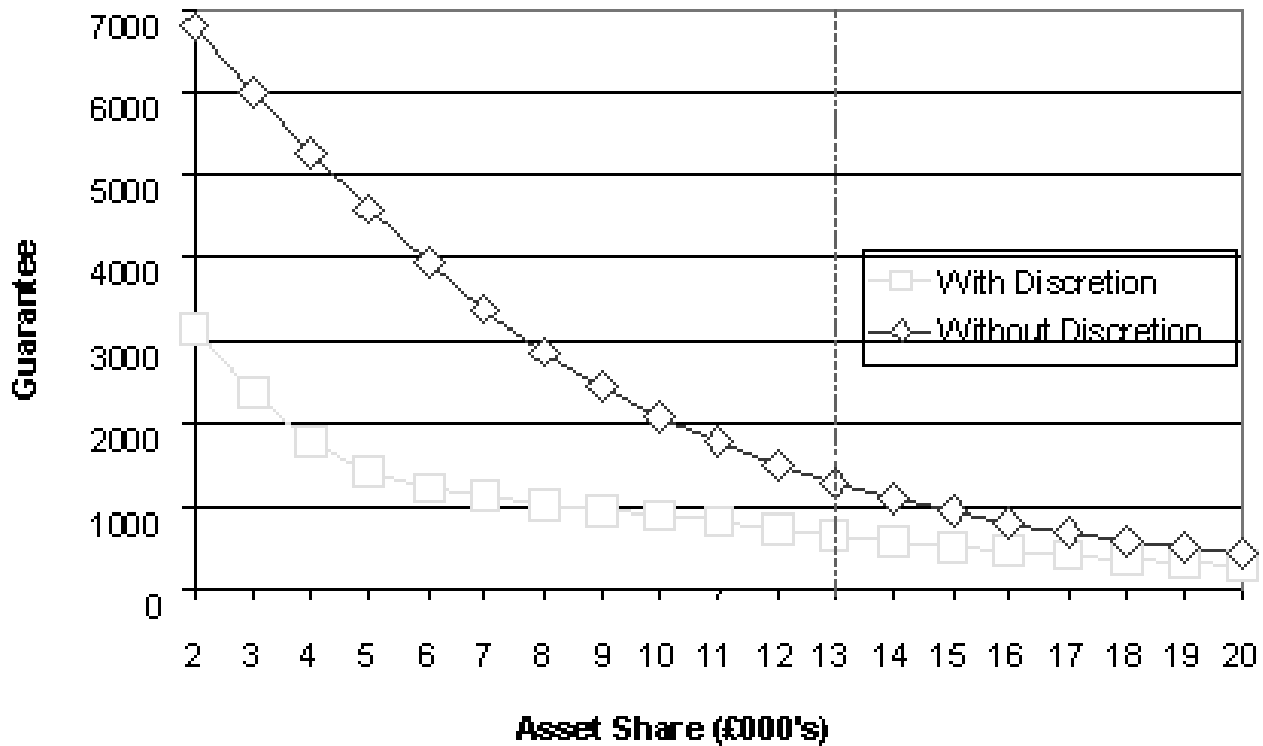


Figure 4.1. Guarantee costs as a function of asset share

of the guarantee fair value is estimated to be -0.08 . The current asset share of £13,000 and a guarantee fair value of £648, imply the following hedge portfolio:

4.3.2 We now use the simulation approach described above to calculate the delta of the guarantee as a function of asset share (remember this is just the rate of change of the guarantee cost with respect to movements in the underlying asset, i.e. the gradient of Figure 4.1). Figure 4.2 shows the results for the example with-profits policy, again with and without discretion being applied.

4.3.3 Firstly, note that the delta of the with-discretion case is almost always of smaller magnitude than the without-discretion case. This is intuitive. The office's ability to limit the impact of very poor investment returns by reducing bonus rates and/or EBRs, reduces the sensitivity of the guarantee cost to movements in asset share. This reduced sensitivity translates into a smaller magnitude of delta.

Table 4.1. Hedge portfolio composition

Parameter	Value
Asset mix of asset share	$-\pounds 1,040 (= \pounds 13,000 \times -0.08)$
Risk-free asset	$\pounds 1,688 (= \pounds 648 - -\pounds 1,040)$
Total	$\pounds 648$

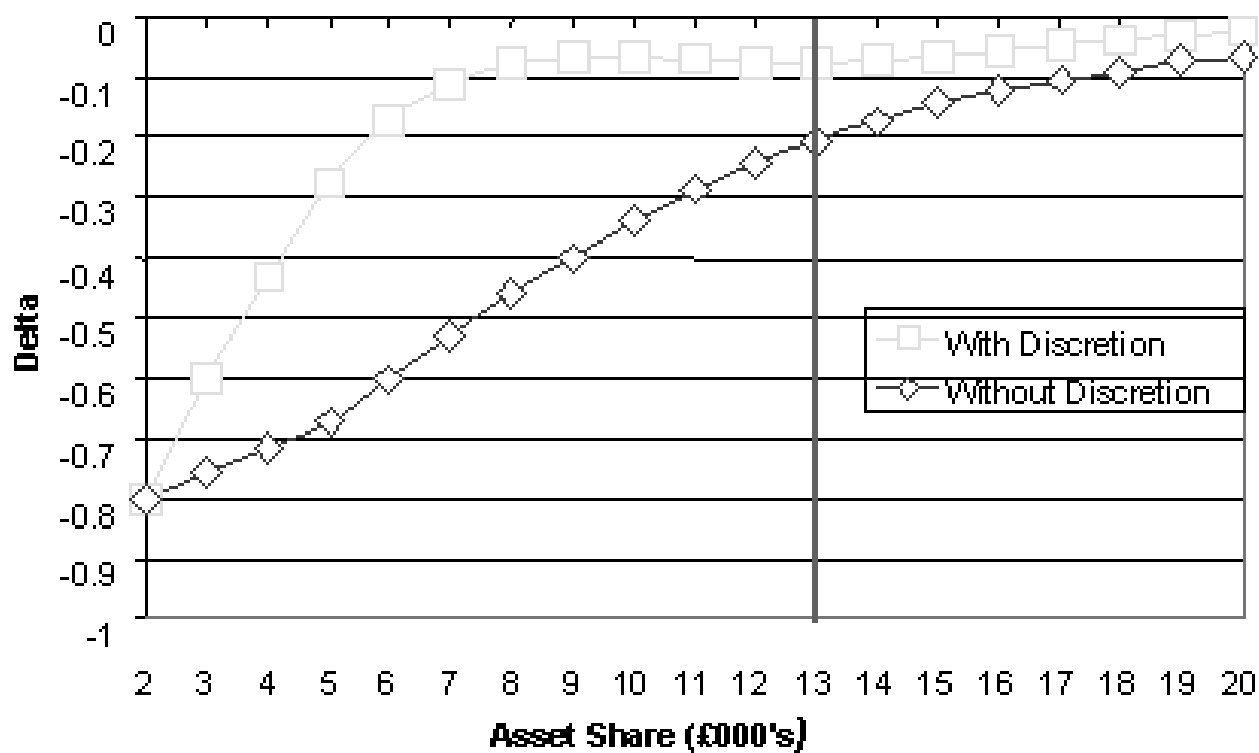


Figure 4.2. The sensitivity of guarantee costs to asset share movements — the delta

4.3.4 The behaviour of the delta in the without-discretion case is very similar to that of the Black-Scholes delta for a plain vanilla put option. This is unsurprising, as the without-discretion guarantee is a plain vanilla put option (at a strike of $\pounds 16,000 \times 1.04^{10} = \pounds 23,684$) subject to the minor complication that policyholder lapse rates are dynamic (policyholders are assumed to be less likely to surrender their policy when the guarantees are valuable).

4.3.5 In the with-discretion case, the delta falls much more slowly as asset share falls. This reflects how the dynamic rules help to mitigate the pain caused by poor asset returns. As a result, asset share could fall from $\pounds 20,000$ to $\pounds 8,000$ without the delta changing very much (indeed, our analysis suggests the with-profits guarantee delta may actually fall in magnitude slightly as asset share falls from $\pounds 13,000$ to $\pounds 9,000$). However, eventually the dynamic rules cannot keep pace with the impact of further asset falls, and the with-discretion delta starts to catch up with the without-discretion delta. By the time that asset share has fallen to $\pounds 2,000$, all the discretion has been 'used up', and the with and without cases have very similar deltas.

4.3.6 The behaviour of the with-profits hedge ratio in the presence of dynamic bonus and investment rules highlights the striking difference between these guarantees and standard plain vanilla guarantees. If hedging programmes are going to be applied, a modified Black-Scholes formula is

unlikely to be very helpful. A more sophisticated estimation technique, such as the Monte-Carlo simulation approach used here, will be required.

5. MANAGING PROFIT VOLATILITY AND CAPITAL REQUIREMENTS

5.1 *Introduction*

5.1.1 Section 4 showed how to measure the sensitivity of the value of a with-profits policy guarantee to movements in underlying asset share. From this we were able to identify a hedge portfolio for the guarantee. In this section we extend this analysis to gain insight into the potential profit variability of a life office which owns with-profits guarantees and into the potential impact of different management strategies on the level of prudential (risk-based) capital held. You may recall from above that this policy's guarantees were valued at £648. As in the previous section, we will use a Monte-Carlo model to develop our analysis. Various 'candidate' strategies for the management of the assets backing the guarantee will be analysed (note that this does not affect the investment policy for the asset shares of the with-profits policy — this is taken as an 'input', and we are now concerned with managing the assets backing the resulting guarantee cost of £648 in a way that minimises the risks to the office).

5.2 *Changes in the Guarantee Fair Value*

5.2.1 Let us begin by reviewing the sensitivity of the guarantee value. Figure 5.1 shows the guarantee fair value at the end of the year as a function of the asset share return earned over that year. In this figure we assume that option-implied volatilities (and so the asset model calibration) are unchanged from the original valuation date (where the guarantee fair value was estimated at £648).

5.2.2 As might be expected, poor asset returns result in the guarantee fair value increasing. The chart also illustrates that the relationship is not quite linear. The curvature relates to the dependence of the guarantee delta on the size of asset share. The remainder of the paper considers how we can hedge the risks created by the guarantee fair value's sensitivity to asset share movements and other factors.

5.3 *Hedging Changes in the Guarantee Value*

5.3.1 In Section 4 we used the fair valuation model to derive the hedge portfolio for the example policy's with-profits guarantees. You might recall from Section 4 that the hedge portfolio consists of £1,688 of the risk-free asset and a short position of £1,040 in the asset mix of the underlying asset share's asset mix. (For the purposes of this analysis, the risk-free asset has been assumed to be cash. In reality, a more appropriate risk-free asset would have a duration equal to the duration of the guarantee.) We now

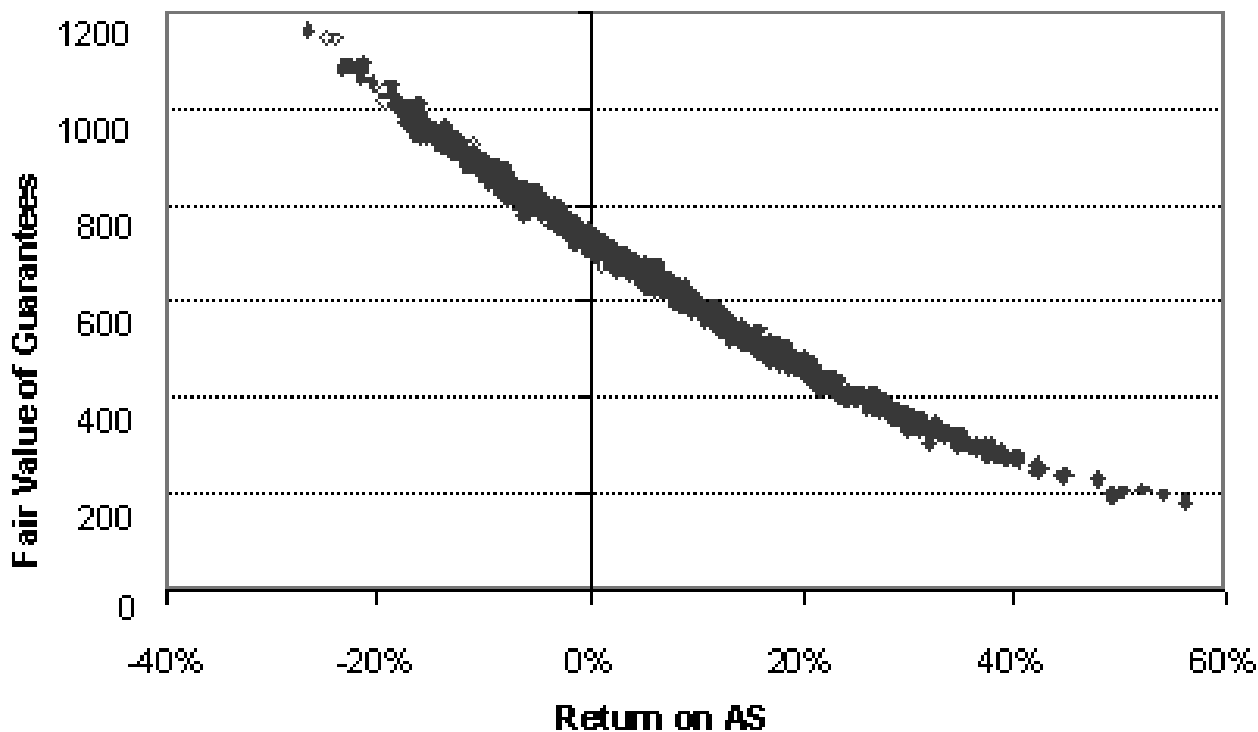


Figure 5.1. End year fair value of guarantees vs return on asset share (with discretion)

consider the hedged fair value guarantee profit, and compare this with the obvious alternatives of investing the assets backing the guarantee fair value in the same asset mix as asset shares or in cash. For ease of modelling, we have assumed that the hedge portfolio is not rebalanced over the course of the one-year projection. (As the delta must be calculated using simulation, to rebalance the hedge portfolio more frequently would require many, many thousands of simulations within each of our simulations of the one-year progress of the fair value guarantee profit.) This makes the hedging results a little poorer than we could otherwise expect — as is discussed in Appendix A, more frequent rebalancing of the hedge portfolio will improve the accuracy of the hedge. On the other hand, we ignore the possibility that option-implied volatilities (and so the end year valuation basis) may change. Dynamic hedging cannot offset the effects of such variability. The consequences of this failing in the hedge are examined later. In the following analysis, the fair value guarantee profit is defined as follows:

$$\begin{aligned} \text{Fair value guarantee profit} = & [\text{Hedge portfolio } (t + 1) - \text{Hedge portfolio } (t)] \\ & - [\text{Fair value of guarantees } (t + 1) - \text{Fair value} \\ & \text{of guarantees } (t)] \\ & - \text{Guarantee cash flow strain } (t, t + 1). \end{aligned}$$

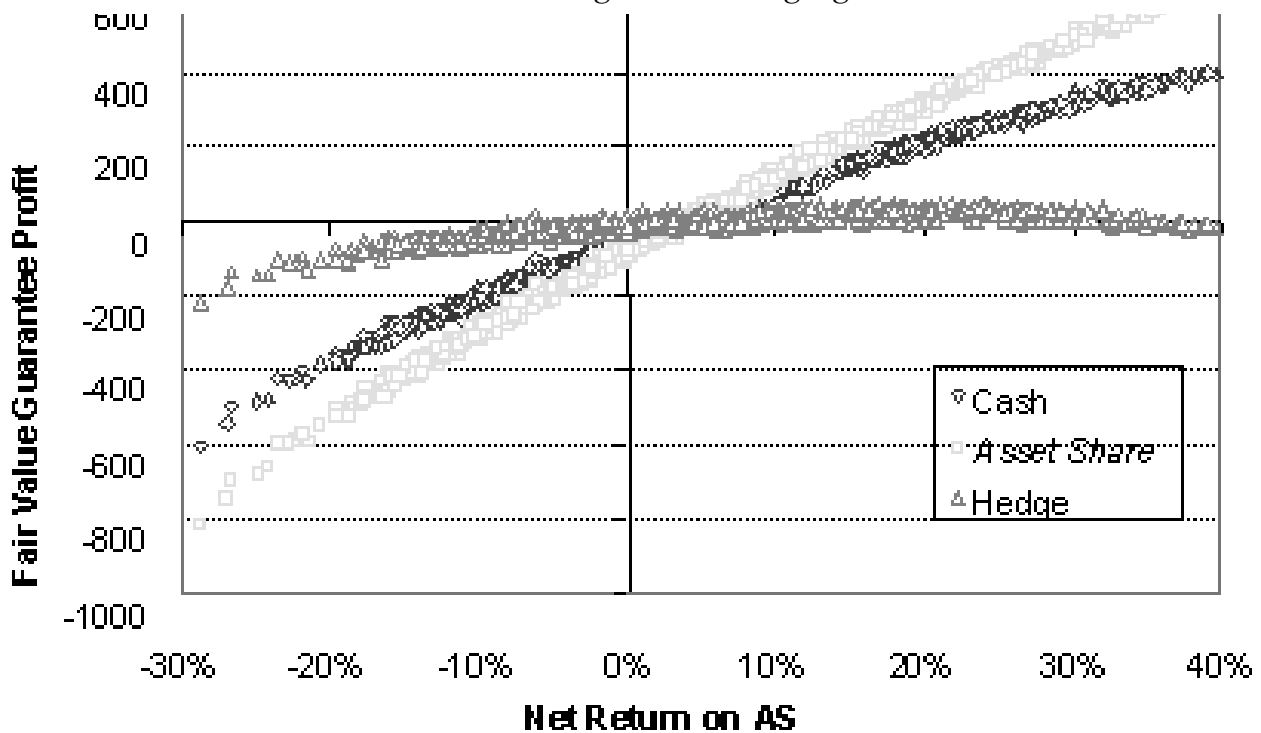


Figure 5.2. Fair value profit as a function of return on asset share and hedging policy

5.3.2 Figure 5.2 shows the fair value guarantee profit emerging over the year as a function of the year's asset share return, where the hedge portfolio is invested in each of the asset mix of asset shares, cash and the dynamic hedge portfolio. (In simulating ahead one year in the 'real world' — as opposed to a risk-neutral world — we have assumed the same level of equity volatility as in that used in the risk-neutral market-consistent projections (i.e. 24% p.a.). This is probably quite prudent, as there are natural reasons to expect option-implied levels of volatility to exceed actual volatility (e.g. the presence of transaction costs, gap risk). We have also assumed an equity risk premium of 4% p.a. in the one-year real world projection.)

5.3.3 Let us work through each of these portfolios in turn. When the assets backing the guarantee are invested in cash, the only significant source of volatility in the fair value guarantee profit is the change in the fair value of the guarantees. As we saw in Figure 5.1, the guarantee fair value is highly dependent on movements in asset share, and this dependence naturally follows through to the fair value guarantee profit. Investing the assets backing the guarantees in the asset mix of the asset shares adds an element of gearing to these profits. In this case, strong asset share returns not only reduce the fair value of the guarantees, they also increase the value of the assets backing them. A similar story applies on the downside — the backing assets fall in value just when the guarantee becomes more burdensome. Finally, we have the hedge portfolio. If this portfolio was rebalanced

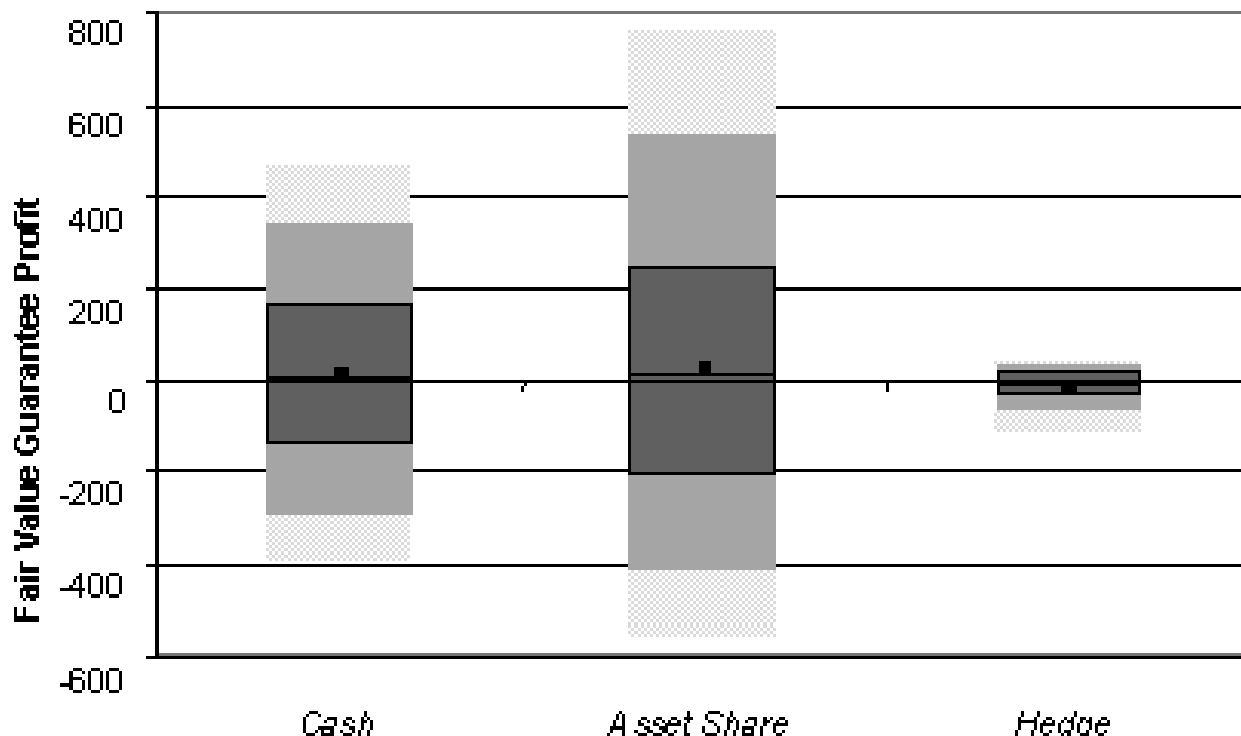


Figure 5.3. Fair value guarantee profit as a function of hedging policy

continuously, the profits would always be zero (ignoring market jumps, transaction costs, etc.). The curvature that is exhibited in the hedged guarantee profits is a result of the simplifying assumption that the hedge portfolio is not rebalanced until the end of the year. As the rebalancing frequency is increased, this line would gradually flatten.

5.3.4 Figure 5.3 summarises the profiles of the fair value guarantee profits emerging at the end of the year. You can see that, even when only rebalancing on an annual basis, the hedged profits are significantly less variable than investing in cash or asset share. We have used our model of the with-profits business to find a hedge portfolio that is consistent with its office-specific features.

5.3.5 Changes in option-implied volatility

In the above analysis, we made the convenient assumption that the volatility parameters of the fair valuation asset model would remain fixed from one year to the next. Unfortunately, in reality this is unlikely to be the case. Option-implied volatilities tend to bounce around over time, reflecting changes in investors' expectations for future asset volatility and the effects of supply and demand. (Few markets exhibit the infinite price elasticity associated with perfect markets).

5.3.6 Changes in the volatility basis for the asset model will naturally impact on the fair value of guarantees. We now introduce a stochastic element to the end year option-implied equity volatility level. (We have

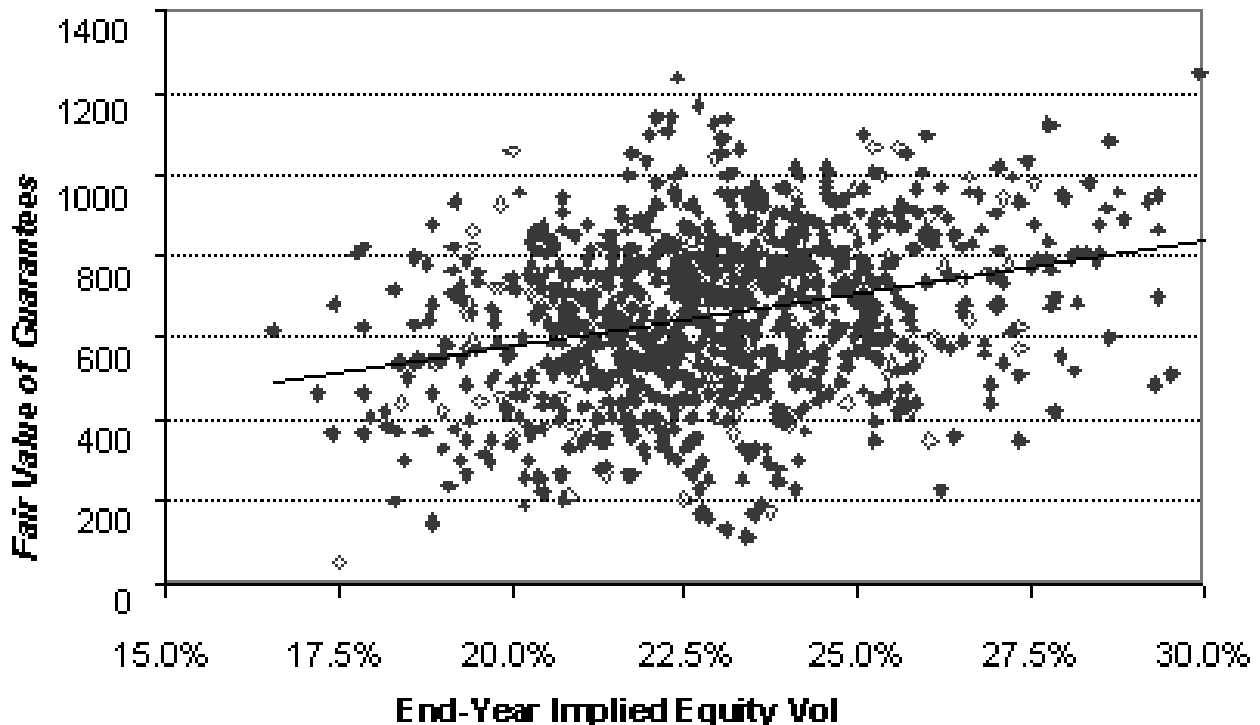


Figure 5.4. End year guarantee fair value as a function of end year option-implied equity volatility

assumed that the end year equity implied volatility level is lognormally distributed with a mean equal to its current value of 23% and a proportional volatility of 10%. This is not an especially sophisticated model of implied volatility, but will suffice for these illustrative purposes. We have ignored possible changes in option-implied interest rate volatility in this analysis. This is another source of hedging error that would arise under the strategy described above. It is likely to be difficult to hedge in practice. In our example policy, it is, however, likely to be of smaller magnitude than the risk created by movements in option-implied equity volatility.) Figure 5.4 shows how the simulated end year guarantee fair value relates to the accompanying level of option-implied equity volatility.

5.3.7 Figure 5.4 illustrates that the option-implied level of equity volatility is an important determinant of the fair value (although it remains second order to the actual asset share return earned in the year). Figure 5.5 illustrates the end year guarantee fair value's dependence on both asset returns and end year option-implied volatility. This figure highlights how the fair value's dependence on volatility changes is greatest when asset returns have been strong (and the guarantee has moved out-the-money).

5.3.8 It is important to understand that it will not be possible to hedge the variability arising from movements in option-implied volatility using a




Figure 5.5. End year guarantee fair value as a function of asset returns and end year option-implied equity volatility

hedge portfolio comprised of the underlying fund and risk-free bonds. Figure 5.6 shows the distribution of the guarantee fair value profit if we allow for movements in implied volatility.

5.3.9 Introducing variation in the end year option-implied equity volatility has increased the variability of the fair value guarantee profit in all three cases considered. However, the impact on the performance of the hedge portfolio is the most marked. There is not much that we can do about these changes in implied volatility if the hedge portfolio is restricted to investments in the underlying asset and risk-free assets. This is highlighted by Figure 5.7, which shows the fair value guarantee profit arising when following the above hedging strategy plotted against the simulated end year implied equity volatility. Whilst Figure 5.2 showed that the dynamic hedge removed most of the fair value profit's exposure to movements in underlying assets, Figure 5.7 shows that the sensitivity to changes in implied volatility remains.

5.3.10 In order to hedge the impact of these volatility changes, the hedge portfolio must include assets whose value depends on the implied volatility of the underlying asset. The following section considers how, in principle, this might be achieved.

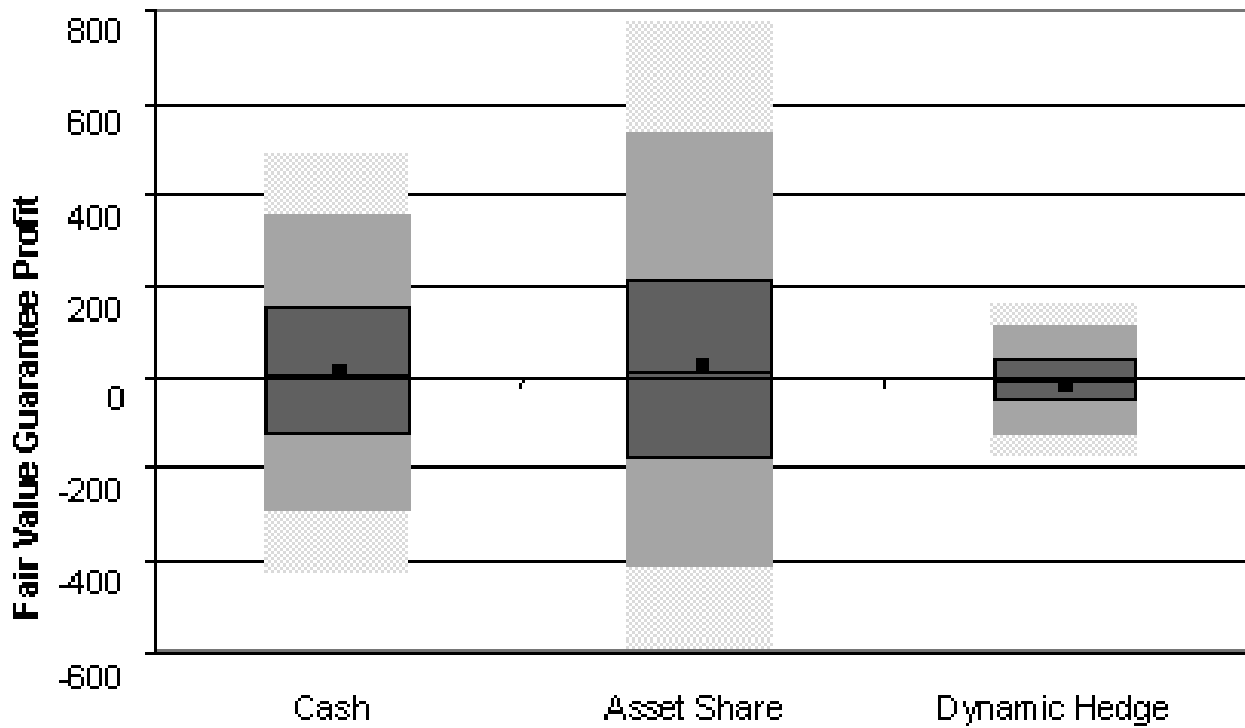


Figure 5.6. Guarantee fair value profit (with stochastic implied volatility)

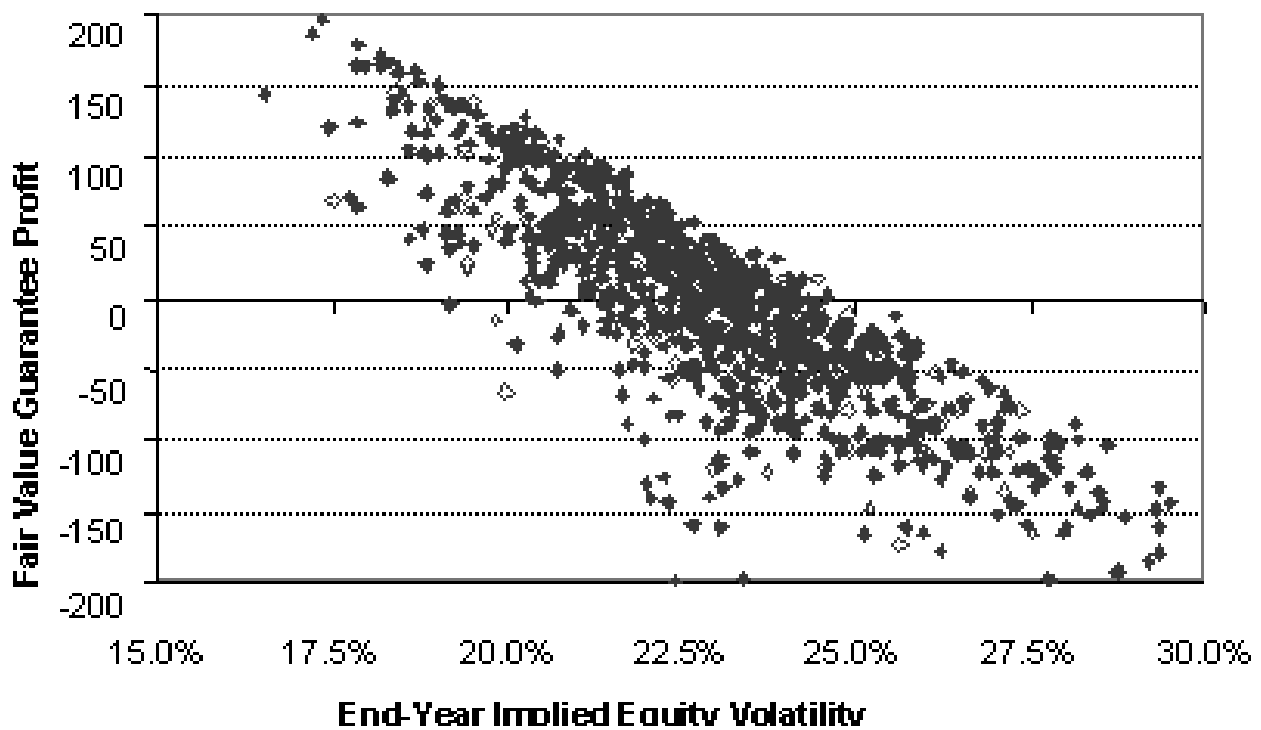


Figure 5.7. Guarantee fair value profit under dynamic hedge (with stochastic implied volatility)

5.4 Other Risk Management Strategies — Developing Derivative-Based Solutions

5.4.1 Dynamic hedging of guarantees using a portfolio of the underlying asset and the risk-free asset can be considered to be the ‘purest’ form of guarantee risk management. Indeed, in theory it delivers the complete removal of market risk. However, as we have seen, there are sources of risk that basic dynamic hedging simply ignores, and there are, inevitably, assumptions underlying the approach that will not always hold in reality. The real world can be a more demanding environment than a computer model. Further, the practical demands of dynamic hedging, and its potential pitfalls, can make this strategy rather daunting to a life office with little or no experience of applying this type of risk management technique. Such offices may feel inclined to ‘sub-contract’ the dynamic hedging (and its costs and risks) to an investment bank by transacting some form of derivative trade. Of course, a bank will naturally require compensation for such a transaction, and the bank’s margin will be an important determinant of the attractiveness of such a trade. Another important consideration will be the extent to which the transaction addresses the office-specific features of the guarantees (Figure 4.2 highlights how important these features are in determining the appropriate hedging strategy for the example with-profits policy).

5.4.2 To illustrate the possibilities in this area, we now review some example derivative transactions that the office might consider as a means of managing its guarantee risk. In all cases, we will consider how the guarantee risks identified above can be managed through suitably selected portfolios of ‘plain vanilla’ options (i.e. standard options to buy or sell an asset with a fixed strike price, expiry date, etc.). Our analysis has shown that matching movements in the guarantee value requires the following characteristics of the hedge portfolio:

- (a) It matches the sensitivity of the guarantee value to changes in the underlying asset value. We have seen that this will ensure that the hedge matches the change in the guarantee value for small changes in the asset value. (In options terminology this is called the option’s delta.)
- (b) It matches the sensitivity of the guarantee value’s delta to changes in the underlying asset value. This will ensure that the hedge matches the change in the guarantee value for large changes in the asset value. (In options terminology this is called the option’s gamma.)
- (c) It matches the sensitivity of the guarantee value to changes in implied volatilities. (In options terminology this is called the option’s vega.)

5.4.3 Our aim is to find a combination of assets which exhibit all of these characteristics. Note that this represents a significant departure from the approach adopted above. In the dynamic hedging case, we effectively create the same derivative (put option) that is attached to the with-profits policy. Since such a complex option is unlikely to be offered by a bank, we will tackle the risk management problem by finding a portfolio of vanilla

options that match — as closely as possible — the characteristics identified above.

5.4.4 A natural starting point is to confine ourselves to a plain vanilla put option. Let us assume that a bank is prepared to write such an option on the current asset mix of the asset share with a maturity date equal to the date of the maturity guarantee. (Some offices' with-profits fund asset mixes may make this assumption very optimistic — especially if, for example, there are significant direct property holdings; but this should not detract from the relevance of this approach, which equips the user with a better understanding of the drivers of liability risk and their magnitude.) In this case, you might expect that simply buying an option with a strike that is in some way consistent with the expected level of the final sum assured might be the best hedging strategy. However, it is important to remember that the analysis of Section 4 showed that the discretionary features of the with-profits guarantee had a significant impact on the nature of the guarantee and its sensitivity to changes in the value of the underlying asset share. So, perhaps the hedge portfolio composition should be more flexible?

5.4.5 Our plain vanilla put option portfolio has (at least) two moving parts. First, we can vary the number of options that we buy. Second, different strike prices can be selected. Table 5.1 shows a candidate option position that has the desired delta of -0.08 .

5.4.6 Interestingly, the strike of this option, £11,644, is around half the projected final level of the guaranteed sum assured ($£16,000 \times 1.04^{10} = £23,684$). Further, our analysis suggests that holding less than one of these option contracts per policy is appropriate for the example policy (in order to obtain the required offsetting exposure to asset share movements). The differences between this strategy and simply buying one put option with a strike equal to the expected guarantee benefits captures the impact that the discretionary and dynamic features of with-profits have on the current sensitivity of the guarantee value to movements in asset share. This option portfolio has an initial cost of £459.

5.4.7 Figure 5.8 shows how the delta of the above put option portfolio compares to that of the fair value of the guarantees (assuming the remaining £189 of the guarantee fair value is invested in risk-free assets).

5.4.8 You can see that the plain vanilla put option delta does not behave like the guarantee. In options parlance, the put option's gamma (i.e. the rate of change of the delta with respect to the underlying asset) is just too big. This is especially true in the event of asset share falls — the

Table 5.1. Candidate plain vanilla put option portfolio

Parameter	Put
Number of options per policy	0.84
Strike price of option	£11,644

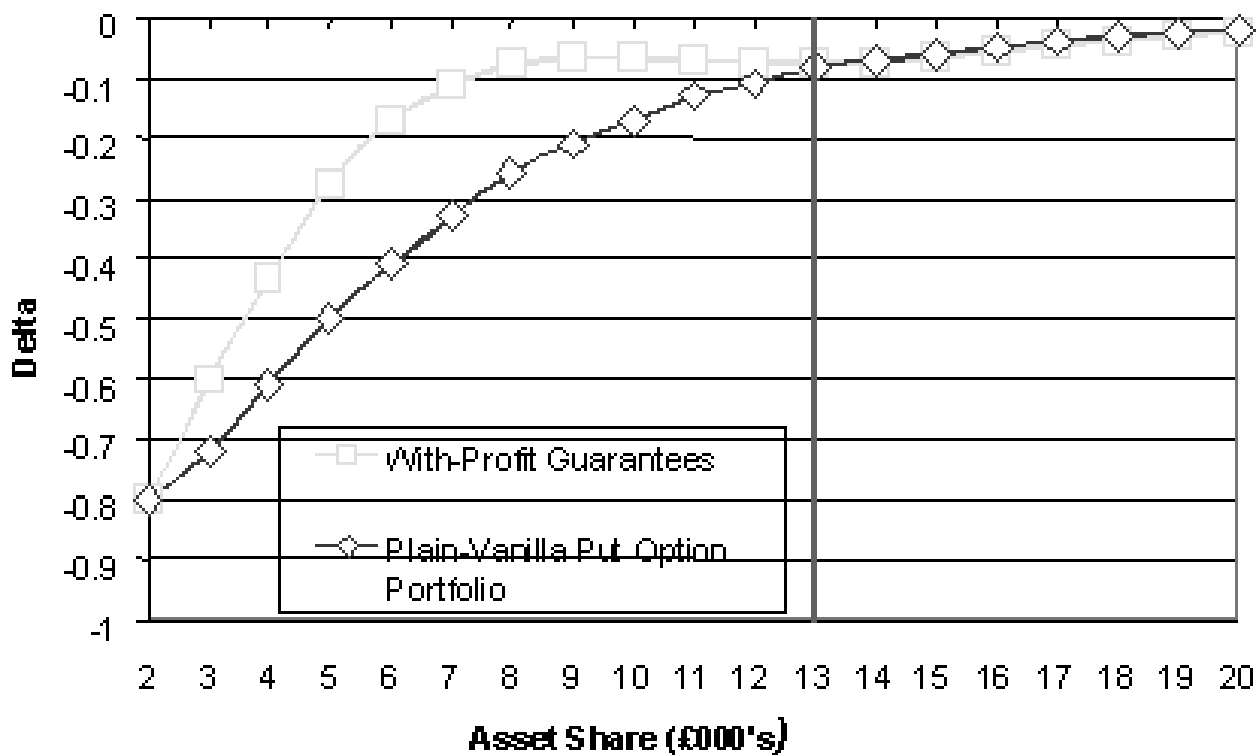


Figure 5.8. Deltas of with-profits guarantees and plain vanilla option portfolio

discretionary adjustments of bonus rates and investment policy that would accompany such a fall help to keep the with-profits guarantee's delta small (in magnitude), and this feature is not captured by the vanilla put option.

5.4.9 So, the gamma of the option portfolio needs to be reduced. Vanilla options (both calls and puts) have positive gammas, so buying more options is unlikely to help. On the other hand, selling some options would reduce the hedge's gamma and could produce a better match to the guarantee. It is therefore worth considering a 'collar' portfolio. Under such a strategy, calls with high strikes are sold whilst puts with lower strikes are bought. In this illustrative example, we again suppose that such options can be bought and sold on the current underlying asset mix of the fund (80%/20% equity/gilts).

5.4.10 With the collar strategy, there are now (at least) four 'dials' that we can turn in determining the composition of these option positions:

- the number of call options sold;
- the strike price of the options sold;
- the number of puts bought;
- the strike price of the puts bought.

5.4.11 Table 5.2 shows a collar strategy which has the same initial delta as the guarantee fair value (-0.08), and a cost of £132. As there are more variables than there are constraints, this is not a unique collar strategy, given

Table 5.2. Candidate collar option portfolio

Parameter	Call	Put
Number of options	-0.10	0.89
Strike price	£35,118	£9,770

the required initial delta. In reality, practical market constraints will dictate which particular strategy is most attractive.

5.4.12 Notice that the call option position is smaller than the put position. This is unsurprising, as the guarantees are a type of put option. However, the negative call exposure does reduce the gamma of the portfolio. This is illustrated in Figure 5.9. You can see that the collar option portfolio provides a closer match to the delta profile of the with-profits guarantee.

5.4.13 We have now identified three candidate risk management solutions — the dynamic hedge, a put option portfolio and a collar strategy. Let us now review the results produced by these strategies. Figure 5.10 shows the fair value guarantee profit produced by these strategies as a function of the return on asset share earned over the year.

5.4.14 You can see that all three strategies show a reasonably similar exposure to the asset share return. In each case, despite initially being delta-neutral, some positive relationship between asset share return and fair value profits remains. This is because asset share movements result in the deltas of

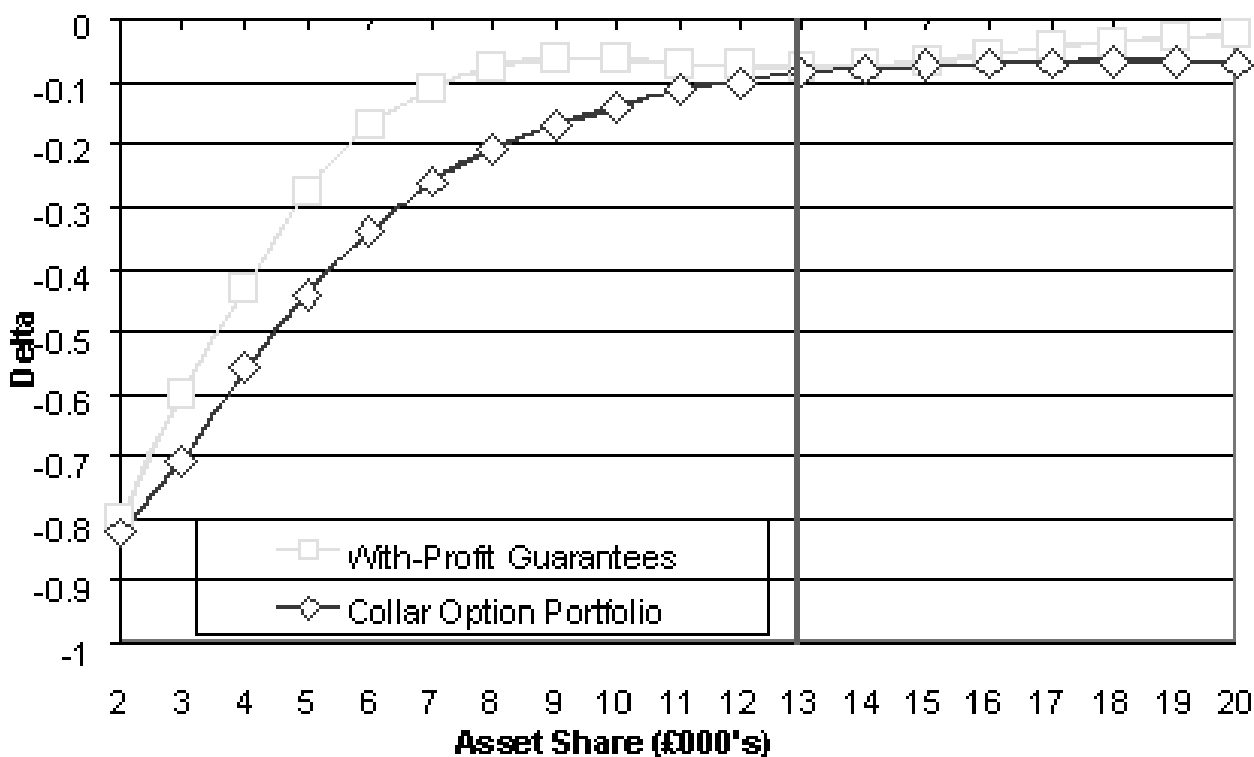


Figure 5.9. Deltas of with-profits guarantees and collar option portfolio

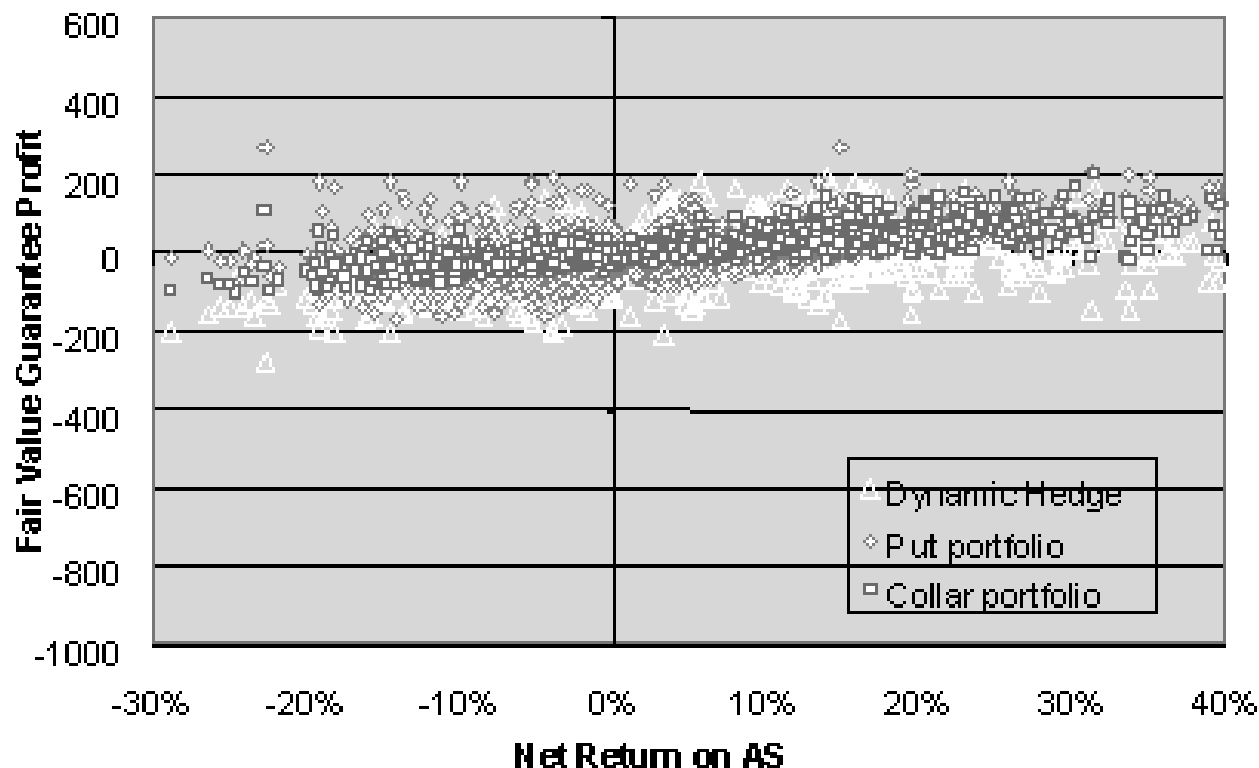


Figure 5.10. Fair value profit as a function of asset share

the hedge portfolio and the guarantees diverging slightly as asset share exposure sneaks back into the net position. However, Figure 5.10 also illustrates a second feature. For a given level of asset share return, there is greatest variation of fair value guarantee profit under the dynamic hedge, and least variation for the collar portfolio. What might be driving this variation in the amount of profit emerging for a given asset share return?

5.4.15 Figure 5.11 shows how the fair value profit relates to the end period option-implied equity volatility.

5.4.16 Interestingly, the different strategies have different exposures to volatility changes. As we have seen, increases in volatility have a negative effect on the fair value guarantee profits when delta hedging. This is natural, as the asset portfolio value does not depend on volatility, whereas the guarantee fair value does — higher volatility means higher guarantee fair value. With the put option portfolio, an opposite relationship is observed. In this case, the put option is more sensitive to changes in volatility than the guarantee fair value. On the other hand, the net position under the plain vanilla put strategy has a positive exposure to volatility movements. The collar strategy has less sensitivity to volatility changes than the put strategy, as the collar strategy involves *selling* some options. As a result, the net position under the collar strategy has the least sensitivity to movements in volatility.

5.4.17 You can see from Figure 5.12 that the three strategies do fairly

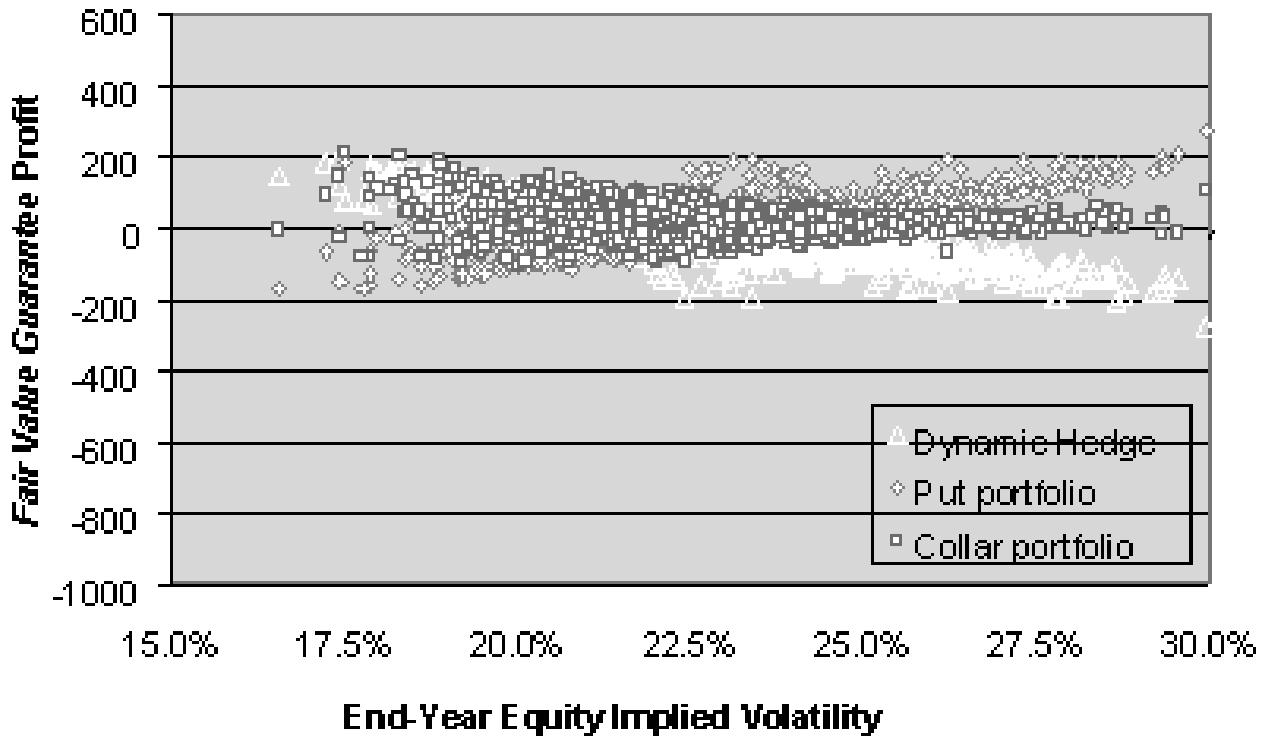


Figure 5.11. Fair value profit as a function of end year option-implied equity volatility




Figure 5.12. Guarantee fair value profit for various risk management strategies

Table 5.3. The greeks of the liability and candidate hedge portfolios

	Liability		Asset: Candidate hedges			
	With-profits guarantee	Dynamic hedge	Put option	Long put	Collar Short call	Net collar
Delta	-0.08	-0.08	-0.08	-0.05	-0.03	-0.08
Gamma	0	0	++	++	-	+
Vega	+	0	++	++	-	+

similar jobs of managing the one-year fair value profit. The collar strategy does a slightly better job than the put option, and also has less downside risk than the dynamic hedge (although our treatment of the dynamic hedge is perhaps a little harsh, as we have not rebalanced the portfolio over the course of the year).

5.4.18 Table 5.3 summarises the delta, gamma and vega exposures of the liability and the three hedges considered above. In managing the office's risk exposures, the aim of the risk manager will be to ensure that the characteristics of the hedge portfolio match those of the liability as closely as possible (all other things — cost in particular — remaining equal). You can see that, in our example, the collar hedge does the best job of replicating the 'greeks' of the with-profits guarantee.

5.4.19 Before we conclude this section, there are a couple of caveats that should be attached:

- The derivative contracts that we have analysed are not readily available. The term and strike of the options (as well as the size likely to be useful to a life office) are not easily available from the financial markets. In practice, banks may be willing to offer such instruments on an 'over-the-counter' basis. However, such trades are likely to incur significant trading costs. The margin taken by the bank could be a vital factor in appraising the attractiveness of these derivative strategies. Nonetheless, many of the benefits of the above strategies could be gained through trading appropriate portfolios of exchange-traded options only.
- Further, we have considered the performance of these strategies over one year only. It is likely that the option strategies would need to be revised at the end of the year — as delta positions move out of shape and the asset mix of asset share is changed. If trading OTC derivatives is costly, this could become problematic.

5.4.20 However, we can still claim to have made some interesting progress in the development of risk management solutions for with-profits guarantees. This analysis highlights how a valuation model for with-profits guarantees can be used for much more than 'merely' estimating a market-

consistent value for the liability. It can be used as a risk management tool, allowing the actuary to understand the dynamic characteristics of liabilities and the types of instruments which might be used to limit those exposures. As well as estimating the accuracy of candidate solutions, the residual risks that remain (and their sources) can be analysed.

5.4.21 Nonetheless, we should not get carried away. The model can be an invaluable tool in determining the appropriate composition of candidate hedge portfolios. Ultimately, however, the choice of which strategy will depend on a number of issues of judgement that lie outwith the jurisdiction of a modelling tool.

5.5 *Implications for Risk-Based Capital Requirements*

5.5.1 The FSA has made it clear that, in the coming years, the regulatory capital regime for insurers will move from the current resilience test basis towards a market-consistent, risk-based approach (see, for example, the FSA's Consultation Papers 136 and 143). The details of this remain sketchy. However, irrespective of the finer points of this regime, under a risk-based approach, hedging should significantly reduce regulatory capital requirements. As an illustrative example, let us suppose that the minimum regulatory capital requirement is defined as the capital sufficient to fund the end-of-year realistic guarantee value with 99% probability. This is slightly different to a more natural actuarial approach of defining capital requirements in terms of the capital required to fund a given percentile of ultimate losses arising from running off the business. The two approaches are actually more similar than they first appear, although they are answering slightly different questions. The run-off approach assumes the current 'hedging' strategy is applied until the business is run off. The approach that we have used assumes it is applied for the next year — the capital is designed to fund the replicating portfolio cost at the end of the year. The differences between these two approaches will be greatest when the assets backing the guarantees represent a significant mis-matched position. In this case, our approach recognises management's flexibility to change this strategy (to the replicating strategy) at the end of the year, whereas the run-off approach assumes that capital must be provided to fund this mis-match position indefinitely.

5.5.2 We can decompose such a capital requirement into the current realistic value and a mis-match reserve. In the theoretical limit, the perfect hedge will reduce the need for any mis-match reserve, and the regulatory capital required would simply be the current realistic value. On the other hand, investing the assets backing the guarantee in equities will create the need for a considerable mis-match reserve. Figure 5.13 plots the hypothetical regulatory capital requirements under the above assumptions for the specimen CWP policy considered throughout this paper. We consider three cases for the investment of the assets backing the realistic guarantee value —

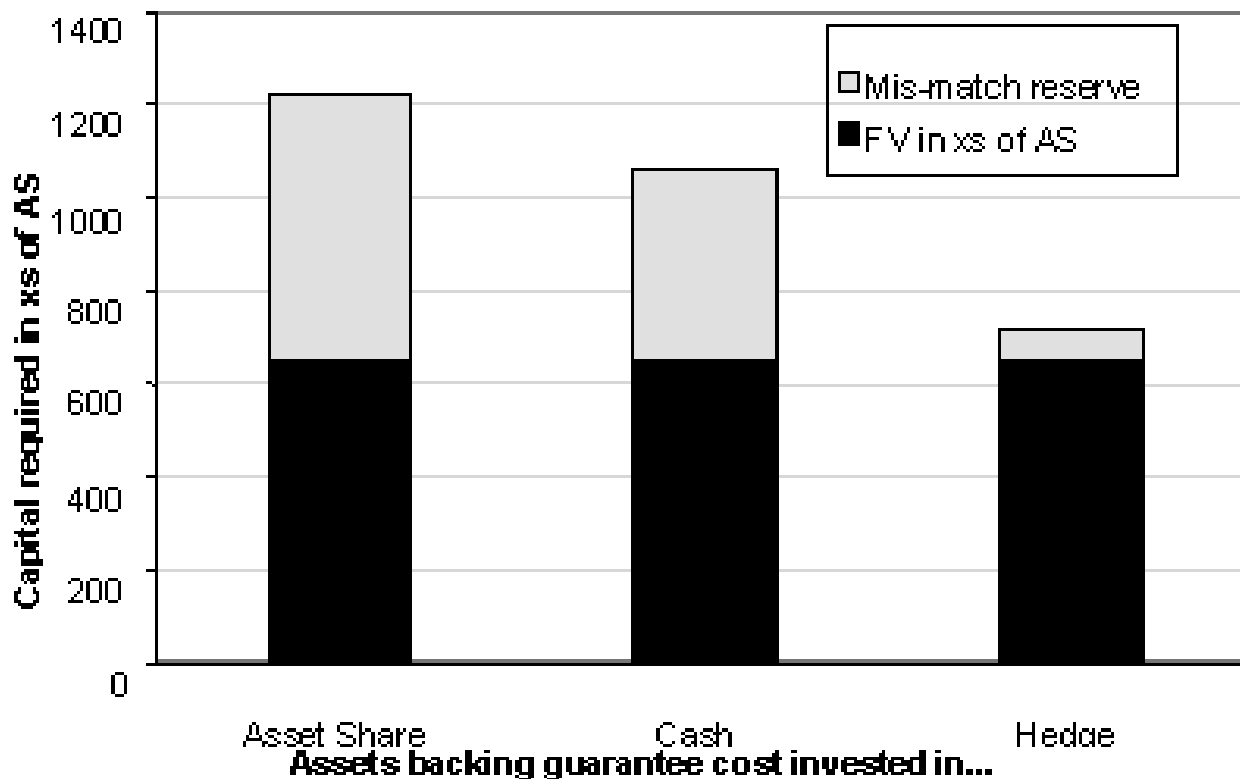


Figure 5.13. Risk-based capital requirements example for the example CWP policy

the asset mix of asset share, cash and the collar hedge discussed in Section 5.3. You can see that, even relative to holding cash, implementing the hedge reduces the capital strain in excess of asset share by almost one third.

6. SUMMARY AND CONCLUSIONS

6.1 The findings and conclusions of this paper can be considered in two parts: those relating to the determination of the fair values; and those relating to their risk management ramifications. The following points summarise the major conclusion from our discussion of fair valuation for with-profits business:

- Asset model calibration for the valuation of long-term guarantees can raise a number of potentially messy issues. There is a lack of liquid, transparent market instruments to which models can be calibrated. Stochastic models will need to be fairly complex if they are to replicate accurately all relevant market prices. A trade-off will therefore arise between model simplicity and pricing accuracy.
- Detailed assumptions will also be required for the behaviour of with-profits liabilities. The fair value of with-profits guarantees could be very sensitive to these assumptions. Indeed, the fair value could be

significantly more sensitive to these assumptions than to differences in the ‘fine-tuning’ of market-consistent asset calibrations.

- The valuation of a specimen conventional with-profits contract was fairly sensitive to the wide range of modelling assumptions required for the behaviour of assets and liabilities and their interaction.
- The fair value of a with-profits liability could conceivably exceed its statutory valuation. The lack of consistency between the two approaches makes it very difficult to make general statements. For example, a year-on-year increase in one of the valuations need not imply an increase in the other. However, the key difference in the treatment of prospective terminal bonus payments is likely to mean fair values often exceed statutory reserves.

6.2 This paper has also discussed approaches to the management of the financial risks created by writing cash guarantees. The basic ideas underlying dynamic hedging are discussed in Appendix A. There, we see that a replicating portfolio can be found for a vanilla put option, and examine its properties in different scenarios. Further, by the principle of no-arbitrage, we know the cost of the replicating portfolio must also be the price of the option.

6.3 Section 4 extended these ideas to the example CWP policy that was valued in Section 3. It turns out that the complexity of with-profits guarantees means that the hedge portfolio cannot be identified through the application of a simple formula. However, using the insights developed by Black-Scholes-Merton, we were able to use a market-consistent valuation model to find the required hedge portfolio. This showed that the dynamic features of with-profits can significantly reduce the exposure of the guarantee value to movements in asset values. As a result, the sensitivity of the hedge portfolio to changes in asset share was significantly lower than for a vanilla guarantee (put option). In effect, the dynamic nature of with-profits does part of the hedging automatically. The valuation model can be used to understand how much risk has been removed by the office’s ability to apply discretion and how much market exposure remains with the office.

6.4 Section 5 considered the performance of the hedge portfolio identified in Section 4 for the CWP guarantee over a one-year period. Whilst the hedge portfolio removed the bulk of the guarantee profit’s exposure to asset movements, it could offer no protection against movements (increases) in the option-implied volatility of the underlying assets. The analysis was then extended to consider other candidate risk management strategies, in particular some derivative trades. Having gained an understanding of the nature of the with-profits guarantees using the valuation model, we were then able to identify suitable candidate (plain vanilla) derivative positions for hedging the guarantee. Over a one-year period, we were able to find derivative positions that did a comparable job of hedging the market risk of the guarantee, whilst potentially reducing the exposure to changes in option-

implied volatility. Finally, as we demonstrated, as well as reducing profit volatility and its exposure to asset movements, implementation of these types of risk management strategies could also have significant implications for the size of an office's risk-based capital. However, the analysis should have the attaching caveats that the derivative positions identified may be expensive and/or incur substantial transaction costs in trading. Further, such positions are likely to require annual revision, as asset values move and with-profits investment policies change.

6.5 Nonetheless, we have developed powerful techniques for evaluating the effectiveness of hedging approaches for with-profits guarantees. As the regulatory and accounting regimes move inexorably towards a market-consistent framework, such arrangements will surely be worthy of life offices' consideration. After all, effective hedging will help to manage profit volatility, solvency risk and capital requirements. The next challenge is for the life industry and the actuarial profession to embrace fully this more rigorous, market-based approach to financial risk management. We hope that this paper can make a contribution to this process.

REFERENCES

- BAXTER, M. & RENNIE, A. (1996). *Financial calculus*. Cambridge University Press.
- BLACK, F. & SCHOLES, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, **81**, 637-659.
- BOYLE, P., BROADIE, M. & GLASSERMAN, P. (1997). Monte-Carlo methods for security pricing. *Journal of Economic Dynamics and Control*, **21**, 1267-1321.
- KEMP, M.H.D. (1997). Actuaries and Derivatives. *B.A.J.* **3**, 51-180 and 381-409.
- HARRISON, M. & PLISKA, S. (1981). Stochastic processes and their applications. **11**, 215-260.
- HULL, J. & WHITE, A. (1994). Numerical procedures for implementing term structure models. *Journal of Derivatives*, **2**, no 2, 37-48.
- HULL, J. & WHITE, A. (2000). Forward rate volatilities, swap rate volatilities, and implementation of the LIBOR market model. *Journal of Fixed Income*, **10**, no 2, 46-62.
- MERTON, R. (1973). Theory of rational option pricing. *Bell Journal of Economics and Management Science*, **4**, 141-183.
- WILKIE, A.D. (1987). An option pricing approach to bonus policy. *J.I.A.* **114**, 21-90.
- WILMOTT, P. (1998). *Derivatives*. John Wiley & Sons.

APPENDIX A

AN INTRODUCTION TO DYNAMIC HEDGING

A.1 *Introduction*

The pricing of guarantees is not a new idea. The major breakthrough in the valuation of options was made in the early 1970s with the seminal work of Black, Scholes and Merton (BSM). They showed how options could be valued and how a replicating ('hedge') portfolio of assets could be found which (if correctly adjusted over time) produced the same contingent payoffs as the option. Academic research in the area has continued ceaselessly for the last three decades. Practical implementation (primarily by investment banks) has developed contemporaneously. This section introduces some of the basic principles that are applied in this area. In particular, we consider how a basic cash guarantee on a portfolio invested in a risky asset can be valued, and the insights that the analysis gives into the composition of the hedge portfolio. The cost of the guarantee is determined by finding the portfolio of assets that will replicate the guarantee pay off. Under the replicating cost argument, the value of the hedge portfolio and the cost of the guarantee are one and the same thing.

A.2 *The Black-Scholes-Merton Analysis*

A.2.1 The classic Black-Scholes-Merton analysis showed that, under a suitable set of assumptions, option payoffs could be replicated by managing a dynamically rebalanced portfolio of the underlying asset and the risk-free asset. This insight is central to the pricing and hedging of derivatives.

A.2.2 Using this approach, Black & Scholes found the value of a European put option on a stock to be:

$$\text{PutOption}(t) = N(-d_2)Xe^{-r(T-t)} - N(-d_1)S_t$$

where:

$$d_1 = \frac{\ln(S_t/X) + (r + \sigma^2/2)(T - t)}{\sigma\sqrt{T - t}}$$

$$d_2 = d_1 - \sigma\sqrt{T - t}$$

where:

$N(\cdot)$ is the cumulative standard Normal density function

S_t is the price of the underlying asset at time t

X is the strike price of the option

T is the maturity date

t is the current date
 r is the continuously-compounded risk-free interest rate (net of dividends)
 σ is the annualised volatility of the underlying asset.

A.2.3 The most significant assumptions in the derivation of this formula are discussed in Section A.3. The above formula, as well as providing a means for pricing guarantees and other options, also provides an insight into how changes in the value of the option can be hedged. Look at the formula. It contains two terms. The first is a term in $e^{-r(T-t)}$ i.e. a quantity of risk-free asset. The second is a term in S_t i.e. a quantity of the underlying risky asset. If we hold a portfolio consisting of $-N(-d_1)$ of the underlying asset and hold $N(-d_2)Xe^{-r(T-t)}$ in cash, then the changes in the option price will be matched by changes in this portfolio. This special portfolio is the hedge portfolio, and the proportion of the underlying asset held in the hedge portfolio, $-N(-d_1)$, is referred to as the ‘delta’ of the option (i.e. delta measures the sensitivity of the option value to changes in the underlying asset price). Dynamically rebalancing the hedge portfolio (to keep the asset exposure in line with the option’s changing delta) is often referred to as delta hedging.

A.2.4 So long as the appropriate hedge portfolio is held throughout the life of an option (i.e. delta of the underlying asset, with the remainder in cash) then — providing BSM’s other assumptions hold — the value of the portfolio at the maturity of the option will be equal to the option payoff. For a life office hedging, say a unit-linked maturity guarantee, changes in the value of the hedge portfolio will exactly offset changes in the fair/realistic value of the guarantee. The office will have removed the fair value profit volatility that arises from writing the guarantee.

A.2.5 Note that the appropriate holdings of the underlying asset and cash will change continuously over time. In order to perfectly replicate the option payoff, it is necessary to continuously rebalance the hedge portfolio. Common sense tells us that this is not feasible. The practical challenges that must be faced in any real world implementation of a dynamic hedging strategy are discussed below.

A.2.6 These insights can be used to create a rule for the management of the hedge portfolio, as follows:

At time t , make a guarantee charge of PutOption (t), and establish the hedge portfolio by holding $-N(-d_1)$ of the underlying asset, and $N(-d_2)Xe^{-r(T-t)}$ of cash.

At time $t + \sigma t$, rebalance the hedge portfolio to reflect the changes in $N(-d_1)$ and $N(-d_2)$ which will result from changes in the underlying asset value and the term to maturity.

Repeat the second step for each subsequent time period of length σt until the option reaches maturity.

A.2.7 Let us illustrate these ideas by considering two example Monte-

Carlo simulations for the risky asset and hedge portfolio. In these examples, let us assume that we are aiming to hedge a ten-year money back guarantee (i.e. a plain vanilla put option) on a unit-linked fund. Rebalancing of the hedge portfolio will take place on a daily basis. For both simulations we start with a fund of 1 unit of currency, and there is a guarantee of 1 after ten years. Under our assumptions, this implies a put option price of around 0.075 and an initial delta of around -0.18 . The initial composition of the hedge portfolio is therefore -0.18×1 of the underlying risky fund, and $0.075 + 0.18 = 0.255$ in the risk-free asset.

A.2.8 In the example projection illustrated in Figure A.1, the underlying asset experience has delivered strong growth over the first four years of the contract's term. This is then followed by long-term poor performance, with the asset value more than halving between its peak and the maturity of the contract. Now consider the behaviour of the delta over this simulation. After four years of the contract's term, the delta has fallen (in magnitude) from its initial value of -0.18 to -0.05 as a result of the fund's strong performance. However, the poor performance of the fund over the second half of its term results in the guarantee moving from being very undemanding to being very valuable. This drives the delta to a higher magnitude, increasing the exposure to the underlying asset as it becomes more likely that the option will mature in-the-money. This increased (negative) exposure to

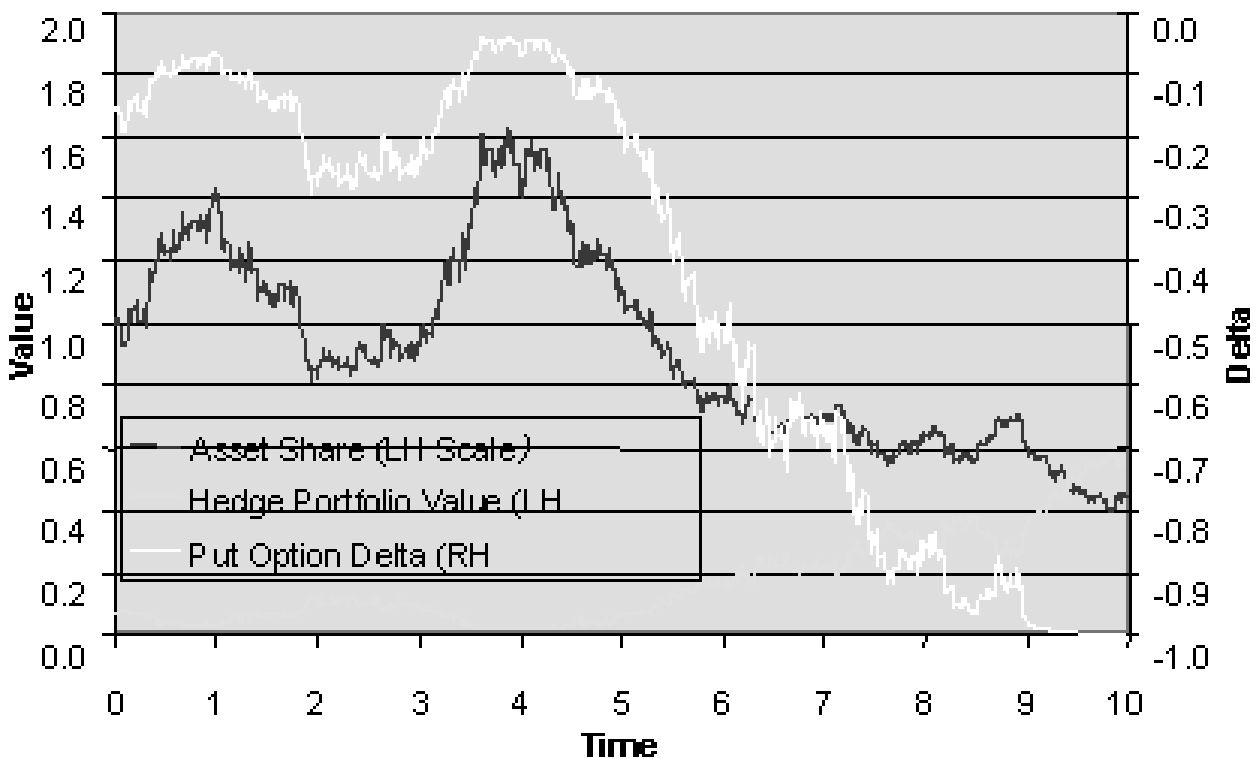


Figure A.1. Example Monte-Carlo projection (1)

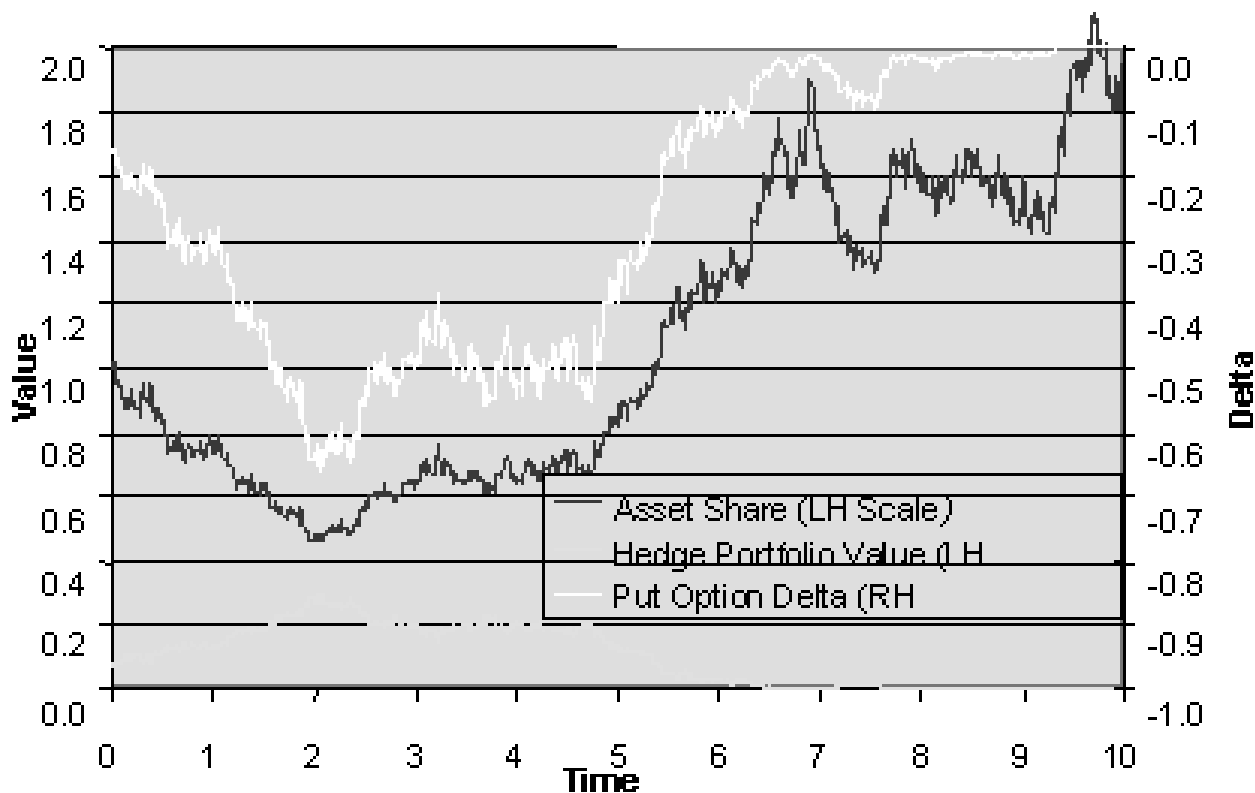


Figure A.2. Example Monte-Carlo projection (2)

the poorly performing underlying asset ensures that the hedge portfolio appreciates strongly in value over the final years of the contract. The value of the hedge portfolio increases from close to zero after 4 years to 0.58 at maturity. Notice that — since the final fund value is 0.42 — the hedge portfolio value neatly offsets the shortfall and ensures that the guarantee can be matched.

A.2.9 In the second example projection, the underlying fund performs poorly over the first couple of years of the contract's life. This results in the delta of the hedge portfolio increasing in magnitude from -0.18 to -0.64 . However, very strong asset performance over the remaining term of the contract means that the guarantee expires worthless. The delta of the option falls to virtually zero as it becomes less likely that the option will expire in-the-money, and the hedge portfolio's value is wiped out by strong underlying asset growth (which the hedge portfolio has negative exposure to). Again, we have a net position of no profit or loss. The option written by the product provider has been matched.

A.3 A Contrast with the Funding Approach

A.3.1 Dynamic hedging is one possible way of using an option premium. It is interesting to review the conventional (life office) approach to the management of guarantees by extending the analysis. We will use Monte-Carlo projections to compare the impact of different hedging

Table A.1. Assumed option parameters

Parameter	Value
Strike price	1.0
Underlying asset	1.0
Risk-free interest rate	5%
Dividend yield	0%
Volatility	20%
AMC	1%

Table A.2. Assumed asset model parameters

Parameter	Value
EBR	80%
Risk-free interest rate	5%
Equity risk premium	3.5%
Long bond risk premium	1%
Equity volatility	24%
Long bond volatility	11%
Equity/bond correlation	0.3%
Transaction costs	0.00%

strategies on the profit/loss arising at the maturity of the contract (i.e. the difference between the final hedge portfolio value and the guarantee cost. Tables A.1 and A.2 list the parameters assumed in this analysis.

A.3.2 The above parameters imply a portfolio volatility of 20%. We assume that portfolio returns are lognormally distributed.

A.3.3 In order to assess the adequacy of the hedge, it is necessary to consider the various strategies open to the hedger in managing risk. One approach would be to set aside the option-based hedging strategy and simply invest the charge made for the guarantee in the underlying asset. (It could be argued that this has historically been the typical life office's strategy, where the free estate is invested in the same asset mix as the underlying asset shares). Figure A.3 shows the profit/loss profile faced by the hedger as a function of the underlying asset value at the maturity date of the policy.

A.3.4 Figure A.3 shows that this is an inherently risky strategy. In this case, the hedger has a geared exposure to the underlying asset — the hedge portfolio simply moves in line with underlying asset values. This means that the net loss increases at a faster rate than the fall in the underlying asset. Writing guarantees under this strategy means accepting a geared position on the future performance of the underlying asset.

A.3.5 Now suppose that the guarantee writer chooses to invest the guarantee charge in the dynamic hedge portfolio discussed above. Further, suppose that the hedge portfolio is rebalanced daily as in the example

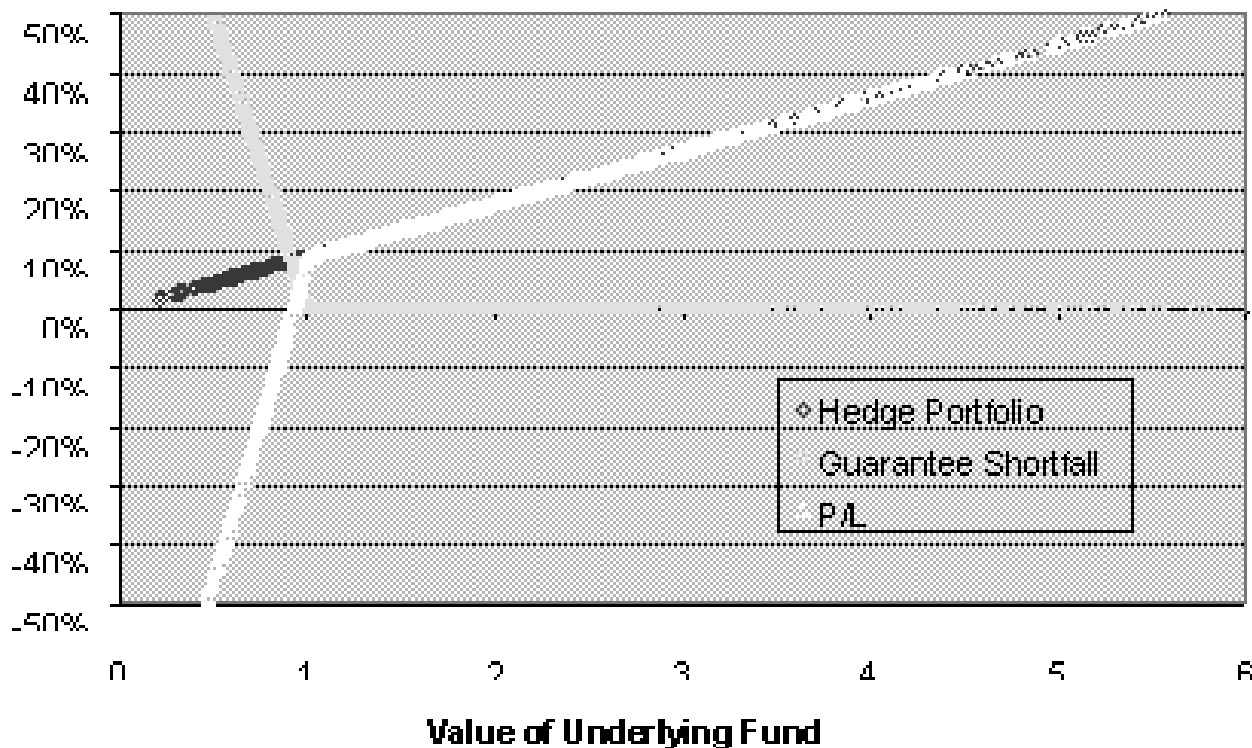


Figure A.3. Profit and loss at maturity with guarantee charge invested in underlying fund

projections shown in Figures A.1 and A.2. Figure A.4 shows how the hedge portfolio compares with the guarantee shortfall at maturity plotted as a function of the final underlying asset value.

A.3.6 You can see that daily rebalancing results in the hedge portfolio providing a very close match to the actual guarantee shortfall (under the assumptions described in Tables A.1 and A.2 and the assumption of a lognormally distributed underlying asset return).

A.3.7 You may recall that the above analysis has ignored the transaction costs that would be incurred in rebalancing a real world hedge portfolio. The existence of transaction costs will naturally constrain the frequency with which the hedger will wish to rebalance.

A.3.8 Figure A.5 shows how the accuracy of the hedge deteriorates as the period between rebalancing is increased from daily to monthly. (It should be noted that when the rebalancing frequency of the hedge portfolio is changed, a small adjustment to the delta should be made in order to maximise the accuracy of the hedge — see Wilmott (1998)). It can be seen that monthly rebalancing still creates a fairly accurate hedge for the guarantee shortfall by comparison with the static strategy shown in Figure A.3.

A.3.9 Figure A.6 plots the distribution of the maturity profit/loss under various rebalancing frequencies, as well as the strategy of investing the

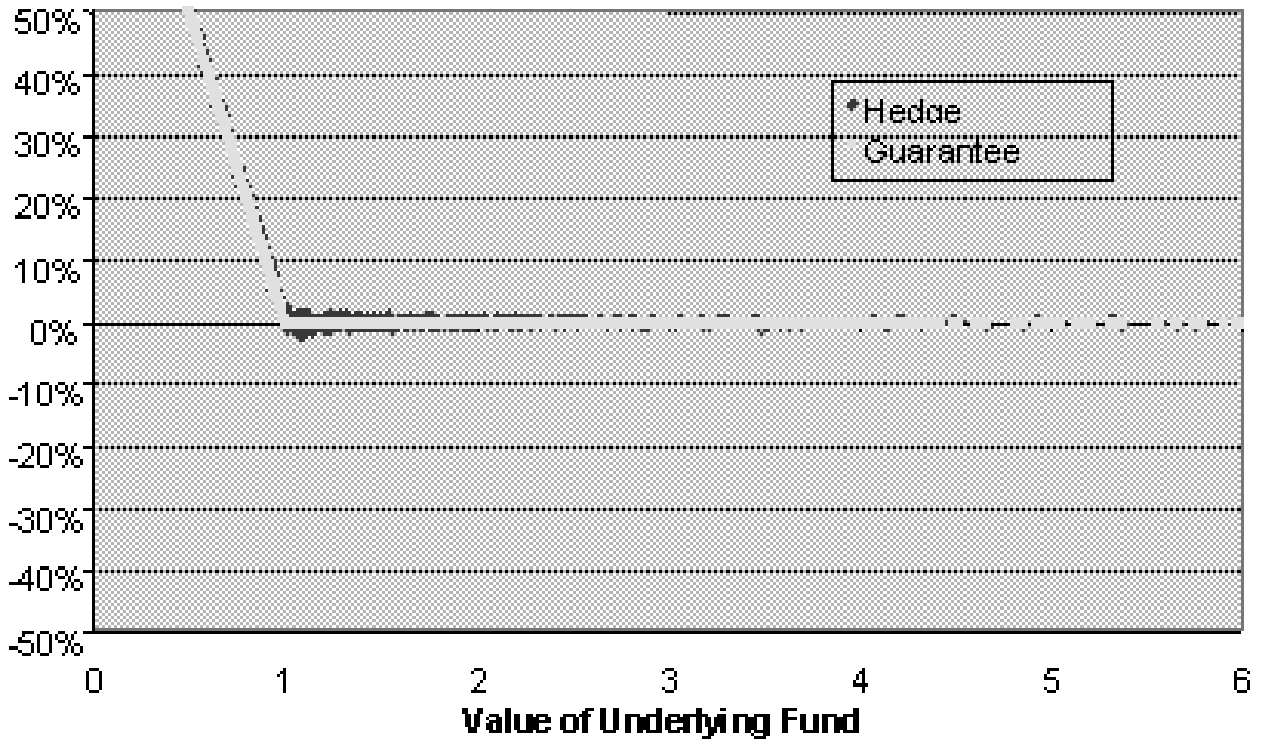


Figure A.4. Profit and loss at maturity with daily rebalancing of the hedge portfolio

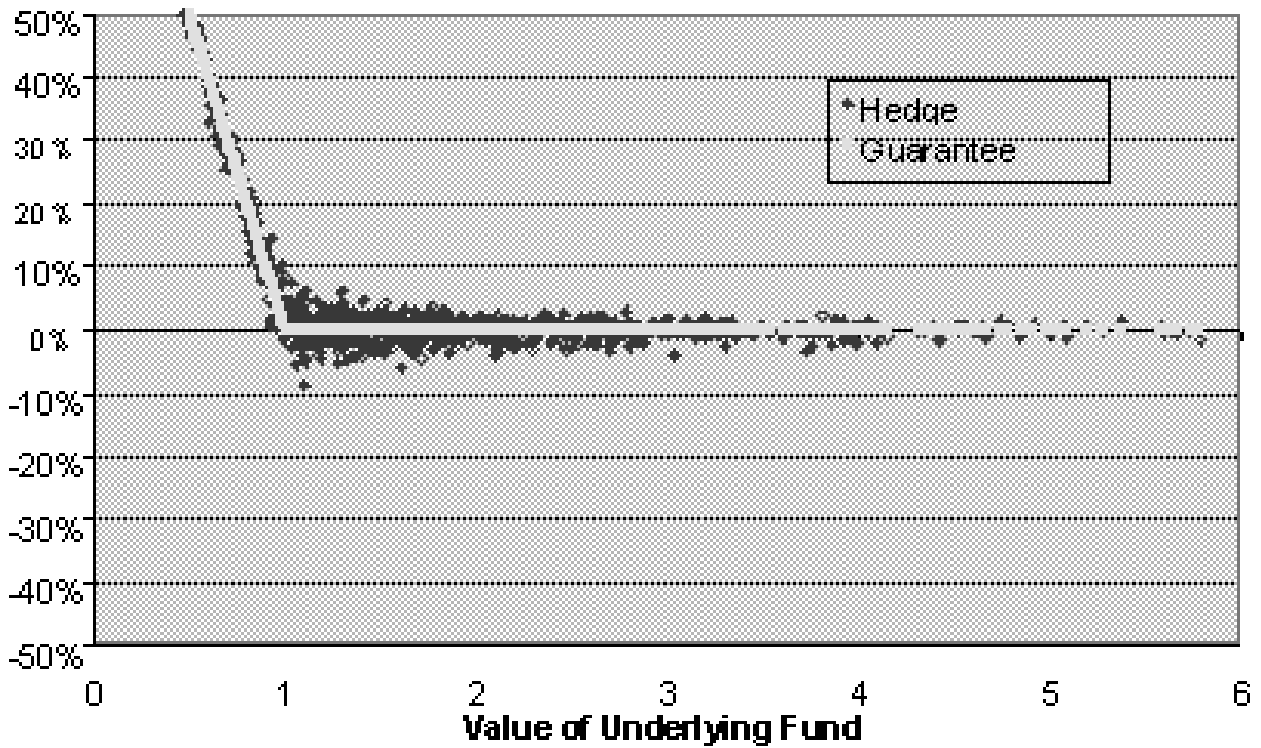


Figure A.5. Profit/loss at maturity, monthly rebalancing of hedge portfolio

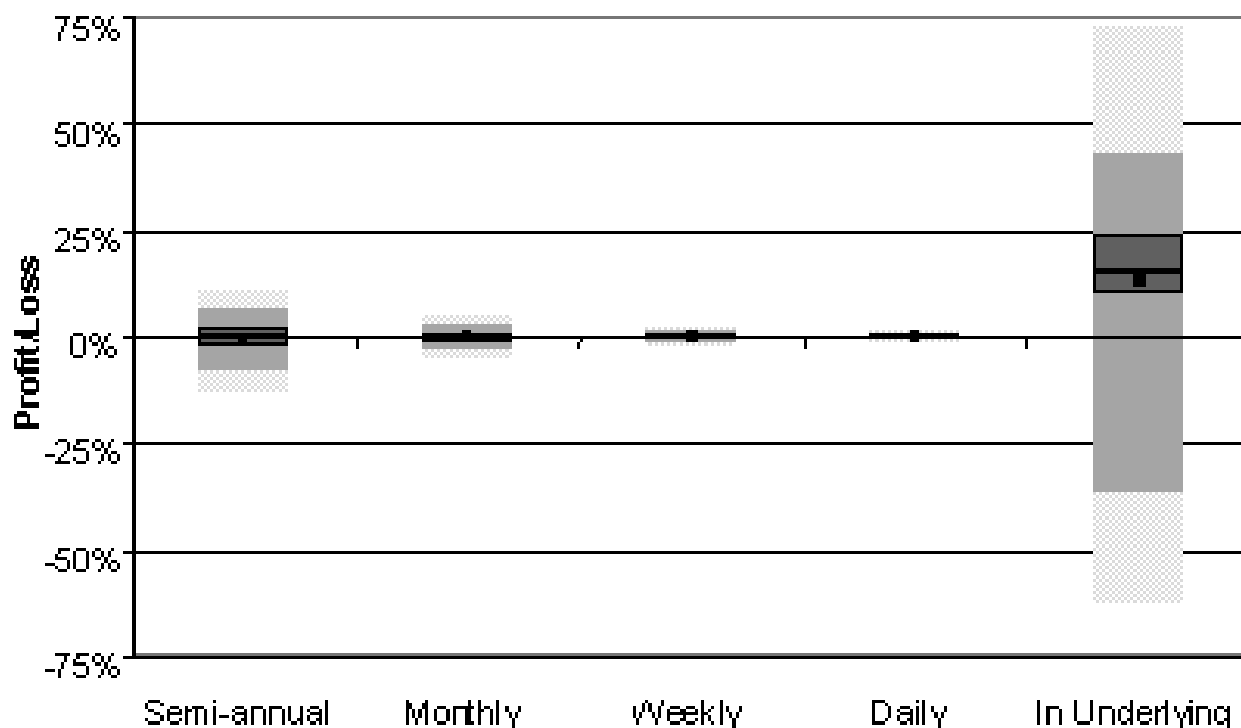


Figure A.6. Comparison of profit/loss distributions

guarantee charge in the underlying fund. In this figure, the darkest segments represent the two central quartiles of the distributions, and the solid box signifies the mean. The next segments out represent the 5th to 25th and 75th to 95th percentiles, whilst the outer segments show the 1st to 5th and 95th and 99th percentiles.

A.3.10 As might be expected, increasing the rebalancing frequency increases the accuracy of the hedge, and therefore reduces the volatility of the profit/loss arising. Investing the guarantee charge in the underlying asset results in a positive expected payoff, although this can only be achieved through exposure to significant downside risk. It is simply the profit/loss profile of a geared equity investment.

A.4 *Some Caveats*

A.4.1 Section A.1 explored how an option can be dynamically hedged in a world that conforms to the assumptions made in the Black-Scholes analysis. Whilst this can give valuable insights into how to hedge financial guarantees, it is important to be aware that some of these assumptions will not hold in reality.

A.4.2 The most significant assumptions in the derivation of the above formula are:

- The underlying asset is lognormally distributed. (This is just a continuous time version of a random walk. It implies that the underlying asset is lognormally distributed, i.e. continuously compounded returns

are normally distributed.) This assumption itself is not very important — the insights from Black-Scholes-Merton allow us to apply their methodology to other stochastic processes for the stock. However, the assumption that the underlying asset follows a continuous process is theoretically necessary in order to construct the replicating portfolio. Gaps in the underlying asset price cannot be dynamically hedged.

- The volatility of the underlying asset is constant. This assumption can be relaxed to allow for deterministic changes in volatility over time. However, the replicating portfolio cannot be extended to cover stochastic volatility — unexpected changes in the volatility of the underlying asset cannot be hedged by a portfolio of the underlying asset and cash.
- Markets are frictionless. In other words, the replicating portfolio can be rebalanced continuously without incurring any transaction costs. This is fundamental to the BSM pricing formula. In practice, hedgers do face significant transaction costs. (It turns out that simple adjustments can be made to the formula that captures the impact of the transaction costs of rebalancing on the replicating cost. The Black-Scholes formula can be adjusted for transaction costs. See Wilmott, 1998.)

A.4.3 Each of these assumptions can have a significant impact on the effectiveness of dynamic hedging. Real world hedging will therefore demand the management of a set of trade-offs between risk and cost, and even the most rigorous dynamic hedging process will still inevitably leave some residual risk.

APPENDIX B

THE TWO-FACTOR BLACK-KARASINSKI MODEL

B.1 In the two-factor Black-Karasinski model used in this paper, the risk-neutral process for the short rate r is:

$$d\ln(r) = a_1[\ln(r_2(r)) - \ln(r)]dt + \sigma_1 dZ_1$$

where u follows the stochastic process:

$$d\ln(r_2(t)) = a_2[\mu - \ln(r_2)]dt + \sigma_2 dZ_2.$$

Here, Z_1, Z_2 are assumed to be independent Brownian motions.

B.2 This popular model has the following desirable characteristics:

- It has two underlying stochastic factors, thus allowing for more realistic yield-curve dynamics than a simple one-factor model.
- The short rate is mean reverting, as observed in historical time series.
- It has nice analytic properties, in the sense that $\ln(r)$ is normally distributed, and we can write down analytic formulae for its mean and variance.
- Interest rates can never become negative.

B.3 The parameters used in the analysis are given in Table B.1.

Table B.1. Parameters in the lognormal model

Parameter	Value
α_1	0.094
α_2	0.059
σ_1	0.1836
σ_2	0.3865
μ	0.0287
r_0	0.0489
$r_2(0)$	0.0571