MARKET CONSISTENT VALUATION OF LIFE ASSURANCE BUSINESS

By T. J. Sheldon and A. D. Smith

[Presented to the Institute of Actuaries, 23 February 2004]

ABSTRACT

In recent years there has been a trend towards market consistent valuation in those institutions for which actuaries have responsibilities. The larger United Kingdom with-profits insurance companies are now preparing realistic balance sheets, both for internal purposes and also at the request of the Financial Services Authority. International accounting standards have been moving to a fair value approach. Pension fund accounting under FRS 17 has also moved in this direction.

In this paper we examine the reasons for the adoption of market consistent valuation and discuss some of the commercial implications and corporate valuation. We consider the methods and assumptions which can be used to develop market consistent valuations of cash flows typically encountered in the liabilities of financial institutions, together with some of the problems inherent in the calibration of models used for the valuation of these cash flows. The volatility assumption is crucial to the valuation of options and guarantees, and we discuss the relationship between historical and implied volatility.

While most insurance companies initially adopted formulae to value their with-profits guarantees, several offices are now using a Monte Carlo simulation approach for their realistic balance sheets. The Monte Carlo approach enables allowance to be made for management discretion in bonus and investment policy, as well as policyholder actions. However, in many cases it is possible to develop analytical formulae for cash flows approximating those payable under insurance contracts.

The valuation formulae have implications for the hedging of embedded guarantees. The authors discuss the construction of hedges for financial risks in with-profits funds, the separate perspectives of policyholders and shareholders, possible funds in which to hold hedging instruments, limitations of capital market hedging tools and the effect of taxation on hedge effectiveness.

KEYWORDS

Market Consistent Valuation; Realistic Balance Sheet; Guarantees; Options; Option Pricing; Black Scholes Formula; Arbitrage; Historical and Implied Volatility; Volatility Smile; Exotic Guarantees; Credit Risk; Calibration; Market Data; Taxation; Monte Carlo Simulation; Hedging

CONTACT ADDRESS

T. J. Sheldon, A. D. Smith, Deloitte, Horizon House, 28, Upper High Street, Epsom, Surrey KT17 4RS, U.K. E-mail: tsheldon@deloitte.co.uk, andrewdsmith8@deloitte.co.uk

1. Introduction

1.1 *Market Consistency*

1.1.1 The term 'market consistent' has become increasingly popular as a description of liability valuations or, more generally, of cash flow valuations. However, there is no widely accepted definition of the term. So, we start with some definitions.

1.1.2 A *valuation algorithm* is a method for converting projected cash flows into a present value. A valuation algorithm may be specified with reference to a set of *calibration assets*. We say a valuation is *market consistent* if it replicates the market prices of the calibration assets to within an acceptable tolerance.

1.1.3 The purpose of a valuation algorithm is to place values on other sets of cash flows, different from the calibration assets. For example, we might look for market consistent valuations of policy liabilities. Alternatively, we might try to decompose the market value of a group share price into the value of each constituent part.

1.1.4 Market consistent valuation can take many forms. Two different models could produce different liability valuations and yet still both be market consistent. Several questions arise. There are four chief decisions:

(1) Which assets are to be used for calibration?
(2) How are model assumptions derived where insufficient market data exist?
(3) Which cash flows are to be valued? For example, what dividend tax credits (if any) are reflected in the market value of an equity share; to what extent are manufacturer expenses and capital costs reflected in option prices?
(4) Algorithms — when are deterministic projections, closed form solutions, numerical integration or Monte Carlo simulations to be used?

We discuss these decisions in the paper.

1.2 *Layout of Paper*

1.2.1 We start in Section 2 with the reasons for market consistent valuation, and identify three motivations: a better understanding of a company's share price; hedging or closing out risks; and the need for objectivity. The idea of hedging and closing out risks is related to the use of market consistent valuation for solvency measurement and the introduction of realistic balance sheets into the United Kingdom regulatory regime for life insurance companies. We also examine the commercial implications of a market consistent framework.

1.2.2 In Section 3 we consider the valuation of fixed cash flows and the determination of the risk free rate of interest. We debate the relative merits of gilt and swap rates, and the impact of credit risk in market consistent valuation. In Section 4 we look at the valuation of simple guarantees by the application of arbitrage to derive possible formulae to value such guarantees, including the well known Black Scholes option pricing formula. We illustrate how market data can be used to determine suitable parameters for the formula and how allowance can be made for life office tax. In Section 5 we discuss the application of more complex formulae to the cost of guarantees in with-profits business, making some allowance for management discretion in setting bonus rates. We also consider regular premium contracts, guaranteed annuity options, allowance for management discretion in investment policy and the impact of interest risk on equity options.

1.2.3 In Section 6 we discuss the relationship between historical and implied volatility, and the problems with the available data on volatility. In Section 7 we turn to the role of market consistency in corporate valuation and profit testing, and the impact of corporate structure on valuation. Section 8 contrasts the method of Monte Carlo simulation with the formula approach, and considers some of the problems of projecting realistic balance sheets of life offices, where it is necessary to evaluate guarantees in these balance sheets on a market consistent basis. In Section 9 we consider the potential for hedging financial risks in a with-profits fund. Our conclusions are presented in Section 10.

1.2.4 While much of the paper is written in the context of market consistent valuation of life assurance liabilities, many of the issues discussed are also relevant to the valuation of cash flows in other financial institutions.

2. Why Market Consistent Values are Useful

2.1 *Three Motivations*

2.1.1 There are different approaches to calculating market consistent valuations; some of these differences may be explained by the different motivations for seeking market consistent valuations. At least three motivations are commonly advanced for examining a firm's operations on a market consistent basis.

2.1.2 We consider the following three motivations for market consistent valuations:

(1) understanding the behaviour of a company's share price;
(2) measurement of solvency relative to a buy-out standard; and
(3) producing comparable valuations which reduce the need for subjective judgement.

2.1.3 We refer the reader to Altmann & Vanderhoof (1998), Hare *et al.* (1999), Hairs *et al.* (2001), Jarvis *et al.* (2001), Abbink & Saker (2002), Blight *et al.* (2003), Dullaway & Needleman (2004), Kipling & Moran (2003), Muir & Waller (2003) for the by now well-rehearsed list of influences leading insurers towards market consistent valuations of assets and liabilities, and for the various uses to which these valuations may be put.

2.2 *Understanding a Company's Share Price*

2.2.1 The first, and perhaps most persuasive, motivation for the use of market consistent valuation is to improve understanding of a firm's own share price. In the past, many companies may have despaired of ever understanding the drivers of their share price. They may have argued that prices were distorted by supply and demand. Fortunately, supply and demand are what modern asset theory is all about. It is when prices are not set by supply and demand that economic theory falters. The motivation, then, is to use an understanding of market price formation to decompose a firm's value into the contributions from different sources. This could include a breakdown by lines of business, or a breakdown between current assets and liabilities and an explanation of franchise value. The economic setting is one of extrapolation, from a known share price to the value of unobservable constituents.

2.2.2 Managers of proprietary companies are appointed to serve the interests of shareholders. The interest of shareholders is clear — they want to see an improvement in their share price. This could be measured in absolute terms or relative to a benchmark. Managers of insurance companies, therefore, seek methodologies which explain how their strategies and actions relate to the market's valuation of an insurance enterprise.

2.2.3 Customer service and compliance with regulatory constraints remain relevant. Managers must understand the demands of many stakeholders. The share price, if only we could understand it, gives a means of expressing the value of each stakeholder's interest in a company, and a metric for trading off their various requirements.

2.2.4 Anecdotal evidence suggests that analysts, too, are influenced by market consistent valuations where these are available. For example analysts have been keen to establish the guarantee costs in life insurers' liabilities, and also the FRS 17 pension liabilities.

2.3 *Closing Out Risks*

2.3.1 The second motivation for examining market consistent values is the notion of hedging or closing out of risks. If an entity is solvent as measured on a market consistent basis, then it might, in theory, be able to transact at those prices and close out its assets and liabilities. Therefore, ensuring market consistent solvency offers some protection to policyholders or creditors.

2.3.2 The realistic balance sheet devised by the Financial Services Authority (FSA) (and described under the realistic peak for with-profits offices in CP195) is designed to capture the cost of guarantees and smoothing on a market consistent basis. Why is this market consistent approach relevant to the measurement and management of non-traded cash flows, such as insurance companies' liabilities? Why not simply set regulatory capital requirements on a run off basis, demanding a sufficiently high degree of confidence that a company can meet its liabilities (including the reasonable expectations of policyholders) as they arise?

2.3.3 The introduction of the realistic balance sheet is, in part, a response to the difficulties that un-hedged guarantees have caused the life industry in recent years. Reliance on long-term solvency tests runs the risk that we overlook more imminent problems, compounded by the use of over-optimistic assumptions and models used to determine capital needs. While there has been some criticism of the transfer of banking techniques to life assurance, the rate of deterioration in life offices' finances over the last three years demands a greater focus on the short term.

2.3.4 The solvency motivation may lead to a definition of market value different to that implied by the shareholder perspective. The differences are due to credit risks and to margins that third parties might require to reflect their institutional costs of bearing risks, and possibly to reflect any additional risks arising as part of the hypothetical transfer. For example, the amount theoretically payable to a third party to take on the guarantees in a with-profits fund might exceed the cost to current shareholders of bearing those risks.

2.3.5 Furthermore, the quotes for closing out may fail at precisely the moment when they are most needed. In normal times, insurers and reinsurers may buy and sell mortality risk on terms which can be observed. However, it is possible that a mortality shock could cause widespread financial distress in the insurance and reinsurance industries. It is in these distress situations where solvency rules based on buyout terms face their harshest tests.

2.3.6 We do not see this as an argument against the use of market consistent values. Rather, it suggests that market consistent valuation has a contribution to make to ensuring financial soundness. Hedging risks is an option available to management, at least for some of the risks run by life offices. Market consistent values capture the cost of this option. Risk management and the quantification of risk based capital need to reflect the variations in costs and the likelihood of significant increases in costs during conditions of financial distress.

2.4 *Objectivity*

2.4.1 The final motivation for market consistent valuations is their relative objectivity. The contrast is made with historic cost accounting, under which firms could alter reported profits by buying and selling investments, or even by notional reallocation of investments between traded and non-traded categories.

2.4.2 Quantitative analysts frequently refer to the 'discipline' of market consistency. The resulting valuations are believed to be more objective and, in general, less subject to manipulation than other valuation methods. As a result, the 'market consistent' label has acquired some kudos. This does not mean that market manipulation or accounting misstatements are impossible.

2.4.3 As an illustration of the kudos effect, consider the British Telecom pension scheme. After some controversy in the *Financial Times*, the Presidents of the Institute and Faculty of Actuaries wrote [*Financial Times*, 16 February 2003]:

> "There are a variety of methods used for funding and for pension cost accounting, and indeed under SSAP 24 many methods are acceptable. Our understanding is that the pension costs reported by BT under SSAP 24 are based on a market value approach and have not used any method that could be the subject of criticism ... "

2.4.4 At 31 March 2003, British Telecom disclosed pre-tax deficits on two 'market' bases; a SSAP 24 deficit of £1.4 bn and an FRS 17 deficit of £9.0bn, relative to assets with a market value of £21.5bn. While this is an extreme example, it highlights how different assumptions can change the results, even when apparently constrained by the 'market' label.

2.4.5 In many ways, the actuarial profession might be considered a natural arbiter of whether a valuation is market consistent or not. Many actuaries believe that evaluation of market consistent models falls on actuarial turf. A few well publicised cases of objective and insightful market consistent valuations by actuaries would help their cause.

2.4.6 This has implications for those seeking to establish a concept of 'market consistent embedded value'. With an unregulated concept, different actuaries could produce very different numbers, perhaps as different as the two BT deficit valuations, and this is already happening in some mergers and acquisitions. The notion of a market consistent valuation is likely to carry more weight if such deviations can be avoided, for example by robust professional guidance.

2.4.7 In the banking world, the introduction of market consistent models immediately generated convergence of reported valuations. This has encouraged regulators and accounting standard setters to apply those same techniques in insurance. However, the convergence has obstinately refused to materialise. Some possible reasons for this are as listed below.

2.4.7.1 Insurance risks contain many elements that are difficult to hedge, such as take-up rates of embedded options and guarantees. Therefore, even if two actuaries agree on market prices of traded instruments, they may differ on a market consistent value, either because they analyse the historical experience differently, or because they hold different views of future experience.

2.4.7.2 While, in the banking world modern finance has displaced all prior methods of valuation, the actuarial world has sought to be more inclusive, perhaps in order to postpone a feared confrontation between traditionalists and financial economists. For example, the SSAP 24 valuation of the BT fund effectively discounts bond-like liabilities using an expected earned rate on risky assets, a traditional actuarial technique which flies in the face of arbitrage-free pricing.

2.4.7.3 The options markets grew up with the Black-Scholes formula in place. Therefore, most options were priced and have been hedged, using modern financial tools. In comparison, many actuarial guarantees were unpriced and unhedged at the time of issue, with financial rectitude only imposed subsequently. This creates transitional issues which are unprecedented in the banking world. There is an understandable temptation to avoid or substantially modify modelling approaches which reveal sudden unfunded liabilities.

2.4.7.4 The valuation approach is most significant when there is a mismatch between assets and liabilities. Banks have traditionally hedged more closely than insurers. As a result, a tweak to a valuation model may dramatically change an insurer's stated net assets, while having little bottom line effect on bank balance sheets. Insurers spend correspondingly greater effort in the search for favourable tweaks.

2.4.7.5 In the banking world, pricing and accounting are closely intertwined. The mark to market for existing products feeds straight into new product pricing in two-way markets, and to deposits under margin agreements. Furthermore, many of a bank's customers are professional investors able to exploit quickly any mispricing by a bank. In contrast, insurers' profit tests are still generally based on statutory valuation methods, rather than reflecting market consistent or fair value principles. Even if an accounting bias were to be reflected in new business prices, an insurer's retail customers may be less able to exploit the mispricing.

2.4.7.6 Banks typically have independent model validation departments who enforce consistency on different models throughout a bank. Risk management processes control the sensitivity to any one parameter. For example, a bank may define a standard tool for measuring historical correlations, and set sensitivity limits for each trader. To change one correlation assumption in his favour, a trader must force a revision in the whole methodology, a change which, as likely as not, will produce losses elsewhere. On the other hand, many insurers find themselves with a correlation position on equities compared to bonds by virtue of guaranteed annuity positions. The choice of this parameter could significantly affect the valuation of liabilities. As yet we are not aware that insurers have robust processes in place for the choice of such parameters, though that is likely to change as results become subject to audit.

2.4.8 The difference between banks and insurers may be overstated, to the extent that banks do not publicise their pricing errors and model inconsistencies.

2.5 *Commercial Implications*

2.5.1 Within a market consistent framework, there is still a range of possible answers. A number of factors may influence how an answer is selected within that range, including financial theory, and professional and regulatory guidance. In addition, commercial considerations may have an effect.

2.5.2 In the context of a proprietary office, managers have a duty to shareholders which includes efficient management of shareholder capital. Where some leeway is available on liability valuation, the higher liability numbers (that is, the stronger valuation bases) may be considered to tie up capital in avoidable margins. We might, therefore, expect management generally to show a preference for weaker liability bases.

2.5.3 This is noticeable, not only with valuations, but also with capital requirements where they are model based. A summary of major modelling developments in the last decade might include:

— better understanding of fat tailed distributions and extreme values;
— use of tail dependency and copulas to measure association of rare events;
— use of unbiased autocorrelation estimators, reducing the apparent mean reversion in capital markets;
— better algorithms for pricing American style options, including embedded policyholder surrender options; and
— integration of process risk and parameter risk, using Monte Carlo Markov Chain algorithms (Hardy, 2003).

2.5.4 It is no surprise that adoption of any, or all, of these tools would increase the stated risk of a business. Implicitly, risks ignored within a model are taken to be zero, so, as we get better at modelling risks, we capture more of them. This, in itself, is not a disincentive for the internal use of new techniques. The tools could still be valuable for the additional insight gained.

2.5.5 However, there is a strong disincentive to the use of new tools in reporting capital requirements to a regulator. Furthermore, some regulators reject model output that is not also used for internal risk management. A cost-benefit analysis of a new modelling approach would include development and implementation costs as well as business benefits, but, in many cases, the cost of additional capital to be put up as a result of the model dominates all other items in the evaluation. In such a competition in laxity, we would not put money on best practice as the likely winner.

2.5.6 At the industry level a different dynamic operates. Regulatory formulae tend to be recalibrated from time to time, and benchmarked against accounting capital ratios internationally. Regulators often have in mind a certain level of acceptable capital for an industry as a whole, and these indications are then re-expressed in terms of model output. It would not, therefore, be possible for the industry as a whole to reduce its need for capital by building favourable models. Instead, an insurer who convinces the regulator of its own lower (regulatory) capital requirements may indirectly increase competitors' regulatory capital requirements.

2.5.7 Furthermore, at the industry level there are other costs associated with different valuation bases. In a market inconsistent framework, an asset may be treated differently if it is a derivative, a bond, a holding in a special purpose vehicle or a reinsurance receivable. In the past, considerable ingenuity has been expended in optimising capital structures to obtain the most favourable regulatory treatment. Although these structures may benefit individual firms, the restructuring is a deadweight expense for the industry as a whole. To the extent that market consistent supervision eliminates these deadweight costs, the industry benefits.

2.5.8 Arbitrary valuation bases also have hidden costs because of their knock-on effects for how a business is run. For example, in a market inconsistent regulatory framework (such as the net premium valuation method coupled with European Union solvency margins), an insurance company wishing to hedge out its financial risks is faced with the choice of hedging either the regulatory balance sheet or its own internal realistic balance sheet. Other solvency regulations, such as the admissibility rules for derivatives held by U.K. insurance companies, further complicate the decision. This combination of rules could prevent the most effective hedging strategy being pursued. A move to a market consistent regulatory framework should mitigate some of these problems and make hedging decisions easier.

2.5.9 While a move to market consistent valuation removes some of the difficulties, problems still remain. A decision to hedge out certain risks has natural implications for the amount of risk-based capital that an insurer requires. Hedging removes some risks from the balance sheet (for example, interest rate risk associated with guaranteed annuity options), but other risks, such as counterparty exposure, are accepted. In future, U.K. insurance companies will be responsible for determining their own risk-based capital needs through the Internal Capital Assessment, though this will be open to challenge from the regulator through the mechanism of Individual Capital Guidance. While the new regulatory regime may make some of the old regulatory rules redundant, in the absence of clear rules or guidance, companies will seek methods and assumptions which reduce regulatory capital requirements. The result is a regime which may still not be consistent with effective risk management.

3. The Risk Free Rate

3.1 *Fixed Cash Flows*

3.1.1 We address firstly the market consistent valuation of fixed cash flows denominated in nominal currency terms.

3.1.2 The simplest valuation model for fixed cash flows is the flat yield curve model, where a constant rate of interest is used to discount all cash flows. In the past, actuaries may have taken the gross redemption yield on a particular gilt as the discount rate. In modern terminology, this would be expressed as calibrating a market consistent valuation model using a selected gilt as the calibration asset.

3.1.3 The flat yield curve approach may be acceptable if most gilts have similar yields. Indeed, this is an important general test of market consistent models. Having chosen some calibration assets, how well does the model price other assets that are not in the calibration set? In this case, the answer will depend on the shape of the yield curve, which, of course, varies from time to time.

3.1.4 If the yield curve does slope, then there are two possible ways to improve the situation. The traditional actuarial route has been to consider carefully the choice of calibration ingredients. For example, we might choose a gilt of similar term to the liabilities, if such a gilt exists.

3.1.5 The other alternative is to build a more complex model, which is capable of calibrating to a larger number of assets. In the case of fixed cash flows, the most obvious generalisation is to assume that the discount factor for a future cash flow depends (only) on the timing of the cash flow, but that this dependency need not decrease geometrically. Instead, we might try to calibrate a function $P(0, t)$, denoting the value at time 0 (the valuation date) of a future cash flow at time t. The quantity $P(0, t)$ is called a zero coupon price or risk free discount factor. Where typographically convenient we will write P_{0t} for $P(0, t)$.

3.1.6 There are many ways to calibrate such a function. All of them involve solving a set of equations where each bond in the calibration set is equated to the sum of the values of its future income and principal cash flows. This gives rise to a proliferation of yield concepts, including spot yields, par yields and forward yields. The papers by Griffiths *et al.* (1997), Feldman *et al.* (1998) and the book by Anderson *et al.* (1996) give several methods for calibrating risk free rates and explain the relationships between different yield concepts.

3.1.7 Variational calculus is a common approach for fitting yield curves. Such formulations tend to give rise to spline solutions. Smith & Wilson (2000) give algorithms for finding zero coupon prices to minimise the expression:

$$\frac{1}{2\alpha^3}\int_0^\infty \left(\frac{d^2}{dt^2}\left\{e^{ft}P(0,t)\right\}\right)^2 dt + \frac{1}{2\alpha}\int_0^\infty \left(\frac{d}{dt}\left\{e^{ft}P(0,t)\right\}\right)^2 dt.$$

3.1.8 f and α are parameters which must be selected by the user. Once these are chosen, it turns out that the optimal fitted yield curve has f as the limiting long forward rate.

Table 3.1. Bonds used in the calibration illustrated in Figure 3.1

Bond	Term (years)	Annual coupon	Redemption yield
Bond 1	2	5%	4%
Bond 2	5	5%	4.5%
Bond 3	10	5%	5%
Bond 4	25	5%	5.5%

3.1.9 If it so happens that all bonds have a common gross redemption yield f, then it is easy to see that the integrand is zero, and the flat yield curve is derived as the best fit. In other cases, we can think of the objective function as punishing deviations from a flat yield curve. The optimisation then provides the closest to a flat yield curve subject to capturing the prices of specified bonds.

3.1.10 The definition of the long forward rate as a mathematical limit makes it unobservable, yet the choice of this parameter is critical. The long forward rate may be hidden inside a chosen family of fitted yield curves, or may be chosen explicitly, as in our example.

3.1.11 To see the possible differences, we show three alternative spot curves, each calibrated to four bonds, as shown in Table 3.1.

3.1.12 We have chosen $\alpha = 0.1$ and long forward rates f of 4%, 6% or 8%. Figure 3.1 shows that many extrapolated curves are possible.

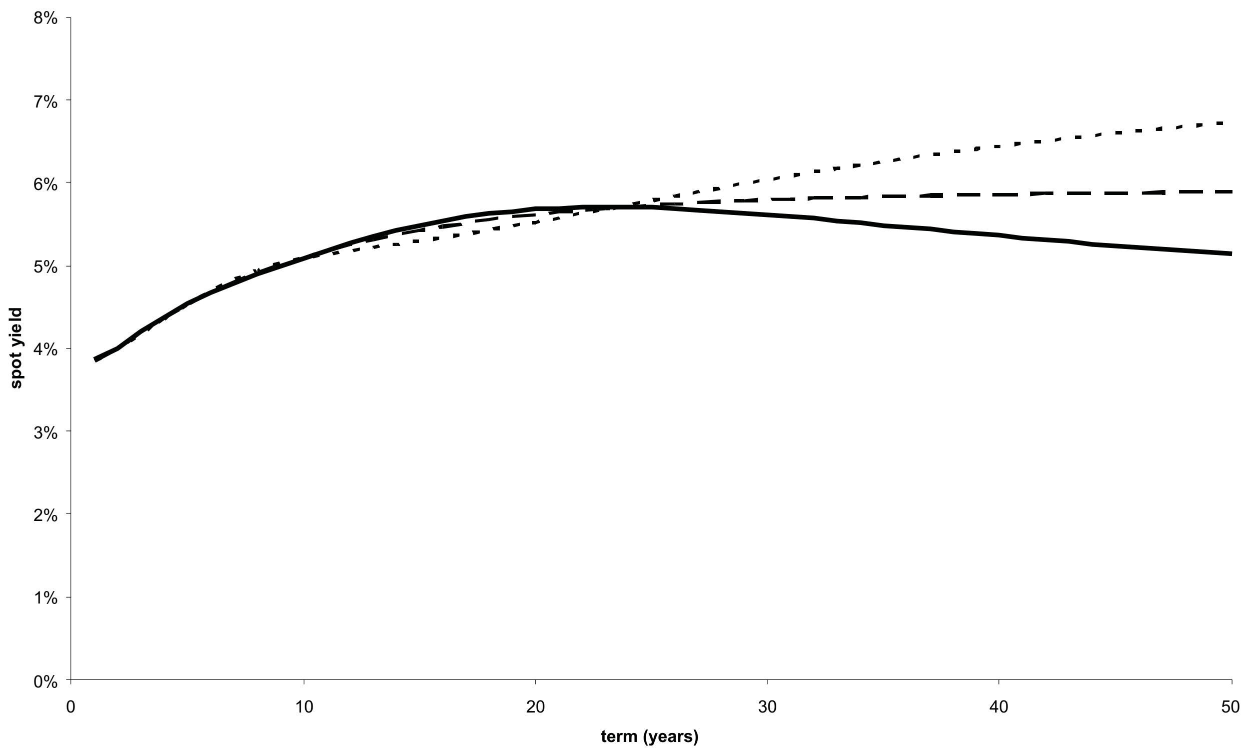

Figure 3.1. Alternative spot curves fitted minimising the formula in ¶3.1.7

3.1.13 The market consistent valuation of fixed cash flows contains a number of other judgemental inputs, in addition to the long forward rate. We now consider the treatment of credit risk, liquidity and tax.

3.2 *Interest Rate Swaps*

3.2.1 There is an extensive inter-bank market in interest rate products. This market is many times larger, and is usually more liquid than the gilt market.

3.2.2 The starting point for these interest rate markets is the fixed rate cash deposit. For example, in the U.K. many inter-bank deals are quoted relative to six-month LIBOR. This rate is the rate at which a bank will accept deposits for a six-month period. There is a small credit risk associated with these deposits, as a result of which six-month LIBOR carries a margin, often between 20 and 100 basis points, over a corresponding gilt yield.

3.2.3 An interest rate swap is a derivative where one party (the *payer*) pays a fixed rate at six-monthly intervals in exchange for six-month LIBOR. At the same time, the other party (the *receiver*) receives the fixed rate and pays a floating rate. This transaction is between two parties, neither of whom need have any underlying deposits. It is necessary for some third party to be making deposits so that LIBOR can be observed in order to settle the cash flows on the swap.

3.2.4 An investor with a series of six-month deposits can enter a receiver swap and turn the deposit interest into a series of fixed payments. This gives a way of constructing synthetic bonds, and yield curves can be derived from these synthetic bonds just as they can from gilts. In some ways the swap calculations are easier, because swap cash flows occur at regular six-month intervals, while gilt cash flows will occur on awkward dates, depending on the maturity date of the particular gilt.

3.2.5 As with any derivative, a swap contract contains some credit risk that a counterparty does not perform its obligations under the contract. This credit risk is mitigated by a number of measures, including the settlement of cash flows on a net basis, the use of master netting agreements and margining (see Kemp, 1997, for a description of margining). As a result, the credit risk on a swap itself is often regarded as very small.

3.2.6 An investor, seeking to replicate fixed cash flows using a deposit and a receiver swap, faces credit risk on the underlying deposit. Equivalently, an investor seeking a low risk variable interest deposit can best do this by holding gilts and a payer swap. For this reason, swap rates are typically higher than government bond yields, to reflect the additional credit risk. Swap yields are typically comparable to yields on AAA bonds, though the gap has widened in recent years, as show in Figure 3.2.

3.2.7 In some jurisdictions, swap rates may fall below government bond yields, particularly if the government credit risk is perceived to be greater than that of international banks.

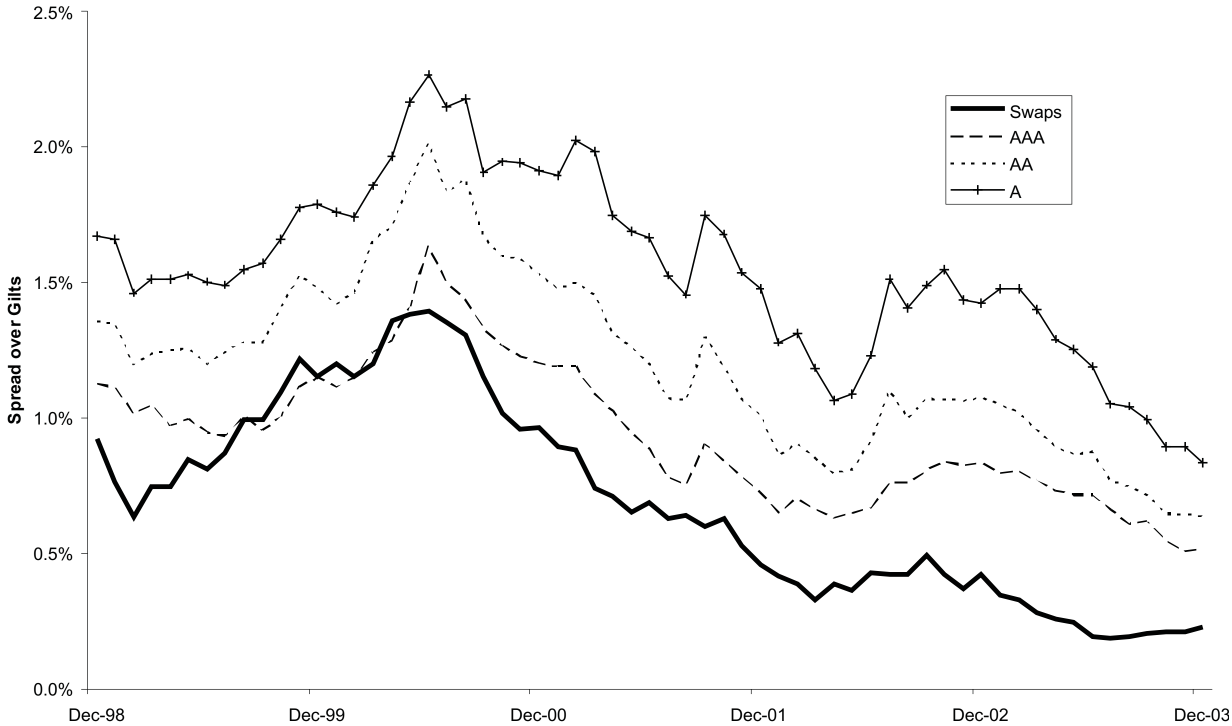

Source: DataStream

Figure 3.2. Historical spreads of swaps and corporate bonds over gilt yields (ten year maturities)

3.2.8 A number of offices have reportedly used swap yields to build a liability discount curve, and asserted that the result is market consistent. In addition to the rationales given by Dullaway & Needleman (2004), this practice does, of course, give a lower liability than gilt yields, which may explain its attraction. However, it does give rise to important consistency questions if gilts are actually held on the balance sheet. If the swaps are truly risk free, then why would an office knowingly accept lower than risk free return on its gilts portfolio? Perhaps it would make sense to take a haircut to the gilt market value to reflect this, although this is an unsatisfactory work around for a problem caused by the pretence that swaps are risk free in the first place.

3.2.9 The use of swaps as a proxy for risk free rates also creates some odd hedging incentives. Stated liabilities could increase, even in the absence of an interest rate move, if swap spreads were to tighten. Tightening swap spreads are associated with improving credit risk, not in the economy as a whole, but specifically in the banking sector. In order to hedge a liability discounted at swap yields, insurers may decide to increase their credit exposure to banks. From a broader economic perspective, it is difficult to argue that such a 'hedge' reduces risk, and could even increase the potential for systemic instability.

3.3 *Credit Risk and Market Consistent Valuation*

3.3.1 The allowance for credit risk in market consistent valuation is contentious. It is helpful to distinguish between *promised* cash flows and *expected* cash flows. The difference is the default option, sometimes called the *limited liability put option*. This option is an asset of shareholders, because it reflects their ability to walk away from a failed company leaving creditors out of pocket.

3.3.2 From a shareholder perspective, the limited liability put option is an asset which has a value. This asset could either be considered separately or as a reduction in the associated liability. The effect for fixed cash flows might be captured by using a discount rate derived from a corporate bond of similar credit standing to the liability owner.

3.3.3 There is an important dynamic here. A shareholding in a limited liability firm can never be a liability. Therefore, after allowing for credit risk, the value of the liabilities can never exceed the assets. As an insurer gets into financial difficulties, the limited liability put option becomes more and more valuable.

3.3.4 From a solvency supervision point of view, it makes little sense to allow credit risk to reduce stated liabilities. The solvency framework is aimed at delivery of promises, not merely delivery of expectations where these are less than the promise. A system which gave regulatory credit for a limited liability put option asset would perversely allow weaker companies to devalue their liabilities, and hence continue to appear solvent, however dire their financial condition.

3.3.5 The argument that short supplies of gilts have depressed yields and created an artificial market may or may not have an element of truth, but supply and demand affect the price of all investments, including swaps. This is discussed in Mehta (1999). Furthermore, market consistency dictates the use of actual market prices, even when those prices are affected by unusual supply or demand effects. A believer in market inefficiency might try to correct market prices to where they ‘truly’ should be, but the result is not then market consistent.

3.3.6 Nevertheless, solvency supervisors recognise that, from time to time, firms do fail. The cost of enforcing government credit quality on the insurance industry would have detrimental cost implications. A compromise is needed. For example, solvency supervision might target a BBB rating or equivalent for an insurer’s overall financial strength.

3.3.7 It is tempting, but wrong, to deduce that such a supervisory regime should be based on a BBB valuation of insurance liabilities. The flaw arises because the credit rating of an insurer is not directly related to the discount rate at which the liabilities are discounted. An insurer which discounts liabilities at BBB yields may still only have a C rating if liabilities and assets are mismatched. Supervision requires a perspective different to that of a shareholder in ¶3.3.2. The regulatory balance sheet is only meaningful when the provision of capital on top of best estimate liabilities is included. In practice this has led insurers’ total provisions, including resilience reserves and solvency margin, to be equivalent to liabilities discounted at a rate significantly below gilt yields.

3.3.8 A further problem can arise in the assessment of capital requirements, if swap rates have been used to discount liabilities. Many companies try to relate capital requirements to default rates on corporate bonds of their own credit rating. CP195 indicates that regulatory capital might be aligned to a 99.5% confidence level over one year, approximating to a BBB rating. There seems to be a risk of double counting in capital calculations if the liabilities already make some, possibly unknown, allowance for credit risk.

3.3.9 It is our view that insurance solvency supervision is most usefully implemented relative to a liability that is as close to risk free as possible. In the U.K., this would mean the use of gilts rather than swaps to calibrate risk free discount factors.

3.3.10 It might be thought that a decision to value liabilities using gilt yields could create solvency strains when a liability has been hedged with a financial institution. For example, the market will value an over-the-counter zero coupon bond at a discount to a corresponding gilt. Some insurers in this situation would prefer to value the policyholder liability by reference to the hedging asset. This would implicitly involve a liability reduction for credit risk — although, in this case, the quantum of the reduction would reflect the credit standing of the derivative counterparty rather than the insurer. Similar issues apply in relation to annuity portfolios backed by corporate bonds. The valuation of fixed liability cash flows should be independent of assets held.

3.3.11 This solvency strain does not create an inconsistency. The insurer is usually obliged to pay out on a policy even if the underlying assets default. A non-government backed zero coupon bond may hedge the interest rate risk of a future liability cash flow, but credit risk remains.

3.4 *Mark to Market of Assets and Derivatives Consistent with Liabilities*

3.4.1 An insurer may decide to use term dependent risk free rates, calibrated to gilt yields. The next challenge is to apply this consistently to other instruments, particularly derivatives which result in a non-insurance liability. For example, should an interest rate swap be revalued using gilt yields when it becomes a liability, and the impact of limited liability be deducted separately?

3.4.2 Our view is that it is appropriate to value swaps at market value, even when a risk free basis has been adopted for liabilities. This is because of the source of the credit spread, which mainly relates, not to the credit risk of the insurer with respect to swap default, but rather to the credit risk inherent in a LIBOR transaction. It is also consistent with a market valuation of all assets.

3.4.3 Other aspects are less clear-cut. For example, if an insurer accepts deposits linked to LIBOR, should the liability calculation include an allowance for the future spread in excess of the risk free rate? To be consistent with our proposals for risk free annuity valuation, such a spread provision would be required, even though in shareholder value terms the spread would be offset by a limited liability put option asset.

3.5 *Rates Net of Tax.*

3.5.1 We now address some initial questions regarding market consistent values net of tax. Let us suppose, for example, that an office pays tax on investment income at a rate of 20%, calculated annually on a mark to market basis and with no allowance for indexation, and that the yield curve is 5% p.a. at all maturities. Simplistically, we might combine these to produce a net discount rate of 4% p.a. for fixed cash flows.

3.5.2 However, when we try to execute this, we have a problem. For a cash flow in ten years' time, we might try to buy a ten-year zero coupon bond as a hedge; but at the end of the first year we have a tax bill depending on market value changes over that year. If we are unlucky, perhaps we could end up liquidating so many bonds to pay tax on an initial capital gain that insufficient remains to meet the liability.

3.5.3 Fortunately, Jensen's inequality offers some reassurance here. It turns out that in this example, if we use the net (of tax) present value to buy a zero coupon bond, then the worst case is that the liability is exactly met. In all other cases, there will be some surplus.

3.5.4 This surplus has a value, and suggests that there may be an argument for a higher net discount rate than 4% p.a. Indeed, this turns out to be the case, though the effect is small. Instead of holding zero coupon bonds of matching maturities, it is possible to devise a portfolio of zero coupon bonds which hedge their own tax liability as well as the original policyholder liability. This results in a slightly lower net of tax liability, and a reduction in the expected surplus (while still avoiding negative surplus). This new valuation is calculated by using forward rates over the term of the liability, netting down each for tax, then recombining to produce a present value.

3.5.5 We can devise a structured derivative whose after tax return is designed to be exactly a fixed number, and hedge out all the tax effects (except, of course, for a change in the underlying tax rate).

3.5.6 These extra surpluses, however, are very small indeed. For example, consider a two-year fixed liability. Let us suppose the that one-year spot rate is 1% p.a. and the two-year spot rate is 5% p.a. Suppose also that the one-year rate in one year's time can take the values of either 6% or 12%. Then, over two years, a simple binomial model can be used to show that:

— netted down two-year spot rate = 4.000%;
— effect of netting down forward rates = 4.012%; and
— theoretical return on structured derivative = 4.018%, assuming credit can be taken for relief of tax losses.

3.5.7 As the total gain from structuring is less than 2 basis points, well below likely structuring costs, offices would be unlikely to pursue so accurate a hedge.

3.5.8 There is one final twist in this tale, relating to investment choice. Different investments are taxed at different rates, with or without indexation or tapering relief, and with different treatments of unrealised gains. This far, we have looked at taxation of matching instruments to a liability. This is a natural choice for a regulatory valuation consistent with the aims discussed in Section 2.2; but if our valuation motivation is shareholder focused, it would be appropriate to examine the actual tax likely payable on the assets held. This could create a lower liability valuation where an insurer holds tax efficient, but mismatching, assets.

3.6 *Variable Cash Flows*

3.6.1 Unit-linked business requires the valuation of variable cash flows, such as future annual fund management charges. These can be valued using actuarial funding principles to discount the future charges. The liability to policyholders is then simply the actuarially funded value, allowing for future decrements.

3.6.2 No reduction for the company's own credit risk is usually made to the value placed on the unit liabilities. This is consistent with the use of risk free gilt yields above. A policyholder may though place a lower value on his benefits to allow for credit risk, though we suspect that few, if any, policyholders consider this in practice.

3.6.3 In bond markets many investors receive cash flows (coupons and redemption proceeds) with no deduction for tax. It is therefore common to assume that the bond market is dominated by gross investors, who therefore determine market prices. We can therefore calibrate to market prices and yields with no adjustment for tax.

3.6.4 In equity markets the position is less clear cut. We need to make an assumption on, for example, tax on dividends. Most investors may no longer be able to reclaim the tax credit on dividends, and that may be the safest assumption for calibration purposes.

3.6.5 The tax assumption on dividends also affects calibration of options with skew and smile volatility characteristics. While insurers need to evaluate guarantees based on total asset returns, most equity option prices are based on capital only indices. To convert from one to the other, an assumption needs to be made on the implied tax rate on dividends.

4. Simple Guarantees

4.1 *Guaranteed Policies*

4.1.1 In this section, we consider simple maturity guarantees. A simple guarantee policy pays the greater of two quantities S_T and K_T at some future date T. For example, S_T might be based on the total return for an investment portfolio and K_T might be a fixed sterling amount, like the benefits on a guaranteed equity product (Dodhia & Sheldon, 1994). We assume that both S_T and K_T must be positive, although their values can be uncertain. We can write the policy benefit $\max\{S_T, K_T\}$ in two ways:

$$\text{Max}\{S_T, K_T\} = S_T + \max\{K_T - S_T, 0\} = K_T + \max\{S_T - K_T, 0\}.$$

4.1.2 The first of these expresses the liability as a unit-linked payoff plus a put option. That has become the usual presentation for realistic balance sheets reported to the FSA. Equivalently, the second presentation shows a fixed payoff plus a call option. The merits of the second presentation are discussed in Dullaway & Needleman (2004).

4.1.3 We assume that a liability of S_T is straightforward to value, and has some value S_0 at time 0. We assume the same for K_T, with value K_0 at time zero. The valuation of the policy, however, is more complicated when it is not known at the outset whether S_T or K_T will be the larger.

4.1.4 We seek a simple formula for valuing a guarantee. We expect the valuation to depend on at least the following:
- the values S_0, K_0 for stand-alone non-guaranteed policies;
- the time T to maturity;
- some measure σ^2 of the level of uncertainty in the guarantee per unit time, that is of the degree of mismatch between the underlying assets representing S and K.

4.1.5 Let us denote a guarantee function by $g(S_0, K_0, \sigma\sqrt{T})$ to place a value on Max $\{S_T, K_T\}$. We would like to use the same function for a variety of policies based on the same underlying assets, but with different levels of guarantee. Arbitrage arguments imply some restrictive conditions on the function g, including:
- g is an increasing function in all three of its arguments (because we prefer to have more than less, and we prefer to commit later rather than earlier);
- g is a convex function in S_0 and K_0, because we prefer two policies with guarantees of 80 and 120 over one policy on the same assets guaranteeing 200;
- $g(S_0, K_0, 0) = \max\{S_0, K_0\}$; and
- $g(S_0, K_0, \sigma\sqrt{T}) < S_0 + K_0$.

4.1.6 Although this is not an arbitrage condition, we might also reasonably expect the function g to be symmetric in its first two arguments, so that $g(S_0, K_0, \sigma\sqrt{T}) = g(K_0, S_0, \sigma\sqrt{T})$. Two possible functions satisfying this equation are as follows:

$$g_{HY}(S_0, K_0, \sigma\sqrt{T}) = \frac{S_0 + K_0}{2} + \sqrt{\left(\frac{S_0 + K_0}{2}\right)^2 - S_0 K_0 \exp\left(-\frac{\sigma^2 t}{2\pi}\right)}$$

$$g_{BS}(S_0, K_0, \sigma\sqrt{T}) = S_0 \Phi\left(\frac{\ln(S_0/K_0)}{\sigma\sqrt{T}} + \frac{\sigma\sqrt{T}}{2}\right) + K_0 \Phi\left(\frac{\ln(K_0/S_0)}{\sigma\sqrt{T}} + \frac{\sigma\sqrt{T}}{2}\right).$$

4.1.7 The function Φ denotes the cumulative normal distribution function. The first of these is a hyperbola, which we think is the simplest possible formula fulfilling the various criteria. The second formula is based on the celebrated Black-Scholes option pricing formula. We have introduced the factor of 2π in the first formula so that the formulae agree when t is small and when $K_0 = S_0$. Subtracting either K_0 or S_0 from these expressions gives formulae for the prices of call options or put options respectively.

4.1.8 The only problem now is to establish appropriate choices for the parameter σ. For a market consistent valuation, we need to establish market values for this parameter σ. This leads us to consider data for which market prices are available.

4.1.9 Figure 4.1 shows these two formulae, together with market prices, of put options on the FTSE 100. These relate to one-year options as at

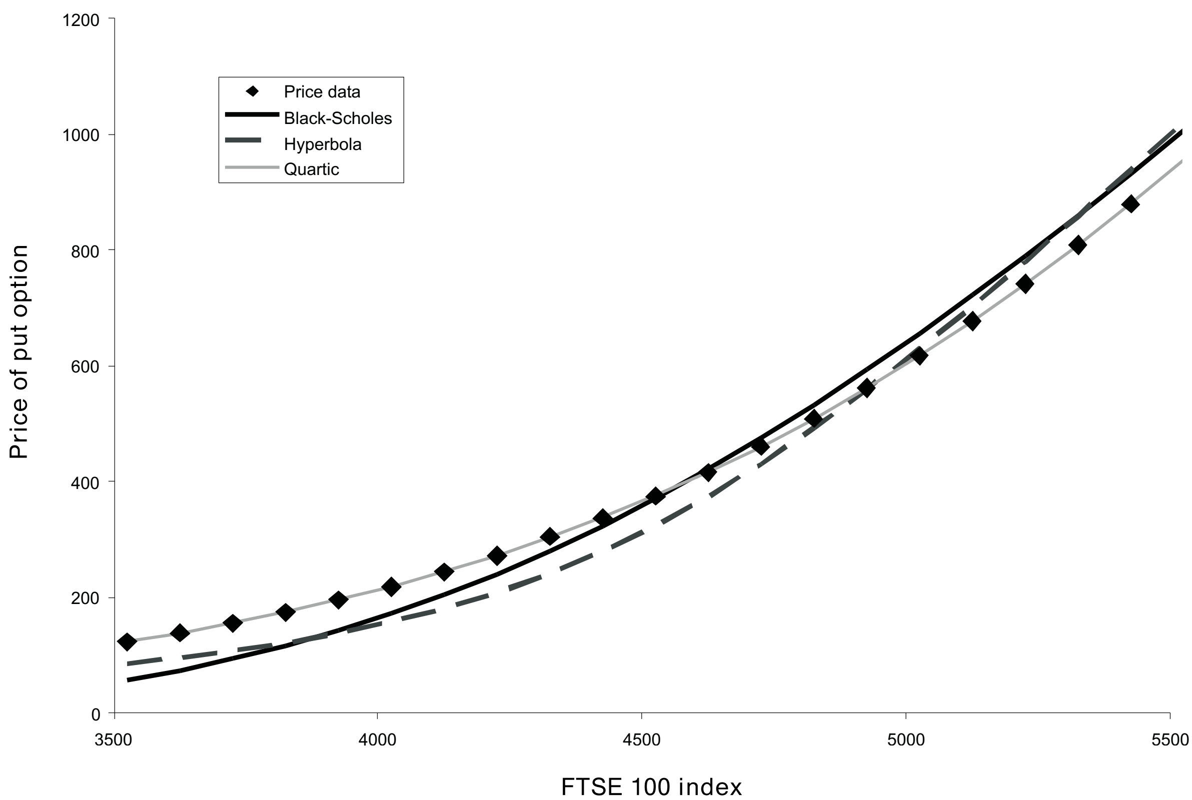

Source of option price data: LIFFE

Figure 4.1. One-year put option prices on FTSE 100 index, and three fitted curves

December 2003. We have also shown a polynomial fit to quoted prices, which naturally provides a very good fit indeed, given that we can choose five calibration parameters.

4.1.10 The idea will be to observe one or more market prices, and then solve for the σ parameter to be used for valuing other options. No model will be able to fit all available option prices. In practice, a balance must be struck between accuracy of the fit for a limited range of options and ensuring a wide enough range of calibration options representative of the guarantees embedded in a company's liabilities.

4.2 *Volatility Calibration*

4.2.1 A market consistent valuation would usually seek consistency, not only with equity and bond markets, but also option markets. This means that a market calibrated model would seek to replicate market prices of selected options.

4.2.2 There are very many options which could be used for calibrating an asset model. Equity options vary according to the strike of the option and maturity. Interest rate swaption prices vary according to the strike, the option term and the term (or tenor) of the underlying swap. This gives a bewildering array of prices; current option pricing models would struggle to fit all of these prices simultaneously. Consideration must be given either to specifying a small number of benchmark options whose prices are used in calibration, or to developing a weighting scheme seeking close approximation to a wider range of instruments. In practice, greater weighting should be given to fitting those options providing the closest match by terms and strikes to the guarantees in the liabilities. While this is a tricky problem for simulation models, it would be possible to derive a matrix of volatilities by term and strike price for use in option pricing formulae.

4.2.3 The method of extrapolation for other option strikes is critical. Figure 4.2 shows the three formulae for one-year put options, extrapolated for lower strikes.

4.2.4 This suggests that for low strikes, the hyperbola is a better fit to market prices than Black-Scholes, at low strikes. The quartic extrapolation is obviously silly, producing negative option prices. This illustrates an important principle, that a close fit to a range of observed prices can create instabilities on extrapolation. While a fit to observed prices is an important part of a test for market consistency, the model that fits market prices most closely may turn out to be the least appropriate for extrapolation to other cash flows.

4.2.5 Although a large number of option prices are potentially available, in practice data may be hard to collect. Traded option prices are often available from the relevant exchange, but these are typically short dated (less than two years in the U.K.), and so are of limited use for most insurance liability valuations. Some online data sources provide historical data on

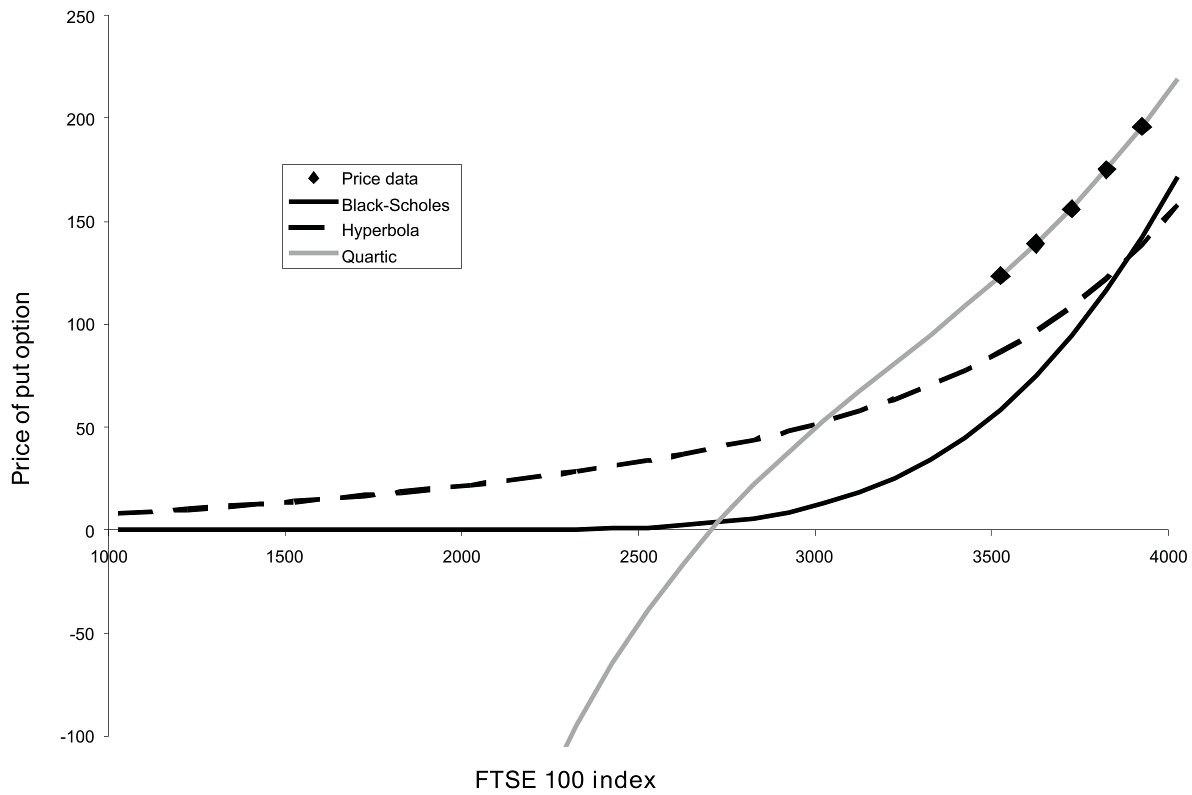

Source of option price data: LIFFE

Figure 4.2. One-year put option prices on FTSE 100 index, and three extrapolated curves

swaptions. We are not aware of commercially available data sources for prices of less liquid instruments such as ten-year equity options. The small group of investment banks which guard these data have been known to open their price databases in exchange for deal referrals.

4.3 *Market Option Data*

4.3.1 Options quotes are available on many underlying assets. These include:

— equity indices, such as the FTSE 100, as well as individual stocks;
— interest rate options, such as gilt options, caps, floors and swaptions;
— foreign currency options; and
— options on commodities, such as gold, oil or gas.

4.3.2 In many cases option quotes are available through specialist financial data providers, directly from an exchange or via friendly banks or brokers.

4.3.3 Options on the FTSE are usually available at a wide range of strike prices, quoted on the exchange with terms out to around one year. Some samples, as at 31 December 2003, are shown in Figure 4.2.

4.3.4 Longer-dated option prices, out to around five years, may be available as over-the-counter quotes from banks.

4.3.5 There are considerable problems with this sort of option data. Although options at the short end trade frequently, the longer options are much less frequently traded. For example, there may be trades in five-year FTSE 100 options once a week or even less. The bid-offer spreads are wide, and latest trade data will suffer from timing problems. For example, put options should be more valuable for higher strikes, but data may not show this if the most recent trades were several days apart and the underlying index level had moved.

4.3.6 In contrast, the data available from commercial data providers often seem suspiciously smooth. Timing problems or other inconsistencies seem rare. The reason for this good behaviour is that the prices listed have often been derived themselves from a model. For example, the closing prices on LIFFE are obtained from real quotes by applying a smooth spline function to implied volatilities.

4.3.7 What appear to be long-dated data are often merely extrapolated from actual trades on shorter-dated instruments, and not real data. If an actuarial model failed to replicate a quoted price for a 20-year option, the correct conclusion may be that the two models use different extrapolation algorithms, and not that one model or the other is market inconsistent.

4.3.8 There are two further difficulties with the use of option prices in market consistent models for insurers. The first is that the indices quoted do not correspond precisely to the portfolios held by many insurers. Some adjustment is required to FTSE 100 implied volatilities to derive volatilities of non-standard portfolios. This issue is much less trivial than it seems.

4.3.9 The second difficulty relates to credit risk. Given call and put option prices for a range of strikes, it is possible to deduce the implied discount factor K_0/K_T. Unsurprisingly, this discount factor reflects inter-bank credit rather than government credit. To get a credit risk free market consistent value, it is necessary to take the market quote and then:

— unwind the effect of inter-bank credit risk; and
— re-state the option price using government credit, or the credit risk of the insurer, according to whether or not the limited liability put option is to be recognised as a reduction in the liability.

4.3.10 The market quote will typically use interbank swap rate or LIBOR both for the discount factor and in the forward rate in the Black Scholes formula for an option. An allowance for credit risk associated with the writer of an option should, however, be based solely within the discount rate applied to the expected payoff evaluated by risk neutral or deflator techniques. The forwards in the Black Scholes formula should not be adjusted.

4.3.11 As a simple example, consider a one-year put option on a non-dividend paying equity share, with current share price 100, strike 102, risk free rate 4.5% p.a. continuously compounded, and a volatility of 17%. The risk free price on these parameters is 5.52.

4.3.12 If we now assume that credit risk is reflected by an additional 100 basis points, and simply replace the 4.5% p.a. by 5.5% p.a. in the formula, we obtain a price for the put option of 5.07. This is too low. The correct adjustment is simply to discount the risk free price by the credit spread for one year (i.e. adjust the discount rate only) to obtain a credit risky price of 5.46.

4.3.13 This adjustment needs to be made in reverse when determining risk free option prices from the quoted credit risky prices. An implicit assumption behind this type of adjustment is the independence of the credit free price for the options and a default occurrence. This is clearly a false assumption, particularly for out of the money put options, where, for example, an equity put option with a very low strike would be worthless if the writer would become insolvent at such low equity prices.

4.4 *Data Availability for Interest Rate Volatility*

4.4.1 Interest rate options are available in a number of forms. These include gilt options and swaptions. For gilt and bond options, the implied volatility will depend on four factors: the term of the bond, the term of the option, the bond's coupon and the strike. Swaption implied volatility depends on three factors: the term or tenor of the swap, the term of the option and the strike. Interest rate models will have their own volatility parameters, which will not necessarily correspond to these observed implied volatilities. For example, a model based on the short rate will have a parameter for the volatility of that rate. All these various forms of volatility rather confusingly get referred to as 'volatility'. For small enough volatility values, they are roughly proportional, and there are convenient formulae available for converting between the various definitions. This does raise a practical issue for stress tests. If an interest rate volatility stress is required, careful definition of the volatility parameter to be altered is required. Further, in the event of a change in interest rates, it is not possible to hold all volatility definitions constant. A decision is therefore required as to which, if any, to keep constant.

4.4.2 The gilt options are potentially the most relevant for calibrating risk free volatility. They tend to mature at awkward dates, reflecting the underlying gilts, and so can be difficult to calibrate. Unfortunately, the traded gilt options tend to be based on baskets of gilts with a 'cheapest to deliver' clause. This means that the gilt options are actually compound options, because the seller of a call (or buyer of a put) can select the underlying gilt (from a list) when the option matures.

4.4.3 For this reason, many analysts prefer to use interest rate swaptions instead. These are options on interest rate swaps. The payoff from a swaption can then be specified by:
— the strike rate k;
— the strike date s;
— the maturity date t; and
— whether a payer swaption or receiver swaption.

4.4.4 A receiver swaption is an option to receive the fixed strike rate and pay the reference floating rate, for a deposit to be made on the strike date and continuing at (for example) half-year intervals until the maturity date. To make the formulae simpler, we will show pricing formulae for cash flows at annual intervals.

4.4.5 Let us consider the swaption at time s. The receiver swaption entitles the holder to receive a fixed stream of payments of k, the value of which is given by the annuity (P denoting zero coupon bond prices):

$$K_s = k[P(s, s+1) + P(s, s+2) + \ldots + P(s, t)].$$

4.4.6 The quantity to be paid in exchange for this fixed stream is a stream of LIBOR payments. These payments can the synthesised (credit risk aside) from a deposit at time s, but rebated at time t. Therefore, the value of the floating rate stream is given by:

$$S_s = 1 - P(s, t).$$

4.4.7 The values of S and K at time 0 are now easily derived from a hedging argument. We find that:

$$S_0 = P(0, s) - P(0, t)$$
$$K_0 = k[P(0, s+1) + P(0, s+2) + \ldots + P(0, t)].$$

4.4.8 We can now use our standard Black formula (or indeed, the simpler quadratic formula) to derive swaption prices. This is, indeed, how swaption prices are quoted, in terms of a Black implied volatility.

4.4.9 Swaption quotes are available, and are reasonably liquid, for short maturities. Unlike equity options where prices are available for a range of strikes, data tend to be available only for at-the-money swaptions. From the equity options prices and their associated volatility smile, it is possible to deduce a distribution of future equity returns. The absence of swaption prices or volatilities at in and out of the money strikes makes it difficult to derive a shape of future interest rate distributions. We understand that some banks do have strike based models which reflect skew and smile effects, but that data are not readily available. However, we can be sure of at least one thing — that a receiver swaption with a strike rate of zero is worthless. As LIBOR must be positive, there is no value in a swaption to pay LIBOR and receive nothing in return.

4.4.10 Calibration to the volatility surface of swaptions is easier in a stochastic model based on swap rates rather than on gilt rates, simply because swaptions are options on swaps, not gilts. However, provided an adjustment made is made to allow for the fact that the strike prices of the swaptions are based on swap rates, a gilt based model, in principle, can be calibrated to swaption volatilities, following the process in ¶4.3.9.

4.4.11 Volatilities derived from swaptions will not, in general, be equal to the corresponding volatility derived from gilt options. The difference arises because the swap rate is a gilt rate plus a credit spread, which is itself variable. The implied volatility of swaptions therefore reflects, not only risk free interest rate volatility, but also the volatility of the credit spread and its correlation to the risk free rate. It often turns out that the spread is negatively correlated with gilt yields. As a result, swaption implied volatilities may lie slightly below the implied volatility for corresponding gilt options.

4.5 *Calibrating a Structural Model to Bond Option Prices*

4.5.1 Calibrating models to bond prices is tricky. The Black formula assumes bond prices are lognormal, but it is not possible for all coupon bonds simultaneously to be lognormally distributed. Therefore, there is no coherent underlying model which implies the Black formula for all bond options.

4.5.2 There are relatively few models for which swaptions have closed form solutions. One such model is the Hull-White model (1990), which we now describe. The Hull-White model is based around a standard Brownian motion W_t. A mean reverting process X_t is then defined by the stochastic integral:

$$X_t = \int_0^t e^{-\alpha(t-u)} dW_u.$$

4.5.3 Equivalently, X_t satisfies the stochastic differential equation:

$$X_0 = 0$$
$$dX_t = -\alpha X_t dt + dW_t.$$

4.5.4 Under the Hull-White model, deflators and term structures have the following form:

$$D_t = \frac{P_{0t}\exp[\lambda W_t + \sigma X_t]}{\mathbf{E}_0 \exp[\lambda W_t + \sigma X_t]}$$

$$P_{st} = \frac{\mathbf{E}_s(D_t)}{D_s} = \frac{P_{0t}}{P_{0s}}\exp\begin{pmatrix} -\sigma(1-e^{-\alpha(t-s)})(X_s - \frac{2\sigma\lambda}{\alpha}(1-e^{-\alpha s})) \\ +\frac{\sigma^2}{2\alpha}(1-e^{-2\alpha s})(1-e^{-2\alpha(st-s)}) \end{pmatrix}.$$

4.5.5 Here, the parameters α, σ have to be estimated. The parameter α describes the degree of mean reversion and σ captures volatility. The parameter λ is related to risk premiums for real world projections, and (as we shall see) cancels out in swaption prices. As X can take any real value, negative interest rates are not excluded. Depending on the parameter choices, the probability of negative rates may, however, be small.

4.5.6 Pricing options on zero coupon bonds is straightforward under the Hull-White model. For example, let us consider a structure paying at time s the greater of:

$$S_s = P(s, t)$$
$$K_s = K.$$

4.5.7 The value of this option is the deflated expectation $\mathbf{E}[D_s \max\{S_s, K_s\}]$. This expectation is given by the Black formula $g_{BS}(S_0, K_0, \eta\sqrt{s})$ as in ¶4.1.6, where:

$$S_0 = P(0, t)$$
$$K_0 = P(0, s)K$$
$$\eta = \sigma\sqrt{\frac{1-e^{-2\alpha s}}{2\alpha s}}[1 - e^{-\alpha(t-s)}].$$

4.5.8 Coupon bonds are more complicated. Jamshidian (1989) gives an exact algorithm. This algorithm is useful for calculating option prices if the model parameters α, σ are known, but it involves the solution of an implicit equation. It is therefore useful for pricing, but inconvenient for calibration. For calibration purposes, we need approximations which are more tractable.

4.5.9 A suitable approximation proceeds as follows. We consider options with a strike K and strike date s on a bond with annual coupons at rate g. Define a function $h(\xi)$ by:

$$h(\xi) = -KP(0, s) + g\left[\sum_{u=s+1}^{t-1} P(0, u)\exp[-\xi(1 - e^{-\alpha(u-s)})]\right] + (1 + g)P(0, t)\exp[-\xi(1 - e^{-\alpha(t-s)})].$$

We define ξ_0 as the root ξ of the equation $h(\xi) = 0$. This depends on known initial bond prices and on the mean reversion parameter α. We can now price bond options using an approximate Black formula, with:

$$S_0 = g \sum_{u=s+1}^{t-1} P(0,u) + (1+g)P(0,t)$$

$$K_0 = P(0,s)K$$

$$\eta = \sigma\sqrt{\frac{1-e^{-2\alpha s}}{2\alpha s}}\frac{\ln(S_0/K_0)}{\xi_0}.$$

4.5.10 With this formula, we notice that both S_0 and K_0 are functions only of known observables. Given an option price, we can solve for the implied volatility η. Given two option prices, we then have two values of η. This gives two equations in the two unknowns α, σ. Fortunately, both contain σ as a multiplicative factor on the right hand side. Therefore we can divide one equation by the other to obtain a single equation in the unknown mean reversion α.

4.5.11 The choice of bond options, or swaptions, to use in the calibration is critical. The degree of mean reversion governs the relative volatilities of:

— options with different strike dates on bonds of the same remaining term (the strike date or option term structure); and
— options with a common strike date on bonds of different terms (the bond term structure).

4.5.12 We take a simple example, assuming a risk free rate of 5% p.a. Table 4.1 shows volatilities typical around the end of 2003.

4.5.13 To analyse the strike date structure, we consider the two options on ten-year bonds, with option terms of three years and one year. Because of mean reversion we would expect the three-year volatility to be below that on the one year, and indeed this is the case, with a three-year bond volatility of 6.00% compared to a one-year volatility of 6.10%. The theoretical ratio is

Table 4.1. Typical bond and swaption volatilities at the end of 2003

Option term (years)	Bond outstanding term	Bond original term	Bond implied volatility	Equivalent swaption implied volatility
1	3	4	2.70%	19.86%
1	10	11	6.10%	15.81%
3	3	6	2.50%	18.44%
3	10	13	6.00%	15.58%

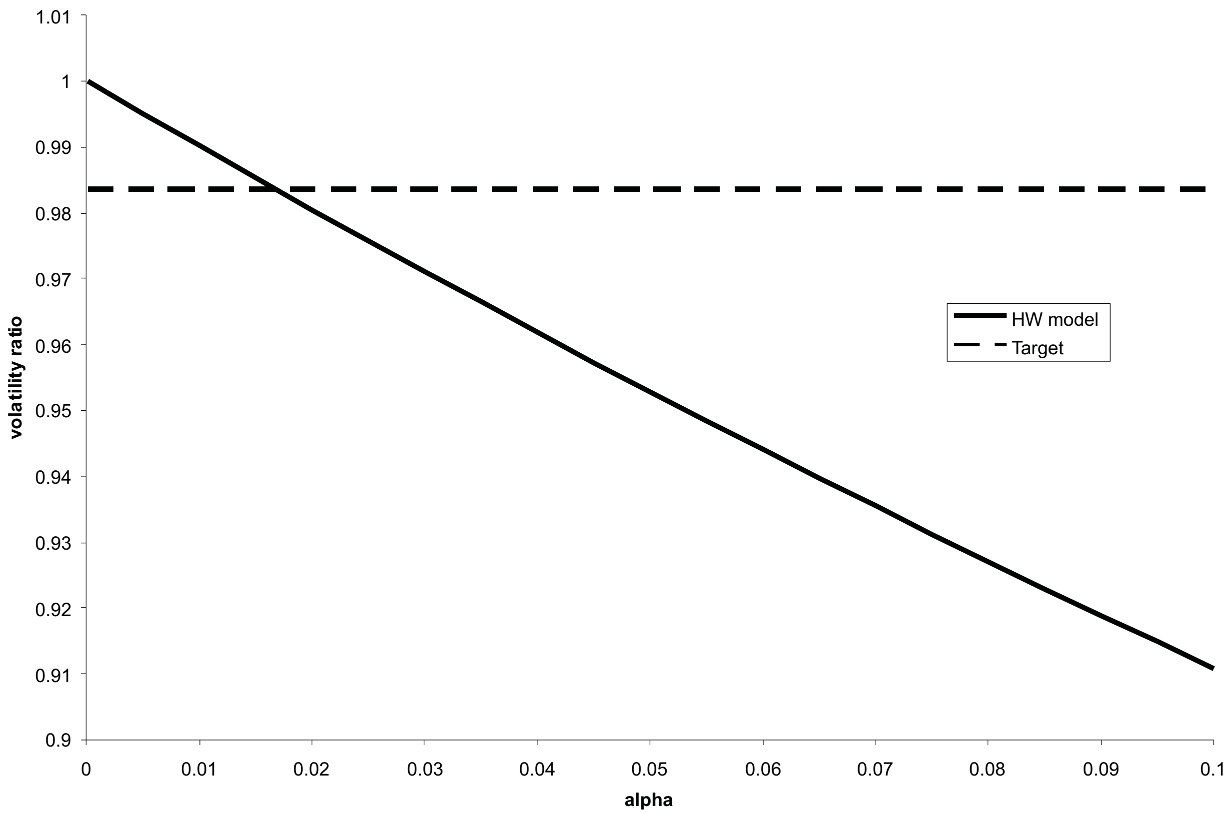

Figure 4.3. Volatility ratio of options on a ten-year bond, strike date of three years vs one year. Solid line is derived from the α parameter of the Hull-White model

a function of the mean reversion parameter α, decreasing from a maximum of 1 at $\alpha = 0$ to a minimum around 0.6 for large α. It is clear from Figure 4.3 that there is a unique value of α around 0.017.

4.5.14 To analyse the bond term structure, we fix the three-year option term and compare the volatility of a ten-year bond to that of a three-year bond. Because of mean reversion, we expect duration calculations to overstate the volatility ratio, so the ten-year bond should not be more than 3.3 times as volatile as the three-year bond. In fact, the ratio is 6.0%/2.5% = 2.4. Under the Hull-White model, the volatility ratio falls from 2.8 when $\alpha = 0$ to 1 when α is very large. Figure 4.4 shows that an intermediate value of α around 0.057 hits the target value.

4.5.15 We notice that no single value of α accurately fits both the strike date structure and the bond term structure. The strike date structure calibration gives a very low value of alpha, equating to mean reversion of 1.7% p.a. Both of these market calibrations give a speed of mean reversion which is slower than most historical time series estimates. These observations are typical for the last five years. Although the numbers have varied, it seems to be systematically the case that the largest mean reversion estimates come from historical data, followed by bond term structure, with bond strike date structures producing the weakest mean reversion.

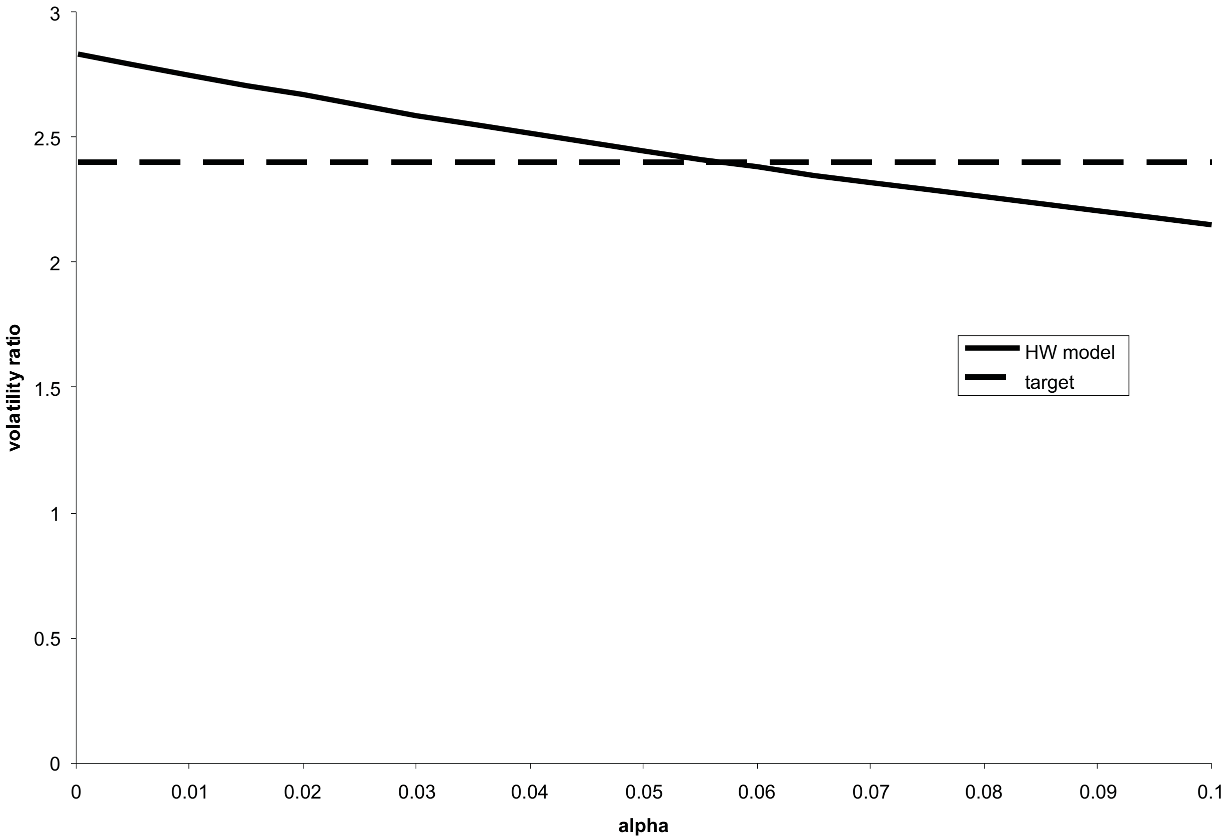

Figure 4.4. Volatility ratio of options struck in three years' time, on a ten-year bond against a three-year bond. Solid line is derived from the α parameter of the Hull-White model

4.5.16 This leaves the modeller with a dilemma — whether to rely on the strike date structure or on the bond term structure of volatility to calibrate a market-consistent model. Another alternative is to build a yet more complex interest rate model. Moving to a more complex model may create more problems that it solves, for example if bond options cannot be priced under the model.

4.5.17 Very slow mean reversion has other implications for liability modelling too. The Hull-White model produces a certain number of negative interest rate scenarios. These cannot simply be discarded, for to do so would result in a model which failed to price the calibration ingredients. This leaves the actuary with a decision of how to model bonus rates, policyholder take-up rates and other dynamic variables, in situations where the model predicts negative interest rates.

4.5.18 In one sense, this is a meaningless and hypothetical question. Life offices reasonably prepare contingency plans for many situations, but even the most meticulously drafted principles and practices would be unlikely to contain a strategy for managing risk in a negative interest rate environment. Unfortunately, using the Hull-White model, this meaningless and hypothetical question can have a large bearing on disclosed value. So, for example, using the strike date structure calibration, we find that over half the value of a 30-year zero coupon bond derives from scenarios where interest rates are negative at some stage. If we assume (as an extreme example) that the life office can cut benefits to zero in the event of negative interest rates, they could apparently halve their market consistent liability.

4.5.19 The most straightforward safeguard against such manipulation is to exclude interest rate models such as Hull-White which give rise to negative values. Alternatively, the Hull-White formula could be used with a matrix of volatilities dependent both on option term and bond term. The relative merits of a simulation approach to modelling and the use of option pricing formulae is discussed in Section 8.

4.6 *Multi-Asset Models*

Options on currencies, commodities and credit are also available. The latter may be relevant to the modelling of U.K. with-profits business where the assets held include corporate bonds, and also to Continental and to American business. There are, though, options not readily available which would be useful to the valuation (and hedging) of insurance liabilities. Ideally, we might like to be able to trade in options on asset share investment mixes, or, failing that, at least on a mix of equities and fixed interest. Even the latter would require a market in the correlation between equities and fixed interest, a market not currently available in the U.K. Similarly, there are no market implied volatilities for property, corporate bonds and inflation.

4.7 *Tax Effects*

4.7.1 The Black-Scholes option pricing formula makes no allowance for tax on investment return. The correct adjustment to make for tax on investment income and gains for basic life assurance business is not obvious, as the following simple example shows.

4.7.2 Consider a one-year contract with 100 invested in equities with annual volatility of 20%, a guaranteed bonus of 2 at the year end, and a risk free rate (continuously compounded) of 5% p.a. With no allowance for tax, the value of the guarantee is 6.45 using the Black-Scholes formula for a put option. Now suppose that equity returns are taxed at 15%, that a tax credit is somehow available on equity losses at the same rate (e.g. through group relief, though some adjustment may need to be made for deferral of losses), and that bonds, from which the risk free rate is derived, are taxed at 20%.

4.7.3 One approach would be to regard tax as a payment from the fund, just as a dividend payment in a fund which only benefits from capital movements. The Black-Scholes formula, adjusted for continuous dividends at the rate of 5% multiplied by the equity tax rate of 15%, produces a value for the guarantee of 6.75. This approach will always lead to a higher cost for a put option, and that seems illogical, since tax relief on losses should limit the extent of asset falls in the fund, thereby reducing the downside risk and making the option less valuable. We cannot therefore simply treat tax as a deduction in the option pricing formula.

4.7.4 A more promising approach is to assume that the dynamic hedging underlying the Black-Scholes model is undertaken in the taxed environment of the life fund. In this case the risk free rate, net of tax at 20%, is 4% p.a. (approx.). The volatility of equity returns in the fund, net of tax, will also be reduced, say by 15%, to give a volatility of 17%. There is no simple relationship between the lognormal distributions of gross and net returns, so this is only an approximation. In practice, the volatility of returns could be determined by simulation allowing for the various tax rates on both losses and gains that the fund incurs. The value of the option is then 5.75.

4.7.5 Another approach is to assume that the fund hedges the guarantee with a one-year option, struck on gross equity returns, and assumed to be priced in the market with no allowance for tax in a Black-Scholes world. The payout from the fund is now $S'' + \max[102 - S'', 0]$, where S'' is the equity value net of tax at the year end. Hence, the value of the guarantee is $\max[102 - S'', 0]$, which equals $\max\,[102 - S + 0.15(S - 100), 0] = 0.85\max[102.353 - S, 0]$, where S is the equity value before tax. Evaluating this using the Black-Scholes formula with the gross of tax parameters gives a value of 5.62. Allowance for tax payable on any gain in the option value (or tax relief on any loss) then needs to be made. The adjustment required for tax is not straightforward. Simply applying a grossing up factor such as $(1 + T)/(1 + vT)$, where T is the tax rate and v the one year discount rate, ignores the further tax liability on the assets held to meet the tax liability on the option. Provided that both the mode and rate of taxation are consistent with the taxation of the assets underlying the option (in this example, tax at 15% on both gains and losses), then we are back in the net of tax dynamic hedging world. It can be shown that the replicating portfolio required for the option and the tax liability has the value of 5.75 calculated in ¶4.7.4. So tax could be allowed for by using net of tax volatility and discount rate. The intricacies of fund level tax and the effects of mismatching also need to be considered.

4.7.6 If the taxation of options is inconsistent with that of the underlying assets, then the above approach will need to be modified accordingly, depending on the action taken to hedge guarantees. There may be potential tax arbitrage in net funds. For example, indexation relief may favour dynamic hedging of equity related guarantees. The relative tax advantages of dynamic hedging, purchasing options or doing nothing may influence the decision of how best to manage guarantees in a net fund. The allowance for tax status is discussed in Mehta (1992).

5. Exotic Guarantee Models

5.1 *Vanilla and Exotic Options*

5.1.1 Call and put options are often described as *vanilla*. These are useful for pricing simple policy guarantees on a market consistent basis. However, there are many more complex derivatives, generically known as *exotic options*. For example:

— A European *put option* gives the right, but not the obligation, to sell an underlying asset at an agreed strike price on an agreed future date.
— An *up-and-in* put option gives the same rights, but only if the price of the underlying touches on or above a specified barrier at least once during the life of the option.

5.1.2 The up-and-in option is an example of an exotic option. Closed form solutions exist for pricing these options; for example, see Hull (2003). The closed form solutions are not the same as the Black-Scholes formula, but they do use the same underlying methodology. There is no corresponding generalisation for our quadratic option pricing formula, as this formula was derived in the absence of any underlying theory. This is a powerful argument for the use of the Black-Scholes model, rather than our quadratic one.

5.1.3 Up-and-in put options are useful for more realistic modelling of bonus policy, allowing for catch-up. For example, let us suppose that a with-profits policy currently has an asset share of 80 and a sum assured with accrued reversionary bonus of 100. We might initially value this as a simple guaranteed policy, but this overlooks the possibility of future bonuses. A crude way to allow for those future bonuses is to assume a one-off reversionary bonus of 10, declared only if the asset share hits 110.

5.2 *Wilkie's Model of Bonus Policy*

5.2.1 The Black-Scholes formula, or some other fitted formula, is convenient for pricing simple guarantees. However, most insurance products are more complex than this. We need a way, not only of extrapolating to different strike prices (as discussed in Section 4.1), but also to different benefit structures.

5.2.2 European call and put options, as priced by the Black-Scholes formula, are examples of *path independent* structures. The value of the option on expiry depends only on the finishing value of the underlying asset. The path taken from start to finish is irrelevant; only the final value counts.

5.2.3 Guaranteed annuity options (GAOs) on unit-linked funds provide another example of a path independent structure. The GAO remains path independent if we add an underlying guarantee to the unit-linked return. In this case the option is more complicated, because it is a function of two unknowns rather than one. We give some formulae for this below.

5.2.4 Other products are more obviously path dependent. For example, consider a regular premium unit-linked policy with a guarantee. In this case, even given a fund return over a period, the policy value depends on the path taken. A series of good returns followed by poor returns will provide a lower policy return than poor returns followed by good returns. This is because the funds under management are expected to increase over the term of the policy, so higher returns are more valuable if they occur later.

5.2.5 With-profits bonds are also path dependent, but in the opposite direction. If low returns are followed by high returns, it is likely that the policy payout could correspond closely to asset shares. If, on the other hand, high returns are followed by low returns, it is likely that the high returns would have triggered reversionary bonuses which cannot subsequently be taken away.

5.2.6 Wilkie (1987) gives an early example of formulae for with-profits policies. His formula treats the bonus build-up as a series of one-year options. The fund return is declared as a bonus, if positive, but is subject to a floor of zero. These returns compound up over the term of the product. The market-consistent liability is then the product of Black-Scholes formulae.

5.2.7 The chief difficulty with Wilkie's formula is that it gives very large liability numbers, and has therefore lacked practical appeal. One reason for the large numbers is that the model gives no allowance for a period of catch-up after poor investment returns. The catch-up period occurs if an office tends to declare low bonuses during the first few good years after poor returns. Higher bonuses would only be triggered after the cumulative returns had caught up to a specified level.

5.3 *Dynamic Bonus Rules*

5.3.1 Reversionary bonuses require a dynamic model because they are often triggered by strong investment performance (or perhaps, more accurately, cut in times of poor performance), yet, once declared cannot be taken away. The cumulative bonus to date is therefore an increasing process, not well described by a random walk. A model bonus policy can be constructed as follows for each group of policies:

(1) An 'asset share' process is calibrated. It is projected into the future using assumptions for volatilities and annual charges.
(2) A 'bonus share' process is calibrated at the fund level. The sum assured plus reversionary bonuses accrued to date are described as the highest ever value of the bonus share process.

5.3.2 Although the bonus share process is a modified random walk, we can think of it as in one of two states. If it has just attained its highest ever value, then a bonus has just been declared, and we say the bonus share is 'active'. On the other hand, under current market conditions some offices have cut bonuses to zero. In this case the bonus share cannot be at its highest ever value, so must have fallen from a previous maximum. The process

Table 5.1. Illustrative scenarios of the progression of asset share, bonus share, sum assured, and resulting maturity value

	$t = 0$	$t = 1$	$t = 2$	$t = 3$	$t = 4$	$t = 5$
Scenario 1						
Asset share	100	110	120	130	140	150
Bonus share	100	105	110	115	120	125
Sum assured	100	105	110	115	120	125
Maturity value						150
Scenario 2						
Asset share	100	110	120	130	120	110
Bonus share	100	105	110	115	110	105
Sum assured	100	105	110	115	115	115
Maturity value						115
Scenario 3						
Asset share	100	90	80	90	100	110
Bonus share	100	95	90	95	100	105
Sum assured	100	100	100	100	100	105
Maturity value						110

is interpreted as being underwater, which means there must be some improvement to the previous maximum before bonuses are again triggered. This dual process permits bonus rates to be dependent on investment returns according to a specified formula.

5.3.3 We illustrate this with a numerical example. An initial premium of 100 is invested for five years. Let us suppose that the bonus share grows at half the increase (or decrease) in the asset share. Three possible scenarios are shown in Table 5.1.

5.3.4 In Scenario 1, the asset share grows by 10 per annum. A bonus of 5 is declared each year, and a terminal bonus of 25 makes the policy proceeds up to the asset share.

5.3.5 In Scenario 2, the asset share rises to 130 and then falls back to 110. The bonus share process rises to 115, and then falls back to 105. As bonuses cannot be taken away, the sum assured (including reversionary bonus) rises to 115 and then stays there. Finally, then, the policy pays no reversionary bonus, as the declared bonuses exceed the asset share.

5.3.6 In Scenario 3, the asset share falls to 80 and then climbs back to 110. The sum assured remains at 100 as the market falls, and also stays there until the market catches up to its original level. Finally, at $t = 5$ the asset share is back to where it started, and bonuses start again. On this occasion, the asset share has risen far enough to allow a terminal bonus too.

5.3.7 The path dependency in this model is clear, because Scenarios 2 and 3 have the same terminal asset share and bonus share, yet result in different policy payouts. Roughly speaking, an ‘up-down’ pattern is more costly to the office than ‘down-up’, because in the former case, reversionary bonuses will have been locked in on the upswing.

5.3.8 The policy benefits depend on the terminal value of the asset share and on the highest ever value of the bonus share. Such payoffs, depending on the maximum of an index, can be valued using a modification of the Goldman *et al.* (1979) model. If the asset share and bonus share are not 100% correlated, then the cumulative normal functions are replaced by a bivariate normal function. The resulting expression is lengthy, so we do not list it here. The implementation is tedious, but straightforward.

5.4 *Dynamic Investment Policy*

5.4.1 Options on a static investment mix can be straightforward to price, with constant fund volatility in the Black-Scholes formula. Many offices have, though, developed dynamic investment policies for their with-profits funds. These policies typically involve a move into less volatile assets such as cash or bonds if solvency is threatened, with some degree of matching by term to the guaranteed liabilities. In good times a higher proportion is held in more volatile assets, such as equities and property. We look for a way of modelling this type of investment policy in the formula approach.

5.4.2 The idea is to model the fund as a proportion which is in the less volatile asset plus the remainder as an exposure to 'super risky' assets. These super risky assets have a higher volatility than the equities (or property) held by the fund, but the proportion of super risky assets is chosen so that the overall volatility of the fund reflects that of the actual asset mix. In reality, the asset mix will be adjusted following a significant fall in the market value of equities or property. In our model, however, there is no dynamic trading, so an upward move in market values is automatically reflected in higher exposure to the super risky asset, and a downward move in lower exposure.

5.4.3 To give a numerical example, we consider the simplest case of deterministic interest rates. Consider a fund as follows:

— 50% equity and 50% cash;
— equity volatility of 20% (and so fund volatility of 10%, with the simplifying assumption of zero volatility for cash); and
— strategy statement that the fund is rebalanced to 35% equity / 65% cash if its market value falls by 10% (and so a new volatility of 7%).

5.4.4 It is interesting to consider what data might be required to support such a dynamic model. One possible requirement on the upside would be to insist that any model successfully explains the high equity backing ratios of the 1990s, given the investment performance at that time.

5.4.5 If we take the rule specification as reasonable, the key attributes that we want to replicate are as follows:

— current fund volatility = 10%; and
— after a 10% fall in value the fund volatility falls to 7%.

Table 5.2. Calibration of 'super equity' model

	Now (before fall)	After 10% fall
Fund value	100	90
Of which: cash	$100(1-\pi)$	$100(1-\pi)$
Of which: super equity	100π	$100\pi - 10$
Volatility	$\pi\sigma = 10\%$	$\frac{100\pi - 10}{90}\sigma = 7\%$

5.4.6 The idea now is to replicate the behaviour of the fund using a static mixture of cash and some proportion π in a new 'super equity', with a different volatility σ to be determined. We can calibrate the model by calculating values before and after the hypothetical fall as shown in Table 5.2.

5.4.7 We can solve these equations, setting $\sigma = 37\%$ and $\pi = 27\%$. We can now price options on the static mix as before, by netting the cash portion against any guarantee.

5.4.8 It is instructive to consider how similar guarantee costs could have been obtained with a simple static model, based on the original equity proportions and with some adjustment to the volatility assumption. Assuming a five-year time horizon, the relevant implied volatilities are shown in Figure 5.1, as a function of the level of guarantee.

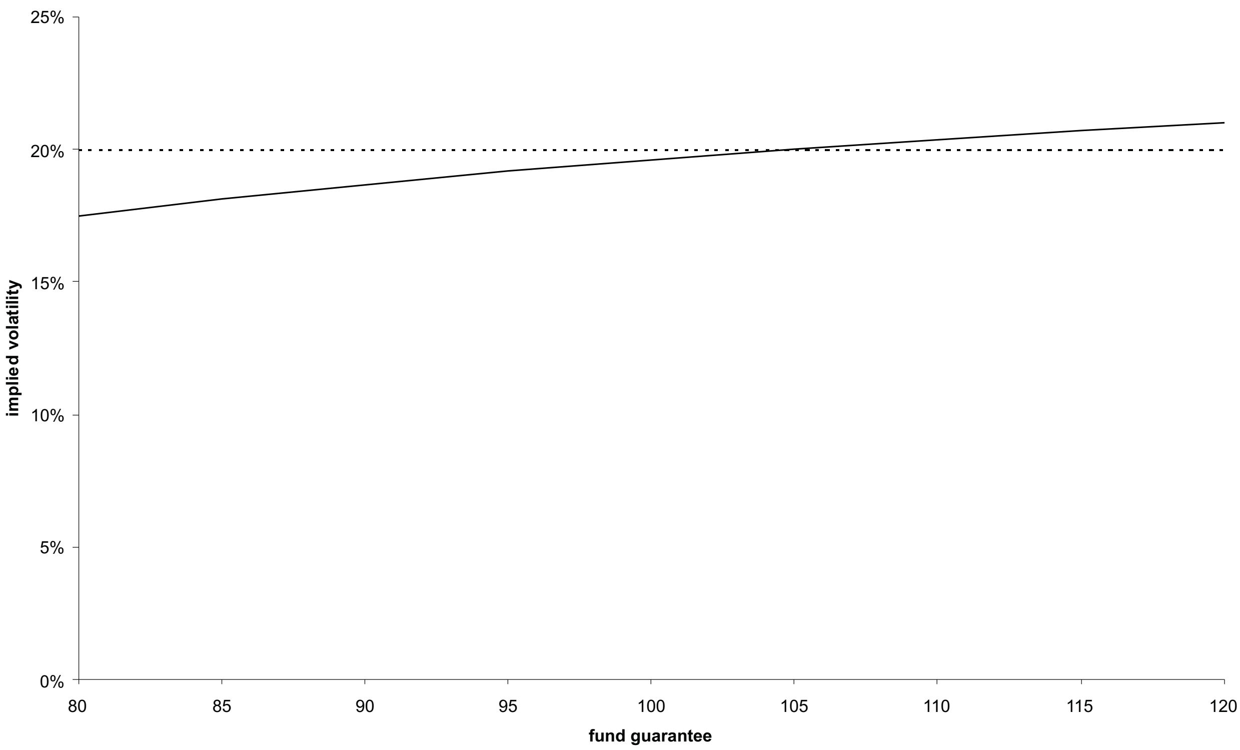

Figure 5.1. Equity implied volatility for static (dotted) and dynamic (solid) investment strategies

5.4.9 We can see that this gives an implied volatility that increases by strike, compared to the flat volatility (dotted line) that arises if we ignore dynamic actions.

5.4.10 The original Black-Scholes model may assume a flat implied volatility as in the dotted line, but we have already seen that real prices fail to respect this model. Indeed, market prices invariably show higher volatility for lower strikes. If we were to take this into account, then the dotted line would have a negative slope. The solid line, allowing for the dynamic strategy would be close to horizontal. In other words, the effect of dynamic investment management should be to flatten out the skew effect when applied to an invested portfolio.

5.4.11 We are aware that many offices are computing market consistent valuations using normal distribution models with constant volatility and, in addition, taking credit for the effect of dynamic actions to mitigate guarantee costs. This effectively double counts the management action. While we can understand the reasons for companies to double count in this way, the rationale is less clear for regulators or auditors to accept this double counting.

5.5 *Regular Premiums*

5.5.1 We return to the question of a regular (or sporadic) premium policy, which we assume is unit linked with a maturity guarantee, but can be extended to with-profits contracts by the approaches in Sections 5.3 and 5.4. Bacinello & Ortu (1993) provide a method for evaluating guarantees on regular premium unit-linked contracts.

5.5.2 Let us take first the case where all the premiums have already been collected. In that case, the structure is again path-independent, because the payoff is a guarantee on the total return between now and maturity. This is straightforward to value using your favourite option pricing model.

5.5.3 The situation is similar if the final premium is due tomorrow. In this case, the terminal value is effectively determined by the total return between tomorrow and maturity. This total return applies to the sum of the current asset share and the future premium.

5.5.4 The situation is similarly simple if the final premium is due the day before the policy matures. In this case, the final premium will be added to the unit value at the same time as the guarantee is applied. Equivalently, the premium is deducted from the future guaranteed amount.

5.5.5 The same idea can be applied for intermediate cash flows. The present value of each future premium is added to the asset share or subtracted from the discounted guarantee, in proportion to the timing of the premium between now and maturity. The final result is then processed using your favourite option pricing formula.

5.6 *Impact of Interest Risk on Equity Options*

5.6.1 A joint model of equity and interest rate markets can still produce guarantee costs consistent with the Black-Scholes formula when the interest rate factors follow a Normal distribution. However, the implied volatility now depends, not only on pure equity volatility, but also on interest rate volatility and correlations with equity markets.

5.6.2 To illustrate this point, let us take a Hull-White model for interest rates, equity movements and dividend yields. The formulae are then as in Section 4.5, but the processes W_t and X_t are now vectors rather than scalars. The corresponding parameters are now $\sigma_C, \sigma_E, \lambda_C, \lambda_E$, the subscripts relating to currency or equity respectively. These too are vectors, with the same number of components as W and X. For simplicity, we suppose that a common mean reversion parameter α applies to both equity and cash.

5.6.3 Using $\bullet$ to denote the scalar product of two vectors, the formulae are now as follows:

$$D_t^C = \frac{P_{0t}^C \exp[\lambda_C \bullet W_t + \sigma_C \bullet X_t]}{\mathbf{E}_0 \exp[\lambda_C \bullet W_t + \sigma_C \bullet X_t]}$$

$$P_{st}^C = \frac{P_{0t}^C}{P_{0s}^C} \exp\begin{pmatrix} -(1-e^{-\alpha(t-s)})\sigma_C \bullet (X_s - \frac{2\lambda_C}{\alpha}(1-e^{-\alpha s})\sigma_C) \\ +\frac{\sigma_C^2}{2\alpha}(1-e^{-2\alpha s})(1-e^{-2\alpha(st-s)}) \end{pmatrix}$$

$$D_t^E = \frac{P_{0t}^E \exp[\lambda_E \bullet W_t + \sigma_E \bullet X_t]}{\mathbf{E}_0 \exp[\lambda_E \bullet W_t + \sigma_E \bullet X_t]}$$

$$P_{st}^E = \frac{P_{0t}^E}{P_{0s}^E} \exp\begin{pmatrix} -(1-e^{-\alpha(t-s)})\sigma_E \bullet (X_s - \frac{2\lambda_E}{\alpha}(1-e^{-\alpha s})\sigma_E) \\ +\frac{\sigma_E^2}{2\alpha}(1-e^{-2\alpha s})(1-e^{-2\alpha(st-s)}) \end{pmatrix}.$$

5.6.4 For reasons explained by Smith & Speed (1998), the equity price index must be the ratio of the equity and cash deflators. For other purposes we are interested in the equity total return index, which is calculated from the cash index by the reinvestment of dividend yields:

$$\text{Price index} = \frac{P_{0t}^E \exp[\lambda_E \bullet W_t + \sigma_E \bullet X_t]}{\mathbf{E}_0 \exp[\lambda_E \bullet W_t + \sigma_E \bullet X_t]} \div \frac{P_{0t}^C \exp[\lambda_C \bullet W_t + \sigma_C \bullet X_t]}{\mathbf{E}_0 \exp[\lambda_C \bullet W_t + \sigma_C \bullet X_t]}$$
$$= \frac{P_{0t}^E \mathbf{E}_0 \exp[\lambda_C \bullet W_t + \sigma_C \bullet X_t]}{P_{ot}^C \mathbf{E}_0 \exp[\lambda_E \bullet W_t + \sigma_E \bullet X_t]} \exp\begin{bmatrix} (\lambda_E - \lambda_C) \bullet W_t \\ +(\sigma_E - \sigma_C) \bullet X_t \end{bmatrix}$$

$$\text{T R index} = \frac{\exp[(\lambda_E + \sigma_E) \bullet W_t]}{\exp[\frac{1}{2}(\lambda_E + \sigma_E)^2 t]} \div \frac{P_{0t}^C \exp[\lambda_C \bullet W_t + \sigma_C \bullet X_t]}{\mathbf{E}_0 \exp[\lambda_C \bullet W_t + \sigma_C \bullet X_t]}$$
$$= \frac{\mathbf{E}_0 \exp[\lambda_C \bullet W_t + \sigma_C \bullet X_t]}{P_{0t}^C \exp[\frac{1}{2}(\lambda_E + \sigma_E)^2 t]} \exp\begin{bmatrix} (\lambda_E - \lambda_C + \sigma_E) \bullet W_t \\ -\sigma_C \bullet X_t \end{bmatrix}.$$

5.6.5 Frequently sampled historical estimates of volatility are estimators of the following:

$$\text{Time series volatility} = \sqrt{(\lambda_E + \sigma_E - \lambda_C - \sigma_C)^2}.$$

5.6.6 By construction, the equity price index starts at 1. Let us consider an option with strike K and strike date t. An option on the equity capital index is priced using the Black-Scholes formula $g_{BS}(S_0, K_0, \eta\sqrt{t})$, with the following substitutions:

$$S_0 = P_{0t}^E$$

$$K_0 = P_{0t}^C K$$

$$\eta = \sqrt{(\lambda_E - \lambda_C)^2 + 2(\lambda_E - \lambda_C) \bullet (\sigma_E - \sigma_C)\frac{1 - e^{-\alpha t}}{\alpha t} + (\sigma_E - \sigma_c)^2 \frac{1 - e^{-2\alpha t}}{2\alpha t}}.$$

5.6.7 In the same way, we can price options on the total return index. We use the same formula, this time with:

$$S_0 = 1$$

$$K_0 = P_{0t}^C K$$

$$\eta = \sqrt{(\lambda_E - \lambda_C + \sigma_E)^2 - 2(\lambda_E - \lambda_C + \sigma_E) \bullet \sigma_C \frac{1 - e^{-\alpha t}}{\alpha t} + \sigma_c^2 \frac{1 - e^{-2\alpha t}}{2\alpha t}}.$$

5.6.8 This gives some insight into the reasons why the Black-Scholes formula has proved so robust, especially when the volatility inputs are calibrated to actual prices. Although the original derivation related to deterministic interest rates, there are many more advanced models where the same principles apply. The one effect which is not easily captured is variation in implied volatility by strike, that is, the smile effect.

5.6.9 These models are also useful for modelling returns on funds containing both equities and bonds. In this case, it is still straightforward to derive formulae for the volatility to be substituted into a Black-Scholes formula. For example, a fund with rebalanced proportion π in a constant maturity bond index of outstanding term τ, and the remainder in equities, the total return index is follows. From this, it is straightforward to develop expressions for the appropriate implied volatility for fund guarantees.

$$\text{T R index} = \frac{\mathbf{E}_0 \exp[\lambda_C \bullet W_t + \sigma_C \bullet X_t]}{P^C_{0t} \exp\left[\frac{1}{2}\begin{pmatrix}\pi\lambda_C + \pi e^{-\alpha\tau}\sigma_C \\ (1-\pi)\lambda_E + (1-\pi)\sigma_E\end{pmatrix}^2 t\right]} \exp\begin{bmatrix}(1-\pi)(\lambda_E - \lambda_C + \sigma_E) \bullet W_t \\ +\pi e^{-\alpha\tau}\sigma_C \bullet W_t - \sigma_C \bullet X_t\end{bmatrix}.$$

5.6.10 The use of joint interest rate and equity models generates complexities in sensitivity testing. A model may have an underlying equity time series volatility parameter, but implied volatilities will also depend on interest rates and on their correlation, and for this reason will necessarily vary by term.

5.6.11 Sensitivity testing is now a difficult problem. The most logical internally consistent test is to consider perturbations in the internal model volatility parameter, but this cannot be specified consistently for different users with different models. A constant shift in implied volatilities at all terms is easier to specify, but may not be consistent with any volatility shape obtainable from the original model.

5.6.12 The notion of long-term equity volatility is also problematic. The value is not unique. We have three different volatility definitions, and all tend to different limits for large t. As we shall see in Section 6, implied volatilities also vary by option strike, and there is no empirical or theoretical reason to believe this effect disappears for long horizons. Disclosures of long-term volatility assumptions lose their meaning if a modeller can select their definition.

5.7 *Guaranteed Annuity Options (GAOs)*

5.7.1 Several papers, including van Bezooyen *et al.* (1998), Wilkie *et al.* (2003), Ballotta & Haberman (2003), give analytical formulae for pricing guaranteed annuity options.

5.7.2 It is perhaps less well known that these formulae can be adapted when the underlying fund size is subject to guarantees. Formulae exist both for simple guarantees and also for some of the path dependent structures such as arise in with-profits business. Example derivatives are described in more detail by Geske (1979).

5.7.3 We consider a simple example. Let us suppose that we have a unit-linked fund with an uncertain maturity value S subject to a minimum guarantee K. The policy maturity value is therefore $\max\{S, K\}$. On maturity of the product the policyholder can either take the money or can buy a annuity at a guaranteed price. We use G to denote the guaranteed price of an annuity paying 1 per annum.

5.7.4 From the perspective of a rational policyholder, the decision is as follows. On retirement, they compare the guaranteed annuity price to the market price M of an annuity of £1 per annum. If $M < G$ then the market annuity is cheaper, and the benefit is simply the policy maturity value $\max\{S, K\}$. On the other hand, if $M > G$ then the right thing to do is to purchase the guaranteed annuity for the price G, giving an annual income of $\max\{S, K\}/G$. The price of buying this annuity in the annuity market would be $\max\{S, K\} * (M/G)$. Taking these together, the overall benefit is $\max\{S, K\} * \max\{1, M/G\}$.

5.7.5 Defining the benefit payable involves more assumptions than appears at first sight. The market consistent value of an annuity will not in general be the same as the premium for a purchased annuity. That is because quoted premiums may contain allowances for cost of capital, allowance for profit and other miscellaneous items. However, from the policyholder perspective, the benefit is defined in terms of annuity prices M and G, and not the market consistent valuation of an annuity. Arguably, the formula $\max\{S, K\} * (M/G)$ involves the office implicitly setting up a provision for its own capital costs and profit margin, in addition to the pure product liability. If the office strips out the margins in M, then the stated GAO liability is significantly reduced. This effectively assumes that other offices cut their margins to zero (including capital costs) when pricing annuities for the open market option.

5.7.6 We now give a suitable option pricing formula for this benefit structure. For simplicity, we ignore expenses and pre-retirement mortality. Let:

S_0 = market consistent value of the maturing unit fund (less charges);
K_0 = present value of minimum fund guarantee, discounted at the risk free rate;
M_0 = present value of a deferred annuity paying 1 per annum from the date of retirement;
G_0 = present value of a payment of G for certain at retirement date;
T = retirement date;
σ_S = annualised volatility for unit fund;
σ_M = annualised volatility of market annuity price; and
ρ = correlation of unit fund and market annuity.

5.7.7 The parameter σ_M should include the volatility of mortality expectations, not only in interest rates. It could be argued that mortality risk is non-systematic, but it still affects the value of options. To disregard mortality risk is to understate the value of the guaranteed annuity option.

5.7.8 The following pricing formula can then be used. This formula applies to the whole benefit, not just the guarantee:

$$K_0\Phi_2\left[\frac{\ln(K_0/S_0)}{\sigma_S\sqrt{T}}+\frac{\sigma_S\sqrt{T}}{2},\frac{\ln(G_0/M_0)}{\sigma_M\sqrt{T}}+\frac{\sigma_M\sqrt{T}}{2},\rho\right]$$
$$+S_0\Phi_2\left[\frac{\ln(S_0/K_0)}{\sigma_S\sqrt{T}}+\frac{\sigma_S\sqrt{T}}{2},\frac{\ln(G_0/M_0)}{\sigma_M\sqrt{T}}+\frac{(\sigma_M-\rho\sigma_S)\sqrt{T}}{2},-\rho\right]$$
$$+\frac{K_0M_0}{G_0}\Phi_2\left[\frac{\ln(K_0/S_0)}{\sigma_S\sqrt{T}}+\frac{(\sigma_S-\rho\sigma_M)\sqrt{T}}{2},\frac{\ln(M_0/G_0)}{\sigma_M\sqrt{T}}+\frac{\sigma_M\sqrt{T}}{2},-\rho\right]$$
$$+\frac{S_0M_0}{G_0}\exp(\rho\sigma_S\sigma_M)$$
$$\times\Phi_2\left[\frac{\ln(S_0/K_0)}{\sigma_S\sqrt{T}}+\frac{(\sigma_S+\rho\sigma_M)\sqrt{T}}{2},\frac{\ln(M_0/G_0)}{\sigma_M\sqrt{T}}+\frac{(\sigma_M+\rho\sigma_S)\sqrt{T}}{2},\rho\right].$$

If K_0 is zero in this formula, we recover the formula for a guaranteed annuity option on a unit linked fund given in van Bezooyen *et al.* (1998).

5.7.9 The function $\Phi_2(a, b, \rho)$ denotes the standard cumulative bivariate normal distribution. This is the probability that $X \leq a$ and that $Y \leq b$ when (X, Y) has a bivariate Normal (0, 1) distribution with correlation ρ. We observe that there is no factorisation here into a GAO portion and a maturity guarantee portion. Separate assessment of these two guarantees may be a first step, but the guarantees interact, so must be considered together within a single model.

5.8 *Smoothing*

5.8.1 The cost of smoothing benefits paid to with-profits policyholders is discussed in both Hibbert & Turnbull (2003) and Dullaway & Needleman (2004). Hibbert & Turnbull point out that smoothing has an economic cost. This can be seen easily in a contract which provides smoothed benefits, but no guarantees. If smoothing is unbiased, that is on average smoothed benefits equate to unsmoothed asset shares, the value of the smoothed benefits will still exceed that of unsmoothed benefits. The Capital Asset Pricing Model would assign a lower discount rate to the smoothed benefits, thus resulting in a higher value. Alternatively, if the smoothed benefits are valued by deflators, higher weights will be applied in adverse conditions, when smoothed benefits would exceed the unsmoothed asset shares.

5.8.2 Dullaway & Needleman comment that the cost (or benefit) of smoothing is not always intuitive. They provide an example where the cost of smoothing reduces when the period of smoothing is extended. That result is likely to be dependent on the smoothing formula used, and may also depend on the level of guarantees. For example, a formula for smoothing incorporating geometric averaging is more likely to provide a benefit (i.e. negative cost), simply because geometric means are lower than arithmetic means. In the presence of guarantees, the additional cost of smoothing would usually be negative, if in adverse conditions when smoothing would provide a benefit to policyholders, the guarantees bite.

5.8.3 In current conditions, when many with-profits offices are paying in excess of asset shares, the major component of the cost of smoothing is likely to be the glide-path until benefits are broadly in line with long-term target payout ratios.

5.8.4 The value of smoothing can be analysed as follows. Let:
— A_T be the unsmoothed asset share at time T;
— S_T be the smoothed asset share at time T, used to determine benefits; and
— G_T be the guaranteed benefits at time T.

Then the liability is the value of Max $[S_T, G_T]$, which can be written as:

$$\text{value } A_T + \text{value Max } [0, G_T - A_T] + \{\text{value Max } [S_T, G_T] - \text{value Max } [A_T, G_T]\}$$

which is the unsmoothed asset share, plus the cost of the guarantee, plus the cost of smoothing. The first term is known and the second is the cost of the guarantee, which is capable of being evaluated using one of the formulae described earlier. The last term is the cost of smoothing, which can be expanded to give:

$$\text{value } S_T + \text{value Max } [0, G_T - S_T] - \text{value } A_T - \text{value Max } [0, G_T - A_T].$$

Again, the third term is known, and the last term can be obtained as above. The second term is another put option on the series S rather than A, and in principle can be evaluated using similar formulae, either with reduced volatility, or potentially through the use of Asian option pricing. (Asian options have payouts averaged over a period.) S is likely to be function of the series A, possibly over several years. The first term, value S_T, could therefore be evaluated as a derivative of A, or, as discussed in ¶5.8.1, by discounting at an adjusted lower rate.

5.9 *Decrement and Take Up Rates*

5.9.1 Allowance needs to be made for decrements — mortality, lapse, paid up and surrender rates — as well as take up rates (e.g. for guaranteed annuity options) in determining market consistent valuations. The uncertainty of these decrements in the future gives rise to additional risks. These risks cannot in general be hedged, though reinsurance is available for mortality risk, and hence market prices for these risks cannot be observed. Some of these risks are non-systematic, that is that they are unrelated to market conditions, and can be allowed for by making best estimates.

5.9.2 Some decrements and take up rates are, or could be expected to be, dependent on market conditions. For example, the take up rate of guaranteed annuity options might depend on interest rates, and surrender rates on MVR free dates on unitised with-profits business might depend on the worth of the guarantee, which, in turn, will depend on asset values. Past experience of these decrement rates is unlikely to provide a good guide to the future, especially if guarantees have not been in the money in the past, but are now.

5.9.3 It is possible to apply formulae to allow for market contingent decrements by combining options at different strikes, with weights dependent on the exercise proportions at each strike. For example, take up rates on guaranteed annuity options could be estimated at a range of interest rates, and then applied to formulae for swaptions and digital swaptions.

5.9.4 To illustrate this approach, consider the general example in Section 4.1, but now suppose that we have a take up rate for the put option of 100% if the guarantee is at least 20% in the money, but only a 50% take up rate if less than 10% in the money, with interpolations in between. The value of the guarantee with these take up rates can then be expressed as the value of:

$$0.5\text{Max}[0, K - S] + \text{Max}[0, 0.9K - S] - 0.05\text{Max}[0, 0.8K - S]$$

which can be found using option pricing formulae.

5.10 *Further use of Exotic Option*

In this section we have given examples of market consistent valuation problems to which analytical solutions exist. Other relevant solutions are given in Briys *et al.* (1998) and Bouwknegt & Pelsser (2001). We have only scratched the surface of what is possible.

6. Volatility — Historical and Implied

6.1 *Historical Volatility*

6.1.1 The value of a guarantee depends on the estimated probability distribution of future asset returns, and the distribution of future interest rates in some cases. The most important parameter affecting guarantee values, as for option prices, is the volatility of asset returns.

6.1.2 Volatility can be measured in two ways — historical and implied. Historical volatility is a statistic sampled from historic data. It is defined as the sample standard deviation of log returns $\log(R)$ over a fixed holding period, divided by the square root of the holding period length. It varies by asset class, accounting currency, data window, sampling frequency and holding period. For example, our data window might be a period of ten years, within which we measure volatility from weekly data (i.e. sampling frequency is weekly) with overlapping three-month holding periods. The resulting historical volatility is dependent on all three choices. If we expect the future to be similar to the past, then it might be appropriate to calibrate a stochastic asset model to historical volatilities (and other characteristics of the distributions of returns, including correlations between different assets).

6.1.3 From the data series, we can identify a historical distribution. It is usually more convenient to identify the so-called *characteristic exponent*. This is a function $\lambda(p)$ defined by:

$$\mathbf{E}[R^p] = \exp[t\lambda(p)]$$

where R is the holding period return factor (i.e. the total return up to time t). The left hand side is most naturally estimated by a sample historic average. The volatility is then the square root of the second derivative at zero:

$$\sigma_{HIST} = \sqrt{\lambda''(0)}.$$

6.2 *Implied Volatility*

6.2.1 Implied volatility is that derived from market prices of instruments such as options whose price depends on volatility, assuming a certain pricing model. For example, implied equity volatilities are usually derived from the standard Black-Scholes option pricing formula. Implied volatility will therefore vary by term and strike of the option. As there is a one-to-one correspondence between volatility and price using this formula, the prices of options are often quoted in terms of implied volatility.

6.2.2 Theoretical models of option prices also exist. The theoretical restrictions can be expressed in terms of the characteristic exponent $\lambda(p)$. For example, the Black-Scholes requires normal distributions, when $\lambda(p)$ is quadratic in p. More general fat tailed distributions can be generated from Lévy process models which simply require that $\lambda(p)$ does not depend on the option maturity date.

6.2.3 Theoretical approaches to option pricing for such models are discussed in Bühlmann *et al.* (1998) and Smith (2003). Lévy process models produce implied volatilities which depend mostly on the annualised strike rate, that is on the option strike expressed as the current asset level growing at a fixed annual rate. Six-monthly data since 1978 generate the theoretical curve for option implied volatility, as shown in Figure 6.1.

6.2.4 In these calculations, we have used an Esccher transform of historical data, as advocated by Bühlmann *et al.* (1998). Saddle point approximations give fast algorithms for implied volatility. Our algorithm is based on ideas in Jensen (1995). Rogers & Zane (1999) evaluate the error in this class of approximations when applied to option pricing problems.

6.2.5 Our algorithm is as follows. Define a function $\zeta(p, q)$ by:

$$\zeta(p,q) = \begin{cases} \dfrac{\sqrt{2}\sqrt{\lambda(p) - \lambda(q) - (p-q)\lambda'(q)}}{|p-q|} & p \neq q \\ \sqrt{\lambda''(p)} & p = q. \end{cases}$$

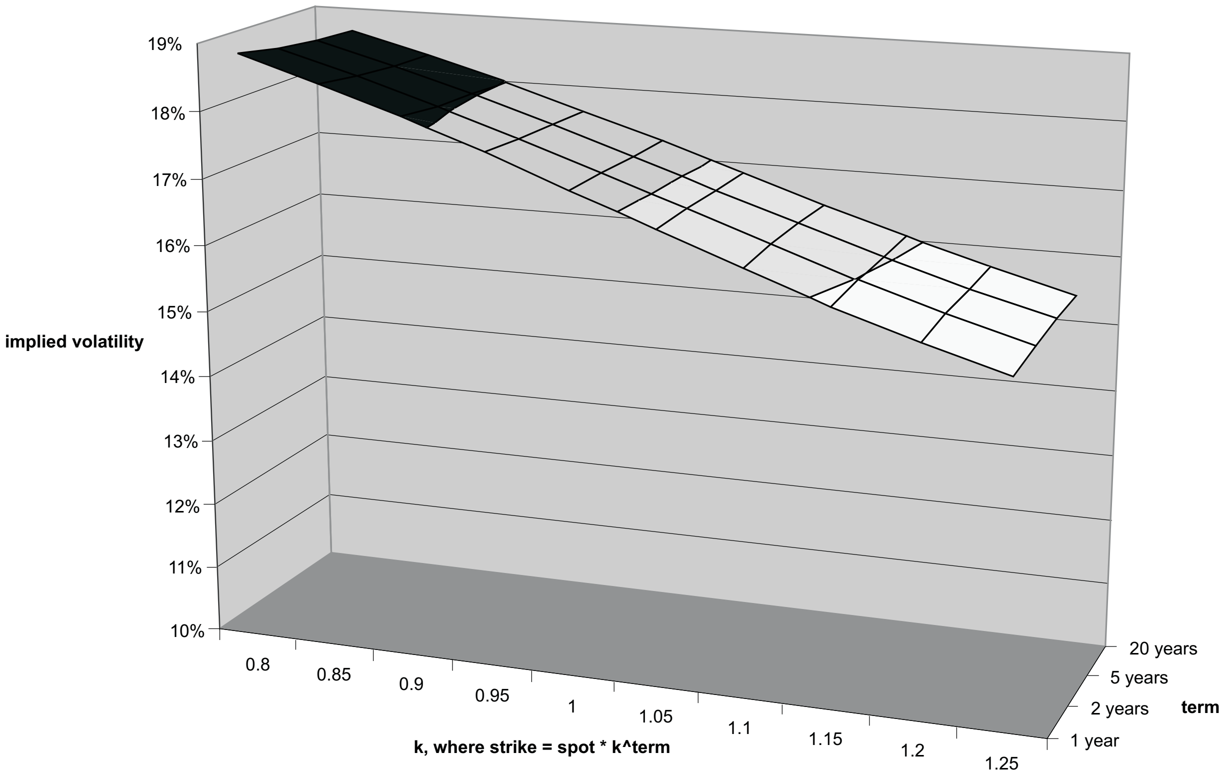

Figure 6.1. Theoretical implied volatilities based on historical six-month returns on the FTSE 100 index

6.2.6 Assuming the initial value of the index is 1, we consider options with term t. Let F denote the forward value of the index, and let K denote the strike of the option. We then find α and β, where:

$$\lambda(\alpha+1)-\lambda(\alpha)=\frac{1}{t}\ln(F)$$
$$\lambda'(\beta)=\frac{1}{t}\ln(K).$$

6.2.7 The option volatility is finally approximated by:

$$\begin{aligned}\eta &= (\beta-\alpha)\zeta(\alpha,\beta)-(\beta-\alpha-1)\zeta(\alpha+1,\beta)\\ &\quad+\frac{1}{t}\left(\frac{1}{(\beta-\alpha)\zeta(\alpha,\beta)}-\frac{1}{(\beta-\alpha-1)\zeta(\alpha+1,\beta)}\right)\\ &\quad\times\ln\left[(\beta-\alpha)\frac{\zeta(\beta,\beta)}{\zeta(\alpha+1,\beta)}-(\beta-\alpha-1)\frac{\zeta(\beta,\beta)}{\zeta(\alpha,\beta)}\right].\end{aligned}$$

6.2.8 These are the first two terms in a saddle point expansion. Higher order terms (of order t and above) have been omitted. It is instructive to consider the limit for large t. The leading term indicates that the variation in implied volatility by strike should still persist even at long horizons.

6.3 *Volatility Term Structure*

6.3.1 Next to the theoretical formulae we can show actual volatilities from option prices.

6.3.2 Implied volatilities are volatile over time, and are usually, but not always, higher than would be expected from historical volatilities. Figure 6.2 shows a history of implied volatilities for one-year options on the London exchange (LIFFE).

6.4 *Explaining Differences between Historical and Implied Volatility*

6.4.1 There is a debate as to whether historical or implied volatilities should be assumed, both for realistic balance sheets and capital assessment. Should the assumption depend, for example, on the approach to risk management? Might it be appropriate to use implied volatilities for balance sheets (or risks) that have been hedged by buying protection, but use historical volatilities for un-hedged balance sheets, or where dynamic hedging is being followed? These arguments seem to ignore the fundamental principle that the valuation of the liabilities should be independent of the

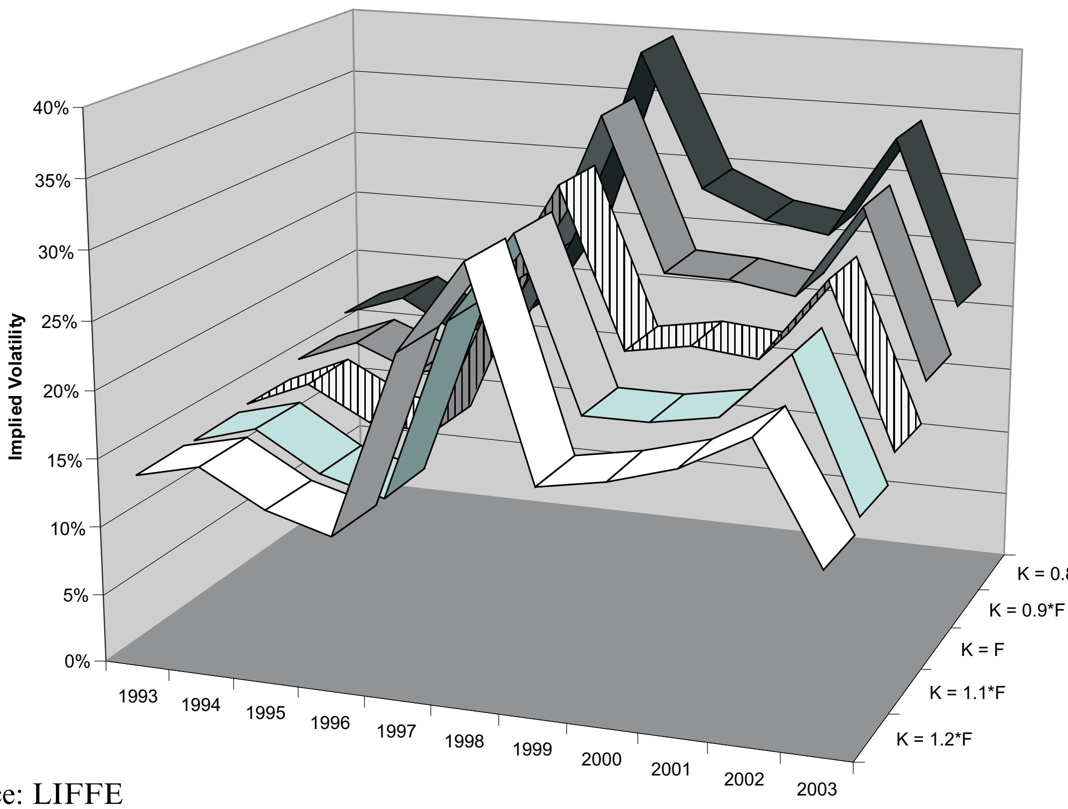

Source: LIFFE

Figure 6.2. Implied volatility of one-year FTSE 100 options

backing assets, at least to the extent that asset performance does not affect the liabilities. The degree of inconsistency that such an approach would introduce therefore seems unwarranted.

6.4.2 The argument regarding the type of model to use for the purpose of capital assessment seems less clear cut. Much depends on the philosophical approach to capital assessment. We simply observe that this is not the only subjective judgement in this field, the choice of equity risk premium and degree of mean reversion being two other examples where subjective assumptions can be expected to have significant impact on the results.

6.4.3 Market consistency would seem to demand the use of implied, rather than historical, volatilities. In practice, implied volatility is usually higher than historical volatility. A possible analysis of the difference between implied and historical volatility is as follows:

(a) The market's estimate of future volatility may well be different to that experienced in the past.
(b) Market prices, and hence implied volatilities, will reflect loadings for the costs of writing options, such as administration and sales costs, and including the expenses of preparing legal documentation and valuations.
(c) Banks will normally try to hedge the risks from writing options, and will therefore incur costs associated with their risk management, such as market spreads, taxation and margining.
(d) Option writers are placing their capital at risk, and will therefore wish to include a charge for cost of capital. That cost could be broken into various components, such as agency costs, taxation, loss of franchise value and financial distress.
(e) Taxation related directly to the writing of options.
(f) Implied volatility may vary with strike and term to reflect fat tailed distributions (see Section 6.2).
(g) The impact of interest rate volatility would be reflected in implied volatility (see ¶5.6.1).
(h) Historical volatility is an unconditional estimate, whereas implied volatility may be conditional on other factors.
(i) Profit margin.
(j) Allowance for the credit risk of the option writer (which would reduce the implied volatility).
(k) The data available from commercial sources may not relate to transactions.

6.4.4 A similar analysis could be made of, for example, annuity rates. Historical annuitant mortality is analysed by the Continuous Mortality Investigation Bureau. A company's estimate of future annuitant mortality will almost certainly be different, and other factors similar to those described in ¶6.4.3, such as expense loadings, are then reflected in the annuity rates actually quoted by the company.

6.4.5 An analysis of the differences between implied and historical volatilities might be of use when trying to decide whether or not to hedge a guarantee through the purchase of a suitable option, although this is likely to require estimates of the various components in a bank's pricing structure. It might be argued that some of the above components should be excluded when evaluating the realistic liabilities, or even that, for un-hedged balance sheets, the use of historical volatilities might be more appropriate. Such arguments would carry weight if we observed large bid offer spreads on option prices, or implied volatilities in bid prices below historical volatilities. In practice, though, we do not observe these features, and the use of mid market prices for calibration removes most of these objections to the use of implied volatilities. If, however, a company was following a dynamic hedging strategy, and specific allowance was being made for the costs of the hedging strategy elsewhere in the balance sheet, then there would be an element of double counting in using a model calibrated to implied volatilities that already made some allowance for these costs.

6.4.6 An argument sometimes made against the use of implied volatilities is that market prices of options, just as for other investments, are themselves volatile. Imbalance between supply and demand in thinly traded markets has caused volatility spikes, particularly in those long-dated derivatives of most relevance to hedging guarantees in life funds. Volatility is though a risk that insurers need to manage. In a hedged fund, changes in volatility will not matter, and in an un-hedged fund the risk needs to be evaluated and managed. While there is no specific volatility stress test required under Pillar I of the new proposed regulatory regime, it is a factor that companies will need to take into account in their internal capital assessments as part of Pillar II.

7. Corporate Valuation and Profit Testing

7.1 *Value of a Financial Firm*

7.1.1 It falls to accountants or regulators to define an insurer's assets and liabilities. This definition establishes criteria for *recognition*: which future cash flows are to be included as assets or liabilities. Accounting standards also specify *measurement*: how future cash flows, once recognised, are to be converted into a present value.

7.1.2 There are very many accounting standards. None has deliberately set out to be unrealistic or to misrepresent economic reality. Differences between accounting standards often reflect genuinely difficult economic questions, such as ownership of future premiums, treatment of policyholder options, standardisation of decrement assumptions and the consolidation of segregated funds.

7.1.3 Having chosen an accounting framework, a firm can calculate its assets and liabilities, and therefore its equity or net assets. However, the market capitalisation of an insurer is different, and usually higher, than its net assets. We refer to the difference as *franchise value*. Therefore, by definition of franchise value the following applies:

$$\text{Market capitalisation} = \text{Accounting net assets} + \text{Franchise value.}$$

7.1.4 If accounting measurement were on a truly market consistent basis, we would still not expect net assets to equal market capitalisation, that is, we would not expect franchise value to be zero. Accounting recognition criteria inevitably exclude some cash flows which the market takes into account when assessing a firm. The excluded cash flows therefore contribute to franchise value. For example, profits associated with future customers are usually excluded from accounting statements, but still contribute to market capitalisation.

7.1.5 We have defined franchise value as a balancing item — the difference between market capitalisation and net assets. As both assets and liabilities are affected by accounting standards, then so will be the franchise value. Measuring assets and liabilities on a market consistent basis is a step forward, but this still does not define which cash flows are to be counted towards an asset or liability valuation for accounting purposes.

7.1.6 The market capitalisation is, by definition, a market consistent assessment of an insurer's future payments to shareholders. If an accounting standard also specifies market consistent measurement criteria, then the cash flows recognised by that standard will be measured on a consistent basis to the business as a whole. In this case, the franchise value may also be deemed market consistent. It represents the value precisely of those cash flows excluded by the accounting recognition criteria.

7.2 *Constituents of Franchise Value*

7.2.1 Life insurance valuations, for example, in mergers and acquisitions, have traditionally been decomposed into three parts — statutory net assets, value of in-force business and goodwill.

7.2.2 The calculation of value of in-force was designed to overcome known conservative margins in the calculation of statutory net assets. If asset and liability bases are market consistent, there is no longer a need for a separate valuation of in-force business.

7.2.3 The idea that the value of a company is the present value of its policy margins relies on assumptions of frictionless institutions and corporate transparency. In practice, there are various frictional costs which reduce the market values of companies.

7.2.4 The three chief areas of frictional costs to be considered are (see Ng & Varnell, 2003) are:

(1) tax payable, including the double taxation effect of assets held in shareholder funds;
(2) agency costs, that is, costs arising from misalignment of interests between shareholders and managers, which might be modelled as a small percentage of net assets; and
(3) financial distress costs, which combine two elements: first the limited liability put option of shareholders to default on corporate obligations; and secondly, and more significant in most cases however, the loss of franchise value in times of financial distress.

7.2.5 One consequence of this approach is the existence of an optimal level of capital. Having too much capital creates agency costs, while having too little increases financial distress costs. Moran, Smith & Walczak (2003) discuss this issue in more detail.

7.3 *Value Impact of Decisions*

7.3.1 The corporate valuation framework has an impact on how profit tests are carried out. In addition to the cash flows of a product, a financial firm should consider the marginal impact of that product on agency and financial distress costs.

7.3.2 Both of these elements are related to the risk of a product in the context of corporate risks. If a product increases corporate risk, it of course increases financial distress costs. It is also likely to increase the optimal amount of capital held, and hence agency costs (while financial distress costs come down a little). These frictional costs are the means for turning a measure of cash flow variability into a measure of shareholder cost. See, for example, Hancock *et al.* (2001) for more details of this approach.

7.3.3 Although banks buy and sell options to other corporate buyers, the dominant trades consist of a corporate buying options or guarantees from a bank. The banking industry in net terms has an exposure to large market moves because of the number of one-way bets the industry has offered to its clients. Therefore, a marginal extra option bought by an insurer is likely to increase the financial distress costs of the bank. Conversely, a bank which buys an option probably reduces its financial distress costs.

7.3.4 As all banks are in a similar position, we can expect that the market price of options reflects the marginal financial distress costs of the issuer as well as the cash flows promised under the option itself. This is the primary reason for option implied volatilities exceeding historical time series volatility estimates.

7.3.5 This type of analysis provides clues as to how options not currently traded, such as those on property, might be valued. Given the diversity of property portfolios, and the absence of futures contracts on property, we would expect the costs of hedging property options to be extremely high. That would be reflected in high option prices. It would not therefore be sensible to obtain an estimate of implied property volatility simply be scaling up historical volatility by the ratio of implied to historical volatility on equities.

7.3.6 In future, we might then hope to see a better class of market consistent asset models, whose volatility forecasts are historically plausible, and where option prices include an element of financial distress costs in addition to the Black-Scholes element. An insurer valuing its own liability could then take account of their own financial distress costs rather than the third party costs implied in market option prices.

8. Comparisons to Monte Carlo Results and Projections

8.1 *The Monte Carlo Approach*

8.1.1 Some offices have prepared realistic balance sheets for with-profits funds using the Monte Carlo simulation approach, and this appears to be the method preferred by the FSA. The main advantage of this method over the formula-based approach is that it permits more accurate modelling of the cash flows. In principle, any cash flow can be valued using this approach, provided that it can be written down mathematically, and the cash flows can be made dependent on a wide variety of parameters. The valuation process then relies on stochastic simulations of investment returns and either the risk neutral or deflator framework in order to price the cash flows. As any builder of such a model will testify, the simulation approach requires little thought, but a considerable amount of effort and time spent in its construction as well as patience in waiting for it to run.

8.1.2 Another advantage of the Monte Carlo method is the ability to allow for the impact of future management decisions at the fund level, policyholder actions on the cash flows, and the links between those actions and economic conditions. Most simulation models will typically allow for dynamic bonus and investment strategies, and will incorporate rules adopted for smoothing final bonuses. See, for example Hibbert & Turnbull (2003) for a description of such strategies. Some models also allow for certain changes in policyholders' behaviour in response to changing economic conditions.

8.1.3 Most valuations, whether performed to support strategic decisions or for internal or regulatory reporting, will require at least some stress and scenario testing. Users of realistic valuations will want to understand the key sensitivities. Here the formula-based approach starts to have the advantage; it is far easier to change a parameter in a formula than rerun a set of simulations. For example, suppose one wants to examine the effect of a change in volatility of asset returns. In most option pricing formulae, this is a single parameter. When running simulations, it will be necessary, not only to rerun the model, but also to adjust the simulations to fit the revised volatility level.

8.1.4 The following are widely believed to be advantages of the Monte Carlo approach:

— Compliance with regulatory requirements for a dynamic stochastic model.
— The ability to model in detail a company's own internal practices for bonus and investment decisions at the fund level.
— Simulation models are easier to explain. Non-technicians readily grasp the idea of running several randomly generated possible outcomes.
— Scalability. The same asset model simulations can be used for each set of liabilities. There is no need to modify the asset model when new liability products are developed.
— Simulations which produce large guarantee costs can be investigated and checked individually. If a rule is not responding as intended to market moves, then the rule can be adjusted and the whole process re-run.
— Testing alternative decision rules and their sensitivity to changes in parameters.
— The ability to isolate scenarios which generate large capital requirements, refining assumed management actions in those scenarios if appropriate.

8.2 *The Closed Form Alternative*

8.2.1 Our closed form solutions can reflect projections which are both stochastic and dynamic. It happens that, with suitable specification, certain dynamic stochastic models have closed form solutions.

8.2.2 There is a limitation to the strategies that can easily be modelled using closed form solutions. For example, we are not aware of closed form solutions for with-profits policies when the bonus strategy is defined using bonus reserve valuations, or where investment mix depends on the free asset ratio.

8.2.3 In practice though, actual bonus and charging strategies have a habit of diverging from model assumptions, however carefully the latter are crafted. There is much to be said for a simple approach to the dynamic part of the model, particularly if this results in valuations which can be computed quickly in closed form.

8.2.4 Until relatively recently, U.K. life offices had high equity backing ratios and fairly high rates of reversionary bonus, and therefore considerable discretion to reduce both in adverse conditions. The value of this discretion is well illustrated in Hibbert & Turnbull (2003), who contrast the shape of the embedded options in with-profits business, both with and without allowance for management actions. However, their example, along with most others, starts with both a high equity backing ratio and high reversionary bonus rates (at least by current standards), so unsurprisingly the option that management has in adverse conditions to reduce both the volatility of asset returns and bonus rates has significant value. Now bonus rates and many funds' equity contents are significantly lower, so the impact of management discretion on market consistent values of liabilities is likely to be much reduced. The benefits of Monte Carlo modelling may therefore not be so great in current conditions, compared with the simpler formula methods.

8.2.5 It is unrealistic to expect a simulation model to fit a large array of option prices, taking into account skew and smile effects, across a wide range of durations. At best, the calibration process can result only in an approximate fit. A formula-based approach, on the other hand, would permit the use of a volatility matrix corresponding to the particular patterns observed in the market.

8.2.6 The closed form approach then provides many advantages, including the following:

— compliance with regulatory requirements for a dynamic stochastic model;
— fast computation, avoiding the need for simulations and also avoiding the sampling error inherent in simulation models;
— transparent calculation trail from assumptions to results, for example from volatility and correlation assumptions for individual asset classes through to fund volatilities;
— easy calculation of sensitivities to key assumptions and re-runs under alternative assumption sets;
— lower risk of implementation errors, as the software runs quickly, and the closed form solutions are tested numerically to ensure compliance with the appropriate Black-Scholes partial differential equation and boundary conditions;
— the closed form solutions split into several terms with intuitive interpretations, enabling the user to separate different components of the liability value;
— the ability to capture term and strike structure of volatilities, by calibrating each model point to appropriate matching derivative prices of similar nature and term;
— easy calculation of market risk sensitivities and identification of appropriate option hedging strategies;
— management actions may depend on the overall financial state of the fund — if that is measured in terms of realistic liabilities, then a formula approach to calculating those liabilities is necessary for practical programming; and
— when the office implements a stochastic model for capital adequacy, they can use the closed form solution to model future realistic balance sheets on a simulation-by-simulation basis.

8.2.7 While we regard the closed form alternative as a very useful, and perhaps underutilised, approach, we would not go so far as to advocate that its use should replace the Monte Carlo simulation method. Rather, there is merit in combining the two approaches. Most companies using the Monte Carlo approach will use some formulae to check the results. Conversely, Monte Carlo simulations could be used to calibrate parameters for the more complex formulae that we have discussed in Section 5. The Monte Carlo approach could also be used to derive and test fund level assumptions and decisions for use in closed form solutions. The difficult step is that of specifying rules for investment strategy, bonus policy and charging for guarantees. Once these are specified, the rules could be implemented either by Monte Carlo simulation or by closed form solutions.

8.3 *The Projection Problem*

8.3.1 Financial condition reports have commonly involved the projection of a fund and its balance sheet over a number of years into the future. The main purpose of these projections has been to examine the future solvency of the fund in a variety of economic and operating conditions. Projections have been performed both by deterministic and stochastic methods. The regulatory balance sheet, constructed in accordance with the current valuation regulations, usually forms the basis for this type of investigation, and since it is formula-based lends itself relatively easily to projection.

8.3.2 For non-profit and much unit-linked business, it would similarly be fairly straightforward to project a realistic or market consistent balance sheet. However, for with-profits business, or business with guarantees or non-linear cash flows, the projection becomes much harder. Companies adopting the Monte Carlo simulation approach face the prospect of handling nested simulations, if they wish to study the evolution of their funds through stochastic methods. If, for example, an office uses 1,000 simulations to determine its realistic liabilities, and wishes to project the balance sheet to a single date in the future under stochastic simulation (in order perhaps to assess capital adequacy to meet the balance sheet at that date), it could require 1,000,000 simulations. This number of simulations for a representative model of a U.K. with-profits fund seems beyond computing power for the foreseeable future, assuming that you wish to see the results before the next FCR is due! Extending that exercise to a range of future dates increases further the number of simulations required.

8.3.3 There seem to be two practical alternatives: either project on a limited number of scenarios, reserving the use of simulations to determine the project market consistent balance sheet; or calculate the projected balance sheet by using formulae, allowing simulations to be used for the projection. The first alternative may be satisfactory for projection to a single date in the future, but is unlikely to be practical if it desired to project the balance sheet to a number of future dates. On balance, we prefer the second of these two alternatives.

8.3.4 Another possible approach to projection involves the construction of a branching process, which attempts to capture the characteristics of nested simulations through a limited number of branches (see Dempster *et al.*, 2003).

8.3.5 The use of a formula-based approach to determining the balance sheets would seem to be almost essential for projection purposes. The challenge is to find suitable formulae for the values of complex cash flows across the full range of economic conditions and consequent states of the fund. Implementation could require specifying the dependence of the parameters in the formulae on those conditions and states.

9. Hedging the Risks in With-Profits Funds

9.1 *The Pooling Principle*

9.1.1 Most with-profits funds in the U.K. have been run as managed funds, with benefits determined by asset shares and a smoothing process. The asset mix of the fund will vary over time, but unlike their unit-linked counterparts, the asset mix in with-profits funds is determined not solely by investment considerations. Solvency constraints and the impact of guaranteed benefits also play a part. It is usual for the asset shares of all policies in a fund to be credited with the same rate of return on assets, and that is often the rate of return earned by the whole fund, irrespective of the type of policy or its duration. One exception is where non-profit business is written in a with-profits fund, when fixed interest assets would usually be allocated to that business, with the remaining assets allocated to the with-profits policies. Some companies allocate a separate block of assets to their estate. The principle that all policyholders in a with-profits fund share equally in its risks and rewards, including the rate of return on investment, is a fundamental one. Ransom & Headdon (1989) provide a clear introduction to the operation of a with-profits fund.

9.1.2 This pooling principle has worked reasonably well in the past, when there were only few types of product with similar characteristics and levels of guaranteed benefits were relatively modest. However, the diversity of products now offered and the increasing value of guarantees, caused by overgenerous reversionary bonus rates, lower interest rates and falling equity values, have placed an increasing strain on this principle. Some companies have started to tackle this problem by allocating different asset mixes to different classes of policy, but few if any have sought to differentiate by duration to maturity.

9.1.3 With-profits funds hold substantial amounts of fixed interest stock. For non-profit business, annuities and guaranteed bonds in particular, the duration of the fixed interest investments are chosen to provide a close match to that of the liabilities. The same principle is often applied to the fixed interest assets in a with-profits fund, and they are usually selected to provide a reasonable match by duration to the guaranteed liabilities. The aim is to achieve some degree of matching between the asset cash flow and the benefits paid to policyholders. Unfortunately, this aim is defeated by the pooling principle of crediting the same rate of return, that on the whole fund, to all asset shares.

9.1.4 The problem can easily be illustrated by considering two policies maturing at different dates in the future, say in five and ten years' time. From a financial risk management perspective, one would like to be able to credit the return on a five-year bond to the five-year policy and that from the ten-year bond to the ten-year policy, rather than the average weighted return from the bonds to both policies. Fluctuations in interest rates will not jeopardise the intention of the first method of allocation, but crediting the average rate of return to both will frustrate the management of the guaranteed benefits when interest rates change. The alternative way of trying to manage the guarantees through reversionary bonuses suffers from the similar problem that reversionary bonus rates do not usually vary by term. In extremely adverse conditions, when the guaranteed benefits are biting for much of the time, matching by term will, of course, help in just the same way as it does for non-profit business, but in the majority of conditions it is of only limited use.

9.1.5 Companies could amend current practices of allocating investment returns to asset shares in order to improve the management of guarantees, but in making such a change PRE would need to be carefully considered, and legal opinions may be required. There would also be implications for systems.

9.2 *Guaranteed Annuity Options and Swaptions*

9.2.1 There has been much activity in recent years in hedging guaranteed annuity options with swaption contracts, but relatively little has been done to hedge the underlying sum assured and reversionary bonus guarantees provided in all with-profits business. That may be because of the high profile guaranteed annuity options have received as well as the greater liquidity in the interest rate options market than in long-dated equity options.

9.2.2 The swaption contracts purchased by companies for the purpose of hedging guaranteed annuity options have been held within the with-profits funds. While we appreciate that there is little alternative for mutual insurers, for proprietary companies consideration needs to be given to whether the options should be held in either the long-term business fund or in the shareholders' fund, and whether the risks being hedged relate to policyholders or shareholders. The latter distinction is not obvious in the traditional 90:10 fund.

9.2.3 The introduction of realistic balance sheets and their use in the management of with-profits funds has led to a convenient segregation of the assets of the fund into those backing the asset shares and those hypothecated to meeting the cost of guarantees and smoothing. Since the purpose of the swaptions is to hedge the guarantee costs, it is logical that they do not form part of the asset mix backing the asset shares. Crediting the return from the swaption contracts to the asset shares would defeat the object. A fall in interest rates would increase both the value of the guaranteed annuity options and the swaptions, and to credit the high return from the swaptions in those circumstances would only inflate the cost of the guarantees. It may though be possible to hold the swaptions in the assets backing the asset shares, but to credit to asset shares only a notional return from the swaptions, not reflecting any variation in interest rates. Such an approach may be necessary in a weak fund.

9.2.4 We deduce from the approach taken by offices to the use of swaptions that it does not appear necessary to credit to policyholders the return on the entire fund, though we appreciate that the legality of this approach may depend on past practice and an office's constitution, as well as being governed by PRE.

9.3 *Hedging other Guarantees*

9.3.1 Hibbert & Turnbull (2003) illustrate the use of put options to hedge the sum assured and reversionary bonus guarantees in with-profits business. While their illustration relies on long dated options on the asset share mix being available, the approach of delta and vega hedging (see e.g. Hull (2003)) can, in principle, be applied separately to the constituents of the asset mix backing the asset shares.

9.3.2 If put options are purchased to hedge the guarantees, their location in the company may depend on the design of the product with the guarantee and the operation of the with-profits fund. While modern with-profits contracts may have a specific charge for guarantees, with the guarantees being met by the shareholders, who would then hold the option hedge, for traditional with-profits funds and contracts, the options are likely to be held in the with-profits fund.

9.3.3 The consideration of whether the options should form part of the assets backing the asset shares is rather different to the situation above concerning guaranteed annuity options and swaptions. A fall in the value of assets backing the asset shares would lead to an increase in the value of the guarantees. There would also be an increase in the value of the put options. Consequently, including the put options in the asset share mix and allocating the return from the options to the asset shares seems reasonable, unlike the example above of the swaptions. However, such an approach suffers from the same drawback as matching the guaranteed benefits by fixed interest securities of similar duration, discussed in Section 9.1. The return credited to an asset share in respect of the put option would ideally need to be related to an option matching the outstanding duration of the policy. That leads to the conclusion that the option hedge should be held outside the asset shares, as indicated by the realistic balance sheet.

9.4 *Charging for Guarantees*

9.4.1 Once guarantees have been quantified, a decision of whether and how to charge for the guarantees needs to be made. Past practice of ignoring these costs and assuming that they are met through the management of the fund no longer appears tenable. Even with a dynamic investment strategy for the with-profits fund, there is likely to be some remaining guarantee cost.

9.4.2 For single premium contracts a charge could be made at inception and deducted from the premium. The charge can then be used to purchase a suitable hedge, if desired. For regular premium contracts, an initial charge or deduction may not be practical, and some form of regular charge, either as a premium deduction or as an annual charge against asset shares, might be made. While the value of such a charge can be reflected in the realistic balance sheet, it represents an intangible rather than a real asset, and cannot be spent on purchasing a suitable hedge. That cost must either be met from other assets in the fund, or by applying actuarial funding to the asset share, enabling part of the asset share to be used to fund the purchase of the hedge.

9.5 *Delta and Gamma Hedging*

9.5.1 There are many ways of selecting a portfolio of put options to hedge the guarantees in with-profits funds, but whatever method is employed, a comparison of the delta of the guarantees and of the options should be made in order to assess the quality of the hedge. The delta of the guarantees or of the option is simply their change in value for a unit change in value of the underlying asset. Since deltas are not linear, but change with the underlying asset prices, some consideration should also be given to matching gamma (the change in delta for a unit change in the underlying asset), vega hedging, which measures the sensitivity to a change in volatility.

9.5.2 These sensitivities can be computed from a Monte Carlo simulation model, but it is far easier to perform the calculations using option pricing formulae. That applies both to the initial determination of the hedge and its subsequent monitoring.

9.5.3 Long-dated options are not widely available, and quoted bid ask spreads are generally large. A company wishing to hedge financial guarantees in its liabilities may therefore consider using shorter-dated traded options. Such an approach may be considered as an alternative to dynamic hedging, and permits some insurance of jump risk.

9.5.4 The calibration of the model used to determine a market consistent value of guarantees will influence hedging decisions. For example, if the model used is calibrated to historical equity volatility of say 20% whereas implied equity volatility is 25%, then the model and the resulting valuation will simply indicate that hedging through the purchase of options appears unattractive, and that it would be better to try to hedge dynamically within the fund. This will not necessarily lead to a faulty decision, since part of the implied volatility reflects the costs to the option writer of its own risk management function and capital structure, as described earlier. Analysis of the decision of whether or not to hedge also needs to include the insurance company's own dynamic hedging costs and the cost of capital required to cover the gaps in the hedging programme.

9.5.5 If implied volatility were lower than historical, use of historical volatility in the model would make the purchase of option protection appear attractive. However, in this case the use of historical volatility might result in an overstatement of an insurance company's own costs of dynamic hedging.

9.6 *Policyholder and Shareholder Risks*

9.6.1 The hedging approach discussed above has focused on hedging risks on behalf of policyholders at the fund level. While such hedging is likely to be beneficial to policyholders in that it enhances security if properly performed, and reduces the risk that benefits will not be paid and guarantees honoured, it is equally important to consider the merits of hedging from the shareholder perspective.

9.6.2 The shareholder perspective is far more complex, and a full analysis requires a robust model of shareholder value. That model needs to include, not only a market consistent valuation of the shareholders' interest in the assets and liabilities of the company, but also quantification of the other factors affecting shareholder value. An analysis of these factors, which include franchise value, agency costs, the limited liability put option and taxation, is provided in Exley & Smith (2003).

9.6.3 Within such a framework it is then feasible to test the sensitivity of shareholder value to equity and interest rate risk, and to formulate potential hedges, and to assess the relative worth of purchasing hedges, performing dynamic hedging and running unhedged positions. In practice this is a significant piece of work, whose starting point is the proper evaluation of the shareholders' interest in the with-profits fund, combined with an assessment of those (hopefully extreme) adverse events where the shareholders are required to support benefits paid to policyholders.

9.6.4 Consideration could be given to hedging other shareholder assets, ranging from deferred tax to franchise value. We leave the debate on the extent to which these could or should be hedged to another day.

10. Conclusions

10.1.1 The move to market consistent valuation of liabilities has the potential to improve transparency of financial reporting, enhance consumer protection and provide useful new tools for financial management.

10.1.2 A large existing literature already addresses many of the technical issues of the calibration of models for pricing cash flows. This includes many closed form solutions, as well as the Monte Carlo simulation approach. Many difficult technical tasks remain however, including an analysis of the impact of taxation and the relationship between historical and implied volatilities.

10.1.3 A logically rigorous approach to market consistent valuation would exclude certain practices which seem currently to be widespread. For example, it is difficult to see how a market consistent valuation could legitimately reflect credit for asset risk premiums (from credit or equity risk) when valuing liabilities. For some purposes, it may be relevant to take account of an insurer's own propensity to default on its liabilities. The terms of reference for a market consistent value should ideally specify such fundamental issues as the allowance (if any) to be taken for liability credit risk.

10.1.4 In addition, there are many elements of a market consistent insurance value where a range of outcomes could legitimately be obtained. These include assumptions for long-dated volatilities, for correlations, for management behaviour, for take-up rates of policyholder options and many more. In many cases, there are little or no real historical data on which to base the judgement, and the chosen parameter is little more than a guess.

10.1.5 The future U.K. supervisory regime places responsibility for valuations on the directors of insurance companies. We hope that this paper will be of practical use to both the actuarial profession and to the industry in their quest for market consistent valuations.

Acknowledgements

This paper could never have come to fruition without input from many friends, clients and colleagues. We especially want to thank Tamsin Abbey, Stephen Carlin, Paul Coulthard, Trevor Maynard, Shyam Mehta, Viktor Mirkin, Ian Moran, Hui Ming Ng, Brendan O'Flynn, Cliff Speed, Elliot Varnell, and two anonymous scrutineers for useful observations and discussions. Any errors remain the responsibility of the authors.

References

Abbink, M. & Saker, M. (2002). Getting to grips with fair value. Paper presented to the Staple Inn Actuarial Society.

Altmann, E. & Vanderhoof, I.T. (1998). *The fair value of insurance liabilities*. Kluwer.

Anderson, N., Breedon, F., Deacon, M., Derry, A. & Murphy, G. (1996). *Estimating and interpreting the yield curve.* Wiley, Chichester.

Ballotta, L. & Haberman, S. (2003). Guaranteed annuity conversion options and their valuation. Working paper, Cass Business School.

BACINELLO, A.R. & ORTU, F. (1993). Pricing equity-linked life insurance with endogenous minimum guarantees. *Insurance: Mathematics and Economics*, **12**, 245-247.

VAN BEZOOYEN, J.T.S., EXLEY C.J. & MEHTA S.J.B. (1998). Valuing and hedging guaranteed annuity options. Paper presented to the Institute and Faculty of Actuaries Investment Conference, September 1998, University of Cambridge.

VAN BEZOOYEN, J.T.S., MEHTA S.J.B. & SMITH, A.D. (2000). Investment strategy and valuation of with profits products. *Investment Conference, Institute of Actuaries.*

BLIGHT, P., KAPEL, A. & BICE, A. (2003). *Market-consistent economic valuations for the wealth management industry*. Institute of Actuaries of Australia.

BOUWKNEGT, P. & PELSSER, A. (2001). Market value of insurance contracts with profit sharing. *AFIR Colloquium, Toronto*.

BRIYS, E., BELLELAH, M., MAI, M.H. & DE VARENNE, F. (1998). *Options, futures and exotic derivatives*. Wiley.

BUHLMANN, H.F.F., DELBAEN, F. EMBRECHTS, P. & SHIRYAEV, A.N. (1998). On Esscher transforms in discrete finance models. *ASTIN Bulletin*, **28**(2), 171-186.

DEMPSTER, M.A.H., GERMANO, M., MEDOVA, E.A. & VILLAVERDE, M. (2003). Global asset liablity management. *British Actuarial Journal*, **9**, 137-195.

DODHIA, M. & SHELDON, T.J. (1994). Guaranteed equity products. Paper presented to the Staple Inn Actuarial Society.

DULLAWAY, D. & NEEDLEMAN, P. (2004). Realistic liabilities and risk capital margins for with-profits business. A discussion paper, *British Actuarial Journal*, (to appear).

FINANCIAL SERVICES AUTHORITY (2003). CP195.

EXLEY, C.J. & SMITH, A.D. (2003). The cost of capital for financial firms. Discussion paper presented to the Finance and Investment Conference.

FELDMAN, K.S., BERGMAN, B., CAIRNS, A.J.G., CHAPLIN, G.B., GWILT, G.D., LOCKYER, P.R. & TURLEY, F.B. (1998). Report of the fixed-interest working group. *British Actuarial Journal*, **4**, 213-384.

GESKE, R. (1979). The valuation of compound options. *Journal of Financial Economics*, **7**, 63-81.

GOLDMAN, B., SOSIN, H. & GATTO, M. (1979). Path dependent options. *Journal of Finance*, **34**, 111.

GRIFFITHS, J., IMAM-SADEQUE, F., ONG, A. & SMITH, A.D. (1997). Actuaries take note — the yield curve is not flat. Investment Conference, Institute of Actuaries.

GROSEN, A. & JORGENSEN, P.L. (2000). Fair valuation of life insurance liabilities: the impact of interest rate guarantees, surrender options and bonus policies. *Insurance: Mathematics and Economics*, **26**, 37-57.

HAIRS, C.J., BELSHAM, D.J., BRYSON, N.M., GEORGE, C.M., HARE, D.J.P., SMITH, D.A. & THOMPSON, S. (2001). Fair valuation of liabilities. *British Actuarial Journal*, **8**, 203-299.

HANCOCK, J., HUBER, P. & KOCH, P. (2001). *The economics of insurance — how insurers create value for shareholders*. Swiss Re Publishing.

HARDY, M.R. (2003). *Investment guarantees: modelling and risk management for equity-linked life insurance*. Wiley.

HARE, D.J.P., DICKSON, J.A., MCDADE, P.A.P., MORRISON, D., PRIESTLEY, R.P. & WILSON, G.J. (1999). A market-based approach to pricing with-profits guarantees. *British Actuarial Journal*, **6**, 143-196.

HIBBERT, A.J. & TURNBULL, C.J. (2003). Measuring and managing the economic risks and costs of with-profits business. *British Actuarial Journal*, (to appear).

HULL, J.C. (2003). *Options, futures and other derivative securities*. Prentice-Hall, Fifth edition.

HULL, J.C. & WHITE, A. (1990). Pricing interest rate derivative securities. *Review of Financial Studies*, **3**, 573-592.

JAMSHIDIAN, F. (1989). An exact bond option formula. *Journal of Finance*, **44**, 205-209.

JARVIS, S., SOUTHALL, F. & VARNELL, E. (2001). Modern valuation techniques. Paper presented to the Staple Inn Actuarial Society.

JENSEN, J.L. (1995). *Saddlepoint approximations*. Oxford.

KEMP, M.H.D. (1997). Actuaries and derivatives. *British Actuarial Journal*, **3**, 51-162.

KIPLING, M. & MORAN, I. (2003). Realistic balance sheets. Presentation to Current Issues in Life Assurance seminar.

MEHTA, S.J.B. (1992). Taxation in the assessment of profitability of life assurance products and of life office appraisal values. *Transactions of the Faculty of Actuaries*, **43**, 484-553.

MEHTA, S.J.B. (1999). Supply and demand imbalances in the government bond market. Investment Conference, Institute of Actuaries, 1999.

MORAN, I., SMITH, A.D. & WALCZAK. (2003). Why can financial firms charge for diversifiable risk? Discussion paper presented to Bowles Symposium, Society of Actuaries (USA).

MUIR, M. & WALLER, R. (2003). Twin peaks — the enhanced capital requirements for realistic basis life firms. Paper presented to the Staple Inn Actuarial Society.

NG, H.M. & VARNELL, E. (2003). Frictional Costs. Paper presented to the Staple Inn Actuarial Society.

RANSON, R.H. & HEADDON, C.P. (1989). With profits without mystery. *Journal of the Institute of Actuaries*, **116**, 301-325.

ROGERS, L.C.G. & ZANE, O. (1999). Saddle-point approximations to option prices. *Ann. Appl. Probability*, **9**, 493–503.

SMITH, A.D. (2003). Option pricing with deflators. In *Asset and liability management tools*, edited by Bernd Scherer, Risk Publications.

SMITH, A.D. & SPEED, C. (1998). Gauge transforms in stochastic models. AFIR Colloquium, Cambridge.

SMITH. A.D. & WILSON, T. (2000). *Fitting yield curves with long constraints*. Financial Options Research Centre, Warwick.

WILKIE, A.D. (1987). An option pricing approach to bonus policy. *Journal of the Institute of Actuaries*, **114**, 21-77.

WILKIE, A.D., WATERS, H.R. & YANG, S. (2003). Reserving, pricing and hedging for policies with guaranteed annuity options. *British Actuarial Journal*, **9**, 263-391.

Printed by Bell & Bain Ltd., Glasgow