A STOCHASTIC ASSET MODEL & CALIBRATION FOR LONG-TERM FINANCIAL PLANNING PURPOSES

John Hibbert, Philip Mowbray & Craig Turnbull

June 2001

Copyright 2001 © Barrie & Hibbert Limited

Barrie & Hibbert Limited is a member of The Securities and Futures Authority

The information in this report is believed to be correct but cannot be guaranteed. All opinions and estimates included in this report constitute our judgement as of the date indicated and are subject to change without notice. This report is intended for information purposes only and is not intended as an offer to buy or sell securities. The company, its clients and officers may have a position or engage in transactions in any of the securities mentioned. This notice does not restrict your rights under the Financial Services Act (1986) or the regulatory system. This document is not intended for the use of Private Customers.

TABLE OF CONTENTS

ABSTRACT

The report provides a specification for a stochastic model for equity returns, inflation and the term structures of real and nominal interest rates together with a discussion of the possible approaches to parameter selection. We contrast the model's output with a typical calibration of the Wilkie investment model.

1. INTRODUCTION

What is it that makes the role of financial intermediaries so special? Surely, some part of the answer to this question is the financial intermediary's objective of pooling and managing risks on behalf of large groups of individuals. In common with other financial intermediaries, life insurance companies are in the risk management business.

The risk exposures accumulated by the shareholders and policyholders of today's life companies come in many different forms. In the past, uncertainty in financial plans has be assessed either by expert intuition or alternatively by developing a small number of handcrafted "*what if?*" scenarios. Such scenario analysis can be extremely valuable in situations where there are only a small number of key sources of risk. Unfortunately, in a situation where there are many sources of risk or where the dimension of the problem is large, building the scenarios by hand becomes impractical.

The risk manager of 25 years ago would be astonished by the range of tools available to his modern counterparts. It is now increasingly common practice to use a computer to generate the scenarios using a stochastic model. Monte-Carlo (MC) simulation techniques can be used to generate very large numbers of scenarios in order to understand the potential behaviour of financial products and other entities in a world of uncertainty and under various strategy options available to the planner. Instead of testing out a handful of possible outcomes, the planner can test out – literally – thousands of possible futures.

In this note we will focus on one specific model out of many interesting possible candidates. We shall tackle the difficult problem of how to simulate consistent future paths for equity returns, dividend yields, inflation and complete real and nominal term structures. It is worth emphasising that there are plenty of models around that deal adequately with the nominal term structure. Here we shall aim to reconcile the behaviour of the inflation rate with both real and nominal interest rates. The model presented has some attractive features for the purposes of analysing certain classes of problem, but we do not claim that it is perfect. We will explore the issue of how to calibrate this model and illustrate two calibrations. The illustrations provided have been judged to be useful by some people. You may prefer a different calibration.

As we shall see, as always, it is necessary to strike a balance between our ambition to make the model as realistic as possible and a need to keep the model simple. Where we strike the balance will depend upon a number of considerations: the specific application of the model as well as the needs and sophistication of model users.

This report is rather longer than we had planned. It would have been straightforward to write down our fancy equations and leave the reader to apply the model. Instead, we have set out to explain *why* we believe modelling is useful at all and what sort of problems are being analysed with stochastic models (section 2). In section 3, background is provided by highlighting some extraordinary changes that have taken place over the past two decades. In the following section we list some properties of 'good models'. Section 5 sets out the specification for the model and the following two sections provide a discussion of calibration strategies and two specific calibrations. Section 8 contrasts the model with the Wilkie investment model. Finally, in sections 9 and 10, we briefly discuss possible extensions and set out some brief conclusions.

Much of the presentation that follows is informal. Our objective is to give the reader insights into the general problem of stochastic model building as well as a description of a practical tool. We do not pretend that the model presented in this report is complete nor that the calibration could not be bettered in some respects. Rather this is our starting point for some stochastic investigations. The ideas presented have been developed over several years and with many mistakes and blind alleys along the way. Let us begin by addressing a basic question: "*why build stochastic models at all?*"

2. WHY BUILD MODELS?

2.1 WHAT LESSONS FROM THE PAST?

Life insurance companies are in the business of managing risk on behalf of large groups of individuals over very long planning horizons. Some of these risks can be diversified away whilst others must be carried by either policyholders or shareholders. The financial products sold by life assurers contain guarantees of numerous varieties. *Life insurance is a risk management business*. Life company managers have experienced a traumatic decade – failing to adequately understand and manage risk across a range of different products. The list is dismal and familiar:

- Some with-profits 'promises' made to policyholders in the late 1980s and early 1990s look hopelessly optimistic in an environment of (normal – but surprising) low inflation.
- Annuity options offered in the 1970s and 1980s have proven costly as a consequence of unanticipated falls in nominal interest rates and surprising improvements in mortality.
- The basic financial position and funding strategy of pension funds (two very distinct properties) remain confused amid a three-way debate among actuaries, accountants and regulators.

It is probably fair to say that – for many life companies – these problems remain unresolved. The quality of understanding of the problems and strategies selected to manage them is variable.

2.2 WHAT CURRENT PROBLEMS DO ACTUARIES & FINANCIAL PLANNERS FACE?

Is this the end of our list? Sadly not. The industry needs to face up to a number of ongoing and new problems over the coming years:

- With-profits products remain under scrutiny – most recently from regulators. Of course, the catch-all 'with-profits' covers a wide range of products with very different guarantees and different strategies for 'smoothing' returns across successive generations of maturing policies. Product providers have struggled (unsuccessfully so far as their end-customer is concerned) to communicate the nature of the product. There exists (at least) two important challenges: first, to make the product (or its replacement) more transparent; second, to manage and price the risks carried by policyholders and shareholders in an appropriate way.

  The recent changes to regulations mean that the capital required to support unitised with-profit business is now more sensitive to the pattern of delivered asset returns.

- Although the move from defined benefit towards defined contribution pensions arrangements (including Stakeholder) shifts the burden of risk-bearing from the sponsor to the individual saver, the risk management challenge remains. The current generation of pension savers will bear far greater risk than their parents.

At present they are poorly equipped to negotiate the complicated trade-off between retirement benefits, contribution levels and the age at which they can finally afford to stop working.

- Pension ‘drawdown’ products and other post-retirement products can expose savers to a potent mixture of investment risk and mortality risk that few are well equipped to understand.
- So-called ‘income’ products come in all shapes and sizes. Many of these products can be expected (in a statistical sense) to reduce a saver’s capital – sometimes by a material amount. As in the past, it seems likely these risk exposures will only be properly appreciated by savers (and issuers and regulators) after a product failure.

The common element in all of these situations is risk. Savings product providers are in the risk management business.

2.3 WHAT FUTURE PROBLEMS?

We can only guess at the risk management challenges to be faced by financial intermediaries in the future. In the light of past experience, it seems likely that many of the problems outlined above will take many years to bring under effective control. So long as the life industry fails to embrace fully the new risk management technologies it is likely that some unanticipated combination of economic, market and demographic change will trigger another round of product failures. As in the past, the result will be damage to provider and regulator reputations and losses of shareholder and policyholder capital.

3. Where Have we Come From?

3.1 Background

It is instructive to highlight two important trends over the past 20 years:

i. First, extraordinary innovation in computer technology has taken place. Slide rules, log tables and punched computer cards have been replaced with unimaginably powerful desktop computers and software. These tools mean that calculations which were unthinkable 20 years ago are now (potentially) routine. Perhaps even more importantly, the means of displaying information is now really only limited by the analyst's imagination.

   It is worth pointing out that some financial institutions – particularly the investment banks – have made substantial investments in this new technology in order to enhance their risk management capabilities.

ii. Second, a huge volume of research has been generated by financial academic researchers and financial practitioners (including actuaries). It is important to understand that this research effort has been motivated by some quite different needs:

   - Traders looking for improved techniques for pricing, trading and hedging a range of new financial instruments. Although options contracts have been around for many centuries, the publication of the *Black-Scholes* model in 1973 provided a spur towards rapid innovation in derivative markets.

     For interest rates in particular, there now exists a vast literature of models of varying degrees of complexity.

   - Economists have developed models largely for the purpose of forecasting and policy-making.

   - Long-term financial planners (including actuaries) must combine multiple sources of uncertainty generally over very long horizons (compared to other users of financial models).

It is probably fair to say that most academic work tends to deal only with *parts* of the problem the actuary is interested in. There is much detailed work on equity price behaviour, on interest rate modelling and on inflation modelling. However, there is very little which puts all of the components together within a consistent framework. The fundamental task of the long-term financial planner is to understand the *joint* behaviour of these variables (and others) on the product or business under scrutiny.

3.2 THE WILKIE MODEL

The Wilkie model[1] was originally developed in the late 1970s against a background of high and volatile inflation and the exceptional UK equity market volatility of 1974/5. It was extended in 1995. Unlike the work of mainstream academics, Professor Wilkie's model *did* tackle the difficult problem of how to put together a model for inflation, real and nominal returns on equities and bonds and their yields. Equity yields – which continue to play an important part in actuarial analysis – are prominent in the model. They are only of passing interest to the economist who tends to focus on *price* and *return* (not how it happens to be packaged). The Wilkie model is relatively straightforward to implement. As a consequence, the model has been widely used by UK-based actuaries over the past two decades and has set a benchmark against which any other proposed approach needs to be judged.

However, we believe that there are some serious problems with the Wilkie model. What is more, the huge development in thinking in mainstream academia means that there are now some (almost) ready-made tools available to fix the shortcomings of the Wilkie model. Our aim is to show how this can be achieved whilst avoiding the mind-boggling complexity that seems to be a characteristic of many of the models proposed in this area.

We do aim to retain one of the primary attractions of the Wilkie model – its ease of implementation. To that end a working version of the model presented in this report will be made available on the Barrie & Hibbert web site (www.barrhibb.com). This model is distributed under the GNU Public licence. Extended versions of the model are available on a commercial basis.

[1] *"Report of the Maturity Guarantees Working Party", (A.D. Wilkie, Journal of the Institute of Actuaries, 107, Part II, No. 435); "A Stochastic Investment Model for Actuarial Use", A.D. Wilkie, Transactions of the Faculty Of Actuaries (No. 268, Vol 39 Part 3); "More on A Stochastic Model for Actuarial Use", (A.D. Wilkie, British Actuarial Journal, 1)*

4. Some Properties of Good Models

The broad objective of the risk analyst is to provide insight across a wide range of problems faced by business managers, product designers, regulators, customers and their advisers into the impact of a range of ‘candidate’ financial strategy choices. For example, when the analyst reviews business written with attaching annuity options, business managers will want to review a range of policy options spanning bonus policy, investment policy as well as hedging and reinsurance solutions. In order to *understand* and *communicate* analysis, a model can be very valuable. The model is intended to be a cut-down, simplified version of reality that captures the essential features of the problem and aids understanding.

It simply is not plausible to argue that there is a single model that can meet the requirements of the risk analyst across all possible problems. Rather the analyst should aim to build a library of models that enable him to tackle the different types of problem that he is faced with.

So, models must be selected, but how? What criteria should the analyst use to pick a particular model? We have listed below some of the attributes that we think are important in ‘good models’. The list is not intended to be complete. Some of the criteria on the list probably overlap with each other. As you might have guessed, it turns out to be very difficult to meet all of the criteria simultaneously. We rarely find models which pass all of the tests.

4.1 Representativeness

The model should aim to provide a good representation of the financial assets contained in the model. The model should “mimic” the behaviour of real-world financial assets by capturing their most important characteristics. If the model is used to generate Monte-Carlo scenarios, we might expect that an expert who scrutinises the model output to be able to say: “*Yes – each of your scenarios looks plausible and the frequencies assigned to particular outcomes look reasonable*”.

This test covers numerous characteristics of asset behaviour – the shape of distributions at different time horizons as well as the relationships between the variables in the model.

4.2 Economic interpretation

The behaviour of assets (within the model) should be consistent with generally-accepted economic principles. The most frequent demand is that a model should be “arbitrage-free”. Since we do not expect to observe systematic opportunities for arbitrage in the real world, it seems sensible to exclude them from models. The “no free lunch” rule seems like a good one for the purposes of modelling. However, it is important to appreciate that implementing the rule often comes at a price. There may be times when the modeller will be prepared to allow some limited arbitrage into a model in exchange for some other benefit.

A further consideration relates to the *joint behaviour* of model variables. Again, we expect to see some consistency between economic principles and model behaviour. It is worth noting that there are some important properties of financial asset behaviour on which there is no clear consensus among economists. The prickly topic of equity market mean reversion is a good example that we will return to.

4.3 PARSIMONY

Keep it simple. Models should be as simple as we can make whilst retaining the most important features of the problem. It is clearly often difficult to judge when complexity is really needed (sometimes some factor can have a big impact on results whilst in other situations it only has a minimal effect). Complexity must be balanced. There is no sense in modelling one aspect of a problem in mind-boggling detail, and then making broad brush assumptions in other areas. The model will stand on its weakest assumption.

There is another reason why complexity should be avoided. A complex model, which tries to mimic as much real-world complexity as the modeller can capture, can create the illusion that we really can model everything. Sometimes this illusion fools the modeller as well as his unlucky audience.

4.4 TRANSPARENCY

Unless we can explain how the basic model works in a few minutes it will be difficult to gain the confidence of non-experts. Success here will depend heavily on the quality of communication. The results produced from the model should be displayed in clear graphical formats wherever possible.

4.5 EVOLUTION

Nothing complex can be designed and built in a single “life”. Anything complicated must be allowed to evolve over a number of lifetimes. This applies whether you set out to build a Boeing 747 aeroplane or a financial model of a with-profits savings contract.

4.6 IMPLEMENTATION TOOLS

Financial models generally combine a set of rules for describing how some payoff or property of interest is determined with a description of the behaviour a set of stochastic variables that determine the payoffs. The models can be implemented in different ways. Implementation tools fall into a number of classes:

i. Analytic calculations where it is possible to find a mathematical function to describe the variables of interest. This is generally only possible for a very limited set of problems.

ii. Historical back-tests are performed by using past data on returns, yields etc.. with the model structure and implicitly assuming that the future will be like the past period selected.

iii. Scenario analysis (deterministic simulation / sensitivity analysis) where the modeller maps out – by hand – a series of scenarios of interest.

iv. Tree-building techniques where the scenarios are built up in the form of a tree. This is generally only feasible where a very small number of stochastic factors influence the problem.

v. "Monte-Carlo" (stochastic) simulation can be used to generate very large numbers of plausible scenarios using a computer[2]. In situations where there is path dependency, Monte Carlo techniques can be particularly valuable[3].

There are advantages and disadvantages associated with each of these approaches. Clearly, if a model can be implemented in different ways it provides increased flexibility to the modeller. In practice, models may be a mixture of Monte-Carlo, analytic and tree-based components.

For the real-world problems faced by actuaries, the flexibility and intuitive presentation offered by Monte-Carlo techniques mean that it will remain the focus of our attention in the remainder of this report.

---

[2] *"Monte-Carlo" simulation can be used in situations where we believe that we can say something sensible about the factors that affect a problem, but we don't know what will happen when we put these factors together. Road traffic engineers have a pretty good idea of how cars and drivers behave (how quickly they brake and accelerate etc..). They know how roads are laid out and the sequences of traffic lights. But it is impossible to capture all of this in a mathematical equation that will predict how traffic will behave. The maths is too complicated. Does this mean that road traffic engineers can't predict how traffic will behave when they fiddle with a traffic light sequence? Not at all. They use computer simulations. The simulations throw up most of the features of real traffic - bottlenecks, queues, sudden empty roads etc.. They allow the engineer to see how changes in some part of the system will affect its overall behaviour.*

[3] *Where results depend on the path taken by financial variables over the planning horizon (not simply by where they end up) we would say that a result is "path-dependent". A good example is a with-profits savings contract where the path of equity returns can affect the final payoff.*

5. AN ALTERNATIVE ASSET MODEL

In this section we describe a stochastic asset model that can be used for long-term financial projections. We believe that the model has a number of attractive properties:

- The model deals explicitly with the economic relationship between inflation, inflation expectations, real interest rates and nominal interest rates.
- It produces a complete and consistent term structure for real and nominal interest rates with rich variation in both the level and shape of the yield curves generated.
- The model for equity returns can be used to generate the negative skewness and kurtosis ('fat tails') that are a characteristic of real-world equity return distributions. Equity yields and dividends are generated in a natural way. The basic set-up of the equity model does *not* incorporate mean reversion. We believe that this is a prudent starting point for long-term financial planning purposes.
- The model is easy to implement with analytic expressions available for discount bond prices.
- It is possible to extend the model relatively easily beyond the basic application presented here. We briefly discuss extensions in section 9.

The model is comprised of a number of component parts that are driven by a set of stochastic drivers. Let us begin by reviewing the model for the term structure of interest rates.

5.1 THE TERM STRUCTURE MODEL

The behaviour of the real-world term structure is reviewed in detail in section 6. At this point it is worthwhile highlighting the challenge ahead. There are many models that allow the analyst to directly mimic the behaviour of the nominal term structure. Here, we aim to deal with both real and nominal interest rates in a consistent way by explicitly modelling the link between inflation expectations and nominal yields.

EXHIBIT 5.1: INFLATION & THE SHORT-TERM INTEREST RATE

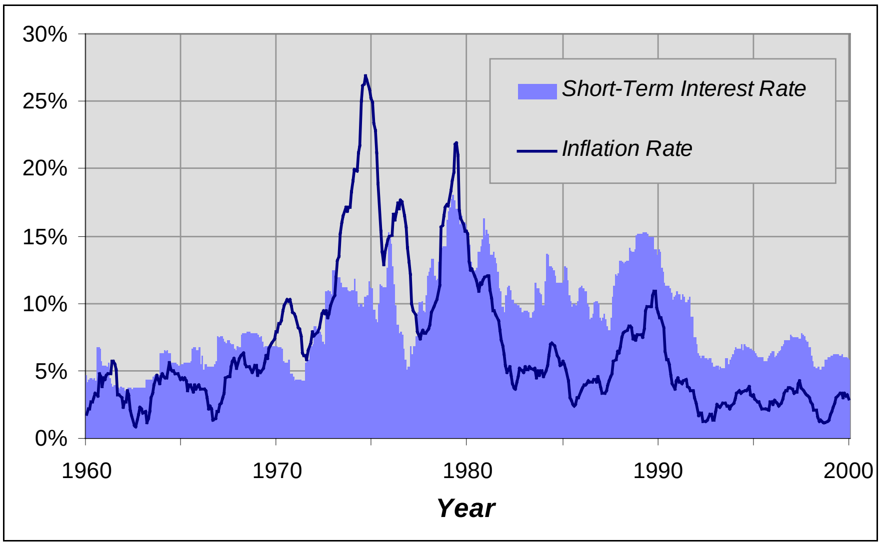

Exhibit 5.1 illustrates the strong linkage between inflation and nominal interest rates observed over the past four decades. In the model we propose, the basic idea is to build a term structure for nominal interest rates from *two* separate components:

- A term structure for real interest rates. You can think of this as the index-linked yield curve. In the model we deal with the yield on real (inflation-protected) discount bonds[4].
- A model for inflation that allows us also to model the inflation *expectations* of investors over different horizons. It is (implicitly) assumed that investors understand the process generating inflation and adjust their expectations for future inflation in a manner that is consistent with the experienced inflation pattern.

The two term structures are combined to form the nominal term structure with allowances for any assumed correlation between inflation and real interest rate changes and for any risk premium associated with bond term and inflation risk. Let us now review each of these components in turn.

5.1.1 REAL INTEREST RATES / 2-FACTOR HULL-WHITE

A huge literature exists on arbitrage-free term structure models. Although much of this work has been motivated by a need to price interest rate derivatives, many of the contributions are directly applicable to long-term planning and actuarial work. Some of the literature is highly mathematical and far from transparent.

We have made use of an extension to one of the first stochastic arbitrage-free term structure models – the *Vasicek* model[5]. The *Vasicek* model specifies a continuous-time mean-reverting stochastic process for the short-term interest rate, and then infers forward rates and spot rates from the expected future path of the short rate (with an allowance for any specified risk premium investors may demand for holding longer-maturity bonds relative to cash). The model can be extended with the addition of a second stochastic factor. This factor allows the mean reversion level for the short rate also to follow a mean-reverting stochastic process. For those who are familiar with the first-order autoregressive process used by David Wilkie to model inflation, our model uses a similar process to model the short-term real rate, but additionally with a mean reversion level which is also autoregressive[6].

---

[4] *We will use the terms '<u>discount bond</u>' and '<u>zero-coupon bond</u>' to mean the same thing. The '<u>T-year spot rate of interest</u>' will mean the annualised continuously compounded interest rate on a T-year discount bond (i.e. –log(P)/T, where P is the discount bond price and T the term in years).*

[5] *See "An Equilibrium Characterization of the Term Structure", O.A. Vasicek, Journal of Financial Economics, 5 (1977)*

[6] *We are grateful to Andrew Cairns of Heriot Watt University who provided assistance in the mathematical development of the model.*

This model we have selected can be shown to be a special case of another published term structure model – a 2-factor model described by Hull & White (1994)[7]. The equations governing the changes in the real short rate are shown below:

$$dr_1(t) = \alpha_{r1}\left(r_2(t) - r_1(t)\right)dt + \sigma_{r1}\,dZ_{r1}(t)$$
$$dr_2(t) = \alpha_{r2}\left(\mu_r - r_{r2}(t)\right)dt + \sigma_{r2}\,dZ_{r2}(t)$$

where:

$r_1(t)$	=	the real short rate at time *t*.
$r_2(t)$	=	the mean reversion level for the real short rate at time *t*.
α_{r1}	=	the autoregressive parameter for the real short rate process.
α_{r2}	=	the autoregressive parameter for the real short rate mean reversion process.
σ_{r1}	=	the annualised volatility (standard deviation) of the real short rate.
σ_{r2}	=	the annualised volatility (standard deviation) of the real short rate mean reversion level.
μ_r	=	the mean reversion level for $r_2(t)$.
g_r	=	a parameter to control the term premium in real bond prices.
$dZ_{r1}(t)$	=	the shock to the real short rate process which is distributed $N(g_r\,dt,\,dt)$
$dZ_{r2}(t)$	=	the shock to the real short rate mean reversion process which is distributed $N(g_r\,dt,\,dt)$
b_{r1}	=	lower bound for the real short rate, $r_1(t)$.
b_{r2}	=	lower bound for the real short rate mean reversion level, $r_2(t)$.

These equations can appear mathematically rather daunting, but they are really just the continuous-time equivalents of two first-order autoregressive time-series processes and, as we will see, can be implemented in a similar way. Note that they imply that real interest rates have a normal distribution allowing the possibility of negative real rates. The potential for an interest rate model to generate negative rates is normally viewed as an inconvenience. However, since *real* interest rates are the subject of the model, it is often argued that the model should be capable of generating some negative rates. It is important to understand that the entire term structure for real interest rates is implied by these equations. They allow us to calculate an expected path for future short rates. This path is naturally related to current forward rates (with an adjustment for any assumed risk premium).

Exhibit 5.2 gives an approximate idea of how the model works. In the example shown in the chart, the initial stochastic mean-reversion level (plotted in green) begins at 10% and is projected to be pulled over time towards its equilibrium level. The real short rate, $r_1(t)$, is projected to be pulled towards the time-varying reversion level and – in this case – traces out a humped path with a peak of nearly 8% after 3 years. Crudely, and for the time being ignoring risk aversion, you can think of this path as the set of instantaneous forward interest rates i.e. the forward interest rate available for a very short time

[7] *See "Numerical Procedures For Implementing Term Structure Models II: Two-Factor Models", Journal of Derivatives, Winter 1994, and B&H Technical Note 2000/024 "2F Vasicek as a Special Case of Hull & White".*

horizon. You can then infer spot rates (zero-coupon rates) for any maturity by combining together the appropriate forward rates. In the case illustrated below you arrive at a humped term structure with a peak at around 10 years.

EXHIBIT 5.2: THE BASIC IDEA ~ 2-FACTOR VASICEK MODEL

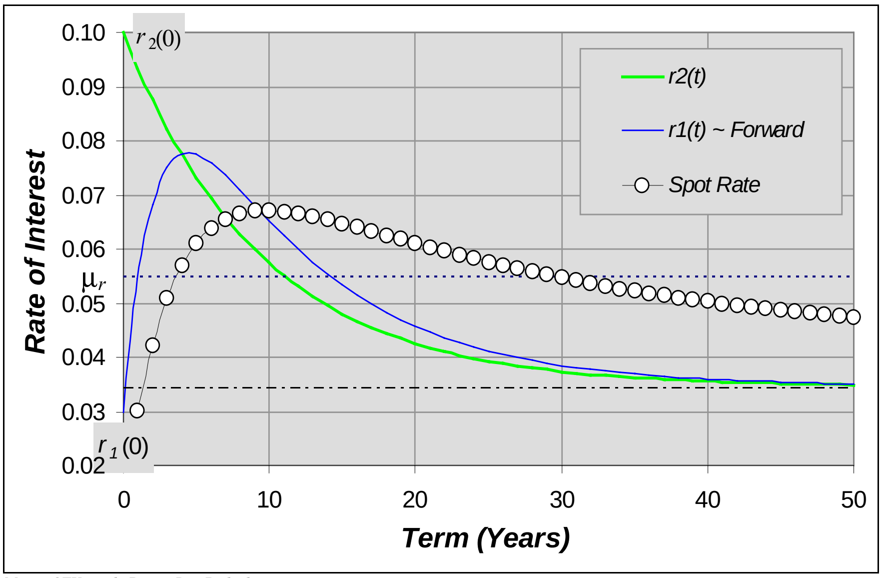

Memo: 2FVasicek_ProjectRatePath.xls

The exhibit gives some idea of how changes in r_1 and r_2 can affect the shape of the curve. The σ_{r1} and σ_{r2} parameters will influence variability of rates. The autoregressive parameters can be seen to have an impact on the *curvature* produced by the model. In practice things are a bit more complicated than suggested by exhibit 5.2.

Although this is a continuous-time model, it can still be applied in discrete time without the need for any approximations – from the above equations it is possible to calculate the expected value and variance of $r_1(t)$ and $r_2(t)$ over any time increment. We can then sample from these distributions to increment the model in discrete time in a manner exactly consistent with the continuous-time model.

An additional parameter, g_r, is introduced to determine the degree to which long-term real bond returns exceed the short real rate i.e. it is possible to introduce a *term premium* into the model. We have to make a small adjustment to how we increment the interest rates when g_r is non-zero. Appendix A describes the equations that are used for incrementing the term structure.

The equations set out above can be used to derive analytic expressions for the real spot rate and real forward rate at any term. The equations determine the entire real term structure. The price of a zero-coupon bond at time t that pays one unit in real terms (i.e. protected from inflation) at time T is given by the pricing equation:

$$P_{real}(t,T) \;=\; exp\;[\;A(T-t) - B_1(T-t)r_1(t) - B_2(T-t)r_2(t)]$$

where :

$$B_1(s) = \left[\frac{1-e^{-\alpha_{r1}s}}{\alpha_{r1}}\right]$$

$$B_2(s) = \frac{\alpha_{r1}}{\alpha_{r1}-\alpha_{r2}}\left[\frac{1-e^{-\alpha_{r2}s}}{\alpha_{r2}} - \frac{1-e^{-\alpha_{r1}s}}{\alpha_{r1}}\right]$$

$$A(s) = \left(B_1(s)-s\right)\left(\mu - \frac{\sigma_{r1}^2}{2\alpha_{r1}^2}\right) + B_2(s)\mu - \frac{\sigma_{r1}^2 B_1(s)^2}{4\alpha_{r1}}$$

$$+\frac{\sigma_{r2}^2}{2}\left[\frac{s}{\alpha_{r2}^2} - 2\frac{\left(B_2(s)+B_1(s)\right)}{\alpha_{r2}^2} + \frac{1}{\left(\alpha_{r1}-\alpha_{r2}\right)^2}\frac{\left(1-e^{-2\alpha_{r1}s}\right)}{2\alpha_{r1}} - \frac{2\alpha_{r1}}{\alpha_{r2}\left(\alpha_{r1}-\alpha_{r2}\right)^2}\frac{\left(1-e^{-\left(\alpha_{r1}+\alpha_{r2}\right)s}\right)}{\left(\alpha_{r1}+\alpha_{r2}\right)} + \frac{\alpha_{r1}^2}{\alpha_{r2}^2\left(\alpha_{r1}-\alpha_{r2}\right)^2}\frac{\left(1-e^{-2\alpha_{r2}s}\right)}{2\alpha_{r2}}\right]$$

We recognise that this is quite a big formula, but you only need to type it into the computer once.

Of course, once we have obtained prices for real discount bonds, it is then possible to price any real coupon bond which can be priced as if it were a package of discount bonds – each element of the package corresponding to one of the coupon or redemption payments of the bond. We can also calculate the continuously compounded yield at time t for maturity T, $R_1(t,T) = -\log\{\,P_{real}(t,T)\,\}\,/\,(T-t)$

5.1.2 THE INFLATION MODEL

Let us now turn our attention to inflation. The idea is to use a model for the behaviour of the inflation rate to generate a term structure for *inflation expectations,* which can then be combined with the real interest rate term structure to build a term structure for nominal interest rates. Exactly the same model structure is used to model the behaviour of the short-term inflation rate as was used in the previous section for real short-term interest rates. (Of course, we may choose to use different model *parameters* when the two models are used in practice).

So, the equations governing the path of the short-term inflation rate are:

$$
\begin{aligned}
dq_1(t) &= \alpha_{q1}\left(q_2(t) - q_1(t)\right)dt + \sigma_{q1}\, dZ_{q1}(t) \\
dq_2(t) &= \alpha_{q2}\left(\mu_q - q_2(t)\right)dt + \sigma_{q2}\, dZ_{q2}(t)
\end{aligned}
$$

where:

$q_1(t)$ = the instantaneous rate of inflation at time t.
$q_2(t)$ = the mean reversion level for the instantaneous inflation rate at time t.
α_{q1} = the autoregressive parameter for the inflation rate process.
α_{q2} = the autoregressive parameter for the inflation rate mean reversion process.
σ_{q1} = the annualised volatility (standard deviation) of the instantaneous inflation rate.
σ_{q2} = the annualised volatility (standard deviation) of the inflation rate mean reversion level.
μ_q = the mean reversion level for $q_2(t)$.
g_q = a parameter to control the inflation risk premium in nominal bonds relative to index-linked bonds.
$dZ_{q1}(t)$ = the shock to the inflation rate process which is distributed $N(g_q\, dt,\, dt)$
$dZ_{q2}(t)$ = the shock to the inflation rate mean reversion process which is distributed $N(g_q\, dt,\, dt)$
b_{q1} = lower bound for the inflation rate, $q_1(t)$.
b_{q2} = lower bound for long-term inflation expectations, $q_2(t)$.

As for the short-term real interest rate, the inflation rate is assumed to be mean-reverting and normally distributed. Using two factors to drive the inflation process, rather than a single factor (as in the Wilkie model) has some major advantages. Firstly, it means that changes in inflation rate expectations at different terms can have a correlation less than 1 (so short-term inflation rates and long-run expectations do not always have to move in lock-step). Secondly, it allows greater control over how the volatility of inflation decays. In a single factor model, it may not be possible to obtain a sensible distribution for long-term inflation without unfeasibly high short-term volatility and vice versa.

The inflation process can be incremented in the same way as for real rates and is described in Appendix A.

5.1.3 INFLATION EXPECTATIONS

In exactly the same way that we use an analytic expression to derive the real interest rate term structure at any time given the values of $r_1(t)$ and $r_2(t)$, a term structure for inflation expectations can be inferred from the current instantaneous inflation rate, $q_1(t)$, and the value of $q_2(t)$. We can therefore use the pricing equation of section 5.1.1 to calculate the value of a discount bond at time t that pays one unit at time T and is discounted only with respect to inflation expectations. We will refer to this quantity as $P_{inf}(t,T)$. The equivalent yield is $R_q(t,T) = -\log\{ P_{inf}(t,T) \} / (T-t)$

Further, in the same way that it is possible to embed a term premium into the real interest rate term structure, the inflation expectations term structure can also incorporate a risk premium reflecting the additional return which may be required by investors to induce them to invest in nominal, rather than index-linked bonds. This is sometimes called the ‘inflation risk premium’.

5.1.4 NOMINAL TERM STRUCTURE

Armed with the real interest rate and inflation expectations term structures, it is now possible to combine them together to obtain a nominal interest rate term structure. Where movements in instantaneous real interest rates and inflation rates are independent, this stage is trivial – the nominal spot rate is simply equal to the sum of the real spot rate and the annualised inflation expectation over the relevant period. Alternatively, the price of a nominal discount bond is obtained as the product of the real and inflation discount bond prices:

$$P_{nom}(t,T) \;=\; P_{real}(t,T)\, P_{inf}(t,T)$$

where $P_{real}(t,T)$ and $P_{inf}(t,T)$ are as described in sections 5.1.1 and 5.1.3.

In practice, innovations in short real rates and inflation may not be independent. For example, we might believe that rises in ‘spot’ inflation are typically associated with increases in real rates as policymakers attempt to squeeze inflationary growth. The model can accommodate an assumption for the correlation between short real rates and inflation through a (small) additional covariance term in the zero-coupon bond price equation:

$$P_{nom}(t,T) \;=\; P_{real}(t,T)\, P_{inf}(t,T) + \rho.sqrt[Var(exp\{-R_1(t,T)\}).Var(exp\{-R_q(t,T)\})]$$

where:

ρ = the correlation between the shock to the real short rate and the instantaneous inflation rate, $dZ_{r1}(t)$ and $dZ_{q1}(t)$[8].

$$R_1(t,T) = \int_t^T r_1(s)ds$$

$$R_q(t,T) = \int_t^T r_{q1}(s)ds$$

[8] *Technically, this correlation must also be applied to $dZ_{r2}(t)$ and $dZ_{q2}(t)$.*

Analytic expressions for the variances in the covariance term are given in Appendix B. Note that the nominal term structure has two separate risk premiums (either or both of which can be set to zero) – a term premium and an inflation risk premium. It is possible to set up the model so that a long-dated index-linked bond will have a higher expected return than a shorter-dated one, and a long-dated nominal bond will have a higher expected return than long real bonds and short-term nominal bonds.

5.1.5 NEGATIVE NOMINAL INTEREST RATES

It must be recognised that the way the model is specified does *not* guarantee that nominal interest rates will always be positive. Since inflation, inflation expectations and real interest rates can take negative values, negative nominal rates of some magnitude will feature in the model. The frequency of these negative rates will depend on the parameters selected.

Faced with this potential problem, the modeller can adopt a number of alternative approaches:

i. Use the negative rates. If the internal mathematical consistency of the model is particularly valued, this may be appropriate.

ii. If the main purpose of the model is to generate plausible scenarios, an alternative approach is to discard the scenarios in which negative rates occur. This appears straightforward, but the discarded scenarios will have an impact on some global characteristics of the model that must be understood and possibly adjusted for.

iii. A similar approach is to constrain the model in a way that guarantees positive rates (or ensures that negative rates appear with only a very low frequency). For the model specified above this might be achieved by setting lower bounds on the value of the stochastic factors.

iv. Find a better model. We have not, yet.

5.2 THE EQUITY MODEL

There is a huge array of models – of varying degrees of complexity – designed to describe the equity returns process. The model that is most widely used (by financial economists) is the *lognormal* model for equity returns. It turns out that this provides a reasonable, but imperfect description of the equity returns process.

EXHIBIT 5.3: END-YEAR & MID-YEAR UK EQUITY ANNUAL EXCESS RETURNS (1900-99)

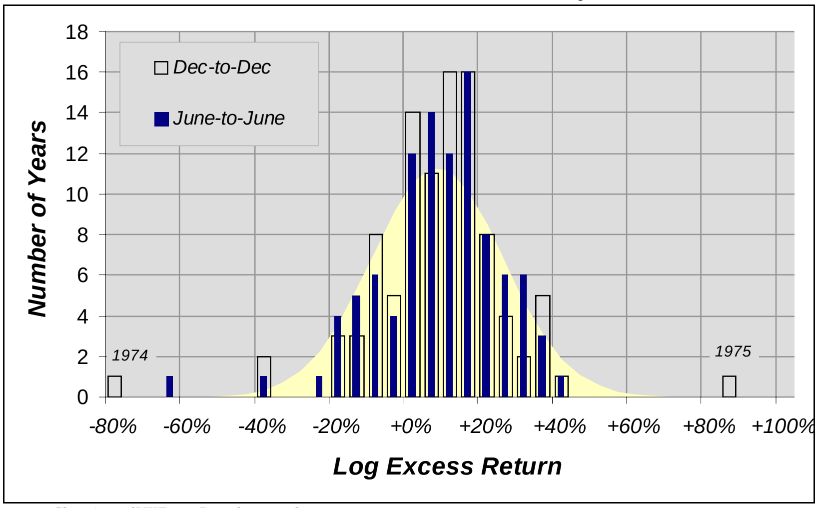

memo: ChartAnnualUKEquityDistributions.xls

The chart above shows the distribution of log excess returns[9] for UK equities since 1900 and a normal distribution with a mean and standard deviation estimated using the data shown *excluding 1974/5*. The picture suggests that the normal distribution does a reasonable job most of the time, but occasionally fails badly. In common with most financial market variables, the distribution for equity price changes shows "fat tails" (statisticians measure this characteristic of the distribution with the kurtosis statistic).

EXHIBIT 5.4: DISTRIBUTION OF MONTHLY EXCESS RETURNS COMPARED TO NORMAL (1900-99)

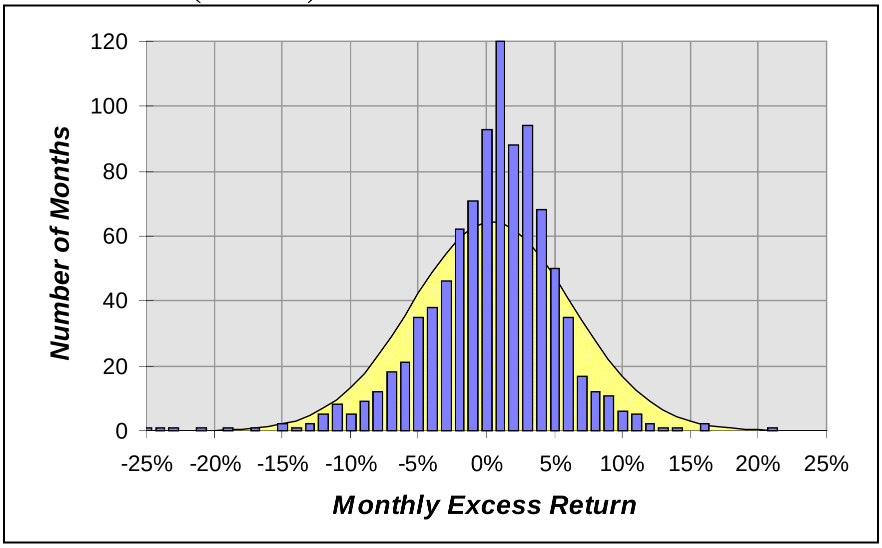

9 The rate of return in excess of the short-term interest rate.

Exhibit 5.4 illustrates the very marked ‘fat tails’ of the monthly excess returns distribution. The kurtosis of the returns distribution tends to increase as the measurement period for returns is shortened. This is not a new conclusion. This feature of financial markets is familiar to most analysts. The widely-used normal only provides an approximation to the real-world behaviour of equity returns. For some purposes this is fine, but for others (particularly when we care about the tails of the distribution) it simply is not up to the job.

EXHIBIT 5.5: EXTRACT FROM DISTRIBUTION OF MONTHLY UK EQUITY PRICE CHANGES

-27.5% to -25.0%	-25.0% to -22.5%	-22.5% to -20.0%	-20.0% to -17.5%	-17.5% to -15.0%	-15.0% to -12.5%	-12.5% to -10.0%	-10.0% to -7.5%	7.5% to 10.0%	10.0% to 12.5%	12.5% to 15.0%	15.0% to 17.5%	17.5% to 20.0%	20.0% to 22.5%	22.5% to 25.0%	25.0% and above
								Sep-97							
								Dec-93							
								Sep-92							
								Apr-92							
								Dec-87							
								Feb-87							
								Jan-87							
								Mar-86							
							Jan-00	Feb-86							
							Sep-90	Aug-84							
							Aug-90	Jan-80							
							Oct-89	Sep-77							
							Mar-80	Jan-77							
							May-74	Nov-76							
							Dec-73	Jan-76							
							Jan-73	Apr-74							
						Aug-98	Jul-69	Nov-72							
						Nov-87	May-62	Oct-67							
						May-84	Jun-61	Aug-62	Feb-91						
						Oct-76	Sep-57	Aug-59	May-90						
						Jun-75	Nov-51	Apr-56	Nov-81						
						Jun-74	Jun-49	Apr-51	Jun-80						
						Sep-72	Feb-48	Jul-40	Mar-79						
						Apr-70	Aug-47	May-39	Aug-77						
						Feb-69	Sep-38	Oct-38	Apr-71						
						Sep-39	May-38	May-33	Mar-68						
						Dec-31	May-32	Aug-32	Oct-39	Jan-89					
					Nov-73	May-31	Apr-32	Nov-24	Oct-31	Aug-75					
		Mar-74		Sep-81	Jul-66	Nov-29	Oct-21	Apr-22	Jun-31	Oct-59		Dec-76			
Oct-87		Jun-40	Aug-74	Nov-74	May-40	Nov-20	Feb-21	Oct-19	Jan-20	Jul-32		Apr-75		Feb-75	Jan-75

Exhibit 5.5 shows the distribution of *equity price changes* for months when the absolute size of the price change exceeded 7.5%. What is interesting about this chart is that the same years seem to appear several times across the picture. Large absolute changes in price do not seem to be generated evenly over time. They appear to be bunched together during periods of market volatility - such as 1931-2, 1940, 1974-5 and 1987. There are long gaps without large absolute returns. One widely-accepted explanation for these results is the notion that the *volatility of returns changes over time*.

The really interesting question is how we respond to the challenge posed by the extreme values plotted in the charts above. Should the equity model be capable of generating equity crashes with the same frequency as this historic data set? Alternatively, should we be prepared to live with an imperfect model for the sake of parsimony? As usual, the answer will depend on what we want to do with the model. If the problem under investigation is very sensitive to the presence of extreme values, it may be worthwhile ensuring that they are properly represented by the model (so long as we can agree on what “properly” means). If the analysis is not especially sensitive to outliers, the log-normal distribution may be perfectly adequate for the job.

5.2.1 THE MARKOV REGIME-SWITCHING MODEL

In order to mimic these important characteristics of equity returns, a *Markov regime-switching model* will be implemented. This type of model is able to generate returns distributions that are consistent with the properties of the empirical data. The basic idea is that returns are not drawn from a single normal distribution; rather there are *two* distributions at work generating the returns observed. The equity returns distribution is assumed to jump between two possible states over time. These states are often referred to as *regimes*. A *transition matrix* controls the probability of moving between states.

We use this regime switching approach to model log equity returns *in excess of the log return on a riskless asset*. In the current implementation, we have used a default-free short-dated discount bond to represent the riskless asset. For the sake of convenience, we set the term of this bond to equal the time increment at which the underlying stochastic variables are updated (for the results presented below this is one month). This also means that the riskless asset will behave very much like ‘cash’, and we will refer to the 'cash return' or 'risk-free return' synonymously. Of course, the choice of riskless asset is arbitrary - although we have chosen a discount bond with a term of one-month, it would be quite possible for us to construct the model using some other numeraire.

Thus the total return on equities in a given period of length Δt, *E(t)*, is the sum of the return on the short term discount bond, plus the excess return on equities, *X(t)*:

$$E(t) = \ln\left\{\frac{1}{P_{nom}(t-\Delta t, t)}\right\} + X(t)$$

The excess return on equities, *X(t)* has a Normal distribution, where:

X(t) has mean $\mu_{E,1}$ and variance $\sigma_{E,1}^2$ if the regime switching model is in *State 1*

X(t) has mean $\mu_{E,2}$ and variance $\sigma_{E,2}^2$ if the regime switching model is in *State 2*

The matrix of transition probabilities which determines how the equity return model switches between the two states can be denoted:

$$P = \begin{bmatrix} P_{11} & 1-P_{11} \\ 1-P_{22} & P_{22} \end{bmatrix}$$

where:

P_{11} = Prob{Model in *State 1* in period $(t, t+\Delta t)$ | Model in *State 1* in period $(t-\Delta t, t)$}

P_{22} = Prob{Model in *State 2* in period $(t, t+\Delta t)$ | Model in *State 2* in period $(t-\Delta t, t)$}

For a typical calibration, in one regime equity returns have a positive mean and low variability. In the second regime equity returns are – on average – lower and exhibit much higher volatility. Returns tend to be generated most often in the benign state, but there are bouts of high variability generated by the second state. Exhibit 5.6 gives an idea of what the two regimes might look like if the model is calibrated to the data shown above in exhibit 5.4. You can see that just over half of returns will be generated from the narrow distribution plotted with a solid line. Negative returns of –10% (per month) and below will be generated from the second high-volatility regime plotted with white circles. One challenge for the analyst who calibrates this model is to decide how long the process will stick in the volatile state: does the system visit the volatile regime frequently (but quickly jump out again) or does it visit only occasionally (and tend to stick there)? The answer to this question will have implications for the distribution of returns at other horizons.

EXHIBIT 5.6: REGIME SWITCHING ~ THE BASIC IDEA

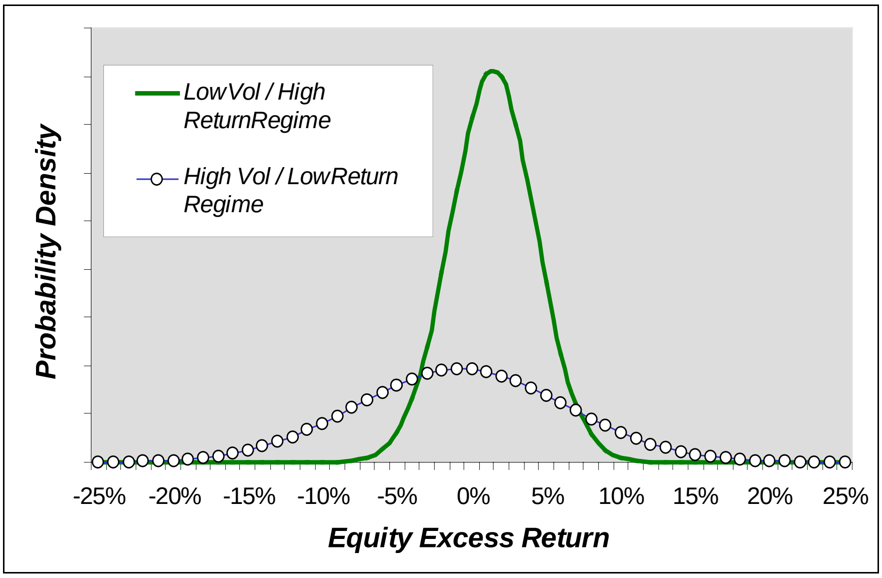

In section 7 a calibration will be presented for a monthly model for equity returns in excess of short-term interest rates. It is possible to calibrate and run the regime-switching model at any time horizon, although it must be appreciated that it is not possible to adjust parameters developed – for example – for monthly returns to a 12-month 2-state model. (Over 12 months there are actually 13 possible states for the 12-month return distribution).

5.2.2 EQUITY DIVIDEND YIELDS

The log of the equity yield is assumed to follow the continuous-time equivalent to a first-order autoregressive process (*AR1*) with long term mean μ_y and drift parameter α_y, so that:

$$d(\log\{y(t)\}) = \alpha_y (\mu_y - \log\{y(t)\})\, dt + \sigma_y . dZ_y(t)$$

where:

$y(t)$ = the equity dividend yield at time t.
α_y = the autoregressive parameter for the (log) equity dividend yield process.
σ_y = the annualised volatility (standard deviation) of the (log) equity dividend yield
$dZ_y(t)$ = is a random shock distributed $N(0,dt)$.

The model says that if $\log\{y(t-\Delta t)\}$ is below the long term average μ_y, then $\log\{y(t)\}$ will increase by (approximately) $\alpha_y\, \Delta t$ times the difference plus a random shock distributed with zero mean and standard deviation of $\sigma_y\, \Delta$t. The standard deviation for the log yield, σ_y equals the equity return volatility, i.e.

$\sigma_y^2 = \sigma_{E,1}^2$ if the equity return regime switching model is in *State 1*
$\sigma_y^2 = \sigma_{E,2}^2$ if the equity return regime switching model is in *State 2*

Suppose we already know the total asset return, $E(t)$, for the time period t-Δt to t and the dividend yield at time t. If the equity price at time t is $S(t)$, we can write :

$$(S(t) + S(t) . y(t). \Delta t) / S(t-\Delta t) = \exp\{E(t)\}$$
$$S(t) = S(t-\Delta t) . \exp\{E(t)\} / (1 + y(t). \Delta t)$$

So the dividend or income at time t:

$$D(t) = S(t) . y(t) . \Delta t$$

So long as a high negative correlation is imposed on the shocks to the yield model and the equity returns model, this specification produces equity yields that move – in the short term – with price changes. Over longer periods *strong* equity returns tend to be *followed by* above-average dividend growth (as the dividend yield reverts to mean) and equity declines will be *followed by* below-average dividend growth or falling dividends. The model is simple and has a natural economic interpretation.

6. Some Alternative Approaches to Calibration

6.1 Introduction

In section 5 a set of models was introduced for the purposes of modelling the behaviour of equities, bonds and inflation using Monte-Carlo methods. The models are only half the story. The scenarios generated by a stochastic model depend on both the model *structure* and the model *parameters*. In this section we explain why calibration is so difficult and then set out some parameter sets which we believe form a sensible starting point for current stochastic investigations.

Calibration is something of a black art. The analyst who faces up to the problem of how to calibrate his fancy new model might consider a number of solutions. Let us consider three alternative approaches.

6.2 Empirical Data

The most obvious approach to calibration is to look at the past behaviour of assets and other variables and to attempt to match this with the model. So, let us begin by reviewing some historic data.

6.2.1 A Long-Term Perspective

Exhibit 6.1 plots UK long-term interest rates since 1725. It shows a profile that falls within a range between 2.5% and 6.5% for over 200 years before the extraordinary experience of the past 50 years[10]. Many investors and financial analysts are still coming to terms with the current low-inflation environment. Headline UK inflation was running at an annual rate of less that 2% pa at the end of 1999. This is a striking contrast to the peak rates of inflation of over 20% pa reached in 1975 and 1980 and the accompanying high levels of interest rates. But it is clear from the chart that this experience is quite different to any other period over the past 300 years. The picture shows us that during the nineteenth century, having started the century at over 6% (towards the end of the Napoleonic Wars) government bond yields never exceeded 4% after 1830 and were never above 3.5% after 1850.

[10] *For interest, note that rates peaked in the eighteenth century at just below 9% in 1712.*

EXHIBIT 6.1: UK LONG-TERM INTEREST RATES (1725-2000) (ANNUAL RANGE OF VALUES)

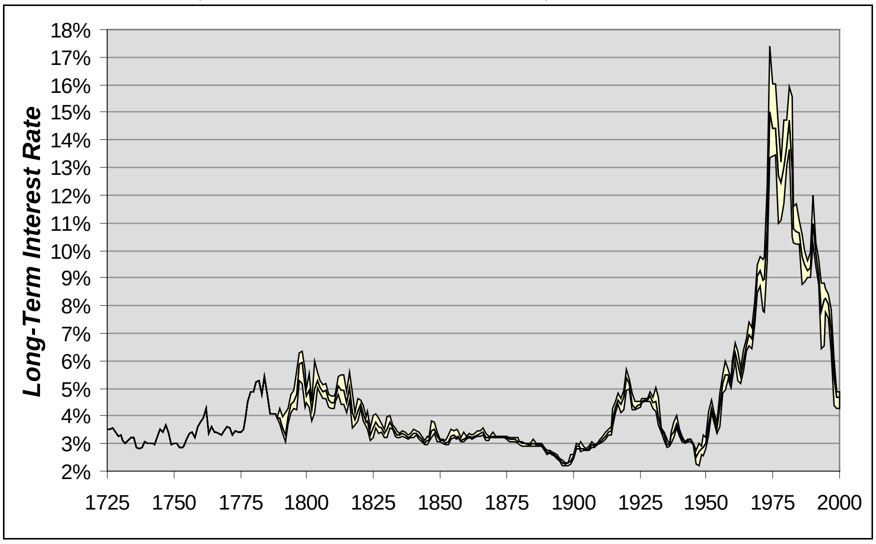

Notes: (a) Source is "A History of Interest Rates", Homer & Sylla, 3rd Edition

(b) The annual range of yields is plotted for years after 1790.

This sort of picture begs an obvious question: *If we calibrate to history, which period should be used – the whole period, the low-interest rate era before 1945 or the 1945/95 high interest rate period?*

Let us make one other point. We are *not* calibrating the model to plan over the next 300 years. Even for very long-term planning problems the horizon rarely extends beyond 30 or 40 years.

An alternative way of viewing the interest rate history is to plot the frequency distribution of rates. This provides a straightforward way of measuring the frequency with which rates have reached different levels over the 275-year period. You can see that the bulk of the distribution falls between 2% and 6%. The frequency with which long-term rates have exceeded 6% is 12.5% (one year in every eight). Rates have fallen between 3% and 4% in over 40% of the past periods analysed.

EXHIBIT 6.2: EMPIRICAL DISTRIBUTION OF UK LONG-TERM INTEREST RATES (1725-2000)

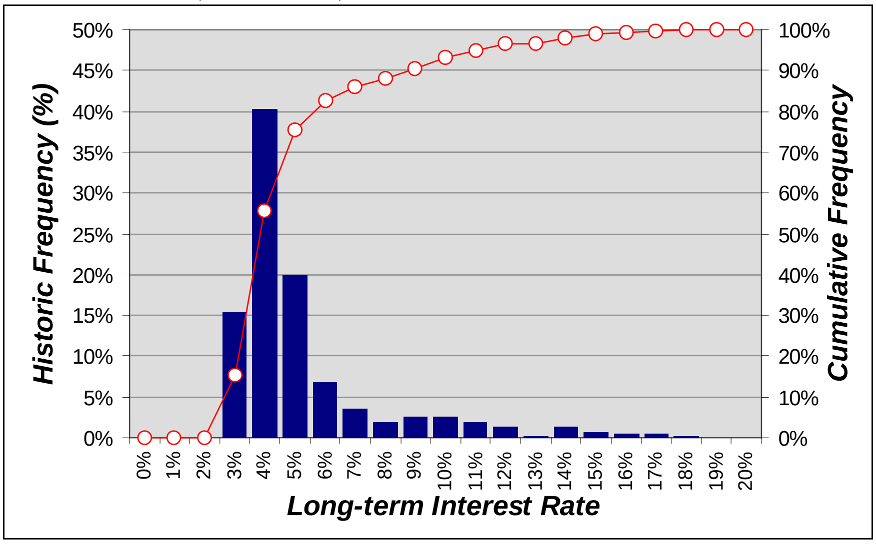

6.2.2 20TH CENTURY INTEREST RATES

Someone might argue that interest rates in the eighteenth century have little relevance today. Indeed, the author's calibration of the widely used *Wilkie* model implicitly suggests that the past century *is* an appropriate period for model calibration. Exhibit 6.3 shows the historical distribution for end-month short rates since 1918. It suggests that the range of plausible rates lies between 1% (the most frequent historical observation) and 18% - the peak for short rates in the 1970s. The distribution looks distinctly bi-modal suggesting that the data was generated from two quite distinct periods[11]. Given the extraordinary post-war inflation and the resulting impact on interest rates and economic policy-making, this should be no great surprise.

[11] *There are techniques for capturing this "regime-switching" effect, but we have not attempted to implement them in the interest rate model calibrations presented in this note.*

EXHIBIT 6.3: EMPIRICAL DISTRIBUTION – SHORT RATES (1918-1999)

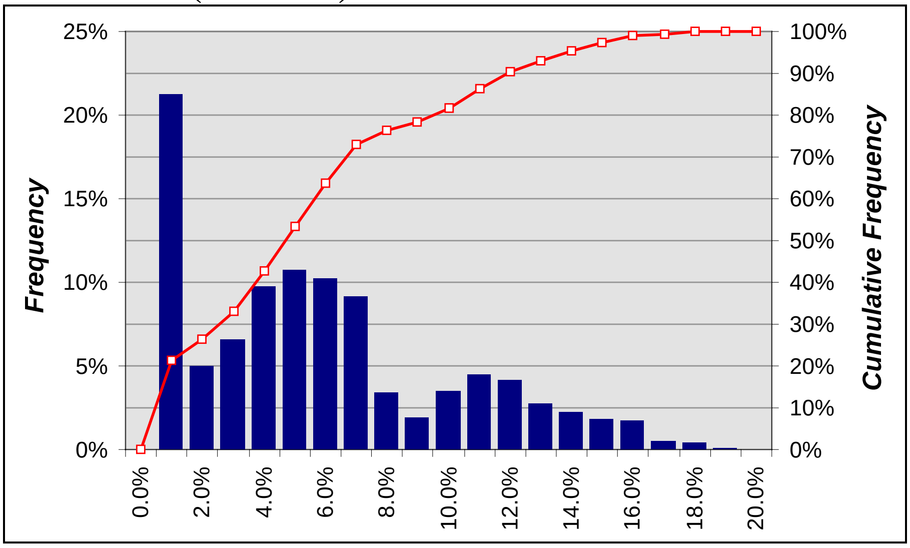

Consider a long-term rate of interest that is plotted in the chart below. For reference we have plotted the distribution of end-month long-term bond yields from 1918 to date. Like the short-term interest rate, there is a suggestion that the distribution is bi-modal (or tri-modal) as a result of the vastly different inflation experiences before and after 1945.

EXHIBIT 6.4: EMPIRICAL DISTRIBUTION – LONG BOND YIELDS (1918-1999)

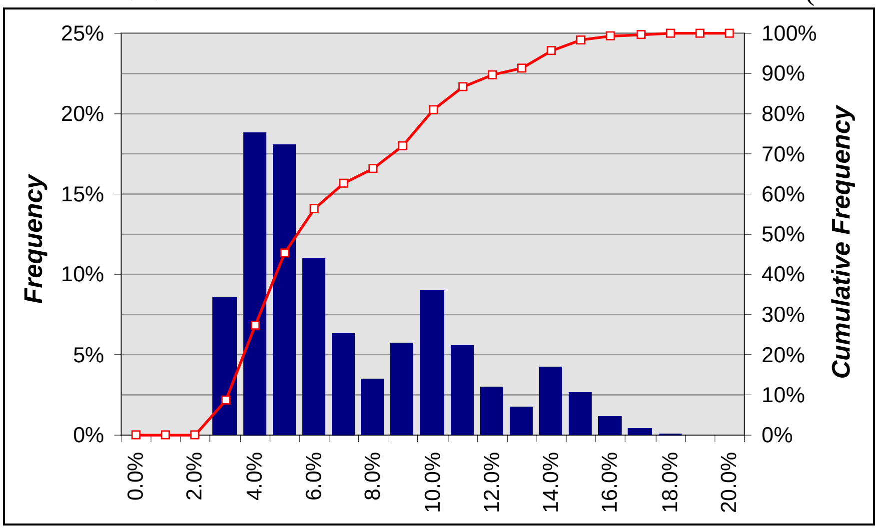

6.2.3 HISTORIC INTEREST RATE VOLATILITY

Aside from the unconditional distribution of rates (our estimate of the probabilities of rates at a long horizon), any model should produce plausible behaviour for rate *changes* over short horizons. Exhibit 6.5 tabulates the historic volatility of interest rates over the past century. Two volatility measures are shown. First, the annualised standard deviations of *absolute* rate changes are shown for 3-month rates and for a long-term bond yield. Short rates have been more variable than long-term rates over this period with the annual standard deviation averaging out at a little above 2% pa. The bottom of the table shows the standard deviation of rate changes measured in proportional terms (which is the conventional way traders use to express interest rate volatility for short-term instruments).

EXHIBIT 6.5: HISTORIC ANNUALISED INTEREST RATE VOLATILITY

	Time Period	*Short Rate (3M)*	*Long-Term Rate (Consols Yield / 20Y Gilt)*
Absolute Volatility	*1900-99*	1.83%	0.95%
	1945-99	2.20%	1.26%
Proportional Volatility	*1900-99*	39%	11%
	1945-99	35%	14%

It is interesting to note that these measures (of volatility) do vary significantly when they are measured over different time periods. Exhibit 6.6 plots the rolling 10-year annualised standard deviation of rate changes for short rates and long-term bond yields. You can see that rate volatility itself varies markedly over time. Again, the peak levels of volatility occurred in the 1970s and 1980s.

EXHIBIT 6.6: ANNUALISED 10-YEAR STANDARD DEVIATION OF RATE CHANGES

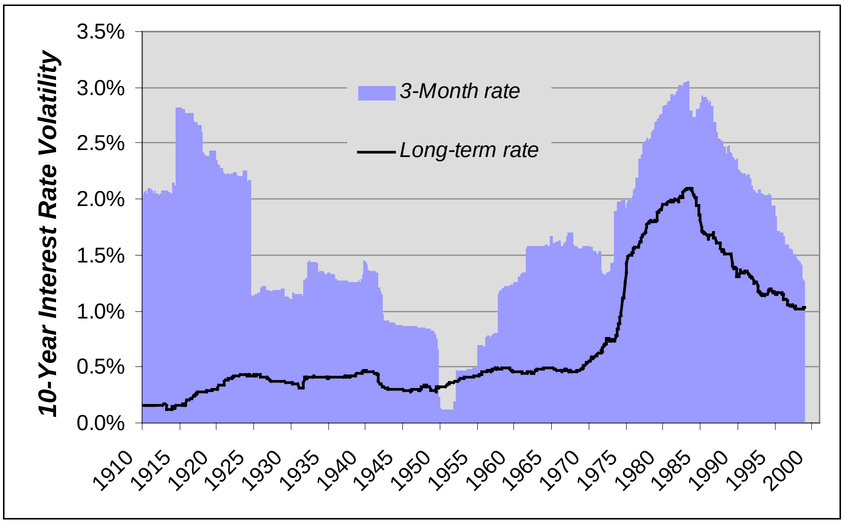

We need to ask whether our chosen model calibration can and should mimic these episodes.

6.2.4 HISTORIC CURVE SHAPES

Finally, when we review past data on interest rate behaviour, it is interesting to examine the *shape* of past term structures in addition to the other properties we have reviewed. As exhibit 6.7 illustrates, the term structure can assume quite a wide range of shapes. Over the 16-year period shown (which really is not very long for our purposes) we can observe upward and downward-sloping curves, humped curves as well as some slightly saucer-shaped curves. When a calibration for the model is selected, we should ask whether it is capable of producing the richness we can see in the real-world data.

EXHIBIT 6.7: HISTORIC SPOT RATE CURVES (1982-98) (END-JUNE & END-DECEMBER)

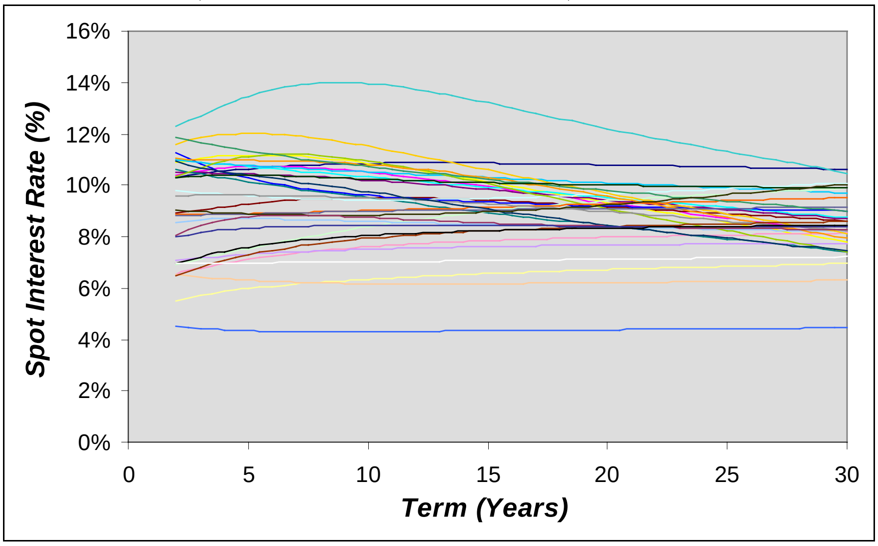

Source: Bank of England

6.2.5 INFLATION

Given the fundamental link between inflation rates and interest rates it is not surprising that we face some very similar issues when the inflation data is examined. Exhibit 6.8 shows the path of inflation rates this century spanning two world wars, the great depression, exit from the gold standard and the post-1945 inflation.

EXHIBIT 6.8: 20TH CENTURY ANNUAL INFLATION RATES

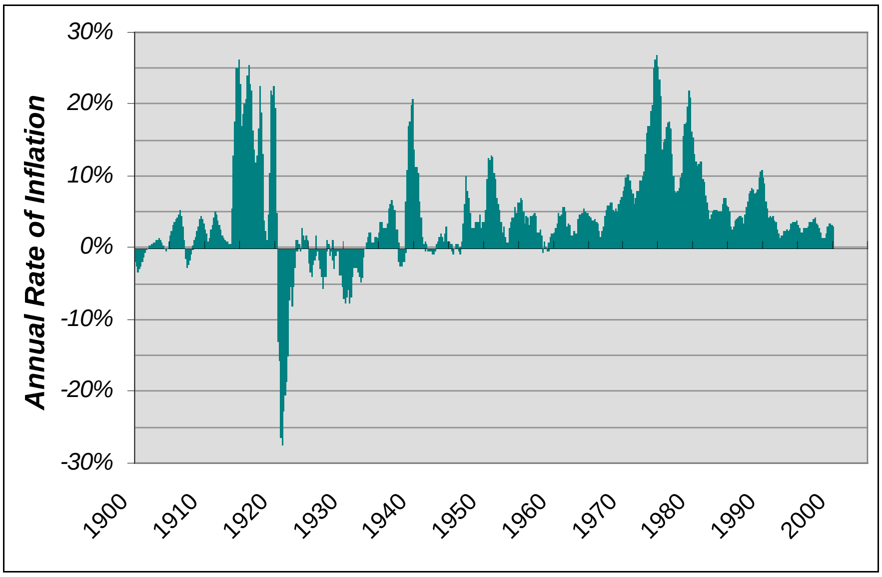

It is interesting to speculate on how much relevance this experience really has for the analyst who is concerned with forecasting the distribution of inflation for the next 20 years. Exhibit 6.9 suggests that there is a wide range of values the analyst could choose for the average rate of inflation, depending on which historic period is selected to calibrate a model.

EXHIBIT 6.9: SOME INFLATION AVERAGES

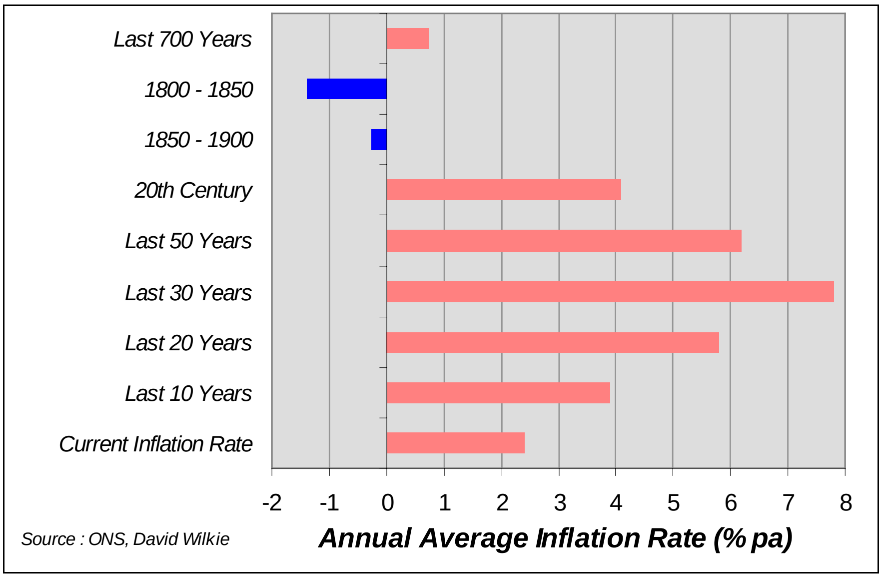

WhatInflationChart.xls

EXHIBIT 6.10: UK INFLATION & LONG-TERM INTEREST RATES
(PRE & POST 1945)

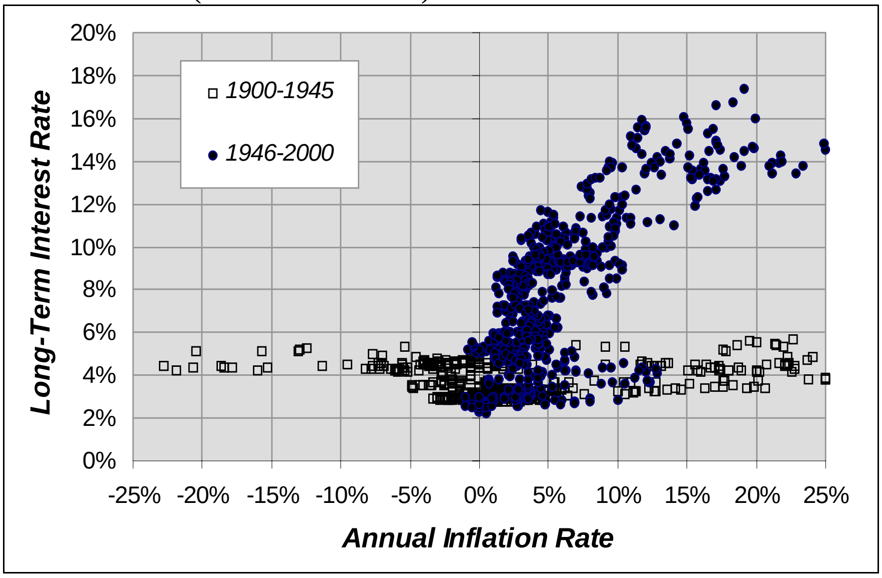

It is worth remembering that we do not simply want to select a long-term average for our model (although its structure means that this is less important than for a simple *AR1* model), it is also desirable to understand the relationship between inflation and other variables. In exhibit 6.10 the past relationship between long-term interest rates and inflation is plotted. You can see that – in the period since 1945 – high rates of inflation have tended to be accompanied by high long-term interest rates as bond investors' inflation expectations are raised. Arguably, the high long-term bond yields of the 1970s and 1980s also included an embedded 'inflation risk premium'.

Inflation also hurts equity returns. Exhibit 6.11 tells us that equities tend to perform relatively poorly in high-inflation environments and when the rate of inflation is accelerating.

EXHIBIT 6.11: UK EQUITY REAL RETURNS & INFLATION
(1919 – 1999)

	Real Equity Return	*Real Bond Return*	*Number of Years*
Falling Prices *Less than -1%*	13.9%	11.1%	9
Price Stability *Between -1% and +1%*	16.6%	14.5%	10
Low Inflation *Between 1% and 4%*	12.4%	9.5%	28
Moderate Inflation *Between 4% and 8%*	5.3%	4.1%	19
High Inflation *Greater than 8%*	-6.7%	-18.4%	15
Falling Inflation Rate	12.0%	6.0%	41
Rising Inflation Rate	3.2%	-1.0%	40

EXHIBIT 6.12: ANNUALISED QUARTERLY INFLATION RATE DIFFERENCES (5-YEAR PERIODS)

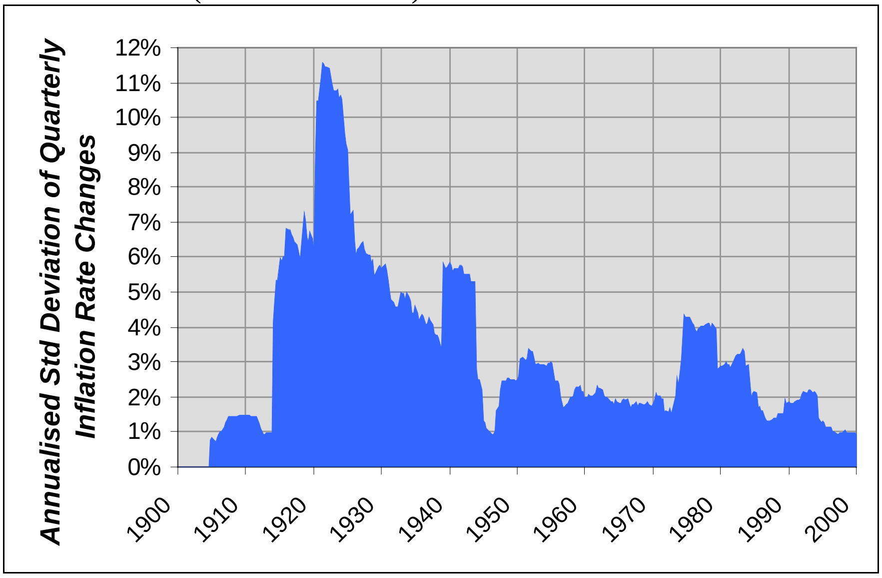

ChartInflationVolatility.xls

The variation in inflation from period to period will feature in the model. The chart above plots the annualised standard deviation of inflation rate changes. It tells us that the variability of the inflation rate (measured in this way) is at its lowest level for over 50 years. It is worth pointing out that the recent average of 1% p.a. is well below the original or revised estimates for use with the Wilkie model for inflation volatility (the QSD parameter).

6.2.6 EQUITY RETURNS

Some important properties of equity returns are illustrated in section 6.2. Exhibit 6.13 summarises the statistical characteristics of excess returns calculated at different frequencies.

EXHIBIT 6.13: SUMMARY STATISTICS FOR EXCESS RETURNS TO UK EQUITIES AT VARIOUS FREQUENCIES (1901-2000)

	Monthly	*Quarterly*	*Annual*	*Annual (excl. 71-80)*
Basic Statistics				
Number of Observations	1200	400	100	90
Mean	0.0038	0.0115	0.0472	0.0471
Standard Deviation	0.0447	0.0821	0.1891	0.1430
Standard Deviation (Annualised)	0.1550	0.1641	0.1891	0.1430
Skewness	-0.219	-0.187	-0.551	-0.231
Relative Kurtosis	10.54	7.45	6.36	0.547
Autocorrelation				
returns	0.104	0.046	-0.094	0.109
absolute returns	0.264	0.174	0.436	0.325
squared returns	0.187	0.065	0.537	0.240

The most notable features of the data are:

- the mean log return is almost 5% pa in excess of short-term interest rates.
- the annualised standard deviation rises from 15.5% pa at a monthly frequency to almost 19% pa at an annual frequency. However, if the returns from 1971 to 1980 are dropped out of the data this annual volatility drops to a little over 14% pa.
- the skewness statistic is used to assess the symmetry of the distributions. A sample drawn from a normal distribution would be expected to show skewness close to zero. All of the estimates are negative.
- the years 1974/5 exert a big influence on the statistics. Notice how the annual autocorrelation estimates change when we drop the 1970s data.
- kurtosis is shown relative to the normal distribution. All of the kurtosis estimates are greater than that expected from a normal distribution. It is very clear that the returns generating process is very different to the normal distribution.
- If the UK equity market really did conform to a normal distribution with an annualised standard deviation of 18% p.a. what would we expect to see? Statistical theory tells us that, over the past 100 years, we would expect to observe perhaps one monthly excess returns of less than -15%. We have actually experienced 7.

6.3 MARKET DATA

There is a second approach that an analyst might adopt in order to calibrate a model. It is to use market information to help find appropriate parameters. A proponent of the market-based approach would argue that market asset and derivative prices contain information that can be used to guide the calibration choice. It turns out that market prices must be interpreted carefully when they are used in the calibration process.

Let us look at some examples of how market data can be used.

6.3.1 UNDERSTANDING THE TERM STRUCTURE

Before we discuss market-implied data for interest rates, it is worthwhile reviewing some of the theories that have been proposed to explain the shape of the term structure. Your view of these theories will bear on how you choose to interpret some of the mixed signals produced by this approach to calibration.

There are a number of "classical" theories of the yield curve. These theories aim to tell us about the determinants of the level and shape of the term structure. The theories earn the "classical" label because they have been around for a long time. The economist's classical approach provides three broad explanations (and lots of variations), as follows:

- The *expectations theory* can be interpreted in a number of ways that are not quite equivalent. One interpretation says that forward interest rates measure market participants' aggregate expectation for the corresponding future short-term rate of interest. The theory suggests that an upward-sloping curve means that short-term interest rates are expected to rise. Conversely a downward-sloping curve implies an expectation of falling short-term rates.

- The basic idea behind the expectations theory is that investors select bonds purely on the basis of their expected returns – they don't demand a premium for accepting risk. Alternatively, the *liquidity premium theory* suggests that, if investors have a preference for stable portfolios compared to volatile portfolios, they will also exhibit a preference for short-dated bonds compared to volatile long-dated bonds. In order to induce investors to hold bonds with long maturities, investors must expect to receive a higher rate of return than on short maturity instruments. This increased expected return is called a "*liquidity premium*" – the extra expected return for bearing the interest rate risk of long-dated instruments. The theory tells us that long-term bonds should, on average, offer higher returns than cash investments. As a result, the "natural" shape of the yield curve should be upward sloping.

  Notice that, if the liquidity premium theory really does hold true, we should expect long-term forward rates to be a biased expectation for future short rates. In other words, if we use forward rates to forecast the short rate we would over-estimate the short rate (on average) by the liquidity premium.

- The liquidity premium theory suggests that investors are averse to variability in portfolio values. But suppose that investors do not just look at the value of their assets, but at the total value of assets *and liabilities*. A pension fund may have long-term liabilities for which the closest match is a long-term bond. The asset that produces the lowest variability in assets plus liabilities may turn out to be a long-term bond rather than cash. You can see that the "preferred" or "natural" maturity of any investor who is averse to total portfolio variability will match the maturity of the investor's liabilities (where the *total portfolio* includes assets and liabilities). Now, in just the same way that the liquidity premium theory rewards long-term instruments with a risk premium that rises in line with term, the *preferred habitat theory* rewards bonds which fall outside the preferred habitat of investors with a premium. Investors must be compensated for moving away from their preferred maturity. Notice that this could mean that the risk premium actually decreases (or becomes negative) with the term of bonds, if bond issuance (of long-dated) paper fails to match a natural appetite among investors.

The three theories outlined above should not be viewed as mutually exclusive. Most commentators (and researchers) believe that all of the factors behind the different theories play a part in the determination of interest rates from time to time.

Now consider exhibit 6.14. It shows a term structure for Sterling government bond yields. The rates plotted with white circles are continuously compounded (log) zero-coupon yields consistent with observed gilt prices. The short rate is a little over 6.5%. The 20-year spot rate is around 4.5% with the 30-year rate at 4%. Note that, although these rates are expressed in terms of notional discount bonds, they are the most convenient way of measuring the term structure. We can think of any coupon government bond as a package of discount bonds. The notional schedule shown below can be used (with the coupon and principal bond cash flows) to replicate fairly closely the observed prices of gilts last October.

The solid line above the zero-coupon curve shows the par yield curve. This is the schedule of yields for notional coupon bonds trading at par. Our analysis shows that a 4.7% 20-year gilt would trade at par under the zero-coupon term structure.

EXHIBIT 6.14: GOVERNMENT BOND INTEREST RATES (END-OCTOBER 2000)

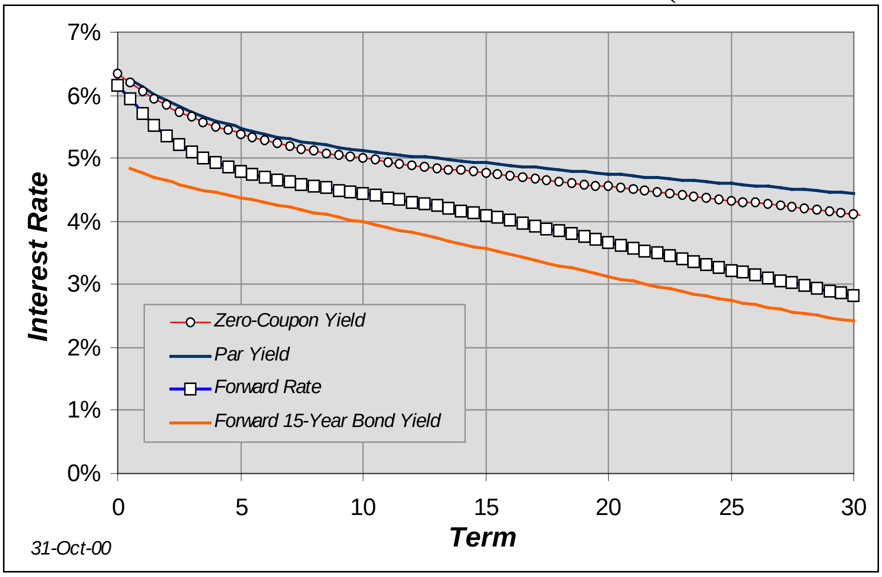

Source: Barrie & Hibbert

The chart also plots *forward rates of interest* for the government yield curve (6-month rates). You can see that, for very short-term rates these are close to spot rates, but for longer maturities they are lower than spot rates. The forward rate at 10 years was 4.5% and at 25 years the forward rate is only a little above 3%. A believer in the *expectations hypothesis* would say that short rates were anticipated (on average) to fall to 3% in 25 years time. Notice that – because the term structure of spot rates is downward sloping – forward rates fall more rapidly with term than spot rates.

Finally, *forward 15-year bond yields* have been plotted. This is the notional yield on a 15-year bond that we could buy forward at various terms. Again, a believer in the *expectations hypothesis* would view these as unbiased expectations of 15-year bond yields at the terms shown. For example, he might say: "*the expected yield on a 15-year gilt at the end of 10 years is 4%. The equivalent figure for 20 years is 3%*". We should add that a believer in the *liquidity premium theory* would view these as over-estimates for future gilt yields – biased upwards by an impossible-to-measure risk premium. On the other hand, the proponent of the *preferred habitat theory* would tell us that the forward 15-year bond yields do not tell us anything useful about the expected level of long-term yields. Rather they simply reflect supply and demand for gilts. The low levels of forward rates might only reflect an imbalance between the demand for long-dated fixed income instruments and supply from the British government and other high-quality bond issuers. Anyone looking to the market for clues to how they should calibrate a model of interest rates should not necessarily aim to match these forward rates in a simulation exercise.

Another reason why you might be suspicious of long-term forward rates as a measure of long-term forward rate expectations is the marked difference across different currencies. Long-term forward rates have been consistently lower in the Sterling bond sector than for Euro- or dollar-denominated assets. It is difficult to explain away these differences with rationale economic arguments. Rather they support the view that Sterling rates are currently biased (as a measure of expected future short rates) by strong demand for long-dated paper.

EXHIBIT 6.15: ALTERNATIVE ASSUMPTIONS FOR THE FUTURE AVERAGE LONG-TERM INTEREST RATES

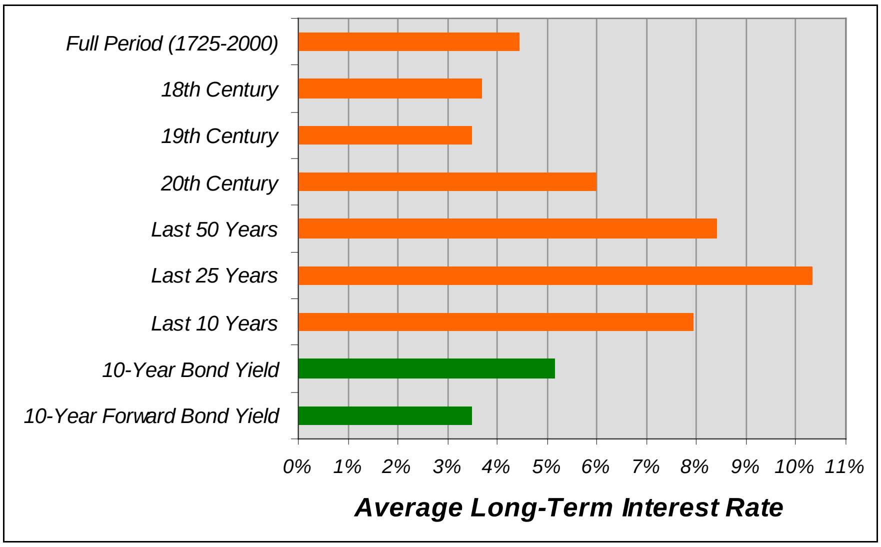

WhatGiltYieldChart.xls

6.3.2 SWAPTION IMPLIED VOLATILITY

The analyst seeking information to calibrate a model looks to the yield curve for investors' expectations for interest rates. By contrast, the options markets offer the enticing prospect of information on the *distribution* of future interest rates. Some analysts believe that it is possible to back out the probability distribution of the underlying asset implied by options by considering how the implied volatility of options varies with the options' strike prices. For example, by considering swaption prices on the 15-year swap rate with a range of strike rates and a common expiry date, it is possible to calculate a probability distribution of 15-year swap rates implied for the expiry date.

This approach can provide insight into market expectations, but we also need to be wary of the assumptions implicit in using option prices in this way – in particular, the approach assumes that options can be delta-hedged without cost or risk, and so option prices are assumed to be solely a function of the market's (risk-neutral) expectations of the future behaviour of the underlying asset.

Exhibit 6.16 below shows the distribution of 15-year swap rates on the 28th February 2002 as implied by swaption prices on 7th November 2000[12].

EXHIBIT 6.16: OPTION-IMPLIED PROBABILITY DISTRIBUTION FOR 15-YEAR SWAP RATES

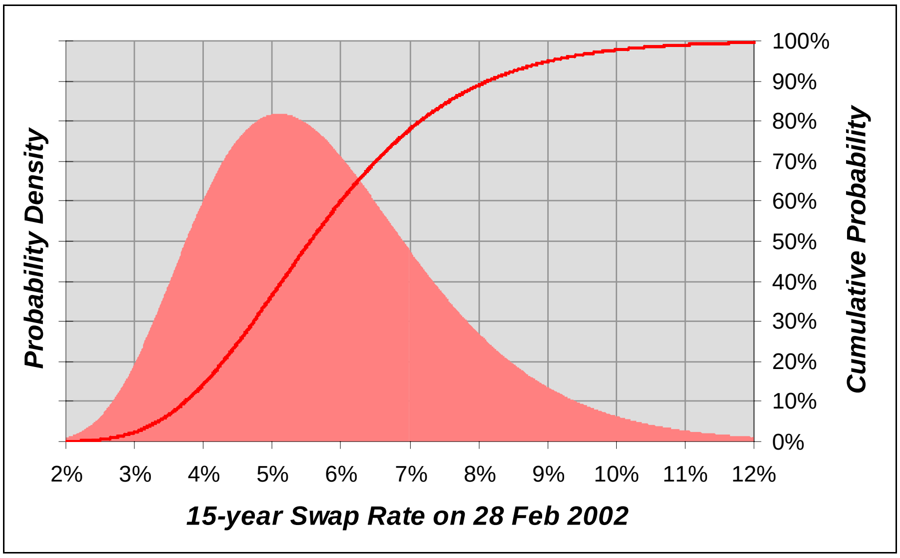

Source: SwapRatesNov_2000_1Yr

The forward swap rate was 5.79% and, by construction, this is equal to the mean of the implied-distribution. However, the chart suggests that swaption prices imply significant potential deviation from this expected swap rate – the standard deviation of the distribution is 1.75%.

---

[12] *Swaptions are not exchange-traded, and so OTC prices have to be used. We are grateful to GenRe Financial Products for supplying us with representative swaption prices.*

6.3.3 EQUITY IMPLIED VOLATILITY

We can also apply this methodology to equity prices, by using traded option prices. Below is an option-implied probability density for the 21st March 2002 as implied by LIFFE traded option prices on the 20th March 2001. Strictly speaking, this is the *risk-neutral* probability density – it assumes the expected return on equities is equal to the risk-free rate. However, we can transform this density into a real-world distribution with a given equity risk premium by making some assumptions regarding investors' utility functions. Such a transform would shift the distribution to the right, and would change the shape of the distribution, though not usually very significantly.

EXHIBIT 6.17: OPTION-IMPLIED PROBABILITY DISTRIBUTION FOR FTSE 100

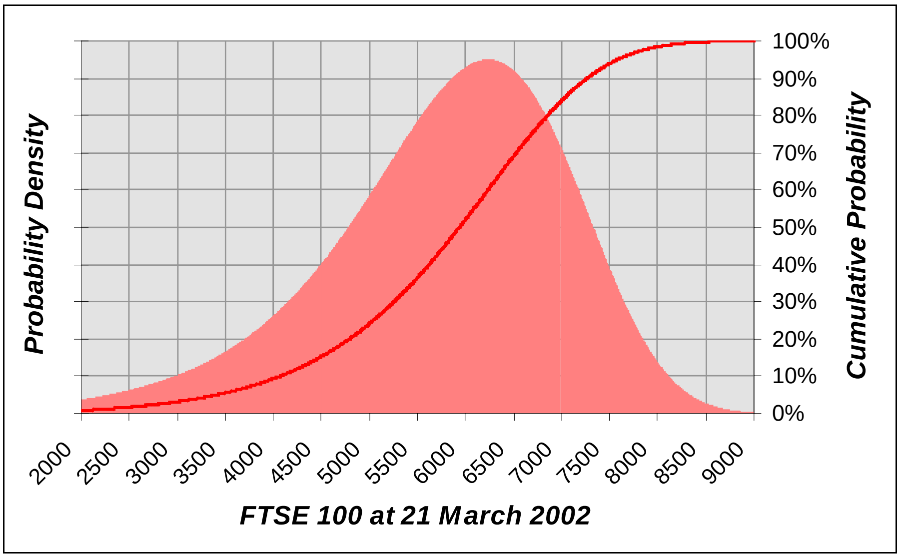

Source: UKEquities0301_0302

Note that, like historical equity market data, the option-implied distribution exhibits negative skew (unlike the lognormal distribution). The left-hand tail is significantly fatter than would be implied by a lognormal distribution, and suggests (to the proponent of this approach) that there is a probability of around 10% that of the FTSE 100 index will fall below 4000 at 21st March 2002.

6.4 EXPERT OPINION

A final potential source of information that can be used to judge the reasonableness of the distributions generated by a model is the judgement of experts. This information can be accessed in different ways:

- Published independent forecasts. This information comes in the form of expected asset returns and, on occasion, distributional information. A good example of this type of data is the inflation forecast distributions published by the Bank of England. The bank estimates a 90% confidence range of for the inflation rate in 2 years time between 0.9% and 3.7% with a median value (in line with the bank's target) of 2.5%. Crudely, this range implies an annual standard deviation for the inflation rate of 0.6% for the next 2 years.
- In-house experts. Most large financial institutions employ economists and forecasters. They are capable of commenting on the location and shape of distributions produced by models.

Although there are no clear rules for how expert opinion is incorporated into the calibration process, it is important to understand that it is a potentially hugely valuable source of insight. As we have already seen, naïve calibration to historic data or market information is likely to produce poor model output. The expert's opinion provides a useful check against this risk.

Let us now focus on two specific calibrations for the model.

7. A CALIBRATION

7.1 A PLAUSIBLE PARAMETER CHOICE

There is a wide range of possible parameter choices for the model specified in Section 5. Here we review two possible choices that we judge to be a reasonable starting point in the light of the analysis presented above. The set-up of the model is shown in exhibits 7.1A and 7.1B below. Exhibit 7.1A sets out parameter values for a 'base case' and a modified set-up for the model where positive interest rates are produced at all times.

EXHIBIT 7.1A: BASE CASE CALIBRATION & WITH REFLECTION

		Calibration Case	
	Parameter	*A. Base Case (Yield Curve Reflection OFF)*	*B. Positive Interest (Yield Curve Reflection ON)*
Real Rates	α_{r1}	0.25	0.25
	α_{r2}	0.05	0.05
	σ_{r1}	0.005	0.005
	σ_{r2}	0.01	0.01
	μ_r	0.025	0.0525
	γ_r	0	-0.125
	$r_1(0)$	0.025	0.025
	$r_2(0)$	0.025	0.0275
	b_{r1}	No	-0.05
	b_{r2}	No	0
	TP(ret)	0	0.0275
	TP(y)	-0.0202	0.0073
Inflation Expectations	α_{q1}	0.3	0.3
	α_{q2}	0.1	0.1
	σ_{q1}	0.008	0.008
	σ_{q2}	0.012	0.012
	μ_q	0.025	0.0433
	γ_q	0	-0.125
	$q_1(0)$	0.025	0.025
	$q_2(0)$	0.025	0.0283
	b_{q1}	No	-0.05
	b_{q2}	No	0
	TP(ret)	0	0.0183
	TP(y)	-0.0076	0.0108
Equity Returns	$\mu_{E,1}$	0.118	0.118
	$\sigma_{E,1}$	0.098	0.098
	$\mu_{E,2}$	-0.136	-0.136
	$\sigma_{E,2}$	0.244	0.244
	$p(1,1)$	0.929	0.929
	$p(2,2)$	0.879	0.879
	$\pi(1)$	0.63	0.63
	$\pi(2)$	0.37	0.37
Yields	α_y	0.25	0.25
	$\sigma_{y,1}$	0.098	0.098
	$\sigma_{y,2}$	0.244	0.244
	μ_y	*log* (.035)	*log* (.035)
	$y(0)$	*log* (.025)	*log* (.025)

Force +ve Rates	No	Yes
Correlation	A	A

Exhibit 7.1B provides one example of a set of correlation coefficients which will control the relationships between the stochastic innovations (the *dZ*'s):

EXHIBIT 7.1B : CORRELATION MATRIX

	Z_1	Z_2	Z_q	Z_μ	Z_E	Z_Y
Z_1	1	0	0.25	0	-0.25	0.25
Z_2		1	0	0.25	-0.25	0.25
Z_q			1	0	-0.25	0.25
Z_μ				1	-0.25	0.25
Z_E					1	-0.95
Z_Y						1

Calibration *A* represents a very simple base case. We have set both g_r and g_q to zero, implying a zero term premium on index-linked bonds and a zero inflation risk premium. In other words, this first calibration should generate scenarios where expected returns on both nominal and index-linked bonds of all terms are the same. Importantly, this first calibration will also produce some negative nominal interest rates. The advantage of starting with this simplistic calibration is that much of the complexity within the model relating to the implementation of risk premia falls away. Using this simple calibration allows us to ensure that the model is producing the results expected.

In calibration *B*, we have incorporated risk premia for both interest rates and inflation, by setting g_r and g_q to equal -0.125. These risk premia have the effect of introducing a term premium to the term-structures for both real interest rates and inflation expectations. In Exhibit 7.1A, we express these term premia in terms of the continuously compounded (log) rate of return, *TP(ret)*, and the zero-coupon yield, *TP(y)*. Furthermore, within calibration *B* we have used separate devices to control the value of real interest rates and inflation expectations, and to force *nominal yields* to remain positive:

- We have imposed minimum barriers on the stochastic variables describing the term-structures of real rates and inflation expectations: $r_1(t) \geq b_{r1}$, $r_2(t) \geq b_{r2}$, $q_1(t) \geq b_{q1}$, $q_2(t) \geq b_{q2}$. In this calibration, we have used b_{r1} and b_{q1} to ensure that our model scenarios exclude both very large negative real rates, and very large negative inflation rates. b_{r2} and b_{q2} ensure that the very long-term real interest rate and inflation expectation do not fall below zero.

- We have also 'reflected' the *nominal* yield curve off zero:

$$\text{if } r_1(t) + q_1(t) < 0, \text{ then set } q_1(t) = -r_1(t) + 0.0001$$

$$\text{if } r_2(t) + q_2(t) < 0, \text{ then set } q_2(t) = -r_2(t) + 0.0001$$

We have touched on the limitations of these adjustments in section 5.1.5. They have the obvious benefit of removing implausible outcomes from the model output. The disadvantage of this approach is the introduction of inconsistency into the model (between its solid theoretical foundation) and the practical implementation. We believe that - for many stochastic simulation applications - these limitations are outweighed by the benefits in terms of generating plausible individual scenarios.

We have used the model to simulate 1000 scenarios over a horizon of 30 years. The simulation trials were built up in time increments of 1 month (i.e. dt = 1/12), although the results were recorded on an annual basis. The choice of time increment and output frequency is entirely flexible. The entire simulation exercise took a few minutes to run on a desktop PC, and produced simulated results for every asset, in every output time period, for each of the simulation trials. As we have already mentioned, generating output in this format means there is huge flexibility in the way the results can be presented.

7.2 SUMMARY STATISTICS : SAMPLE MEAN RETURNS & STANDARD DEVIATIONS

First, it is worthwhile considering some summary statistics that describe the distributions of asset returns generated from the model. These summary statistics are provided for the two alternative calibrations:

EXHIBIT 7.2A: SAMPLE MEANS & STANDARD DEVIATIONS (REFLECTION *OFF*)

Asset	*Log Return (% pa.)*	*Real Log Return (% pa.)*	*Ordinary Expected Return (% pa.)*	*Std. Dev. (% pa.)*	*10Y Historic Std. Dev.*
Equities	7.5	5.0	9.2	18.9	18
Cash	5.0	2.5	5.0	3.5	1
Constant Maturity Coupon Bond	4.2	1.7	5.0	10.6	9
Constant Maturity Coupon IL Bond	4.5	2.0	5.0	8.9	7
Underlying Inflation	2.5	-	2.5	2.3	na

EXHIBIT 7.2B: SAMPLE MEANS & STANDARD DEVIATIONS (REFLECTION *ON*)

Asset	*Log Return (% pa.)*	*Real Log Return (% pa.)*	*Ordinary Expected Return (% pa.)*	*Standard Deviation*	*10Y Historic Std. Dev.*
Equities	8.4	5.5	10.1	18.8	18
Cash	5.9	3.0	5.9	2.8	1
Constant Maturity Coupon Bond	6.5	3.6	7.0	9.5	9
Constant Maturity Coupon IL Bond	6.3	3.4	6.6	8.3	7
Underlying Inflation	2.9	-	3	1.9	na

Firstly note that in calibration A, the expected ordinary returns on all the fixed interest assets is equal to 5%, the expected rate of return on cash. Calibration B produces a term premium on 20-year bonds of about 1%, and on 20-year index-linked bonds of about 60 basis points (i.e. for 20-year bonds there is an inflation risk premium of 40 basis points). These results show that the effect of excluding negative nominal rates, as described above, is to increase the expected returns on the various assets whilst producing similar standard deviations. It is quite possible to offset this effect by reducing the μ_r and μ_q parameters.

7.3 EQUITIES

Exhibit 7.3 shows one way of plotting results for log (continuously compounded) annual equity returns and comparing them to the historic and market-implied distributions. Note that the solid red band in the centre of the chart shows the spread from 25th to 75th percentile for annual log returns. The outer pink bands plot the 5th/25th and 75th/95th ranges. Notice that the alternative model calibrations (*A* & *B*) produce very similar distributions for equity returns, and this distribution appears to lie somewhere between historic experience and the distribution implied by current market prices (as of end-March 2001).

EXHIBIT 7.3: DISTRIBUTION OF 1-YEAR EQUITY RETURNS

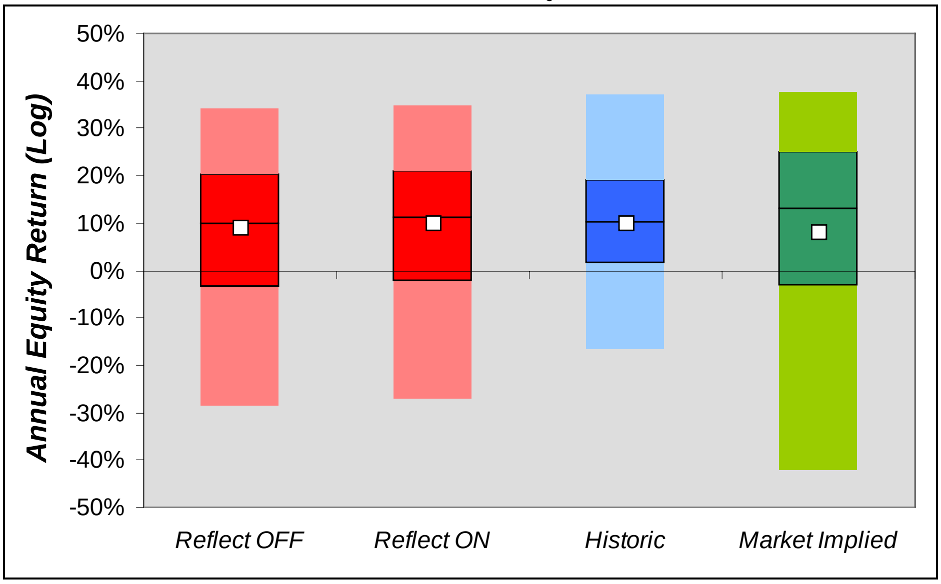

Note: Historic period=1900/2000; Market-implied data for FTSE options @ 31/03/2001

EXHIBIT 7.4: UNCONDITIONAL DISTRIBUTION OF UK EQUITY RETURNS

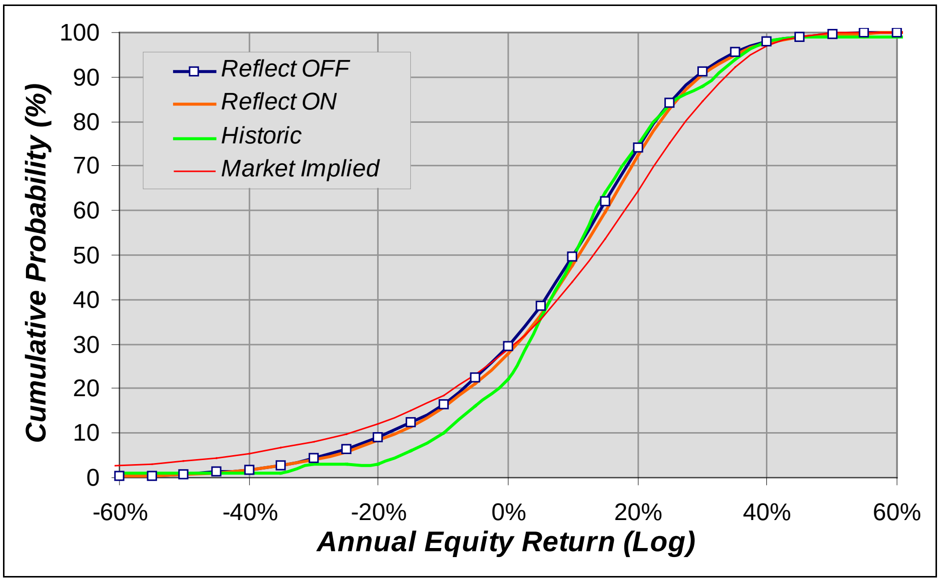

Exhibit 7.4 illustrates the entire cumulative probability distribution, rather than just five selected percentile points. Again, you can see that the distributions produced by the model appear to fall somewhere between the historic distribution and the current market

implied distribution. An important feature highlighted by this chart, which was not evident in exhibit 7.3, is the very heavy downside tail in the current market implied distribution of equity returns for the period 31/3/2001 to 31/3/2002. According to option prices at end-March, the market appeared to be assigning a probability of roughly 5% to a return of -40% over the following 12 months. Although, the current calibration of the regime-switching model for equity returns does not capture the size of this tail, it does very much better than a simple lognormal assumption. Furthermore, you might get much closer to the market implied distribution by adjusting the parameters of the regime-switching model if judged appropriate.

7.4 SHORT-TERM INTEREST RATES & CASH RETURNS

In exhibit 7.5 we consider the distribution of short-term nominal interest rates simulated by the model, and compare this distribution with the distribution of interest rates over two historic periods. The yellow distribution is positioned well to the right of the green distribution, reflecting the fact that post-war interest rates have been very high relative to pre-war rates (and the rates that have been observed more recently). The distributions of short rates produced by both calibrations generally fall between the two historic distributions, with the exception that calibration *A* produces short rates below zero in just under 10% of the simulations. These negative rates are removed in calibration *B*, where we force nominal rates to be positive.

EXHIBIT 7.5: UNCONDITIONAL DISTRIBUTION FOR SHORT RATES

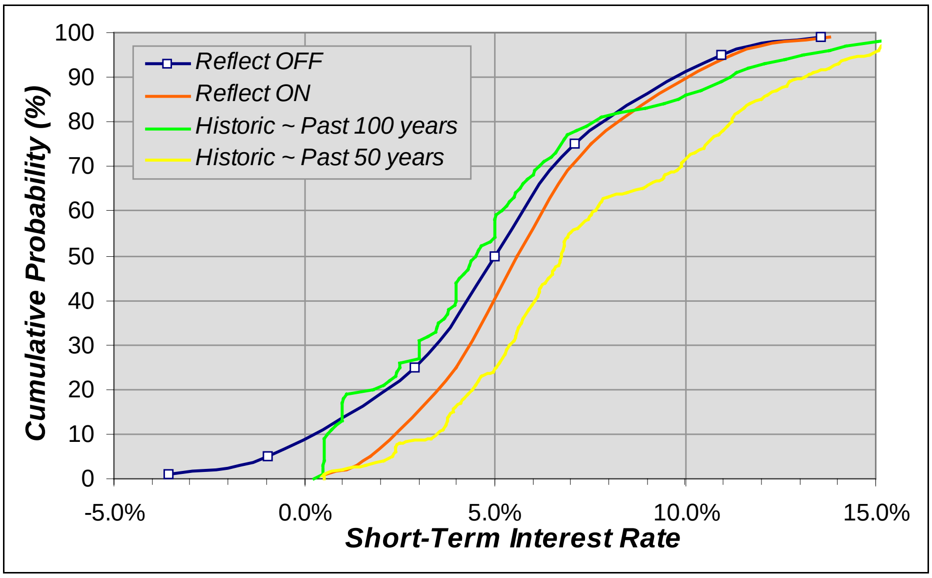

Exhibits 7.6A and 7.6B illustrate the distribution of the path of short rates over the course of the 30-year simulation horizon. These charts illustrate the percentiles of a distribution in a manner similar to the bar charts in exhibit 7.3, the only difference being that we are now interested in the change in the distribution as we progress through the 30-year simulation horizon, rather than a single distribution across all time periods (an *unconditional* distribution).

In calibration *A*, we have initialised the model for real and nominal interest rates in its equilibrium position, meaning that $r_1(0)$ and $r_2(0)$ equal μ_r, and $q_1(0)$ and $q_2(0)$ equal μ_q. One consequence of this is that the centre of the short rate distribution (the median) remains at 5% throughout the course of the 30-year horizon. In the two charts, we can see how uncertainty in movements in the short rate creates a distribution around this central value - a ‘funnel of doubt’. The spread of the distribution around its median increases over the first few years, but the effect of mean reversion is such that the distribution stabilizes after about 20 years or so. Importantly, in this basic first calibration, short rates become negative in approximately 10% of the 1000 simulations. This feature is likely to be viewed as an undesirable property of this calibration.

EXHIBIT 7.6A: DISTRIBUTION OF PATH OF SHORT RATE OVER 30-YEAR HORIZON (REFLECTION *OFF*)

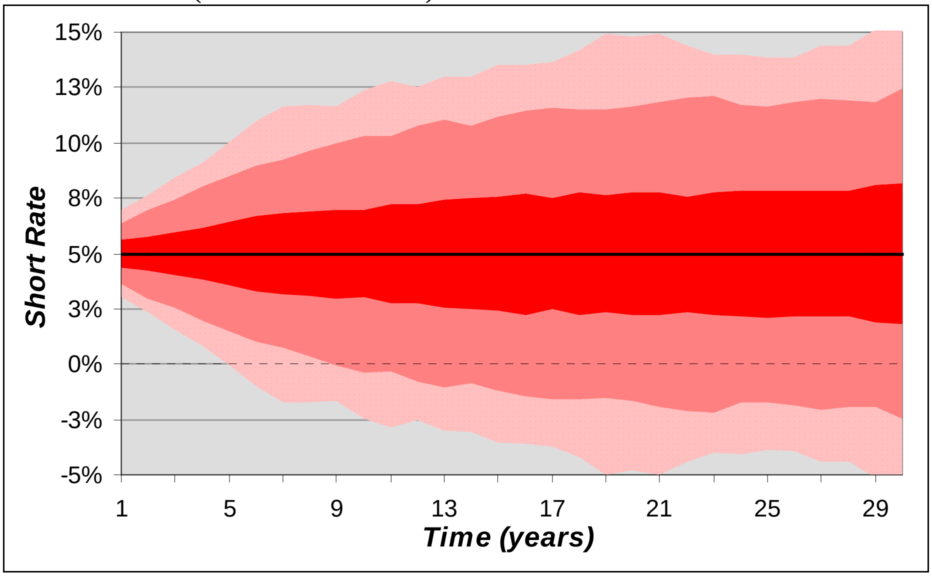

In exhibit 7.6B, you can see that forcing nominal interest rates to be positive means that the lower percentiles of the distribution remain positive. An awkward consequence of the trick used to ensure positive rates is that the expected short rate tends to drift up over time. This effect can be partly offset by reducing the value of the μ_r and μ_q parameters.

EXHIBIT 7.6B: DISTRIBUTION OF PATH OF SHORT RATE OVER 30-YEAR HORIZON (REFLECTION *ON*)

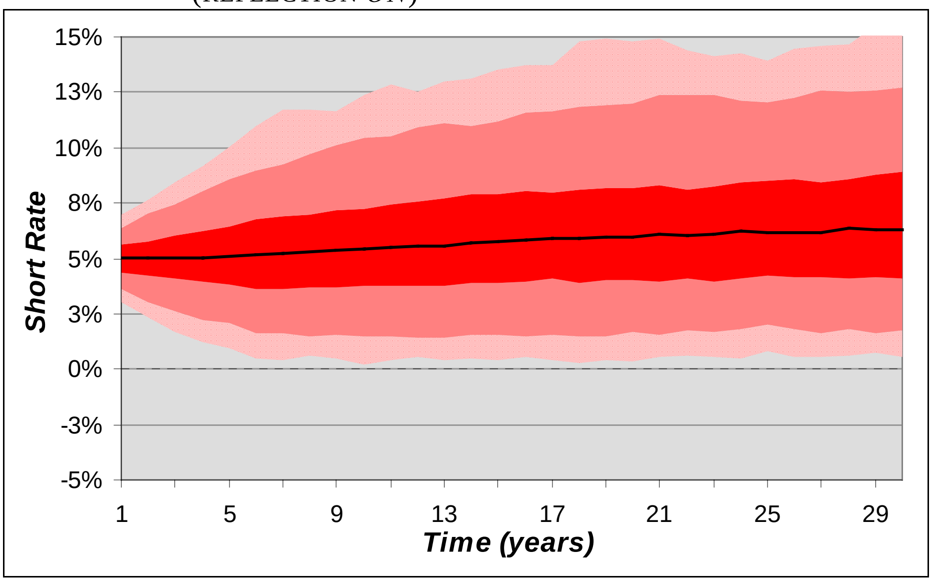

In Exhibit 7.7 we illustrate the distribution of returns on cash, and compare the distributions produced by the two calibrations with the historic distribution from the last century. Our second calibration produces a distribution which looks similar to the historic distribution, although the upper percentiles of the cash returns generated by the model are somewhat less the historic values. This is a direct consequence of this particular calibration, which assigns lower likelihood to the very high interest rates that have been observed historically. It is sensible to ask what we believe the chances are of observing a short-term interest rate of 15% over the course of the next 30 years. Our second calibration assigns a probability of just less than 1%.

EXHIBIT 7.7: DISTRIBUTION OF ANNUAL CASH RETURNS

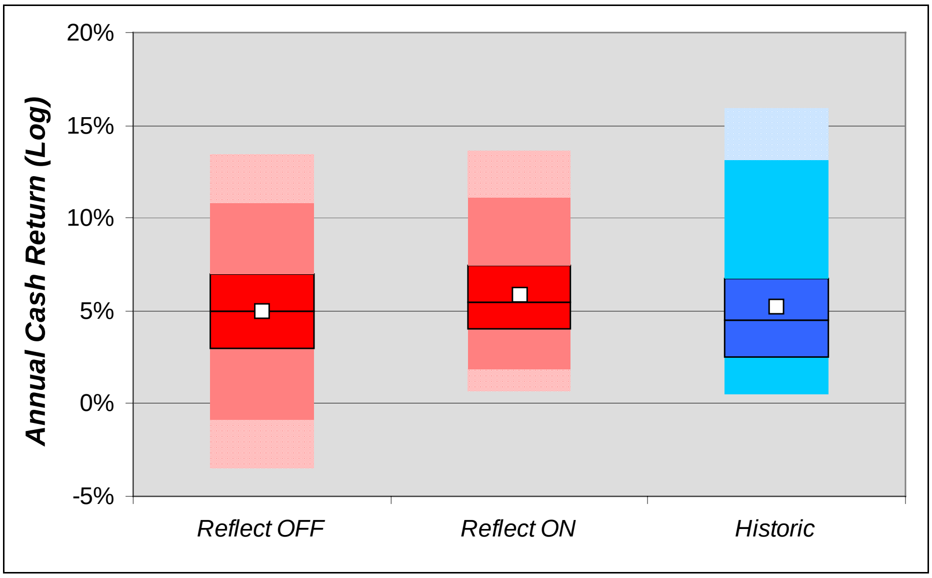

7.5 15-YEAR CONVENTIONAL BOND YIELDS

Exhibit 7.8 shows the resulting cumulative frequency plot for long-term bond yields. For comparison, we have also plotted cumulative plots for historic long-term yields and market-implied rates. The chart tells us what we already know – that there is a range of plausible distributions. Expert opinion suggests that the immediate inflation and interest rate future will look nothing like the last 50 years and that we should assign much lower probability to 10% bond yields than the 30% frequency over the past 50 years. The two distributions generated by the model look reasonable, although a refinement would be to extend the left-hand to tail to assign greater probability to Japanese-style bond yields.

EXHIBIT 7.8: UNCONDITIONAL DISTRIBUTION OF 20-YEAR COUPON BOND YIELDS

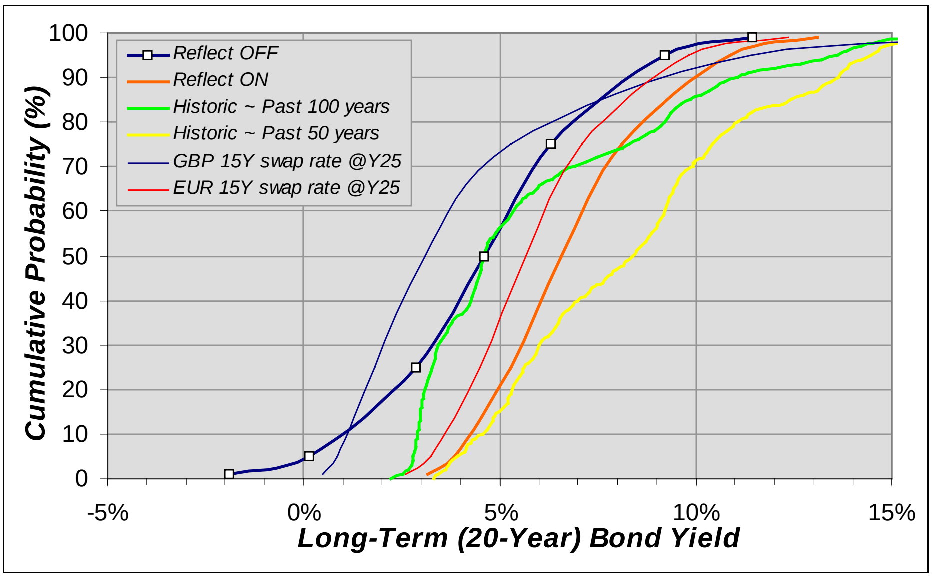

Exhibits 7.9A and 7.9B are analogous to the funnels of doubt for cash returns illustrated in exhibits 7.6A and 7.6B. Again, we can see a funnel that grows over the first 10-15 years of the simulation, and then flattens out as mean reversion takes effect. Another obvious effect of mean reversion is that the spread of the distribution of long-term yields is less than that for the short rate. The central 98% of the distribution for the 15-year coupon bond yield in exhibit 7.9A covers the range -2.5% to 12.5%, compared to -5% to 15% for the short rate in exhibit 7.6A.

Exhibit 7.9B shows that the second calibration ensures that no negative bond yields are generated. In fact, bond yields never appear to fall below about 2.5%. Again, this is a feature of a particular calibration. There are many other calibrations we could use to reduce (or remove) this implicit lower bound on bond yields. There are also other devices we could use which would ensure nominal rates remained positive, but which did not preclude very low bond yields.

EXHIBIT 7.9A: DISTRIBUTION OF PATH OF 15-YEAR COUPON BOND YIELDS OVER 30 YEAR HORIZON (REFLECTION *OFF*)

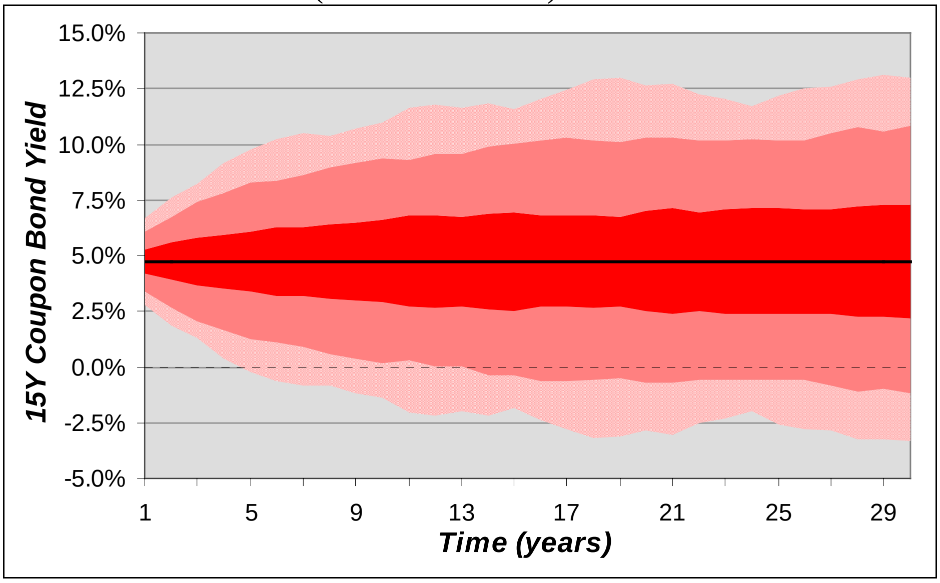

In exhibit 7.10, we show the distribution for yields on 15-year coupon bonds, and compare the distributions from the two model calibrations with both the historic and market implied distributions. Notice that the distribution produced by the second calibration is somewhat narrower than the historic experience. As for short rates, in using this particular set of parameters, we have assigned lower probabilities to very high long-term nominal interest rates than the frequencies been experienced over the last century. Notice that the second calibration produces a distribution that is rather similar to the distribution for long-term bond yields implied by current market prices for long-dated swap contracts.

EXHIBIT 7.9B: DISTRIBUTION OF PATH OF 15-YEAR COUPON BOND YIELDS OVER 30 YEAR HORIZON (REFLECTION *ON*)

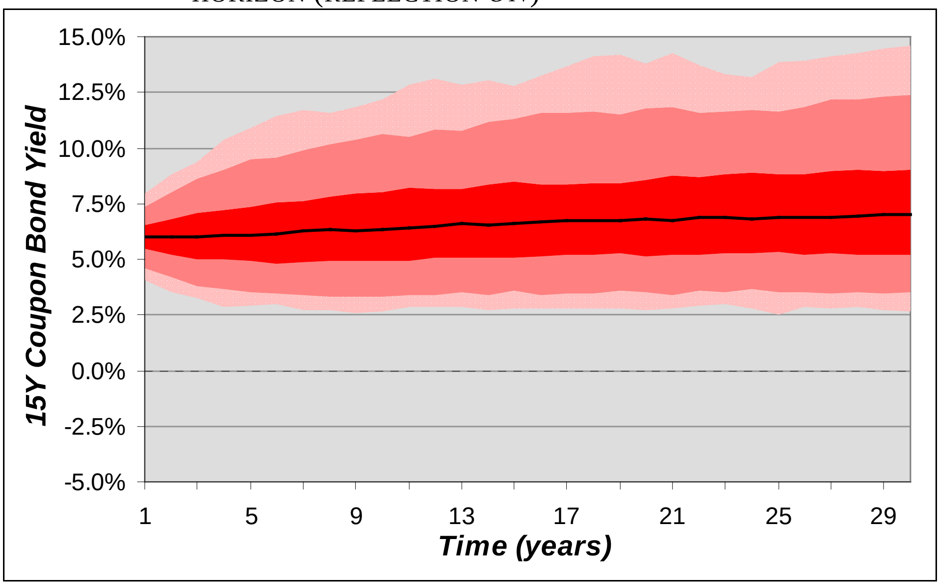

EXHIBIT 7.10: DISTRIBUTION OF COUPON BOND YIELDS

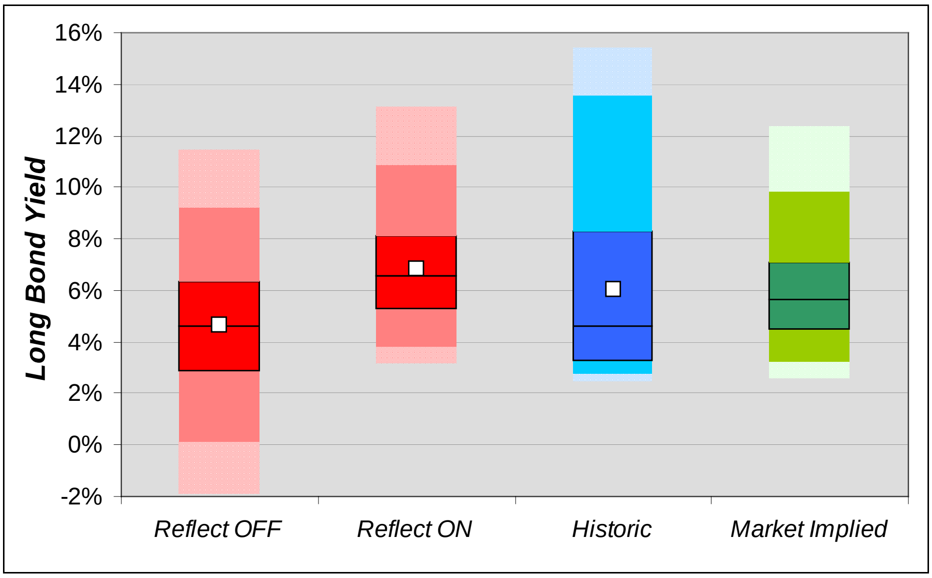

7.6 INDEX-LINKED BOND YIELDS

Now let us consider yields on index-linked bonds. In the first calibration, we can see from exhibit 7.11 that negative index-linked bond yields are generated with a probability of just over 10% of the simulation trials. When the minimum barriers are applied to the values of $r_1(t)$, $r_2(t)$, $q_1(t)$, $q_2(t)$, and the nominal yield curve reflection is activated in the second calibration, the lowest simulated index-linked bond yield is just over 1%. The median of the distribution for index-linked yields also increases from 2.5% to over 3%. This effect is evident when we compare the distribution of the path for index-linked yields under the two calibrations in exhibits 7.12A and 7.12B. In the second calibration, the term premium in real interest rates will mean that the expected return on index-linked bonds will increase with term.

EXHIBIT 7.11: UNCONDITIONAL DISTRIBUTION OF 20-YEAR INDEX-LINKED BOND YIELDS

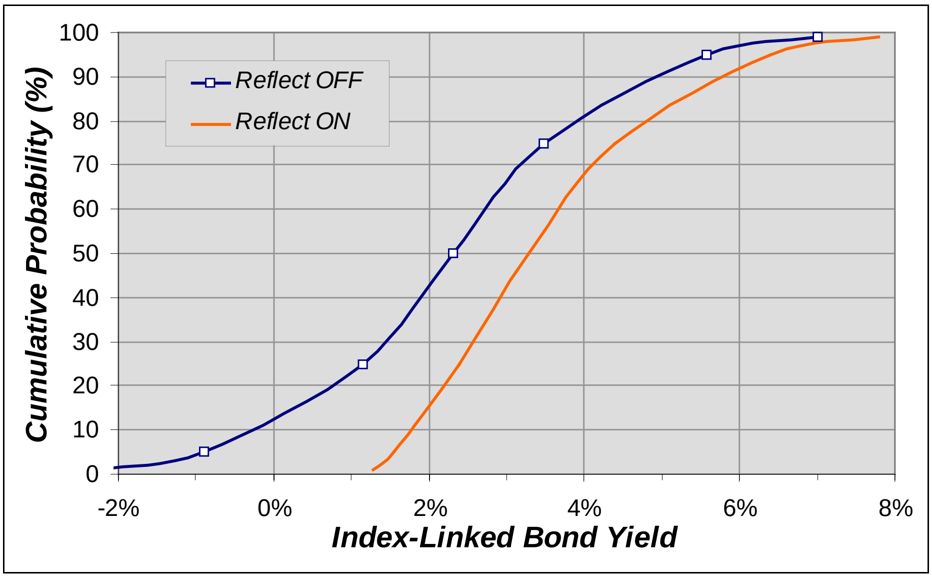

EXHIBIT 7.12A: DISTRIBUTION OF PATH OF 20-YEAR INDEX-LINKED YIELDS (REFLECTION *OFF*)

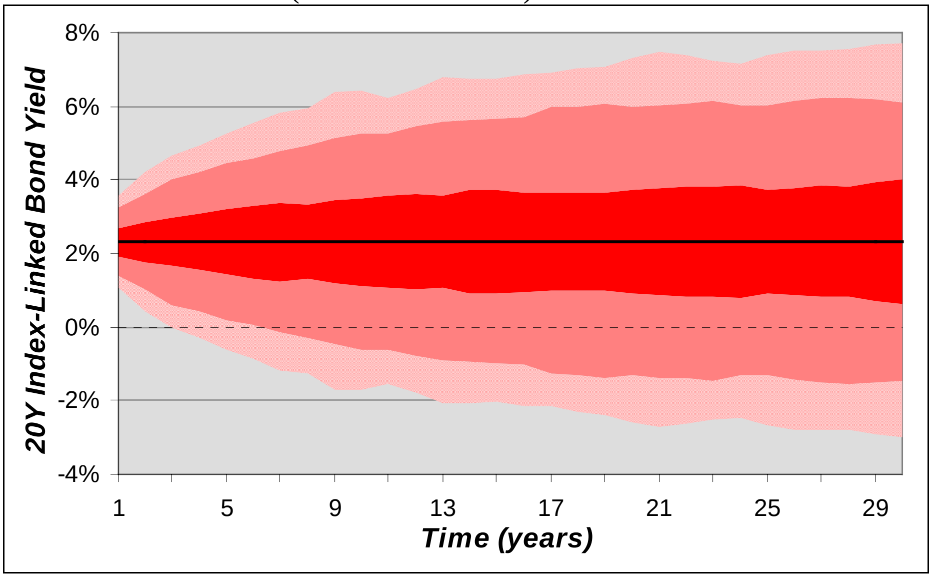

EXHIBIT 7.12B: DISTRIBUTION OF PATH OF 20-YEAR INDEX-LINKED YIELDS (REFLECTION *ON*)

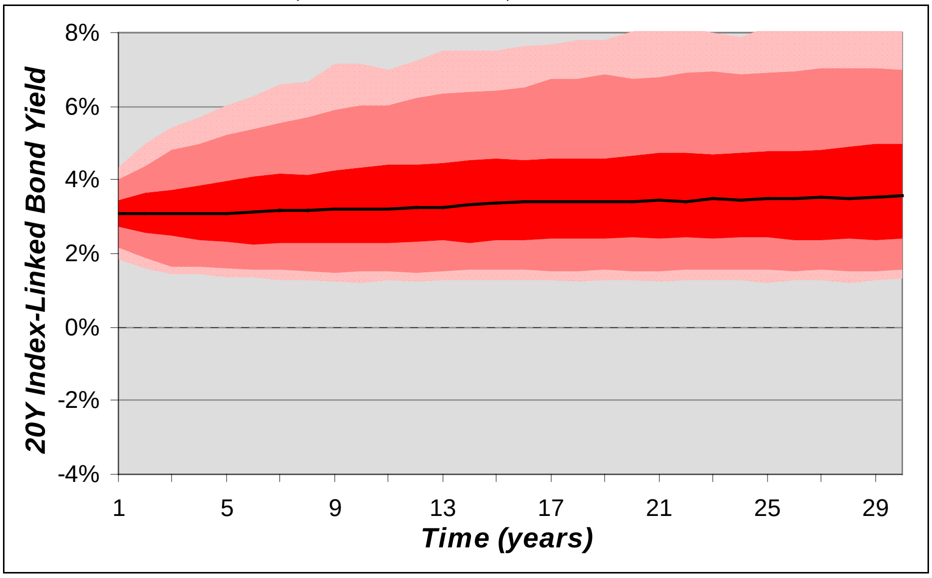

7.7 NOMINAL INTEREST RATE TERM-STRUCTURE

Rather than considering the distribution of nominal and index-linked yields on bonds of particular maturities, it is important that the model captures movements in the *entire* term-structure in a realistic way. The value of a portfolio of assets is affected by movements in the level and shape of the yield curve, and so it can be useful if the model can produce a representative range of yield curve shapes (and changes in shape). The following two charts illustrate a couple of rather different scenarios for changes in the yield curve over the course of individual 30-year simulation trials, from its 'flat' starting position of 5%. Exhibit 7.13A illustrates a scenario where, for much of the 30 years, the nominal yield curve lies below its starting position, including two or three years where zero coupon yields for many maturities fall below 1%. Exhibit 7.13B illustrates a scenario where yields have generally increased over the course of the 30-year simulation.

EXHIBIT 7.13A: SIMULATED PATH OF NOMINAL YIELD CURVE (SIMULATION #5)

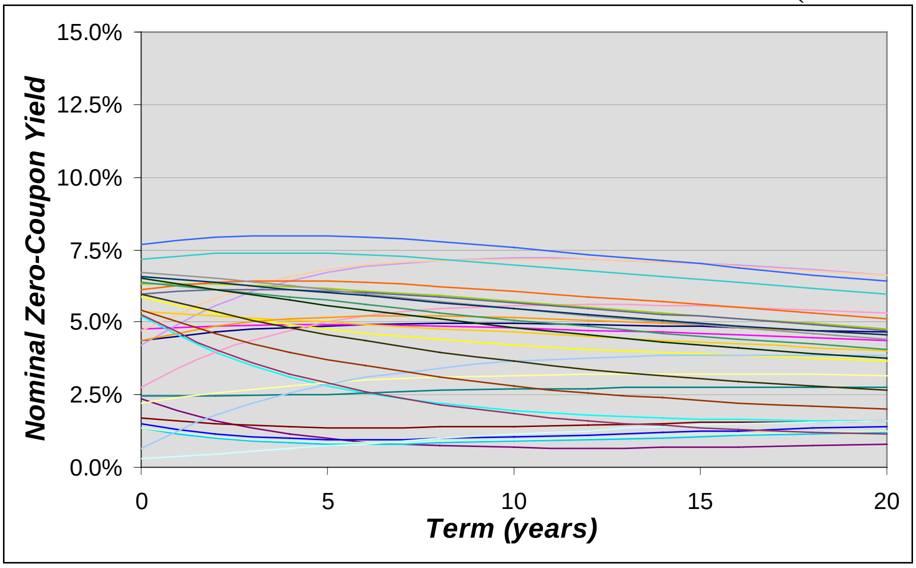

EXHIBIT 7.13B: SIMULATED PATH OF NOMINAL YIELD CURVE (SIMULATION #47)

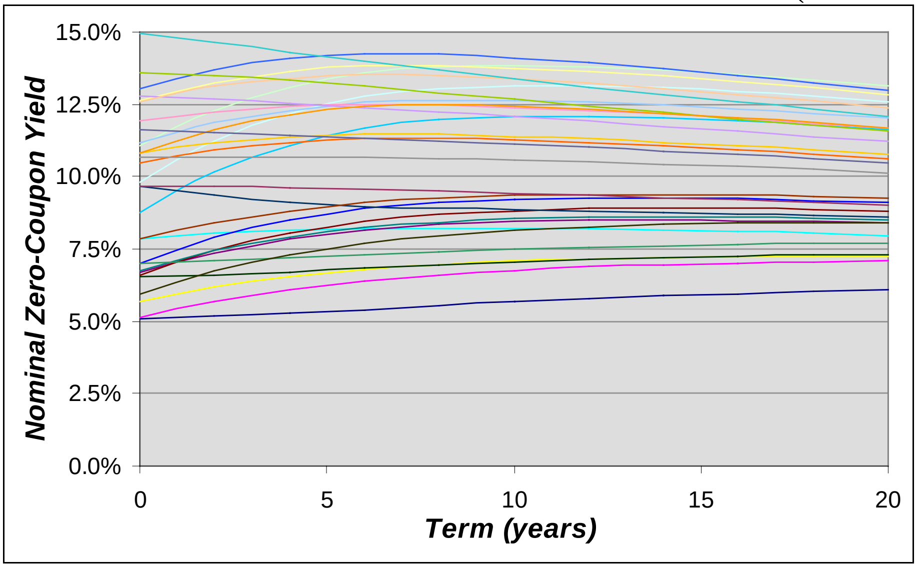

As well as looking at the behaviour of the yield curve within individual simulation trials, it is useful to understand the way in which zero-coupon yields are distributed as the model is run out over longer terms (the *unconditional* distribution of the nominal term structure). Needless to say, the distributions illustrated in exhibits 7.14A and 7.14B look rather different. For the first calibration, the distribution of nominal zero-coupon yields across all terms are centred at slightly below 5%[13]. This chart certainly

[13] *Although the expected instantaneous rate of return at all terms is exactly equal to 5%, the expected yields fall slightly with increasing term due to Jensen's inequality [Bulletin of the Australian Mathematical Society, 1997, 55, 185-189.]*

highlights the problem of negative nominal interest rates inherent in the first calibration: over 5% of nominal zero-coupon yields fall below zero. When we use the reflection adjustment to the model to force positive nominal rates in the second calibration (exhibit 7.14B), the entire distribution is shifted up, and we experience no very low long-term nominal rates. These issues have been mentioned previously and it is important to re-iterate that such features are a consequence of this particular calibration. Users who believe that such a distribution for long-term nominal rates was implausible should investigate other parameter choices.

EXHIBIT 7.14A: UNCONDITIONAL DISTRIBUTION OF TERM-STRUCTURE OF NOMINAL INTEREST RATES (REFLECTION *OFF*)

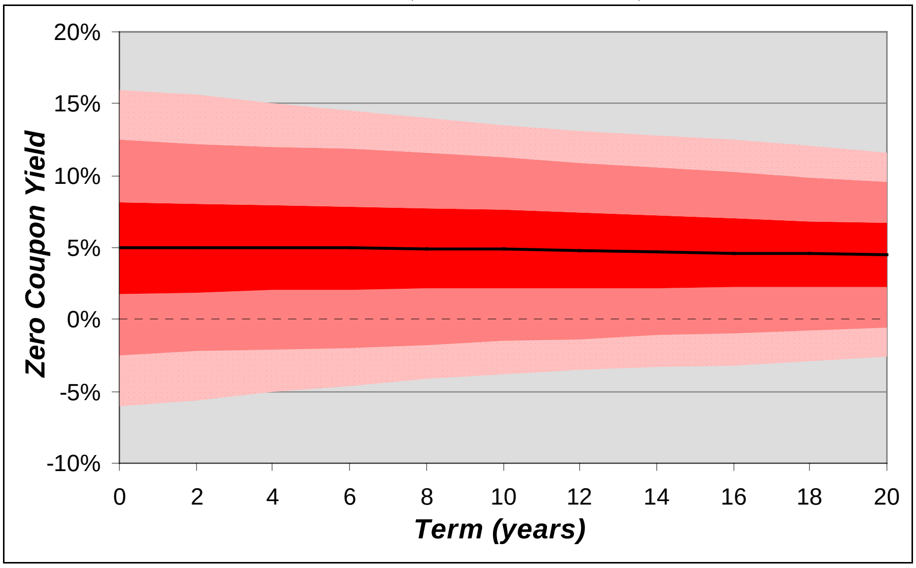

EXHIBIT 7.14B: UNCONDITIONAL DISTRIBUTION OF TERM-STRUCTURE OF NOMINAL INTEREST RATES (REFLECTION *ON*)

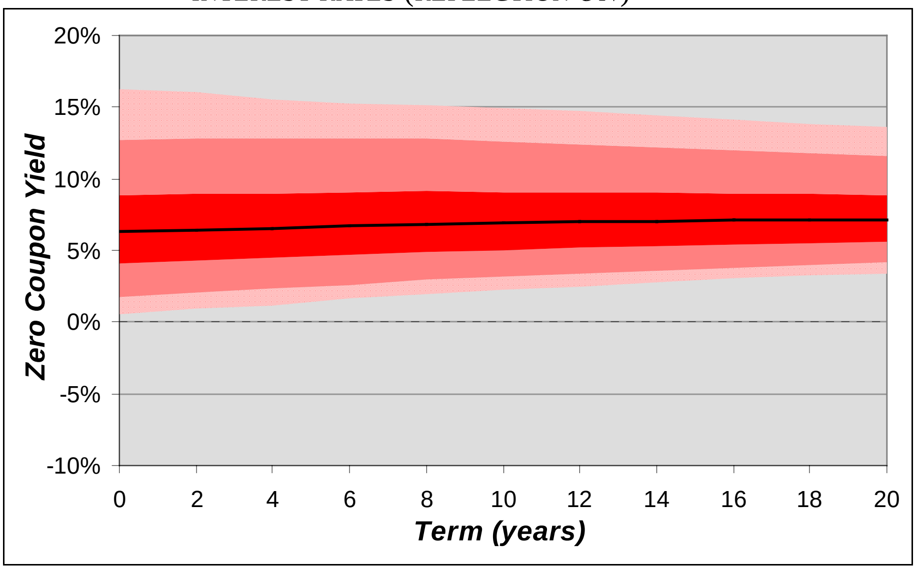

7.8 INFLATION

In this section we will illustrate the behaviour of inflation rates and the term-structure for inflation expectations as we have done for bond yields and the nominal yield curve. Exhibit 7.15 shows the entire cumulative probability distribution for the instantaneous rate of inflation under the two calibrations, and compares these distributions with historic inflation rates from the last 100 years. The median value generated by both model calibrations is approximately 2.5% per annum. The most striking feature of this chart is that the historic distribution has much greater spread than the distribution generated by the model. We can look back to exhibit 6.8 to confirm that the UK inflation rate has indeed exceeded 15% for roughly 10 of the last 100 years. We have chosen the model parameters to reflect a view that the likelihood of the inflation rate reaching 15% at some point over the next 30 years is very much smaller than observed over the course of the 20th century. Exhibit 7.15 shows that our both calibrations assign a 10% probability to the inflation rate exceeding about 6%, rather than 15%!

EXHIBIT 7.15: UNCONDITIONAL DISTRIBUTION OF INFLATION RATE

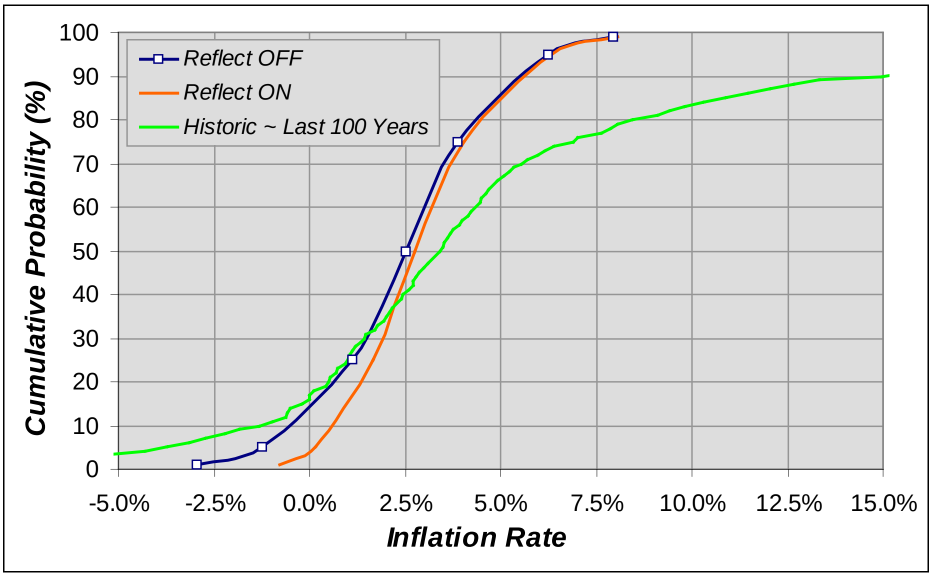

Exhibits 7.16A and 7.16B show the distribution of the path of inflation over the course of the 30-year simulation horizon. Very like the distributions for bond yields and interest rates, the mean reversion within the inflation model means that these ‘funnels of doubt’ spread out over the first few years of the simulation before stabilising after about 15 years. Notice that, even when we apply the minimum barriers to $q_1(t)$ and $q_2(t)$, and force nominal rates to be positive, deflationary scenarios are still generated in about 5% of the simulations.

EXHIBIT 7.16A: DISTRIBUTION FOR PATH OF INFLATION RATE OVER 30 YEARS (REFLECTION OFF)

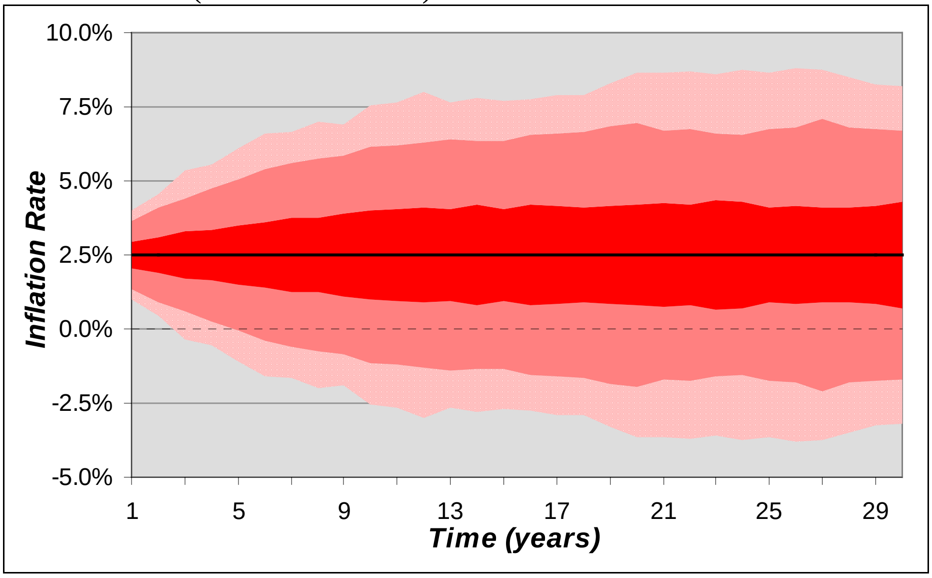

EXHIBIT 7.16B: DISTRIBUTION FOR PATH OF INFLATION RATE OVER 30 YEARS (REFLECTION ON)

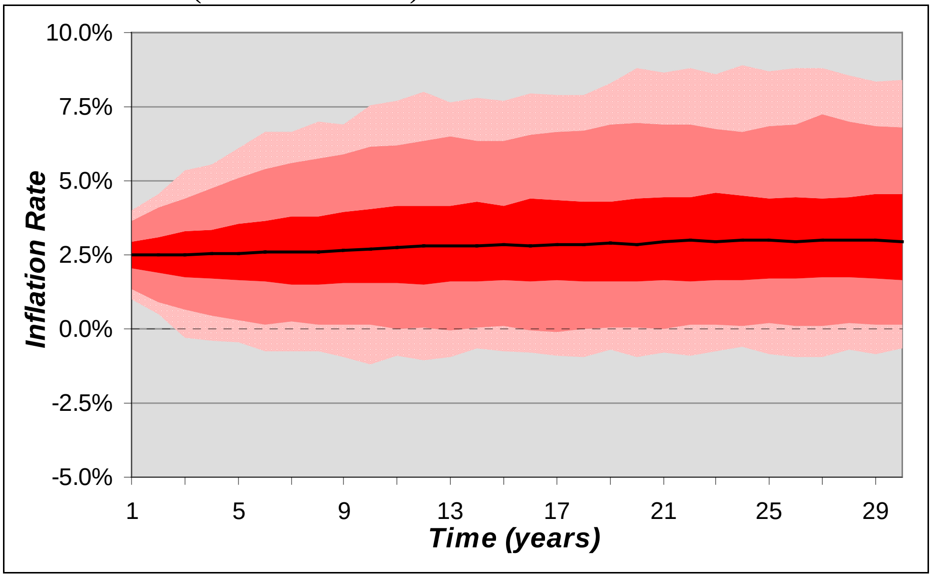

7.9 INTER-RELATIONSHIPS BETWEEN INFLATION, BOND YIELDS & EQUITY RETURNS

Although our model may produce individual asset scenarios and distributions that appear quite sensible when compared with empirical data, when we are considering entire portfolios consisting of a range of assets and liabilities, it is important to ensure that the model generates plausible inter-relationships between different asset classes. For instance, generating a significant number of scenarios where high interest rates coincide with stable, low inflation rates would seem rather unreasonable.

Firstly we look at the relationship between the simulated inflation rate and 15-year coupon bond yield at a particular point in time (year 25), in each of the 1000 simulation trials. Exhibits 7.17A and 7.17B demonstrate that the model generates quite strong correlation between the simulated inflation experience and nominal bond yields. Lower inflation scenarios tend to coincide with lower bond yields, and vice-versa.

EXHIBIT 7.17A : INFLATION VS BOND YIELD (REFLECTION OFF)

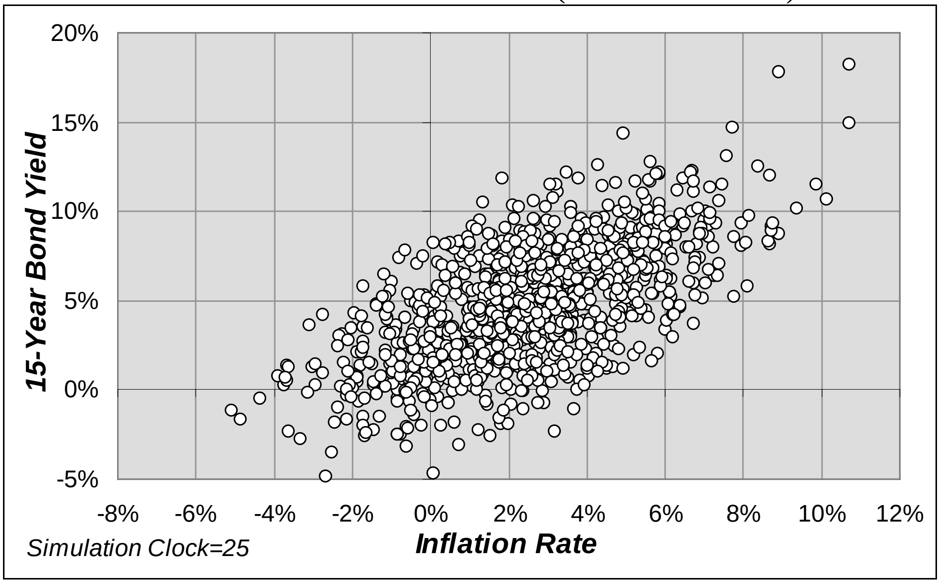

EXHIBIT 7.17B : INFLATION VS BOND YIELD (REFLECTION ON)

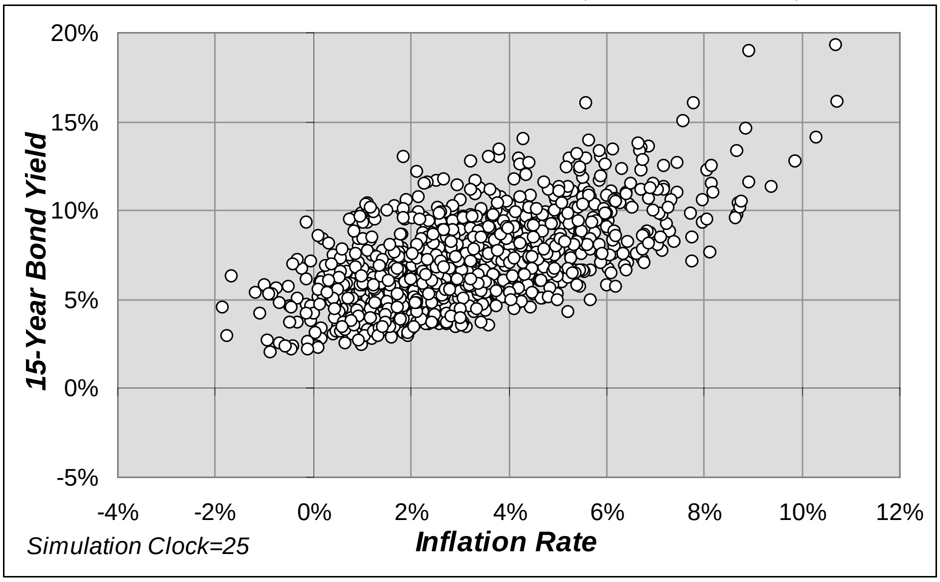

In exhibits 7.18A and 7.18B we look at the relationship between simulated scenarios for rolled-up equity returns and the corresponding level of the inflation index. Generally speaking, high inflation has tended to hinder equity performance, particularly in real terms. This effect can be seen (moderately) within the scenarios generated by this model. High inflation scenarios are often associated with lower rolled-up equity returns. This is particularly true when we consider equity returns in real, rather than nominal, terms.

EXHIBIT 7.17A:EQUITY ROLL-UP VS INFLATION (REFLECTION *OFF*)

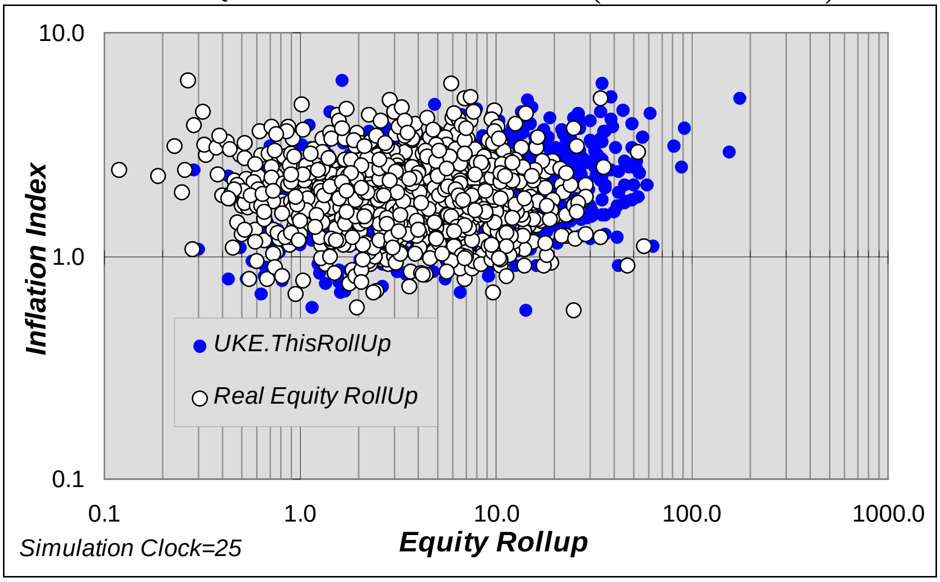

EXHIBIT 7.17B: EQUITY ROLL-UP VS INFLATION (REFLECTION *ON*)

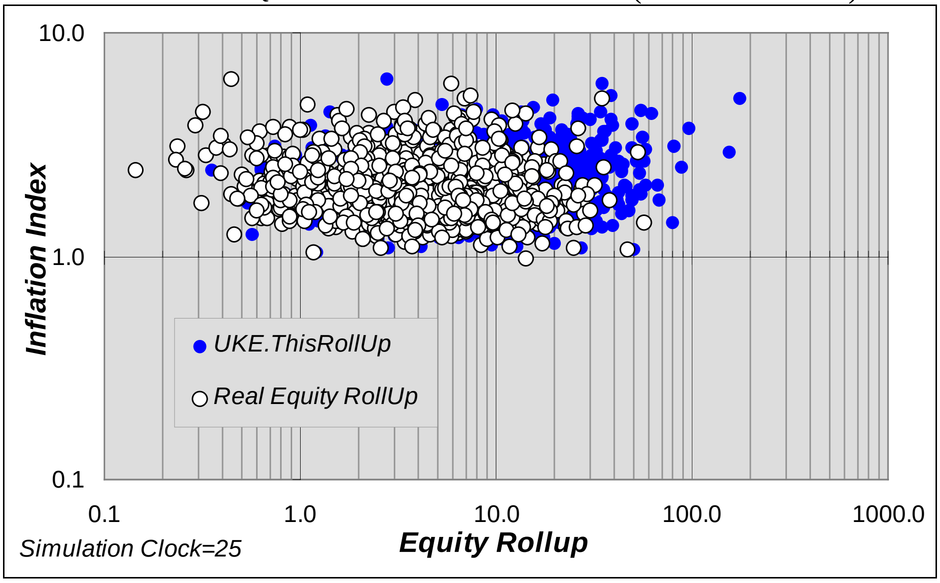

8. A Comparison with the Wilkie Model

Someone might – quite reasonably – ask: "*Who needs another model? Surely the Wilkie model is good enough for my purposes?*" As we have already observed, any model will be fit for some purpose but not others. It would be difficult to argue that the *Wilkie* model is not fit to analyse certain problems (providing the modeller holds a particular set of beliefs about equity market behaviour). On the other hand, our experience suggests some serious failings.

We now compare our proposed model and calibration (which we will refer to as the *B&H* model) with the *Wilkie* model. Of course, there could be as many calibrations of the *Wilkie* model as there are model users. Professor Wilkie has certainly encouraged users to try out different parameters. We will present a *Wilkie* calibration that we believe is fairly representative of current parameter selections (see Appendix C for the parameters used in this section). The comparison between the two models is undertaken both in terms of comparing model outputs and making some general observations with regard to differences between the two models.

8.1 Some General Observations

The two models adopt very different approaches to the challenge of modelling the long-term behaviour of financial variables. In the *Wilkie* model, statistical time-series relationships are developed for a number of observable variables – inflation, short-term interest rates, consol yields, dividend income and the dividend yield. Market prices are then derived from these processes. For example, in the *Wilkie* model, equity prices are inferred from the ratio of dividend income to dividend yield, whereas in the *B&H* model, the total return on equities is modelled separately from the dividend yield, and the process for the equity price is therefore transparent and parsimonious.

Whilst the structure used by *Wilkie* may seem simpler to those with an aversion to maths, it has some major drawbacks. The structure of the model is rather convoluted and lacks transparency. In a statistical analysis of the properties of the *Wilkie* model, Huber[14] found that the model did not provide a good representation of historical data and was over-parameterised, with a number of parameters being effectively redundant: he argued that the complexity of the model structure does not add to the effectiveness of the model. Whilst calibration is a challenging problem in almost any financial modelling, the poor statistical fit may seem slightly more surprising in the case of the *Wilkie* model given the complex structure of the model seems to have been driven more by consideration of fitting observed patterns in the data rather than by any reference to building structures consistent with economic theory.

We believe that the structure of the model presented in section 5 is simpler and more economically logical. We model real interest rates, inflation rates and the equity return in excess of the nominal interest rate. This structure ensures a consistency in the joint asset behaviour generated at any given time that is lacking in the *Wilkie* model.

[14] *"A Review of Wilkie's Stochastic Investment Model", P. Huber (1995).*

The following sections review some output of the *Wilkie* model. We highlight some of the potential problems with the output and make comparisons with output generated by the *B&H* model.

8.2 REPRESENTATION

As has been mentioned above, the *Wilki*e model does a poor job of representing certain features of the real world. For example:

- The relationship between simulated inflation outturns and long-term bond yields simply is not plausible. The model (using a typical parameter set) generates frequent scenarios of very low average inflation accompanied by high bond yields (and high inflation coupled with low bond yields). Two examples of typical - but rather surprising - joint paths for inflation and bond yields are plotted below.

EXHIBIT 8.1A: A WILKIE SCENARIO

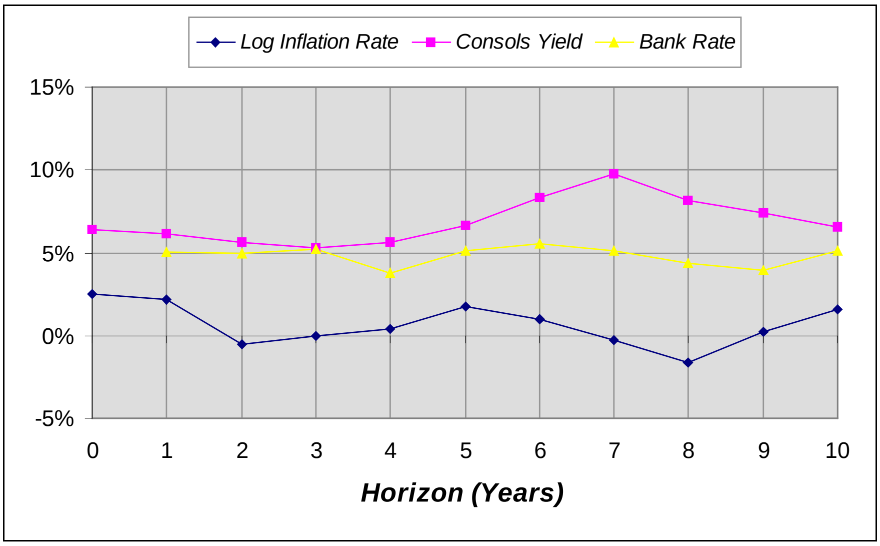

EXHIBIT 8.1B: ANOTHER WILKIE SCENARIO

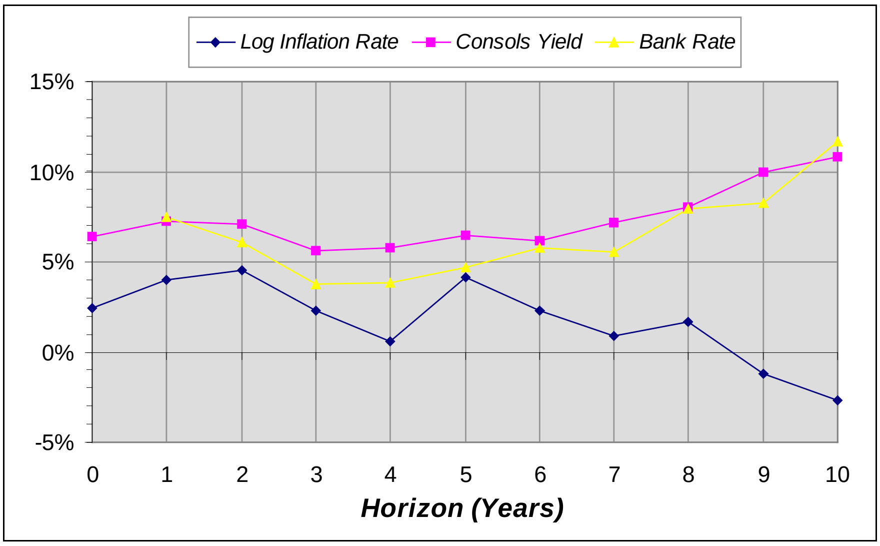

EXHIBIT 8.2: WILKIE INFLATION RATES & CONSOLS YIELD (10-YEAR HORIZON)

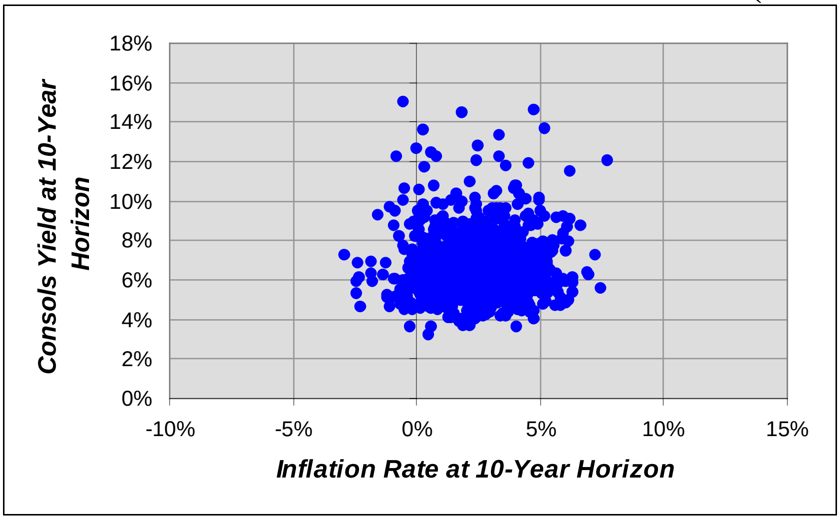

Exhibit 8.2 illustrates the relationship between the inflation rate at a 10-year horizon and the resulting Consols yield. There is no apparent correlation between these two quantities. For an analyst who uses the model to investigate the properties of conventional bonds in different sorts of inflation environment, the model fails to capture the fundamental link between inflation and nominal yields.

➢ The model (using a typical parameter set) generates plausible variability in equity returns over short horizons, but very narrow distributions at long-term horizons. The probabilities assigned to equity market declines over long-term horizons look implausibly low given 20th century experience. Since one of the primary purposes of the model is often to show users the potential impact of low-probability outcomes, this is particularly surprising. Exhibit 8.3 illustrates three 20th century episodes of 10-year year equity index declines – one for each of the UK, US & Japanese markets. There are other examples for these markets. Indeed, equity claims were wiped out altogether in Japan in 1945. Note that the scale of the UK decline is increased dramatically if the illusory impact of inflation is removed by plotting the *real* price of equities.

EXHIBIT 8.3: 20TH CENTURY BEAR MARKETS

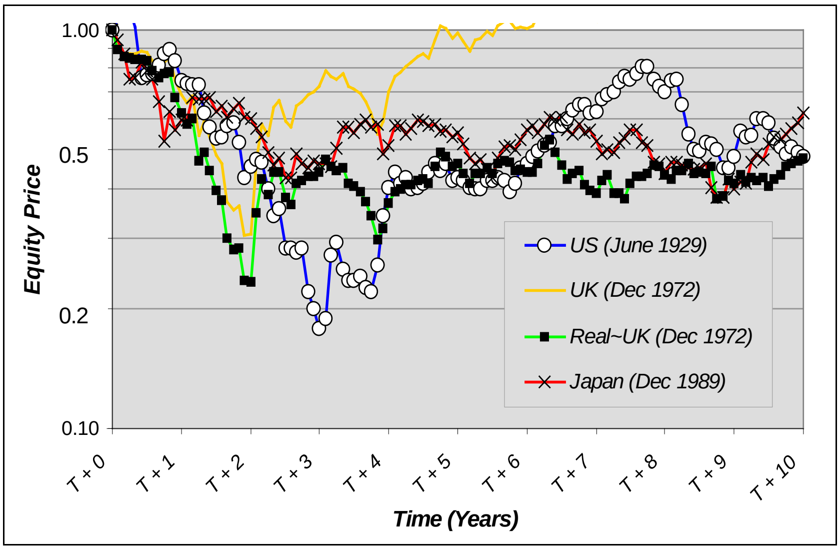

So, very crudely, you might guess that since each of these markets has experienced a 10-year 50% price decline (somewhat less in terms of total return and approximately the same in terms of excess returns) at least once in the past 100 years, the probability assigned to this sort of scenario would be something like 10%. The mean reversion which is an important part of the Wilkie model means that the probability assigned to long-term declines of this magnitude is much lower. The jury is still out in the (complicated) mean reversion debate. Whatever you might think, it seems reasonable that the modeller should begin by *excluding* mean reversion from models (a prudent assumption?) unless there is compelling evidence to support its existence[15].

It is possible to increase long-term variation in Wilkie equity returns by raising short-term variability. Unfortunately, this has the effect of producing implausible variability in short-term equity returns.

- The mean reversion feature means that short-term returns generated from the model can be sensitive to the way the model is initialised. This requires extreme care from the model user.
- Even with a fairly dramatic reduction in the parameter describing the variability of inflation (QSD), short-term inflation variation looks too high (by comparison with almost any economic forecast) and the distribution at long horizons looks too narrow.

8.3 MEAN REVERSION

The structure of the *Wilkie* model means that – with typical parameter choices – the model generates mean-reversion in equity returns. Whilst the extent to which equity markets *actually* mean-revert is the subject of much debate, we would argue that assuming mean-reversion to the extent generated by the *Wilkie* model is perhaps unwise. Such mean-reversion means that simple re-balancing rules can increase returns whilst at the same time reducing risk. It could be argued that taking credit for the on-going existence of this supposed free lunch going forward is a rather imprudent starting point for making long-term equity projections. Though there is some statistical evidence for mean-reversion in equity markets, its statistical significance is dubious, particularly when we consider the excess return (i.e. the return in excess of the prevailing cash rate).

Given the structure of the *Wilkie* model, whereby a stochastic process for the dividend yield is used to derive equity prices, mean-reversion seems difficult to remove from the model. It could be argued that deriving prices from dividend yields is really a case of the tail wagging the dog. In the *B&H* model, equity prices and dividend yields are modelled as two separate (but highly negatively correlated) processes. As in the *Wilkie* model, the dividend yield is modelled as a mean-reverting process. However, the *B&H* model has

---

[15] *There is large literature on mean reversion. For a recent example see "Mean Reversion in Stock Returns: Evidence & Implications", L Summers & J Poterba (Financial Markets Group discussion paper, LSE (1998))*

quite different implications for the behaviour of dividend income. For example, the *B&H* model implies that when markets crash, yields will immediately rise and then gradually fall back to their long-term mean, without any systematic mean-reversion in equity *returns*. This implies that following a market crash, the rate of expected dividend growth is lower – that is, the market has fallen due to a reduction in expected dividend growth. Put simply, the expected return in excess of cash is assumed to be constant, so any change in dividend yield implies a corresponding change (in the opposite direction) for dividend growth. Note that we are considering the *excess* return here – the expected nominal return on equities in the *B&H* model will be higher when short-term interest rates are higher.

This is fundamentally different from the *Wilkie* model – here the expected return on equities depends on the level of the dividend yield relative to its assumed mean, as expected future dividend growth is not a function of the current dividend yield (except insofar as they both are a function of inflation). So when yields are high, expected returns are high – that is, expected dividend growth has not fallen to the extent that it offsets the higher yield. In this case, the model is implying that equities have fallen by more than is implied by the reduction in expected dividend income, so equities then mean-revert back to their 'fair' value.

So the *B&H* model is consistent with a world where equity prices are based on rational expectations of future dividend growth, whilst the *Wilkie* model assumes that dividend income is far more stable than would be suggested by the volatility of the dividend yield, and equity markets go through periods where the expected return on equities can vary very significantly (and note that a change in inflation or interest rates is not necessary for such a change in expected return to occur).

We now turn our attention to how these differences in the model affect the distributions of equity returns generated.

8.4 EQUITY MODEL

The differences in the structure of the two models perhaps have the greatest impact on the simulated distributions of equities, particularly at longer horizons. The mean-reversion inherent in the *Wilkie* model implies far lower equity risk at longer time horizons than the *B&H* model – the mean-reversion results in annualised equity volatility decaying much faster than it otherwise would. In a model with no mean reversion, we would expect the annualised volatility to decay at rate of 1/sqrt(t). Let us call the annualised volatility multiplied by the square root of time the standardised annualised volatility. Exhibit 8.2 plots this statistic as simulated for the two models over a ten-year horizon.

EXHIBIT 8.2: STANDARDISED ANNUALISED VOLATILITY

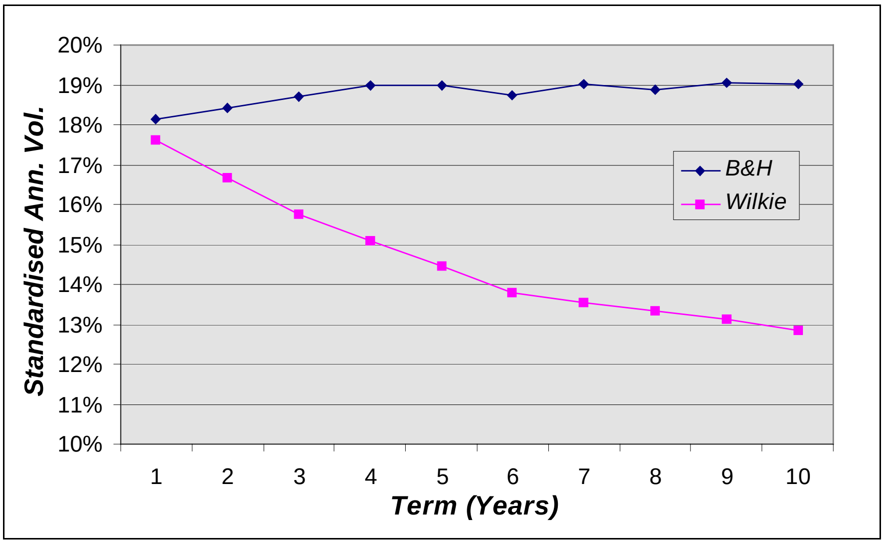

Exhibit 8.2 suggests some striking differences in the simulated long-term equity behaviour of the two models. Note that the one-year volatilities are very similar. But as the time horizon is extended, the Wilkie model suggests that volatility falls very significantly, whereas in the B&H model the volatility remains constant (ignoring sampling error). What does this imply for the distributions of equity returns? Exhibit 8.3 plots the simulated cumulative distributions for the annualised equity returns generated by the two models over ten years.

EXHIBIT 8.3: CUMULATIVE DISTRIBUTIONS OF 10-YEAR ANNUALISED EQUITY RETURN

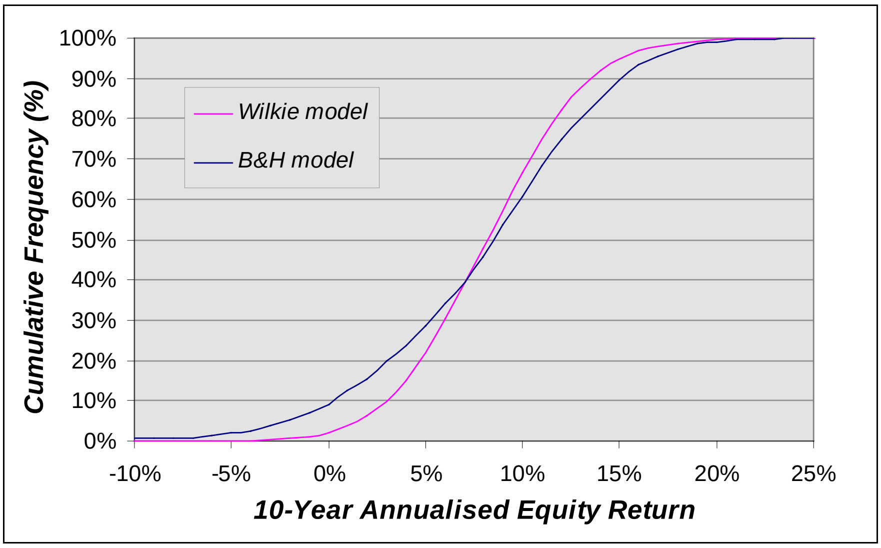

You can see that the *Wilkie* model attributes a very low probability to equities generating a negative return over a 10-year period: the 99th percentile is -0.7% p.a., whilst the corresponding *B&H* value is –6.0 %p.a. The rolled-up equity values over the ten-year period are 88 % and 53%. Whilst you might think the *B&H* value is very bearish, as was mentioned above there is plenty of evidence of such poor returns in equity markets.

8.5 TERM STRUCTURE MODEL

The *Wilkie* model does not generate a full term structure of interest rates. This means that the analyst who wishes to model the returns on a bond which is not a consol or cash will need to make some difficult assumptions, e.g. some form of interpolation of yields. The lack of a full term structure makes analysis of some forms of interest rate risk very difficult.

The two models can be configured to generate similar distributions for short and long-term interest rates. But this, to an extent, misses the point – the *B&H* model generates an economically consistent relationship between the short and long-term rates (the long-term spot rate can be regarded as the expected path of the short-term rate plus any specified risk premium), whereas the *Wilkie* model has no such economically meaningful relationship between the rates as there is no model for the term structure. This feature, and indeed the way the model is built from a series of statistical relationships rather than being based on any notion of rational economic expectations, can lead to some very odd joint paths for interest rates and inflation easily occurring, as we saw above.

8.6 INFLATION

The structure of the inflation models used in *Wilkie* and *B&*H are actually quite similar – in both models, the rate of inflation is normally distributed and is modelled as a mean-reverting process. Whilst the *B&H* model works in continuous time and the *Wilkie* model works in discrete time, there is no reason why this should have a significant impact on simulated inflation rates. However, there is one very important difference: the *Wilkie* model is a *single factor* model, where as the *B&H* model has two stochastic variables – as well as the rate of inflation being able to vary, so too can the rate to which it is pulled at any moment in time. This allows more varied shapes of implied future inflation expectations to be generated (which can be important in generating nominal yield curves in the *B&H* models). Further, as we shall see below, it also permits greater flexibility in the pattern of short-term and long-term distributions of the inflation rate. Exhibit 8.4 illustrates the distributions of the inflation rate simulated for a 10-year horizon by the two models.

EXHIBIT 8.4: UNCONDITIONAL DISTRIBUTION OF THE INFLATION RATE

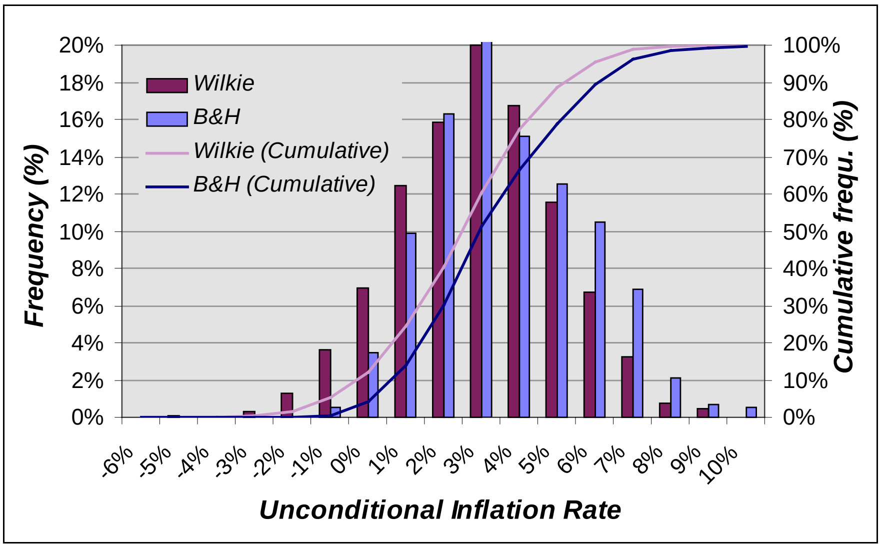

It can be seen that the two models produce broadly similar shapes of unconditional distributions for inflation (with our selection of parameters for the *Wilkie* model). Although the mean parameters of the two models are assumed to be equal (at 2.5%), the *B&H* simulated mean is actually higher at 3.2%. This is because we the simulations have been run with 'reflection' turned on when generating the nominal term structure. This has the effect of producing a higher mean than when the model is run without reflection. Clearly we can quantify this effect and adjust the parameters of the model to offset the increase in the simulated mean if we wish.

Now take a look at exhibit 8.5 below, which shows the distributions of next year's inflation rate.

EXHIBIT 8.5: DISTRIBUTION OF NEXT YEAR'S INFLATION RATE

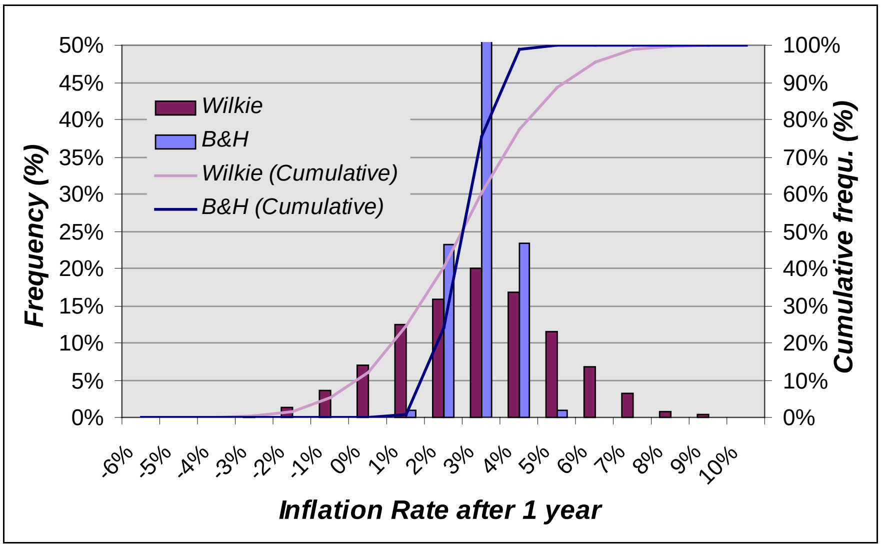

Interestingly, the differences between these distributions are far more significant – the *Wilkie* model generates unfeasibly high short-term inflation volatility. In the *Wilkie* model, the structure of the model makes such high short-term volatility necessary in order to generate longer-horizon distributions with sufficiently high volatility. The use of a second stochastic factor in the *B&H* model means realistic distributions can be generated for both short and long-term distributions of the inflation rate.

Advocates of the *Wilkie* model might argue that the short-term distribution of inflation is not really important – the model is designed for analysing long-term outcomes. However, the direct dependence of asset returns on prevailing and past inflation rates suggests the path taken by inflation could be very important to the model output.

8.7 DERIVATIVE PRICING

The indirect approach to modelling equity prices and the lack of an arbitrage-free term structure for interest rates makes both equity option and interest option pricing and modelling very difficult within the *Wilkie* model. In pricing equity options, finding a closed-form equity option pricing formula as implied by the equity model is very difficult. We could of course simulate option pay-offs (under a risk-neutral set of assumptions for equity returns). We would find that option prices implied by the Wilkie model are lower than market prices as market participants anticipate higher long-term equity volatility than that implied by the mean-reverting equity return behaviour of the *Wilkie* model.

The lack of a full term structure puts the task of pricing/modelling almost any interest rate derivative beyond the scope of the *Wilkie* model. Generally, the analysis of derivatives both in terms of pricing and in terms of their behaviour within a portfolio will require a more sophisticated stochastic asset model, and one which is more consistent with the economic principles of efficient, rational markets.

9. Extensions to the Basic Model

We believe that the model described above is useful in situations where we seek plausible *joint* paths for equities, inflation and both real and conventional bonds. As always, new problems will arise which require extensions to the basic model. We have set out below some very brief comments on potential extensions to the model – some of which are relatively straightforward and some for which further research is required.

9.1 Foreign equity & property

The addition of equity-type assets and yields is fairly straightforward if – as for the domestic equity asset class – returns are specified in excess of short rates and the modeller is prepared to make all equity-type assets switch regimes at the same time. Correlation among equity excess returns is imposed through the correlation structure specified between the impulses to the various components of the model.

9.2 Credit risk & corporate bonds

This is another area in which academic models can be used to extend the basic model to mimic the behaviour of credit spreads in general and the credit behaviour of individual assets. The *Jarrow-Lando-Turnbull* credit model[16] can be used to extend our model (or any other model of the default-free term structure) to allow the analysis of credit-risky bonds within a Monte-Carlo framework.

9.3 Foreign exchange & foreign term structures

The extension of the basic model to foreign currency and foreign inflation adds considerable complexity to the model to ensure:

- Exchange rates respect purchasing power parity relationships at long horizons.
- Inflation patterns can accommodate a global component (given the apparent synchronisation of past inflation patterns).

Whilst these features can be added to the model, we have come across few situations (to date) where they add significantly to the insights gained from modelling a reduced set of asset types.

9.4 Equity mean reversion

Our starting point has been to ignore mean reversion because the statistical evidence – whilst suggestive – does not compel us to include it in the model. Of course, it is possible to impose mean reversion on the model if we choose to. Certainly, there are plenty of people who *do* believe that mean reversion is a strong feature of equity markets and expect to see it represented in models.

---

[16] *Jarrow RA, Lando D, Turnbull SM, "A Markov Model for the Term Structure of Credit Risk Spreads", The Review of Financial Studies 1997, Vol 10, No 2, pp. 481-523.*

9.5 MORTALITY

Errors (in the statistical sense) in forecasting mortality improvement rates over the past 20 years have caused problems to life assurers, contributing to the magnitude of annuity option losses. Again, it is possible to incorporate a stochastic mortality feature into the model to mimic future mortality uncertainty.

9.6 DERIVATIVE PRICING / CONTINGENT CLAIMS VALUATION

Our model can be used to estimate the value of certain types of derivative or contingent claim using the risk-neutral pricing methodology routinely used by investment banks and academics. Mean returns on all assets are set equal to the risk-free rate of interest and then a mean estimate for the expected cash flow from the derivative is calculated using Monte-Carlo methods. This cash flow is discounted at the risk-free rate to obtain the economic value of the derivative or contingent claim.

The risk-neutral pricing methodology makes many assumptions that are violated in the real world (and in part by our model if implemented as described above). As a consequence, these estimated values must be interpreted carefully. Nevertheless, they do give a meaningful benchmark and are often markedly different to estimates for contingent claims made using conventional actuarial techniques. For that reason alone, they deserve the attention of the actuary.

10. CONCLUSION

Life companies are in the risk management business. The risks carried by them come from many different sources. Some risks can be diversified and others must be borne by shareholders and policyholders. Over the past 20 years there has been an enormous increase in the computing power available to the financial planner who sets out to build financial models. In parallel with this technological innovation there has also been rapid development of models by academics and practitioners. The new technology and technical know-how offers the opportunity to address old problems in new ways.

Our report presents our ideas on what constitutes a good model and then sets out a single example of a stochastic model (out of many interesting potential candidates). We give some background information to the messy task of calibrating the model and provide some sample calibrations. Again, there are many interesting alternatives. The output of the model is illustrated with the specific calibrations presented. Finally, we contrast the model with the Wilkie model and very briefly discuss extensions.

The model we have presented is far from perfect – no model ever is. However, we do believe that its relative parsimony, ready economic interpretation and its ability to mimic some important features of financial markets means that it deserves the attention of analysts seeking to model jointly the behaviour of inflation, interest rates and equity markets.

APPENDIX A – INCREMENTING THE TERM STRUCTURE

Given the current values $r_1(t)$ and $r_2(t)$, we can calculate the expected values and variances of $r_1(T)$ and $r_2(T)$. They are as follows:

$$E\left(r_1(T)\middle|r_1(t),r_2(t)\right)=\mu_r+e^{-\alpha_{r1}(T-t)}\left(r_1(t)-\mu_r\right)+\frac{\alpha_1}{\alpha_1-\alpha_2}\left(e^{-\alpha_{r2}(T-t)}-e^{-\alpha_{r1}(T-t)}\right)\left(r_2(t)-\mu_r\right)$$

$$E\left(r_2(T)\middle|r_2(t)\right)=\mu_r+e^{-\alpha_{r2}(T-t)}\left(r_2(t)-\mu_r\right)$$

$$Var\left(r_1(T)\middle|r_1(t),r_2(t)\right)=\frac{\sigma^2{}_{r1}}{2\alpha_{r1}}\left(1-e^{-2\alpha_{r1}(T-t)}\right)+\frac{\sigma^2{}_{r1}}{2\alpha_{r1}}\left(1-e^{-2\alpha_{r1}(T-t)}\right)+\ldots$$

$$\ldots\left(\frac{\alpha_{r1}\sigma_{r2}}{\alpha_{r1}-\alpha_{r2}}\right)^2\left[\frac{1}{2\alpha_{r1}}\left(1-e^{-2\alpha_{r1}(T-t)}\right)+\frac{1}{2\alpha_{r2}}\left(1-e^{-2\alpha_{r2}(T-t)}\right)-\frac{2}{\alpha_{r1}-\alpha_{r2}}\left(1-e^{-(\alpha_{r1}+\alpha_{r2})(T-t)}\right)\right]$$

$$Var\left(r_2(T)\middle|r_2(t)\right)=\frac{\sigma^2{}_{r2}}{2\alpha_{r2}}\left(1-e^{-2\alpha_{r2}(T-t)}\right)$$

As we know $r_1(t)$ and $r_2(t)$ are normally distributed, so all we need are the above moments in order to sample from the distributions and increment the term structure. That is:

$$r_1(T)=E\left(r_1(T)\middle|r_1(t),r_2(t)\right)+\sqrt{Var\left(r_1(T)\middle|r_1(t),r_2(t)\right)}Z_1(T)$$

and

$$r_2(T)=E\left(r_2(T)\middle|r_2(t)\right)+\sqrt{Var\left(r_2(T)\middle|r_2(t)\right)}Z_2(T)$$

where $Z_1(T)$ and $Z_2(T)$ are independent standard normal deviates.

However, if we have a non-zero term premium parameter, g_r, these equations are adjusted as follows:

$$r_1(T)=E\left(r_1(T)\middle|r_1(t),r_2(t)\right)+\sqrt{Var\left(r_1(T)\middle|r_1(t),r_2(t)\right)}\left(Z_1(T)+g_r\sqrt{T-t}\right)$$

and

$$r_2(T)=E\left(r_2(T)\middle|r_2(t)\right)+\sqrt{Var\left(r_2(T)\middle|r_2(t)\right)}\left(Z_2(T)+g_r\sqrt{T-t}\right)$$

The term premium parameter has the effect of adjusting the evolution of the short-term interest rate path, so that the shape of the yield curve is not an un-biased estimate of the path of future short-term interest rates, but instead has a 'loading' which reflects investors' risk preferences. For example, if investors require an additional return to invest in longer-term bonds, then g_r is negative, and the expected value of a rolled-up cash account will be lower than the expected value of an account invested (and continuously-rebalanced) in, say, 10-year bonds. The term premium can be expressed in terms of expected zero-coupon bond yields, or expected returns. The term premiums are as follows (where the term premium is defined as the difference between the expected yield/return on an infinite-maturity zero-coupon bond and the expected yield/return on a instantaneously-maturing zero-coupon bond):

$$Term\,Premium(Yields) = -g_r\left(\frac{\sigma_{r1}}{\alpha_{r1}} + \frac{\sigma_{r2}}{\alpha_{r2}}\right) - \frac{1}{2}\left(\frac{\sigma_1{}^2}{\alpha_1{}^2} + \frac{\sigma_2{}^2}{\alpha_2{}^2}\right)$$

$$Term\,Premium(Returns) = -g_r\left(\frac{\sigma_{r1}}{\alpha_{r1}} + \frac{\sigma_{r2}}{\alpha_{r2}}\right)$$

Note that a positive return term premium requires a negative g_r.

In the inflation model, the risk premium parameter reflects the additional return available on a nominal bond relative to index-linked bond, and works in exactly the same way as in the real interest rate model.

APPENDIX B – CALCULATING COVARIANCE TERM IN NOMINAL TERM STRUCTURE

To calculate the covariance term, we need to evaluate the following variance:

$$Var\left(\exp\left(-R_1(t,T)\right)\right) = \exp\left\{-2E\left(R_1(t,T)\right) + Var\left(R_1(t,T)\right)\right\}.\left(\exp\left\{Var\left(R_1(t,T)\right)\right\} - 1\right)$$

where:

$$E\left(R_1(t,T)\right) = \mu_r.(T-t) + x_1 + x_2$$

$$Var\left(R_1(t,T)\right) = y_1.(T-t) - 2y_2y_3 - 2y_4y_5 + \frac{1}{\alpha_{r1}}y_6y_7 - 2y_8y_9 + y_{10}y_{11}$$

$$x_1 = \left(r_1(t) - \mu_r\right)\frac{\left(1 - \exp\{-\alpha_{r1}(T-t)\}\right)}{\alpha r_{r1}}$$

$$x_2 = \frac{\alpha_{r1}}{\alpha_{r2} - \alpha_{r1}}\frac{\left(1 - \exp\{-\alpha_{r2}(T-t)\}\right)}{\alpha_{r2}} - \left(r_2(t) - \mu_r\right)\frac{\left(1 - \exp\{-\alpha_{r1}(T-t)\}\right)}{\alpha_{r1}}$$

$$y_1 = \left(\frac{\sigma_{r1}}{\alpha_{r1}}\right)^2 + \left(\frac{\sigma_{r2}}{\alpha_{r2}}\right)^2$$

$$y_2 = \left(\frac{\sigma_{r1}}{\alpha_{r1}}\right)^2 - \frac{\sigma^2{}_{r2}}{\alpha_{r2}\left(\alpha_{r1} - \alpha_{r2}\right)}$$

$$y_3 = \frac{\left(1 - \exp\{-\alpha_{r1}(T-t)\}\right)}{\alpha_{r1}}$$

$$y_4 = \frac{\alpha_{r1}\sigma^2{}_{r2}}{\alpha^2{}_{r2}\left(\alpha_{r1} - \alpha_{r2}\right)}$$

$$y_5 = \frac{\left(1 - \exp\{-\alpha_{r2}(T-t)\}\right)}{\alpha_{r2}}$$

$$y_6 = \sigma^2{}_{r1} + \frac{\alpha^2{}_{r1}\sigma^2{}_{r2}}{\left(\alpha_{r1} - \alpha_{r2}\right)^2}$$

$$y_7 = \frac{\left(1 - \exp\{-2\alpha_{r1}(T-t)\}\right)}{2\alpha_{r1}}$$

$$y_8 = \frac{\alpha_{r1}\sigma^2{}_{r2}}{\alpha_{r2}\left(\alpha_{r1} - \alpha_{r2}\right)^2}$$

$$y_9 = \frac{\left(1 - \exp\{-\left(\alpha_{r1} + \alpha_{r2}\right)(T-t)\}\right)}{\alpha_{r1} + \alpha_{r2}}$$

$$y_{10} = \frac{\alpha^2{}_{r1}\sigma^2{}_{r1}}{\alpha^2{}_{r2}\left(\alpha_{r1} - \alpha_{r2}\right)^2}$$

$$y_{11} = \frac{\left(1 - \exp\{-2\alpha_{r2}(T-t)\}\right)}{2\alpha_{r2}}$$

An analogous expression exists for this variance with regard to the inflation term structure. By taking the product of these variances and the correlation, we find the covariance term that applies to the zero-coupon nominal bond price.

APPENDIX C – WILKIE PARAMETERS USED IN SECTION 8

EXHIBIT C.1: WILKIE MODEL PARAMETERS

Inflation		*Dividend Yields*	
qmu	*2.5%*	*yw*	*1.35*
qa	*0.58*	*ya*	*0.6*
qsd	*1.7%*	*ymu*	*0.025*
i0	*2.5%*	*ysd*	*0.155*
		yn0	*Ln(0.025)*
Dividends		*Consols*	
dw	*0.58*	*cw*	*1*
dd	*0.13*	*cd*	*0.045*
dmu	*0.0316*	*cmu*	*0.039*
dy	*-0.175*	*ca1*	*1.2*
dsd	*0.07*	*ca2*	*-0.48*
dx	*0.42*	*ca3*	*0.2*
		cy	*0.06*
		csd	*0.14*
Short-term interest rates			
bmu	*0.198*		
ba	*0.74*		
bsd	*0.18*		
bd0	*0.198*		