DESCRIPTIVE BOND-YIELD AND FORWARD-RATE MODELS FOR THE BRITISH GOVERNMENT SECURITIES' MARKET

BY ANDREW J.G. CAIRNS [1 2 3]

[Presented to the Institute of Actuaries, 20 October 1997,
and to the Faculty of Actuaries, 17 November 1997]

ABSTRACT

This paper discusses possible approaches to the construction of gilt yield indices published by the Financial Times. The existing method, described by Dobbie & Wilkie (1978) splits bonds into high-, medium- and low-coupon bands and fits separate yield curves to each. This method has been identified as susceptible to 'catastrophic' jumps when the least-squares fit jumps from one set of parameters to another set of quite different values. This problem is a result of non-linearities in the least-squares formula which can give rise to more than one local minimum. A desire to remove the risk of catastrophic changes prompted this research which is being carried out as part of the work of the Fixed Interest Working Group.

Recent changes in the taxation of bonds has, further, prompted the need for a review of the yield indices. Significantly, since the announcement of the new tax regime, the old coupon effect has been removed. This has made the use of a single forward-rate curve appropriate for the first time.

A particular form of forward-rate curve is proposed as the basis for a revision of the gilt yield indices. This curve appears to give a significantly better fit than the present yield-curve model. It is also argued that the risk of catastrophic jumps has been reduced significantly.

KEYWORDS

Yield curve; Forward-rate curve; Catastrophic jumps; Least squares; Maximum likelihood; Forward-inflation curve

[1]Department of Actuarial Mathematics and Statistics, Heriot-Watt University, Edinburgh, EH14 4AS, UK
[2]E-mail: A.Cairns@ma.hw.ac.uk
[3]WWW: http://www.ma.hw.ac.uk/∼andrewc/

1 INTRODUCTION

1.1.1 This paper investigates two approaches to the modelling of fixed-interest bond prices and the term structure of interest rates. The work described here makes up part of that of the Fixed Interest Working Group of the Institute of Actuaries and the Faculty of Actuaries.

1.1.2 The model described by Dobbie & Wilkie (1978) has been in use, now, for many years. However, certain problems with their approach have been identified since the model was originally described. The investigation described in this paper was, therefore, stimulated by a need to find a new approach which avoided these problems. In Dobbie & Wilkie's model, parameters are estimated by minimising a least-squares function. It is now known that the function which is being minimised can often have two or more local minima giving rise to unexpected jumps in parameter estimates from one day to the next. This problem is illustrated in Section 2.

1.1.3 The use of forward rates in this piece of work rather than yields was motivated by Gwilt (1982, and in more recent discussions). The initial stages of the work were carried out before the government announced in 1995 that the taxation of government bonds was to change. Prior to the announcement it was not clear whether or not capital and income could be valued on the same basis. The forward-rate approach described in this paper, it is hoped, avoids the problem of multiple minima. It also brings the actuarial profession more into line with current thinking in the field of financial economics. Furthermore, the timeliness of the investigation of this new method was fortunate given the announcement of the tax changes.

1.2 *Approaches to modelling*

1.2.1 There are a number of different types of interest-rate function which one can estimate. These are (for example, see Anderson *et al.*, 1996):

1. The yield curve: The name 'yield curve', in fact, means different things to different people. Here, it is taken to refer to the gross-redemption yield curves estimated by Dobbie & Wilkie (1978) and described in Section 2 of this paper. Although such curves are rather imprecise, practitioners have a good feel for what they are. Indeed since 1995 one can plot the gross redemption yields for individual stocks against maturity and immediately see a relatively well formed curve without the need for any fitting.

2. The spot-rate curve: This is the curve of gross redemption yields on zero-coupon bonds.

3. The (instantaneous) forward-rate curve: This is the curve of implied short-term interest rates in the future. It can be used to price (in a riskless way) forward bond contracts. The use of this curve is developed from Section 3 onwards.

4. The par-yield curve: This curve specifies the interest rates at which new gilts should be priced if they are to be issued at par.

1.2.2 Note that each of the spot-rate, forward-rate and par-yield curves precisely defines the others. This cannot be said of the Dobbie & Wilkie (1978) yield curve unless a precise coupon rate (rather than a band) is specified.

1.2.3 This paper concentrates on the use of *descriptive* models for the forward-rate curve. They are not necessarily intended for use in cheap/dear analysis (although Wiseman, 1994, does advocate their use in this context) or for the pricing of derivative securities. Instead the intention is that they should be used to give a good indication of the interest rates which are implicit in market prices. The model therefore gives us a snapshot of the market and not a movie: that is, it does not describe a stochastic model for the evolution of the term-structure of interest rates. The model can be used in conjunction with the framework of Heath, Jarrow & Morton (1992) which would allow interest-rate derivatives to be priced. This application is described in the paper by Feldman (1998).

1.2.4 A good overview of descriptive models is given in Chapter 2 of Anderson *et al.* (1996). This includes discussion of models for prices, yields and forward rates using parametric curves (as described here), non-parametric curves and semi-parametric spline fitting methods.

1.2.5 As with many curve fitting problems any curve will fit well (that is, have a relatively low standard error) in the middle of the range of dates and less well at the end points. In particular, the curves which have been fitted fit less well at the short end (that is, with a remaining term of less than one year). This problem is exacerbated by the fact that yields on the shortest dated stocks are susceptible to rounding in the prices of stocks.

1.2.6 The curve-fitting process excludes certain stocks with special features, such as double-dated bonds and convertible bonds, which may give rise to prices which are apparently out of line with the market. So any term structure which is produced is giving us information about what interest rates are implicit in the prices of what one might describe as 'normal' stocks.

1.3 *Changes in the Gilts Market*

1.3.1 During the first few months of existence of the working group in 1995 the British government announced that it would be changing the taxation structure of the market. In particular, since April 1996 income and capital gains (for corporate and institutional investors) have been taxed on the same basis, thus removing the advantages of low-coupon stocks for tax payers. (Small, private investors will continue to be taxed on income only. It is felt that the impact of this group on market prices is minimal.)

1.3.2 One reason for this move was to allow the introduction in 1997 of a zero-coupon bond market created by the 'stripping' of a limited number of existing coupon-paying bonds into a series of individual cashflows (see Appendix A).

1.3.3 A particular side-effect of these changes in taxation was that the use of forward-rate curve models (for example, that described in Section 3 of this paper) became much more appropriate for application to the UK. A forward-rate model does not distinguish between income and capital and would not have given an adequate description of the market before the government's announcement. Before 1995, a three-dimensional model taking account of coupon as well as term was necessary (for example, see Dobbie & Wilkie, 1978, Clarkson, 1979 and Feldman, 1977).

1.3.4 When gross redemption yields are plotted one can still see a small coupon effect. This, however, is due to the fact that the gross redemption yield is, in some sense, a weighted average of the forward rates. High-coupon stocks place more weight on lower durations so that if the forward-rate curve is falling then high-coupon stocks will have slightly higher gross redemption yields than low-coupon stocks with the same term. This ordering will be reversed if the forward-rate curve is rising.

1.3.5 The disappearance of the coupon effect is apparent when gross redemption yields are plotted. Whereas previously there would have been a fairly wide band of yields at all terms, there is now quite a well defined curve even before any model-fitting has taken place (for example, see Feldman *et al.,* 1998, Figure 1). An objective view of the decline is plotted in Figure 11 of this paper.

1.4 *Plan for the rest of the paper*

1.4.1 In Section 2 of this paper we consider the yield-curve approach described by Dobbie & Wilkie (1978). The discussion includes an analysis of recent data and illustrates the problem of having multiple minima.

1.4.2 Section 3 introduces the forward-rate curve model which, it is proposed, should replace the Dobbie & Wilkie (1978) model in the construction of the published indices. We define $f(t,t+s)$ to be the instantaneous forward-rate curve observed at time t for payments to be made at time $t+s$. The form of the curve proposed is

$$f(t,t+s) = b_0(t) + b_1(t)e^{-c_1 s} + b_2(t)e^{-c_2 s} + b_3(t)e^{-c_3 s} + b_4(t)e^{-c_4 s}.$$

1.4.3 The curve is a flexible model with four exponential terms and nine parameters in total. However, four of these parameters (the exponential rates) are fixed which reduces the risk of multiple solutions. Rather than take a least-squares approach, as is common in this field, here we describe fully a statistical model which underlies the prices of the stocks. This allows a rigorous process of estimation. The method is, however, shown to be equivalent to a reasonable version of least squares. The section concludes by taking a look at the models described by Nelson & Siegel (1987) and Svensson (1994) and the occurence of catastrophic jumps is demonstrated and also discusses the use of splines as an alternative to parametric curve fitting.

1.4.4 Section 4 presents the results from fitting the model monthly data from January 1992 to November 1996. The monthly data allow us to examine how the coupon effect has disappeared and how stable the forward-rate curve is. Section 5 continues the analysis by carrying out some statistical tests on the results of the fitting exercise – the significance of the coupon effect; correlations between residuals – and the calculation of standard errors of and correlations between various interest rates. Section 6 considers the possibility of catastrophic jumps if the new model is adopted. It is concluded that although multiple maxima could, in theory, exist, catastrophic jumps should be much less frequent if they occur at all.

1.4.5 In Section 7 we consider similar approaches to the pricing of index-linked gilts. The method involves taking the forward-rate for nominal bonds as given and deriving an implied inflation curve. The approach is found to work well even when the implied inflation curve is tied down or restricted at the short end in order for it to tie in with the last year's inflation.

2 THE EXISTING YIELD CURVE APPROACH AND PROBLEMS

2.1 *The current FT-Actuaries approach*

2.1.1 Dobbie & Wilkie (1978) proposed the model

$$y(t) = A + Be^{-Ct} + De^{-Ft}$$

where $y(t)$ is the gross redemption yield on a bond with maturity date time t from now.

2.1.2 In the late 1970's (and right up until 1995) there was a marked *coupon effect* whereby high-coupon gilts had higher gross yields than low-coupon gilts. (In 1978, high-coupon stocks consisted of those with a coupon of more than 7%). Such an effect existed because coupons were taxed as income whereas capital gains were tax-free (see Feldman *et al.,* 1998, Section 3.2).

2.1.3 The precise relationship between coupon and yield was, at the same time discussed by Feldman (1977) and Clarkson (1978). Feldman (1977) felt that, for a fixed term, there should be a linear relationship between coupon and yield. Clarkson (1978) used economic reasoning to argue that the relationship was convex rather than linear. Empirical evidence since these papers for the precise form of this relationship has been mixed.

2.1.4 Dobbie & Wilkie (1978) used a least squares approach to the estimation of the parameters of the yield curve. We, therefore, minimise the function:

$$S(A,B,C,D,F) = \sum w_i[Y_i - y(t_i)]^2$$

where Y_i is the observed gross redemption yield on bond i with maturity t_i. The weights, used in the Dobbie & Wilkie (1978) approach were proportional to the market capitalisation of each stock. In the analysis described in this paper, the w_i were all set to 1 and there was no subdivision of the stocks according to the size of coupon. (This was in the interests of simplicity, since the aim was to illustrate the frequency with which multiple minima can occur, rather than to match, precisely, the the past development of the index through time.) Currently the calculations carried out for the Financial Times place an upper limit on C and F of 3, although such a restriction was not included in this analysis.

2.1.5 For the purposes of this investigation the model was reparametrised in the following way:

$$y(t) = A + Be^{-Ct} + D\frac{e^{-Ft} - e^{-Ct}}{C - F}.$$

This ensures that the estimates of B and D have sensible limits as C and F get closer and closer together. In the original form of the model, as C converges to F, estimates for B and D tend to infinity. With the reformulated model, the limit as C tends to F is

$$y(t) = A + Be^{-Ft} + Dte^{-Ft}$$

and the estimates for B and D do not diverge.

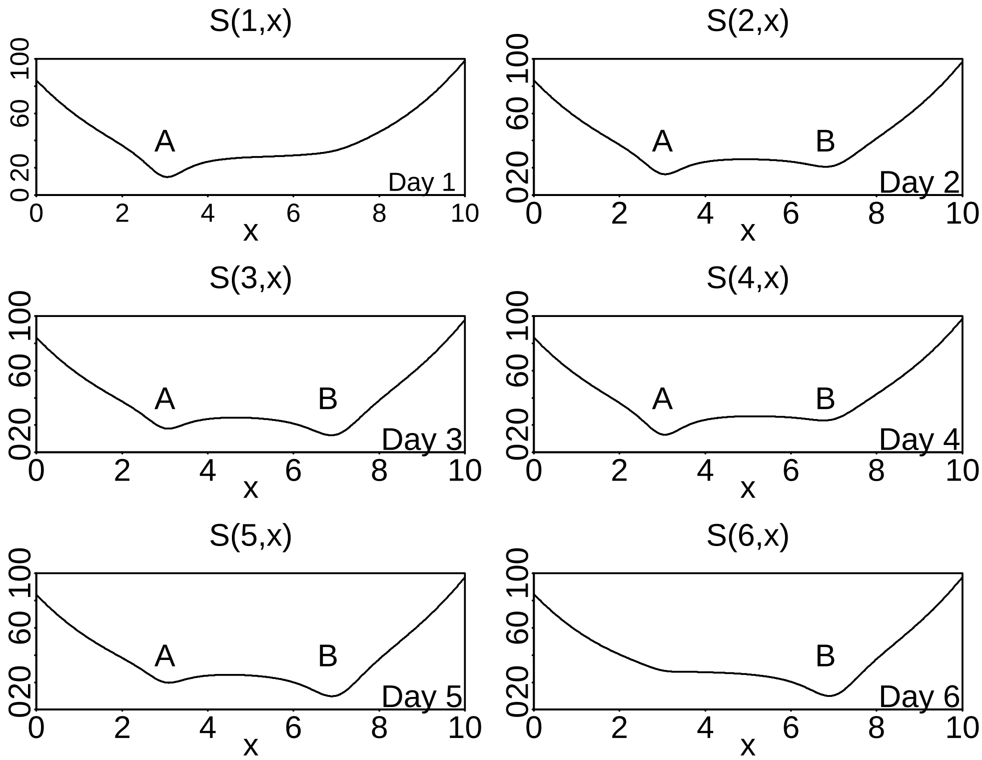

Figure 1: Example of a problem with the minimisation algorithm. A catastrophic jump occurs between days 5 and 6.

2.1.6 Similar curves have been used as a model for forward rates by Nelson & Siegel (1987) and Svensson (1994).

2.1.7 Now $y(t)$ is linear in A, B and D so that it is straightforward to write down the values of $\hat{A} = \hat{A}(C,F)$, $\hat{B}$ and $\hat{D}$, given C and F. This leaves us with the problem of how to minimise the function

$$\hat{S}(C,F) = S(\hat{A},\hat{B},C,\hat{D},F)$$

2.1.8 It is now well known that $\hat{S}(C,F)$ can often have two or more local minima (for example, see Figure 3, further on). This gives rise to occasional jumps from one minimum to another. Such a '*catastrophic*' jump can sometimes go unnoticed. Occasionally, however, such a jump will occur during a period of relatively low volatility in prices, and the resulting change in the shape of the curve will be noticed by at least some practitioners. (Unfortunately the dates of these jumps have not been recorded and their reconstruction is prevented by the lack of detail in the published indices. However, the frequency of jumps is thought to be about one or two per year.)

2.2 *A simplified explanation*

2.2.1 A stylised 1-dimensional version of this problem is illustrated in Figure 1. Suppose that we wish to minimise the function $S(t,x)$ with respect to x. This function varies randomly but continuously through time, t (measured in days). The sequence of plots in Figure 1 represents six snapshots of the evolution of this curve. On day 1, the curve has a unique minimum at A. On days 2 to 5 there are two minima at A and B, but the minimisation algorithm starts at the previous minimum and stays at A (which varies slightly from one day to the next). On days 2 and 4 A represents the global minimum while on days 3 and 5 the chosen minimum at A is only a local minimum and not the global minimum, which is at B. On day 6 the algorithm jumps to to what is now the unique minimum at B. This type of discontinuity between days 5 and 6 is often referred to as a '*catastrophic*' jump and has a branch of mathematics devoted to it (for example, see Poston and Stewart, 1978, Chapter 5). It is not until day 6 when the algorithm finds the global minimum at B as a result of the disappearance of the local minimum at A. If the time sequence had been reversed then the solution to the algorithm would have been at B from time 6 down to time 2 before jumping to A at time 1 (the so-called '*hysteresis*' effect). At the time of the catastrophic jump there may be an identifiable shift in the shape of the fitted yield curve, which would be particularly noticable during quiet trading periods when actual yields had not shifted by any considerable margin.

2.2.2 From the above description it will be seen that the minimisation algorithm (the Nelder-Meade simplex algorithm [different from the better known Simplex algorithm in linear programming]) starts at the previous day's solution and tries to find the nearest local minimum which may or may not be the global minimum. This has the advantage that the frequency of catastrophic jumps is not as high as would be the case if a global optimisation algorithm was used. On-the-other-hand, a disadvantage of this algorithm is that it may stick at a local minimum which gives a significantly worse fit than the global minimum.

2.2.3 In practice catastrophic jumps might occur more frequently than would strictly be the case from the above description. First, we only take a daily snapshot, so that the new starting point for minimisation (that is, yesterday's minimum) might be sufficiently far from today's minimum that the optimisation routine would find the other minimum rather than to continue to track the previous solution. Second, the minimisation algorithm might step inadvertently outside the current minimum's domain of attraction. Both cases have the result that the other minimum might be found despite the fact that the previous minimum still exists.

2.3 *Recent history*

2.3.1 A sample period of 100 consecutive trading days during 1995 is given in Figure 2.

2.3.2 The program which generated this output took a slightly different approach to the problem from the framework of Dobbie & Wilkie (1978). First the data were not split into high- medium- and low-coupon stocks (for convenience). Second the weights (the w_i) were all set equal to 1. Third the program located and tracked as many local minima as it could track down (that is, some minima might have been missed out). In Figure 2 we plot the progress of the local minimum estimates of C, F and of the Root-Mean-Squared Error (RMSE). Different numbers of

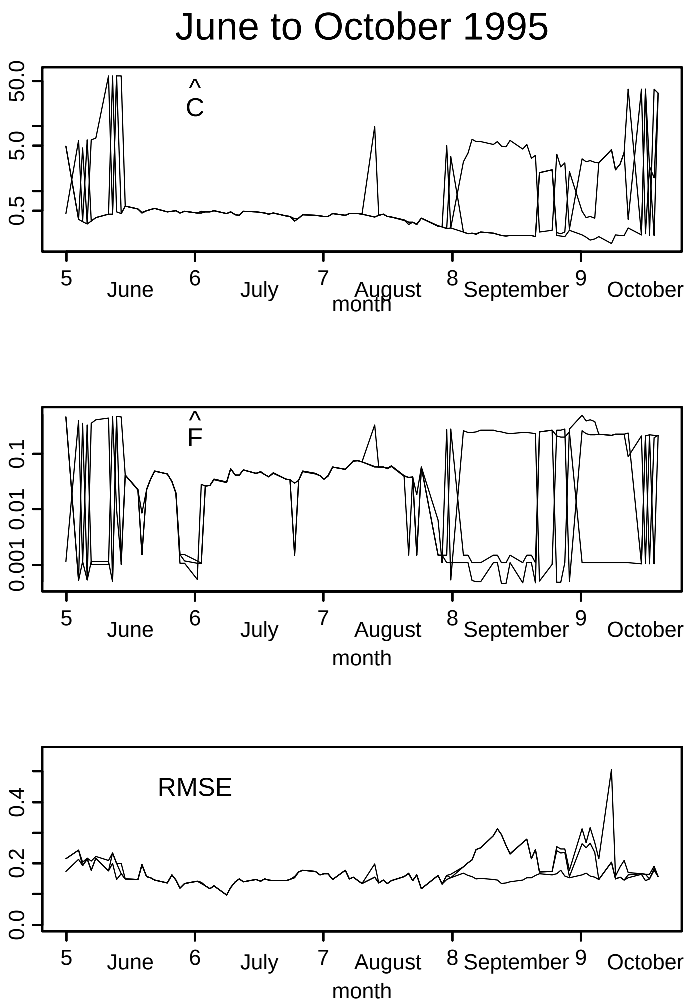

Figure 2: Evolution of the estimates for $\hat{C}$ and $\hat{F}$, and of the Root Mean Squared Error (RMSE) in the estimate of the gross redemption yield during the period June to October 1995. (A figure of 0.2 equates to a standard error of 20 basis points in the estimate of the yield.) On each date the algorithm located and tracked one, two or three local minima. All one/two/three are plotted.

minima were observed at different times: for example, around the middle of July there appeared to be a unique minimum; around the end of July there were two minima; and early in October there were three minima each giving noticably different yield curves. In September, note that the global minimum gives a substantially better fit than the other local minimum. In June, on-the-other-hand, the two minima both fit the data almost equally well with the location of the global minimum switching between the two regularly.

2.3.3 Similar results were observed when then program was run using monthly data from 1992 to 1996.

2.3.4 Figure 3 gives three yield curves on 27 September 1995 and a contour plot of $\hat{S}(C,F)$. (C and F have been rescaled in the contour plot to $\sqrt{C}$ and $\sqrt{F}$ in order to illustrate the situation more clearly.) P, Q and R were the three local minima identified on that date and mentioned above. The root mean squared errors are 18, 23 and 24 basis points, so each curve fits reasonably well. The best fit, P, succeeds by fitting the short end very well and adequately elsewhere. At terms 0 to 10 years Q and R are very similar and quite different from P. Between 10 and 15 years they are all similar, and after that P and R are similar and lower than Q. A catastrophic jump can occur when prices are relatively stable. From the example given here, such a jump from P to Q would be noticable.

2.3.5 Suppose, then, that we started at the global minimum at the start of June. There is always a minimum with $\hat{C}$ in the range 0.1 to 0.5 and it is suspected that the original algorithm would have tracked this all the way through without any catastrophic jumps. However, it is necessary to track simultaneously $\hat{F}$ to check for this continuity. The lowest values of $\hat{F}$, however, are slightly unreliable as many of the small values of $\hat{F}$ are combined with very small values of $\hat{D}$ and so have very little impact on the shape of the yield curve. When this happens the minimisation routine is not able to improve upon the previous estimate of F due to rounding errors.

2.3.6 If there were no catastrophic jumps then a potentially serious problem would have arisen during the early part of August. The old algorithm would have picked a value of $\hat{C}$ of between 0.1 and 0.5: a local minimum. The global minimum during a good part of this period gave a much higher value for $\hat{C}$ and, more importantly, a much better fit, and the algorithm would have completely missed this feature.

2.3.7 The combination of possible catastrophic jumps, of an algorithm which might track a poor local minimum, and of the recent changes in the taxation of gilts means that a fresh approach to the yield curve is now appropriate.

3 A FORWARD-RATE CURVE APPROACH TO THE PRICING OF GILTS

3.1 *Forward rates*

3.1.1 An alternative to the use of yield curves, and which is popular in the field of financial economics, is the use of forward-rate curves. In a traditional actuarial context the forward rate at time s is sometimes referred to as the force of interest $\delta(s)$ (for example, see McCutcheon &

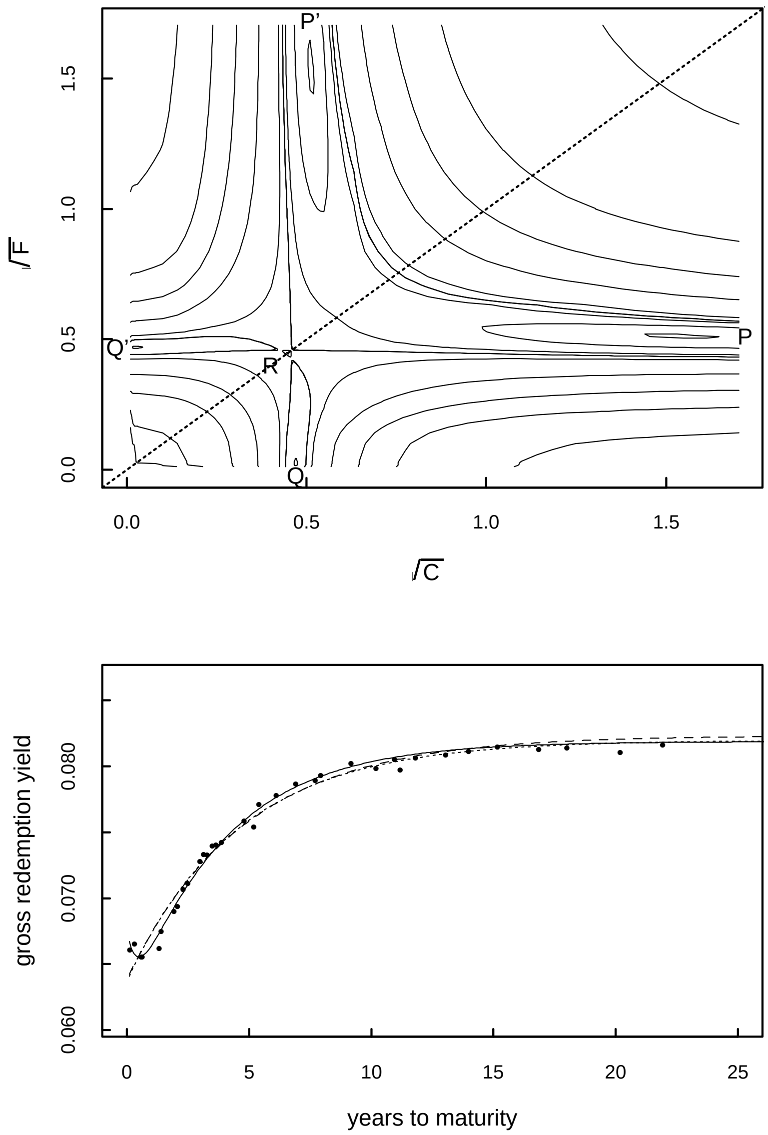

Figure 3: 27 September 1995. (Top) A contour plot of the function $\hat{S}(C,F)$ for this date showing the unusual feature of three local minima at P, Q and R (plus the reflections of P and Q: P' and Q'). R is on the diagonal (dotted line). The locations of these local minima are typical for other dates. The corners of the plot are all 'high' points. (Bottom) Three yield curves corresponding to the local minima at P (solid curve), Q (dotted curve) and R (dashed curve). The dots represent the gross redemption yields on the gilt market at the end of the day.

Scott, 1986). The force of interest is, however, a one dimensional deterministic process whereas the forward-rate curve is two dimemsional and evolves stochastically. This two-dimensionality reflects the difference, for example, between an investment in an n-year zero-coupon bond and an investment in the short-term money market for *n* years.

3.1.2 In the financial economics literature forward-rate curves play a central role in the pricing of bonds. In particular, the general framework described by Heath, Jarrow & Morton (1992) takes the initial forward-rate curve as input (almost regardless of its shape) and describes how the curve must evolve thereafter in order to maintain an arbitrage-free market. Other models such as those proposed by Vasicek (1977) and Cox, Ingersoll & Ross (1985) take only one or two variable factors as input (such as the short-term interest rate) and estimate other fixed parameters by getting the best match to the observed forward-rate curve (for example, see Chaplin, 1998). Such models have necessarily been relatively simple in order to facilitate the pricing of derivative instruments. As a consequence the models often fit available data relatively poorly or do not provide an accurate picture of how the forward-rate curve actually evolves through time.

3.1.3 The forward-rate curve used here is more complex and does not, on its own, provide an arbitrage-free framework within which derivatives can be priced. It does provide a much better fit than the existing, simpler equilibrium models. As mentioned above, the forward-rate curve described here could provide the input for a stochastic model within the Heath-Jarrow-Morton framework. This application, however, was not considered to be essential as it was not designed as a means of pricing derivatives: instead it is designed to give an indication of what interest rates are currently being implied by the market.

3.1.4 Let $f(t,t+s)$ be the forward-rate curve observed at time t for $0 \leq s < \infty$. That is, $f(t,t+s) = -\partial \log P(t,t+s)/\partial s$ where $P(t,t+s)$ is the price at time t of a zero-coupon bond maturing at time $t+s$. (Note that $P(t,t+s)$ is to be distinguished from P_i, $\hat{P}_i$ and $\hat{P}_i(b)$ which are used later for the prices of coupon-bearing bonds.)

3.1.5 The model for $f(t,t+s)$ used in this paper is:

$$\begin{aligned} f(t,t+s) &= b_0(t) + \sum_{i=1}^{4} b_i(t) e^{-c_i s} \\ &= b_0 + \sum_{i=1}^{4} b_i e^{-c_i s} \text{ for brevity.} \end{aligned}$$

If the value of c_i is small then the relevant value of b_i affects all durations whereas if c_i is large then the relevant value of b_i only affects the shortest durations.

3.1.6 Although this is a nine-parameter model, the model is structured in such a way that c_1, c_2, c_3 and c_4 are regarded as fixed parameters and that we optimise over the b_i only. As is described in Section 4 and in Section 5.2 a choice of $c = (0.2, 0.4, 0.8, 1.6)$ was found to be appropriate for the period 1992 to 1996. With this set of values for c, a wide variety of shapes for the forward-rate curve can be generated by using different combinations of values for the b_i. In particular the forward-rate curve can have 0, 1, 2 or 3 turning points. Furthermore, b_1 will have more influence over the long term, while b_4 will have more influence over the short term. b_0, clearly, sets the overall level of the forward-rate curve.

3.1.7 This particular model in fact belongs to the more general class of *Restricted Exponential Polynomial* curves described by Bjørk & Christensen (1997). For consistency in the Ho & Lee and Hull & White models they show that it is a necessary condition that the c_i remain constant through time. (Further constraints are discussed in paragraph 5.2.2.)

3.1.8 This model was favoured over an alternative five-parameter model

$$f(t,t+s) = b_0(t) + \sum_{i=1}^{2} b_i(t) e^{-c_i(t)s}$$

where optimisation occurs over b_0, b_1, b_2, c_1 and c_2. This preference arises out of the conjecture that the former model cannot give rise to catastrophic jumps while the latter is known to have multiple maxima in the likelihood function even for the simple case where there are only zero-coupon bonds (when the optimisation problem becomes the same as that considered by Dobbie & Wilkie, 1978, and discussed in Section 2). Similar models for the forward-rate curve have been used by Nelson & Siegel (1987) –

$$f(t,t+s) = b_0(t) + b_{10}(t)\exp(-c_1(t)s) + b_{11}(t)s\exp(-c_1(t)s)$$

– and by Svensson (1994) –

$$f(t,t+s) = b_0(t) + b_{10}(t)\exp(-c_1(t)s) + b_{11}(t)s\exp(-c_1(t)s) + b_{21}(t)s\exp(-c_2(t)s),$$

– and these are (as shown in Section 3.5) also subject to catastrophic jumps. The Nelson & Siegel (1987) and Svensson (1994) models only allow for a maximum of one and two turning points (respectively), so that the alternative 5 out of 9 parameter model has a much richer range of curves available for fitting the data.

3.1.9 The curve is equivalent to the polynomial approach to the fitting of forward rates:

$$\begin{aligned} f(t,t+s) &= b_0 + \sum_{j=1}^{4} b_j x^{c_j/c_1} \\ \text{where } x &= e^{-c_1 t}. \end{aligned}$$

3.1.10 More normally, however, the polynomial approach would use the powers 0, 1, 2, 3, 4 instead of the 0, 1, 2, 4, 8 used here. Since we restrict ourselves to the range $0 < x < 1$, each set of polynomials can generate an equally rich family of curves (in particular, up to three turning points).

3.1.11 The discount function at time t for payments at time $t+s$ is

$$\begin{aligned} P(t,t+s) &= \exp\left[-\int_0^s f(t,t+u)du\right] \\ &= \exp\left[-b_0 s - \sum_{i=1}^{4} b_i \frac{1-e^{-c_i s}}{c_i}\right] \end{aligned}$$

3.1.12 In the estimation problem we have N fixed interest stocks. Stock i has coupon g_i per annum per £1 nominal, with coupons payable at times $T_i = \{t_{i1}, \ldots, t_{in_i}\}$ from time t_0 (so that, $0 < t_{i1} \leq 0.5$, $t_{ij+1} - t_{ij} = 0.5$ and $\tau_i = t_{in_i}$ is the time to redemption). The first payment may be zero if the stock is currently ex-dividend. Given a particular parameter set b ($b = (b_0, b_1, b_2, b_3, b_4)$) and corresponding forward rate curve $f(t, t+s)$ the theoretical price of stock i is

$$\begin{aligned}
\hat{P}_i &= \hat{P}_i(b) \\
&= \sum_{i=1}^{n_i} c_{ij} P(t, t+t_{ij}) \\
\text{where } c_{i1} &= \begin{cases} 0 & \text{if the stock is ex-dividend} \\ \frac{g_i}{2} & \text{otherwise} \end{cases} \\
c_{i2} = \ldots = c_{i,n_i-1} &= \frac{g_i}{2} \\
c_{in_i} &= 1 + \frac{g_i}{2}
\end{aligned}$$

3.1.13 Now the $\hat{P}_i$ are theoretical prices, whereas the actual prices P_i are subject to errors. It is appropriate, therefore, to formulate some form of statistical model.

3.2 *Statistical model*

3.2.1 Three sources of error taken into account here.

3.2.2 First, published prices are subject to rounding to the nearest $1/32$ per £100 nominal (described here as 'Type 1' errors). This affects primarily the short-dated stocks and can lead to large errors in the gross redemption yield. Short-dated stocks in an appropriate model therefore will have standard deviations of about $1/32\%$ or, equivalently, the logarithm of the price has a standard deviation of $1/3200$ (and, for the sake of argument, a normal distribution). It can readily be argued that this standard deviation should be less than this, in the sense that the 'true' midmarket price should lie in a band of width $1/32$. The value taken here is, therefore, more cautious. As described below, an important aspect of this error structure is that it simulates weightings which tend to zero as the term of a stock approaches zero.

3.2.3 Second, (ignoring temporarily the $1/32$ error) we postulate that the differences between the actual and expected yields on independent stocks are independent and identically distributed normal random variables with mean 0 and standard deviation σ_d (Type 2 errors). This part of the model is consistent with the statistical model which is contained implicitly in the methods of least squares adopted by Dobbie & Wilkie (1978) and by Phillips (described in Feldman *et al.*, 1998).

3.2.4 Third, a limit was placed or the magnitude of the price errors which can be experienced by long-dated stocks (Type 3 errors – or rather restriction of errors). For example, if the second type of error described above used $\sigma_d = 0.0005$ (5 basis points) (an appropriate value since 1995) then a stock with a duration of 10 years would have a standard deviation of 0.5% of the price (or about 50 pence). This can be considered to be too large.

3.2.5 A precise description of the error structure is given in Appendix B. In summary, however, the statistical model is as follows:

$$\log P_i \sim N\left(\log \hat{P}_i, \frac{\sigma_d^2 d_i^2}{w_i}\right) \tag{1}$$

where d_i is the duration of stock i and w_i is the weight attached to stock i. (A larger weight is equivalent to a lower standard deviation. Hence the form of equation 1 above.) The errors are assumed to be independent. This assumption relies on the hypothesis that the forward-rate curve being fitted does, indeed give a good description of the market. In practice, the model is, of course, only an approximation to reality. One can then argue that where the theoretical curve deviates from the true curve errors in the estimated prices will be correlated. Nevertheless, it is demonstrated in Section 5 of this paper that such correlations, if they do exist, are relatively small.

3.2.6 In this paper, we have not considered the time-series aspect of the data. Each date has been taken as a snapshot and the curve fitted without reference to previous dates. In effect this assumes that the errors in the prices of individual stocks are independent from one day to the next. In reality the underlying curve depends upon previous values and individual price errors are correlated through time. A full analysis taking account of this time-series aspect would involve a very complex Kalman-filtering routine. While the number of stocks is significantly larger than the number of parameters and while price errors are small, the marginal gain from such an exercise is small.

3.2.7 In practice there will exist other minor features which could be added to this model. For example, each stock, i, could have a bias term, μ_i, to account for special characteristics, such as benchmark status or if it has special taxation concessions. (For example, $8\frac{3}{4}$% 2017 had benchmark status until May 1996 before which it attracted a premium of as much as 4 basis points. This premium immediately disappeared at the time of the introduction of 8% 2021.) In the fitting of the forward rate curve described below, we have specifically excluded double-dated bonds and convertible bonds which have a clear bias attached to their prices which results from their inherent option characteristics.

3.2.8 The log-likelihood function for the data is thus

$$g_1(P|b) \quad = \quad -\frac{1}{2}\sum_{i=1}^{N}\left\{\log(\sigma_d^2 d_i^2/w_i) + \frac{(\log P_i - \log \hat{P}_i)^2}{\sigma_d^2 d_i^2/w_i}\right\}$$

3.3 *Equivalent least squares approach*

3.3.1 Maximisation of the likelihood function with respect to b can easily be shown (see Appendix C) to be equivalent to minimisation of the weighted least-squares function

$$S^2 \quad = \quad \sum_{i=1}^{N} w_i \frac{(P_i - \hat{P}_i)^2}{d_i^2}$$

$$\text{where } w_i = \text{the weight attached to stock } i \text{ (see Appendix B)}$$
$$\rightarrow \begin{cases} \infty & \text{as } d_i \rightarrow \infty \\ 0 & \text{as } d_i \rightarrow 0 \end{cases}$$

3.3.2 If the Type 3 (restriction of) errors are removed then $w_i \rightarrow 1$ as $d_i \rightarrow \infty$.

3.3.3 This least squares approach is very similar to that suggested by Phillips (see the description in Feldman *et al.*, 1998) who uses $w_i = \min\{d_i, 1\}$. This is more familiar to many practitioners and it is reassuring to know that there is a sensible statistical model underlying the least-squares approach.

3.4 *Avoiding negative forward rates*

3.4.1 In the early stages of this investigation it was found, occasionally, that, for example, some predicted rates could be negative; or that the long term forward rate implied by fitted market prices (that is, estimated rather than actual prices) could tend to infinity. (This could occur if the c_i are being optimised as well as the b_i.) Such outcomes were considered to be undesirable: the first clearly leads to arbitrage possibilities and would not exist in practice; and in the second example it seems very unlikely that the market would expect interest rates to increase without bound. (In the latter case, however, some plausible autoregressive models which have increasing risk premia do allow forward rates to drift off to plus infinity. The author feels less strongly about this constraint.)

3.4.2 It was therefore felt appropriate to introduce some way of constraining the solution to forward-rate curves which are bounded and non-negative.

3.4.3 One way to do this would be to maximise the log-likelihood function $g_1(P|b)$ subject to the constraint that $f(t,t+s)$ is bounded and non-negative. This could still lead to solutions in which the forward rate was equal to zero for some value of s. From a subjective point of view this still seems to be unlikely since if $f(t,t+s) = 0$ for some $s \geq 0$ we create an arbitrage opportunity of order ds. That is, we could go long in zero-coupon bonds with term s and short in zero-coupon bonds with term $s+ds$. The difference in price would be $o(ds)$ since $f(t,t+s) = 0$ whereas investing the proceeds of the s-dated bond would create a profit of the same order as $r(t+s).ds$ (where $r(t)$ is the risk-free rate of interest). Since $r(u) > 0$ with probability greater than 0 (under any reasonable model) this means we have a guaranteed profit of $O(ds)$.

3.4.4 We therefore adopt a Bayesian approach which incorporates this subjective view-point. This is done in the form of saying that before we look at the data we would say that it is more likely that the forward rate t years hence, $f(t,t+s)$, is more likely to be, say, 4% than 2% which itself is more plausible than, say, a rate of 1%. This is done by introducing some prior distributions for the parameters underlying the forward-rate curve.

3.4.5 Now we do not want the incorporation of this subjective view to swamp the data. The prior distribution was therefore chosen to reflect this view in a weak way. There are various ways of doing this, and the following joint prior density function for $\delta = f(t,t)$, $f(t,t+10)$ and $f(t,\infty)$ was selected:

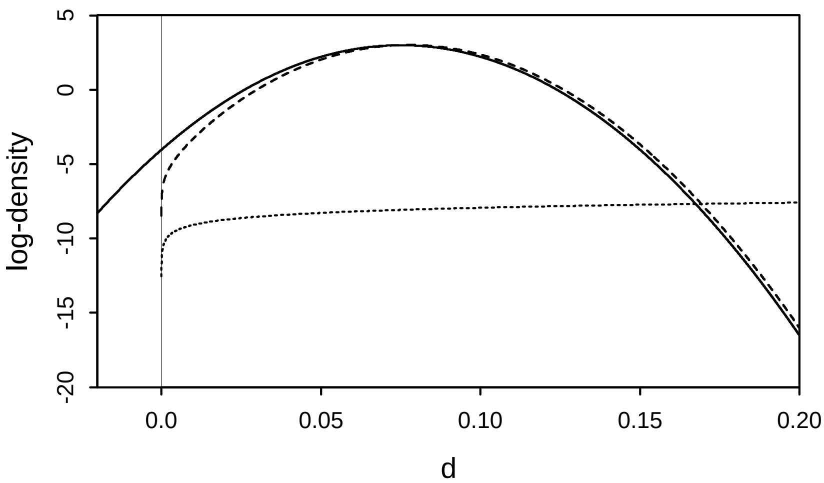

Figure 4: Possible log-likelihood (solid curve), log-prior density (dotted) and log-posterior density (dashed) curves for a parameter d. (For example, $d = f(t,t+10)$.)

$$\begin{aligned} \text{prior-density} &= \exp[g_2(f(t,t))] \times \exp[g_2(f(t,t+10))] \times \exp[g_2(f(t,t+\infty))] \\ g_2(\delta) &= \text{individual log-prior-density function} \\ &= (\alpha-1)\log\delta - \beta\delta + \text{constant} \end{aligned}$$

where $\alpha = 1.5$ and $\beta = 0.01$. (The prior thus assumes that $f(t,t)$, $f(t,t+10)$, and $f(t,t+\infty)$ are independent.) β needs to be positive to ensure that we have a proper prior density function (that is, one which will integrate over its full range to 1), and to inhibit the possibility that the forward rate tends to $+\infty$. β could equally well have been set at 0.1 or 0.001 without significantly altering the results of the estimation process. α must be greater than 1 to inhibit the forward rate from getting too close to zero. If α is greater than 2 then there is, perhaps, too much of a push away from zero interest rates. The choice of 1.5 for α is the chosen compromise: again, though, the precise value has little impact. The prior distribution is very diffuse relative to the likelihood function so it has very little impact on the posterior distribution except when some points on the forward rate curve get too close to zero.

3.4.6 A stylised, one-dimensional representation of this is provided in Figure 4. The three curves are log-likelihood (solid curve), log-prior density (dotted curve) and log-posterior density (dashed curve) for some quantity d (for example $f(t,t+10)$). It can be seen that the main effect of the introduction of a prior distribution is to ensure that the log-posterior tends to $-\infty$ as d tends to 0. The maximum log-likelihood occurs at $\hat{d} = 0.075$ while the maximum log-posterior occurs at about $\hat{d} = 0.078$. This shift is relatively small. In reality, however, the difference would be even smaller given that the log-likelihood will normally be concentrated over a much narrower range.

3.4.7 The log-posterior density is thus

$$g_3(b|P) = g_1(P|b) + g_2[f(t,t)] + g_2[f(t,t+10)] + g_2[f(t,\infty)] + \text{constant}.$$

3.4.8 In practice the the estimates of parameter values were not found to be sensitive to the choice of values for α or β, unless the unconstrained maximum likelihood estimate was predicting negative forward rates.

3.4.9 The parameter estimates are $\hat{b}$ where

$$g_3(\hat{b}|P) = \max_b g_3(b|P)$$

3.4.10 The very diffuse nature of the prior distribution means that these estimates will be very close to the maximum likelihood estimates.

3.5 *The Svensson model*

3.5.1 Nelson & Siegel (1987) proposed the following model for forward rates

$$f(t,t+s) = b_0(t) + b_{10}(t)\exp(-c_1(t)s) + b_{11}(t)s\exp(-c_1(t)s).$$

3.5.2 This was extended by Svensson (1994) who included an extra polynomial-exponential term

$$f(t,t+s) = b_0(t) + b_{10}(t)\exp(-c_1(t)s) + b_{11}(t)s\exp(-c_1(t)s) + b_{21}(t)s\exp(-c_2(t)s),$$

bringing the total number of parameters to be estimated to 6. In order to improve the estimation process the model was reformulated as

$$\begin{aligned} f(t,t+s) &= b_0(t) + b_{10}(t)\exp(-c_1(t)s) + b_{11}(t)s\exp(-c_1(t)s) \\ &\quad + b_{21}(t)s\frac{[\exp(-c_2(t)s) - \exp(-c_1(t)s)]}{(c_1 - c_2)}. \end{aligned}$$

3.5.3 In some cases the estimates of c_1 and c_2 get very close to one another. In the new formulation, as $c_1 \to c_2$:

$$f(t,t+s) = b_0(t) + b_{10}(t)\exp(-c_1(t)s) + b_{11}(t)s\exp(-c_1(t)s) + b_{21}(t)s^2\exp(-c_1(t)s)$$

and $b_{21}(t)$ does not get arbitrarily large. In the original formulation, as $c_1 \to c_2$, $b_{21}(t)$ tends to infinity.

3.5.4 The Svensson (1994) model was fitted to price data using the same statistical structure described in Section 3.2. It is possible to find more than one maximum (if more than one exists) by using the following procedure:

3.5.5 Take a relatively crude grid of values for c_1 and c_2 (these are the only non-linear parameters in a zero-coupon bond market). For each pair of values (c_1, c_2) maximise the log-posterior density function over the remaining (linear) parameters $b_0, b_{10}, b_{11}, b_{20}$. If there is

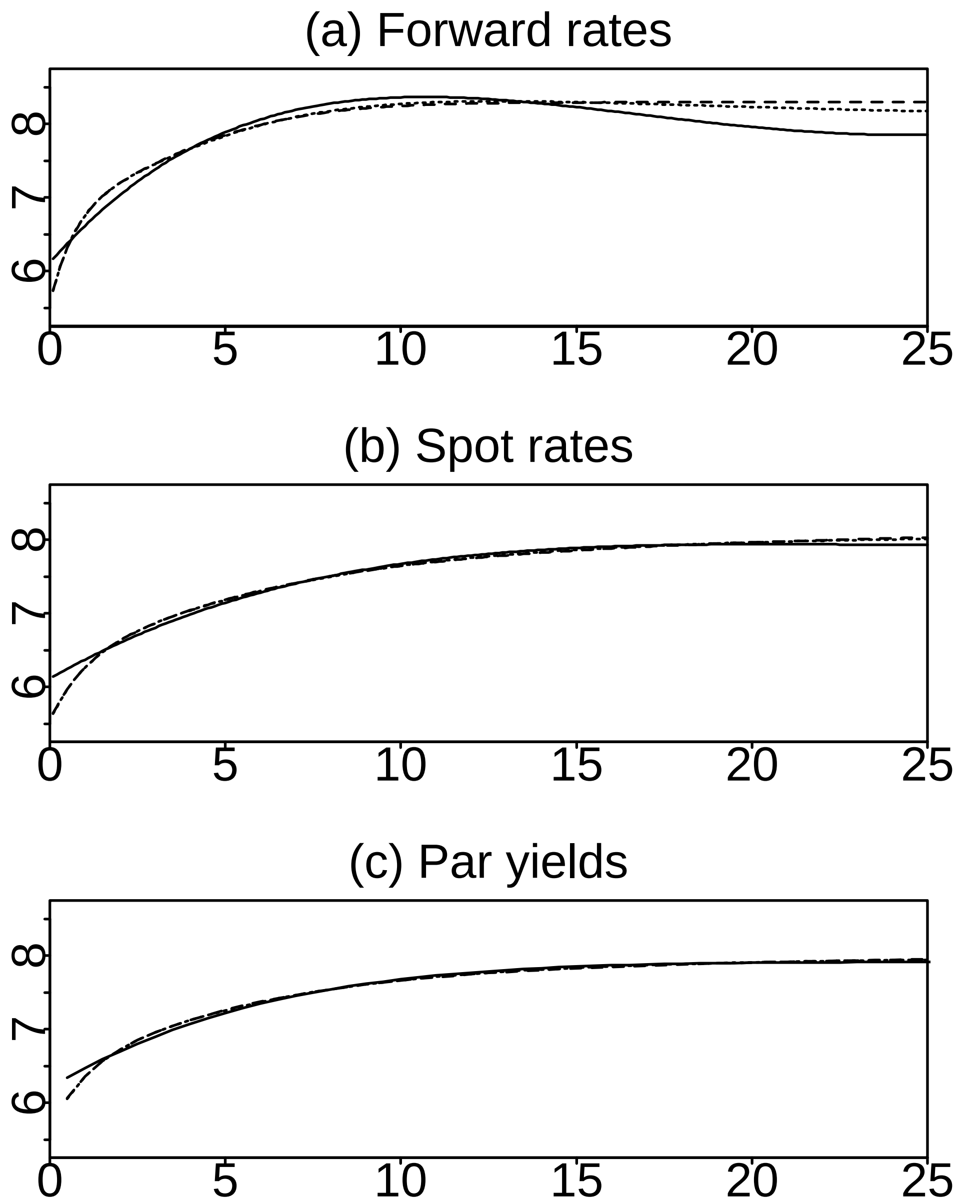

Figure 5: Forward-rate, spot-rate and par-yield curves for the Svensson model on 1 November 1996. $(c_1, c_2) =$ solid curve (0.0369,0.0267); dotted curve (0.1754, 0.9282); dashed curve (0.6847, 0.6847).

more than one maximum then these will show up in a contour plot (similar to Figure 3) of the maximum log-posterior function when plotted as a function of c_1 and c_2. The true local maxima can then be found by carrying out unconstrained maximisation starting from the local maxima on the grid.

3.5.6 As an example the model was fitted to prices published on 1 November 1996. It was found that there were (at least) three local maxima. The three fitted forward-rate curves are plotted in Figure 5(a) and their equivalent spot rates in Figure 5(b). Importantly, the three maxima have very similar log-posterior values: in other words, they all fit the data almost equally well.

3.5.7 It can be seen that the three fitted curves are significantly different. The differences are more marked for forward rates, but they are still noticeable when we consider the spot rates or par yields. The dotted and dashed curves are relatively similar throughout in comparison with the solid curve (which gave the best fit of the three curves).

3.6 *Splines*

3.6.1 The models described, so far, in this paper all come under the heading of parametric curves. A popular alternative to this is the use of splines. In the context of bond pricing this was pioneered by McCulloch (1971, 1975) with subsequent work relevant to the present discussion by Mastronikola (1991) and Deacon and Derry (1994) for the Bank of England.

3.6.2 In the context of the working group the use of splines was rejected on the grounds that a term structure estimated using splines is too highly parametrised. It was considered important that complete information about the term structure can be published in the Financial Times: this increases significantly the usefulness of the yield indices. The fitted curve, also, is intended to give a broad picture highlighting the main features of the market rather than a detailed version, warts and all.

3.6.3 Other points considered were as follows:

1. Splines can produce a more lumpy or wiggly curve than a parametric model because they are more highly parametrised. Furthermore, in a heterogeneous market where some gilts have special features, splines will, it is suspected, overcorrect for ‘specials’ (for example, strippable or benchmark gilts) more than a parametric curve which will average out over a broader range of gilts in the same neighbourhood. However, where there is a preponderance of specials (for example, long-term gilts) both approaches will be heavily influenced by the relatively high density of specials.

2. The new market in zero-coupon bonds will be homogeneous and the application of splines to this market is entirely appropriate for some applications.

3. The Bank of England working paper by Deacon and Derry (1994) indicates that fitted spline curves are sensitive to the number and positions of knots right across all durations. The parametric curve described in this paper only exhibits sensitivity at the extremes of the maturity range.

4 RESULTS

4.1 *General description*

4.1.1 Monthly data running from 1992 to 1996 (59 separate dates) was used, representing a typical period straddling the announcement of the tax changes. The announcement early in 1995 that the taxation of gilts was to change means that yield curves could almost be drawn by hand from that date onwards. Certain stocks were excluded from the fitting process: in particular, convertible stocks, those with optional redemption dates (including the two stocks which will continue to be taxed under the old regime), all undated stocks (except for 4% Consols and 2.5% Treasury) and the stocks 3% 1992, 6% 1993 and 6% 1999. (The last three stocks were observed, in general, to be well out of line with the rest of the market. However, they were reinstated when the time series of errors in the prices of individual stocks were considered in Section 5.)

4.1.2 Considering the 1992-1996 data (Figure 6), it was found that $\sigma_d = 0.001$ to 0.002 (i.e. 10 to 20 basis points in the yield) was about right for the pre 1995 data. For the 1995 and 1996 data (Figure 6) $\sigma_d = 0.0005$ was more appropriate. This difference is not surprising as prior to 1995 there was a marked coupon effect, which the present model for the forward-rate curve was unable to take account of, given that it does not distinguish between income and capital. Nevertheless, it was perhaps surprising that the value of $\sigma_d = 0.001$ was so low. (Further discussion of the coupon effect is given in Section 5.)

4.1.3 The effect of the choice of σ_d^2 was also considered. (It is intended that, like c, σ_d^2 should be fixed at the outset.) If there is no prior distribution then the value of σ_d^2 has no effect on the estimate of b. If prior distributions are included, there is an effect on the estimate of b but this is very insignificant.

4.1.4 Results for 1992-1996 are plotted in Figure 6. (The estimated instantaneous spot-rate $f(t,t)$ and the 3-month par-yield $y(0.25)$ are given to allow comparison with money-market rates and treasury-bill prices.)

4.1.5 It can be seen that there was considerable variation over this period, but this variation appeared to be relatively smooth. Possibly the least variation was observed in the 10-year forward rates and yields.

4.1.6 Several choices for the vector $c = (c_1, c_2, c_3, c_4)$ were considered. Of these, $c = (0.2, 0.4, 0.8, 1.6)$ was found to be best over both the 1992-1996 period investigated. Other choices were comparable in terms of fit and an example of one alternative ($c = (0.1, 0.2, 0.4, 0.8)$) is given in Figure 6. However, $c = (0.1, 0.2, 0.4, 0.8)$ leads to much greater instability in the long-term forward rate $f(t, t+\infty)$ as can be seen in Figure 6 (dotted curve). Although this gives some cause for concern, it is unlikely that this problem could be removed through the use of an alternative model, because there will always be a range of possible parameter sets which are almost as good as the optimum and which give a wide variety of log-term forward rates.

4.1.7 In Figure 7 we compare the estimated 3-month spot rate with the 3-month spot rate for Treasury Bills (TB's) and for Commercial Bills (CB's). This figure includes a 95% confidence band for the estimated 3-month spot rate calculated along the lines described in Section 5. This band is relatively wide and reflects the fact that we are looking at one end of the yield curve rather than the middle (which has lower standard errors).

4.1.8 The rates for TB's and CB's have been taken from *Financial Statistics* published by

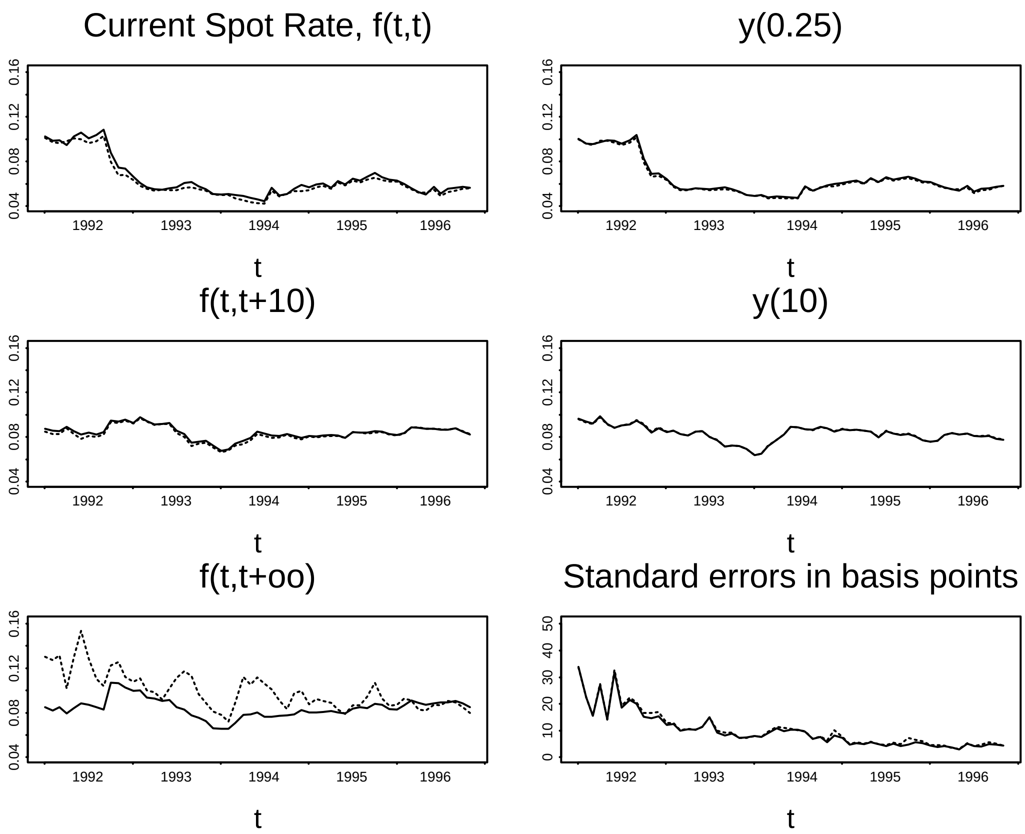

Figure 6: Evolution of various points on the forward-rate curve and on the yield curve during the period 1992 to 1996. The solid lines represent values for $c = (0.2, 0.4, 0.8, 1.6)$ and the dotted lines represent values for $c = (0.1, 0.2, 0.4, 0.8)$.

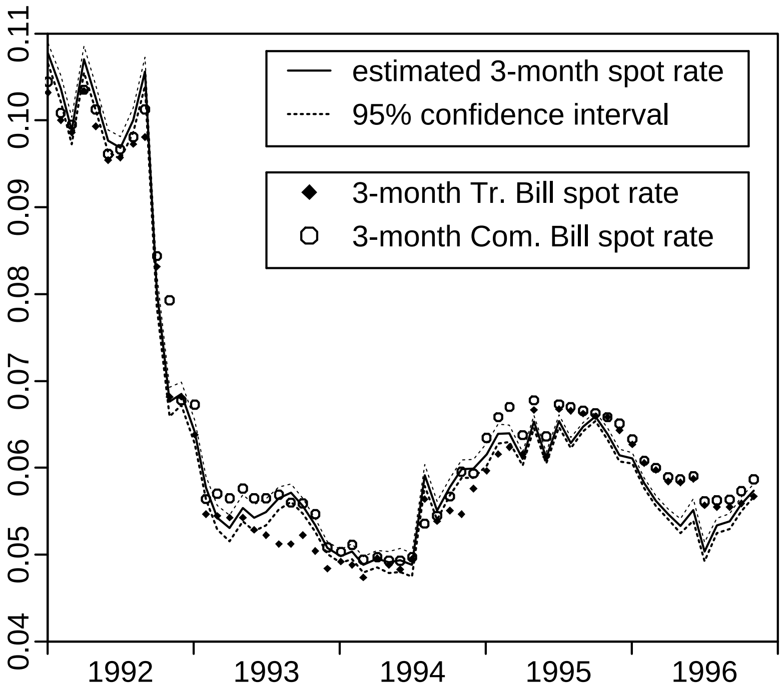

Figure 7: The fitted 3-month spot rate compared with market spot rates for Treasury Bills and for Commercial Bills (converted from discount rates).

HMSO and apply to the last Friday in the preceeding month. Thus these rates are sampled up to 4 trading days before the gilts prices were taken (on the last trading day of each month). (In *Financial Statistics* the CB's have now taken the name 'eligible bills').

4.2 *Further remarks on the analysis*

4.2.1 On further investigation of Figure 6 we note:

1. The long-term forward rate jumps up by about 2% in October 1992 reflecting the UK's exit from the ERM, but then drifts down as the market decided that the government was serious about keeping inflation low in the long term.

2. The standard errors plotted in Figure 6 are approximately equal to (but slightly less than) the standard errors in gross redemption yields (expressed in basis points). (That is, if Δy_i is the error in the redemption yield for stock i then the standard error plotted is S where $S^2 = \frac{1}{n}\sum_{i=1}^{n} w_i \Delta y_i^2$ and the w_i $(0 < w_i < 1)$ are weights which are mostly very close to 1 except for short-dated stocks. The precise formula is given in Section 3.)

   It can be seen that the standard error is relatively high initially. This reflects the coupon effect which was in existence before 1995. After 1995 the market had factored out the proposed changes to the taxation of fixed interest stocks which will came into effect in 1996 and under which interest and capital gains are taxed on the same basis. (Allowance would be made for taxation under the old regime of immediate payments.)

   A gradual reduction in standard errors in the run up to 1995 reflected a narrowing of the differences between high and low-coupon stocks. It is possible that the market was anticipating the tax changes announced in 1995.

3. Although TB's and CB's both follow the fitted curve (Figure 7) very well there are significant differences: TB's are too low through some of 1993 and both TB's and CB's are too high through much of 1995 and early 1996 (often the confidence interval and the bid/offer spread are far apart).

   The 3-month CB rate for the end of October 1992 at about 7.9% is substantially higher than the estimated rate (6.9%) and the TB rate (6.8%). Possibly this was a result of the use of different sampling dates.

   The use of these (or some other) published rates to constrain the short end of the forward-rate curve was considered. This was not found to improve results substantially and the approach was rejected.

4. During the first quarter of 1994 duration 10 forward and spot rates rose by about 2% from a low of 6%. This was caused by the United States Federal Reserve which raised its rates by 0.25%. This sparked off a world-wide bout of selling of bonds which lasted for several months.

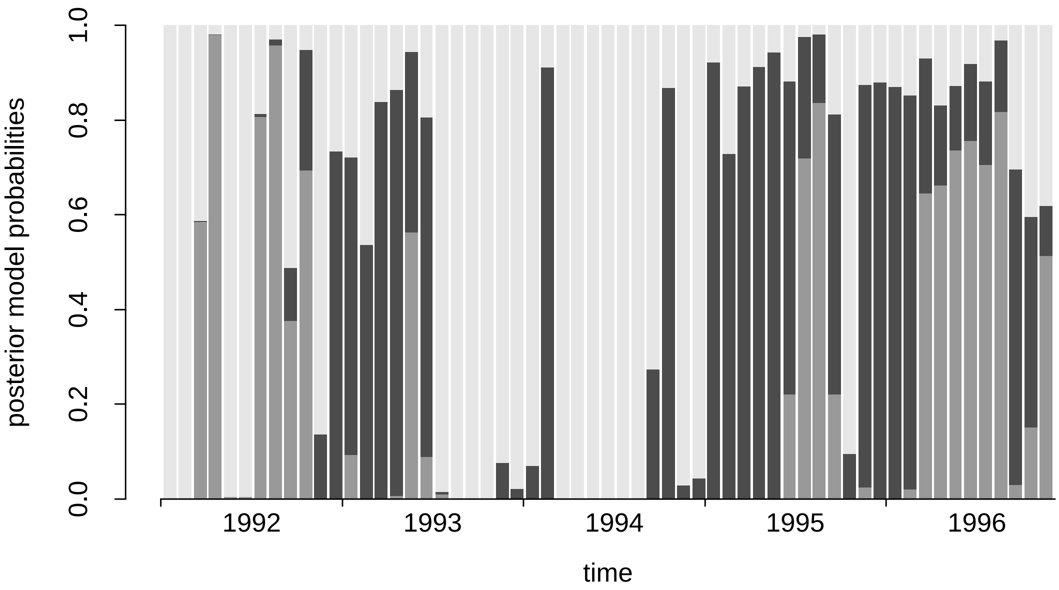

Figure 8: Posterior model probabilities. On each date the bar is divided up into three sections adding up to a total probability of 1: lower section (dark grey) posterior probability for $k = 3$, $PP(3)$; middle section (black) $PP(4)$; top section (light grey) $PP(5)$.

5 STATISTICAL ISSUES

5.1 *Model selection*

5.1.1 It is necessary to justify the use of a 5-parameter model rather than 3 or 4. There are a number of ways in which to choose the best model. Chaplin (1998) uses hypothesis tests with the null hypothesis, for example, that $b_4 = 0$ to choose between 4 and 5 parameters. Here we describe the use of model selection criteria.

5.1.2 The two most popular model selection criteria are the Akaike Information Criterion (AIC) and the Schwartz Bayes Criterion (SBC or BIC). These can be written in the form (for example, see Wei, 1990, or Cairns, 1995)

$$\begin{aligned} AIC(k) &= g_1(P|\hat{b}) - k \\ SBC(k) &= g_1(P|\hat{b}) - \frac{1}{2}k\log(N) \end{aligned}$$

where $\hat{b}$ is the maximum-likelihood (or Bayesian) estimate, k is the number of parameters being estimated and N is the number of stocks in the dataset.

5.1.3 The effect of the second term in each criterion is to penalise over-parametrised models. Put another way: under the AIC it is necessary for the maximum-likelihood to increase by at least 1 for it to be worth adding in one extra parameter. Except for small values of N,

$\frac{1}{2}\log(N) > 1$. Therefore, the SBC tends to select models with fewer parameters than does the AIC. It is known, in fact, that the AIC tends to select an overparametrised model. In particular, where the true model is contained within the class of models being considered, as more and more data is collected the AIC will converge on a more highly parametrised model which contains the true one. The SBC, on-the-other-hand, will converge to the true model. (For example, see Wei, 1990.) It should further be noted that neither the AIC not the SBC was designed for sequential model selection. This point is discussed by Chaplin (1998).

5.1.4 A useful adjustment of the SBC which aids comparison of models is to let $PP(k) = \exp[SBC(k)]/\sum_j \exp[SBC(j)]$. The function $PP(k)$ gives (to a good approximation) the posterior model probabilities.

5.1.5 In Figure 8 we plot the PP for the 3-parameter ($b_3 = b_4 = 0$), 4-parameter ($b_4 = 0$) and 5-parameter models. It can be seen from this figure that all of three models would have been selected on a range of dates.

5.1.6 From a consideration of the most relevant period in 1995 and 1996 the 3 and 4-parameter models dominate. It is tempting, therefore, to use the 4-parameter model for use in the construction of the yield indices. However, 5 parameters have proved to be necessary once in 1995 but most significantly during 1993 and 1994. It may be argued that there was, to a certain extent, an unusual shape during that period with unusual market forces acting. This, however, is not an argument against the use of 5 parameters: it is quite likely that, in the future, other different but still unusual circumstances might arise which require the use of 5 parameters.

5.1.7 In summary the 5-parameter model was selected as the most appropriate model for use in the construction of the yield indices for the following reasons:

1. It has given the best fit on a number of dates in the past.
2. The yield indices need a model which will fit well almost all future Gilts markets without the need to vary the number of parameters. It is likely that, if 5 parameters have been needed in the past, 5 parameters will be needed in the future. Adopting a strategy which varies the number of parameters periodically would result in catastrophic jumps on those days: a problem which the Working Group was set up to eliminate.

5.2 *Choice of the fixed parameters c*

5.2.1 In Section 4 there was some mention of the possibility that different values for c might be used. In Figure 5 it was demonstrated, first, that $c = (0.2, 0.4, 0.8, 1.6)$ and $c = (0.1, 0.2, 0.4, 0.8)$ both fit the data equally well on all dates but, second, that the latter gave much more volatile long-term forward rates. In general the choice of c does not affect substantially the goodness of fit. For example, Chaplin (1998) uses $c = (0.23, 0.46, 0.69, 0.92)$ and achieves an equally good set of fits, although he excludes stocks with a term of less than one year to maturity. Some sets of values, however, such as $c = (0.1, 0.2, 0.4, 0.8)$ (Figure 5) give rise to undesirable characteristics such as volatility in the long-term forward rates whereas others with much larger values of c might give rise to excessive volatility in the very short-term rate. In both cases the volatility occurs outside the range of the observed data. From a subjective point of view such rates should be less volatile than the rates predicted by, for example, $c = (0.1, 0.2, 0.4, 0.8)$.

The choice of $c = (0.2, 0.4, 0.8, 1.6)$ satisfies the desire for relative stability. It is not claimed, however, that this set of values is optimal in any sense.

5.2.2 Bjørk & Christensen (1997) discuss constraints on the forward-rate curve for consistency within the Ho & Lee (1986) and Hull & White (1990) models for the pricing of interest-rate derivatives. They demonstrate, first, that the c must remain fixed and second that there are further constraints on the relationship between the b_i given an initial forward-rate curve. For example, in the Vasicek (1977) model (which is a special case of the Hull & White, 1990, model) forward rates will have the form $f(t, t+s) = b_0(t) + b_1(t)\exp(-c_1 s) + b_2(t)\exp(-c_2 s)$. This curve initially has four degrees of freedom (the only constraint being that $c_2 = 2c_1$, but, thereafter there is only one degree of freedom (which can be taken as the risk-free rate of interest, $r(t)$). b_0, c_1 and c_2 must remain fixed while $b_1(t)$ and $b_2(t)$ depend solely upon $r(t)$. Bjørk & Christensen (1997) generalise this to cover more complex models.

5.3 *The coupon effect*

5.3.1 A relevant question at this stage is whether the coupon effect has genuinely disappeared. The discussion so far has commented that the coupon effect was very obvious in 1992 but had disappeared by the time of the announcement in 1995 of the tax changes.

5.3.2 It is an important feature of the model described in Section 3 that income and capital are valued in the same way: that is, no coupon effect is allowed for. It is therefore important to demonstrate that, since, 1995, no significant coupon effect exists.

5.3.3 Two forms of investigation were carried out. First, on each date the sample correlation, R^y, between coupon and yield errors was considered. For each stock i let g_i be the coupon, y_i be the gross redemption yield and $\hat{y}_i$ be the predicted yield. The sample correlation coefficient is

$$\begin{aligned} R^y &= \frac{\sum_{i=1}^{N}(y_i - \hat{Y}_i)(g_i - \bar{g})}{\sqrt{\sum_{i=1}^{N}(y_i - \hat{Y}_i)^2 \sum_{j=1}^{N}(g_j - \bar{g})^2}} \\ \text{where } \bar{g} &= \frac{1}{N}\sum_{i=1}^{N} g_i \end{aligned}$$

5.3.4 (An alternative approach calculates a sample correlation coefficient for the three subsets of short, medium and long-dated stocks. This would be more likely to pick up an effect if it varied significantly according to term. On-the-other-hand, it reduces the power of the test described below.)

5.3.5 The sample correlation coefficient was calculated for each of the 59 dates in the 1992 to 1996 data, giving values $R_1, R_2, \ldots, R_{59}$. The following null hypothesis was tested:

H_0: There is no correlation between the g_i and the $y_i - \hat{y}_i$.

versus

H_1: The g_i and the $y_i - \hat{y}_i$ are positively correlated.

5.3.6 The null hypothesis is rejected on an individual date (that is, without reference to other dates) if $R_t > R_{min}$, where R_{min} depends on the significance level, α, and on the number of stocks, N. Over the period N ranged from 28 to 35. Sample values for R_{min} are given in Table 1 (these values can, for example, be found in Neave, 1978, Table 6.2).

R_{min}	$\alpha = 4\%$	$\alpha = 10\%$
$N = 28$	0.437	0.374
$N = 36$	0.386	0.329

Table 1: Critical values for sample correlation coefficients.

5.3.7 Values of R_t^y are plotted in Figure 9(a), along with $R_{min} = 0.4$. (0.4 is representative of the values in Table 1. For a more rigorous test a more precise value would be used.) Also plotted in Figure 9(c) are the sample correlation coefficients between coupon and price errors (which correspondingly should be zero or negative). In Figures 9(b) and 9(d) the fitted slopes (using a simple linear regression to link the yield errors with coupon and the price errors with coupon respectively) are plotted.

5.3.8 It is clear from Figure 9 that there is still some sort of coupon effect. The coupon effect was at its lowest during 1995 but it has increased again in 1996. However, when Figures 9(b) and 9(d) are considered it can be seen that even though there is still possibly a small coupon effect: about one tenth of what it was in 1992.

5.3.9 It can be seen that the levels of correlation with coupon have fallen very substantially (Figures 9(a) and (c)). If individual dates are considered then there is insufficient evidence to reject the hypothesis that there is no coupon effect. However, since the correlation coefficients are continually positive in the case of yield errors and negative in the case of price errors over the whole period from 1995 onwards, one might conjecture that there is still a coupon effect of sorts, albeit much reduced. (If there was genuinely no coupon effect then we would observe a mixture of positive and negative correlation coefficients.)

5.3.10 The reduction in the size of the coupon effect can be seen by looking at the reductions in the regression slopes plotted in Figures 9(b) and (d). From these graphs it can be conjectured that, if a coupon effect still exists, then it is about 10% of what it was in, say 1992.

5.3.11 Figure 9 includes the 3% and 6% stocks which were excluded from the estimation process (for completeness). If these three stocks are excluded from the correlation analysis then Figure 9 would be largely unchanged. However, if we exclude those stocks with a coupon of 15% or more (which appeared to be somewhat erratic) as well as the 3% and the 6% stocks (that is, all stocks with a coupon of 6% or less) then, perhaps strangely, the levels of correlation become more marked.

5.3.12 This again suggests that the evidence for the removal of the coupon effect is inconclusive. It is easy to look at scatter plots and to identify and remove outliers (for example, stocks above 15%, or some of the strippable stocks) and to improve the situation in favour of one hypothesis or another. This is the worst form of what is called ”data mining” whereby the data and model are continually modified to make things look better than they perhaps are.

5.3.13 There is a notable discontinuity during 1996. This coincides with the introduction of 8% 2021. This particular stock (which replaced 8.75% 2017 as the longest dated fixed interest gilt on the market) seems to trade at a substantial premium and may have had a distorting effect.

5.3.14 No attempt has been made to test residuals for dependence jointly on term and

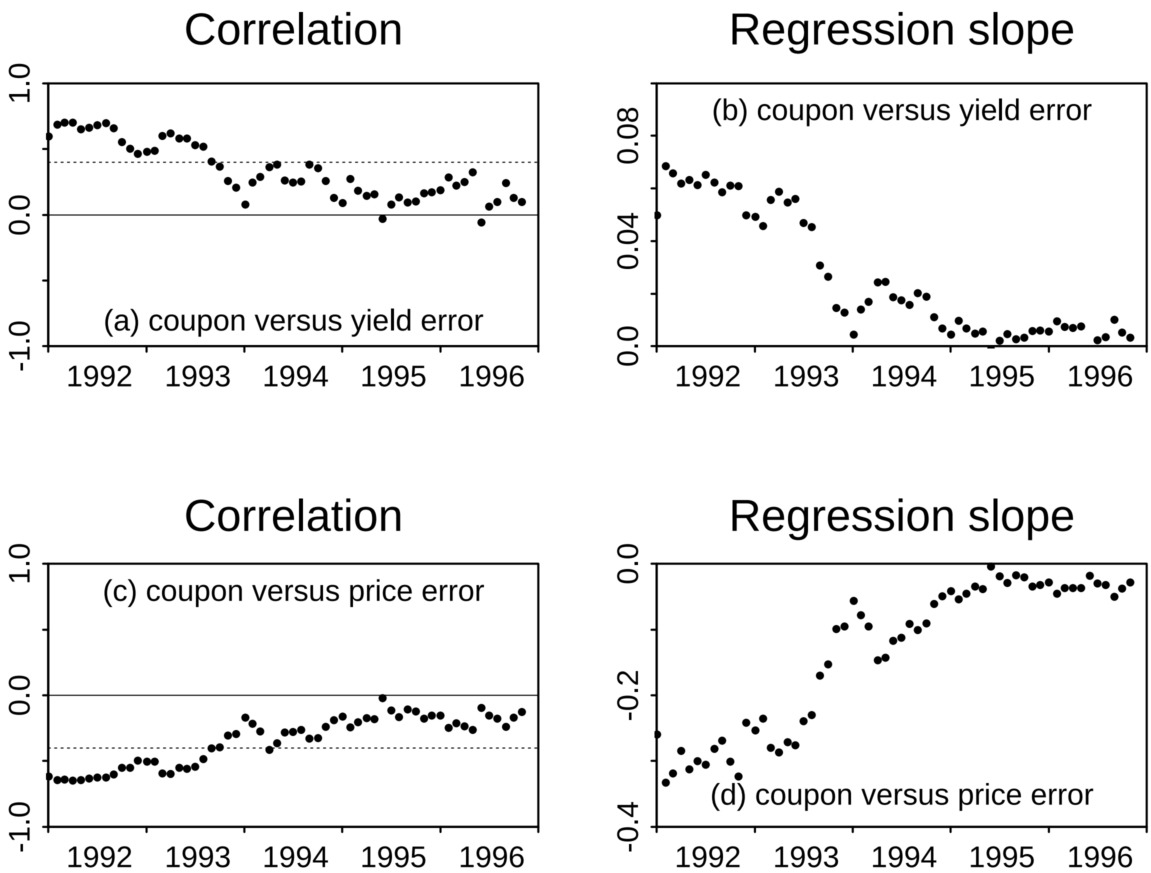

Figure 9: Sample correlation coefficients between: (a) coupon and yield errors; (c) coupon and price errors. Regression slope linking: (b) coupon with yield errors; (d) coupon with price errors. (That is, a modified model might be of the form $y_i = \hat{y}_i + \alpha_y(g_i - \overline{g}) + \varepsilon_i$ where α_y is the slope and ε_i is the error. For example, $\alpha_y = 0.04$ means an increase of 4 basis points in the yield for every 1% increase in the coupon.)

coupon.

5.3.15 In a second, non-rigorous test, the sequence of 59 errors in the price of each stock were plotted. Eight examples are given in Figure 10. We note:

1. 9% 2008 was a 'run-of-the-mill' stock with no large or systematic errors over the whole period.

2. 12% 1998 (a high-coupon stock) was always cheap (most obviously during 1993). Since 1995 it has consistently traded 10 to 20 pence below its theoretical price. This might be evidence of the continued existence of a coupon effect. However, neither of 7.25% 1998 and 15.5% 1998 have exhibited a possible coupon effect since 1995. Similarly 7% 2001 appeared to be expensive while 6.75% 2004 was cheap over the period 1995 and early 1996 but expensive later in 1996.

3. Benchmark and strippable status does seem to command a premium. 8.75% 2017 was relatively expensive during 1994 and early 1995 and when it ceased to hold benchmark status on the issue in 1996 of 8% 2021 there was a marked drop in price. All of the strippable bonds, except for 8.5% 2005, appear to trade at a premium (for example, 8% 2015 is plotted in Figure 10).

5.4 *Independence of errors*

5.4.1 The Runs Test was applied to the results to test for independence of the residuals on a specific date. The stocks were ordered according to term to maturity ($i = 1$ representing the shortest-dated stock, and $i = n_1 + n_2$ representing the longest-dated stock). In this test we considered the residuals $z_i = y_i - \hat{Y}_i$. Let n_1 be the number of positive and n_2 be the number of negative residuals. In the construction of the likelihood function it was assumed that the price errors (that is, the difference between the theoretical and observed prices) would be independent. The purpose of the Runs Test is to assess whether or not stocks with similar maturity dates have similar price errors. If the estimated curve provides a poor fit then we may, for example, observe that all stocks which mature in, say 8 to 12 years have yields which are below their estimated values, and the number of runs would be quite low.

5.4.2 As an example, suppose $z = (1.1, 0.6, -0.3, -0.7, -0.2, 0.7, -0.1, 0.6)$. We first translate this into +'s and -'s: that is, $(++---+-+)$. A sequence of +'s (minimum length 1) then counts as one run. So then the number of runs here is 5.

5.4.3 If the residuals are independent with $Pr(z_i > 0) = n_1/(n_1 + n_2)$ (that is, the null hypothesis) then we can derive the distribution of the number of runs. If the observed value is *consistently* too low then the model may be considered to have provided a poor fit. If the number of runs calculated each day occasionally looks too low then this is not a problem, since random variation will dictate that this will happen occasionally.

5.4.4 The Runs test was applied to each of the 59 dates included in the 1992 to 1996 data and the results are plotted in Figure 11. The plot gives percentage points under the null hypothesis for the observed number of runs using the Normal approximation to the number of runs. The following observations are made:

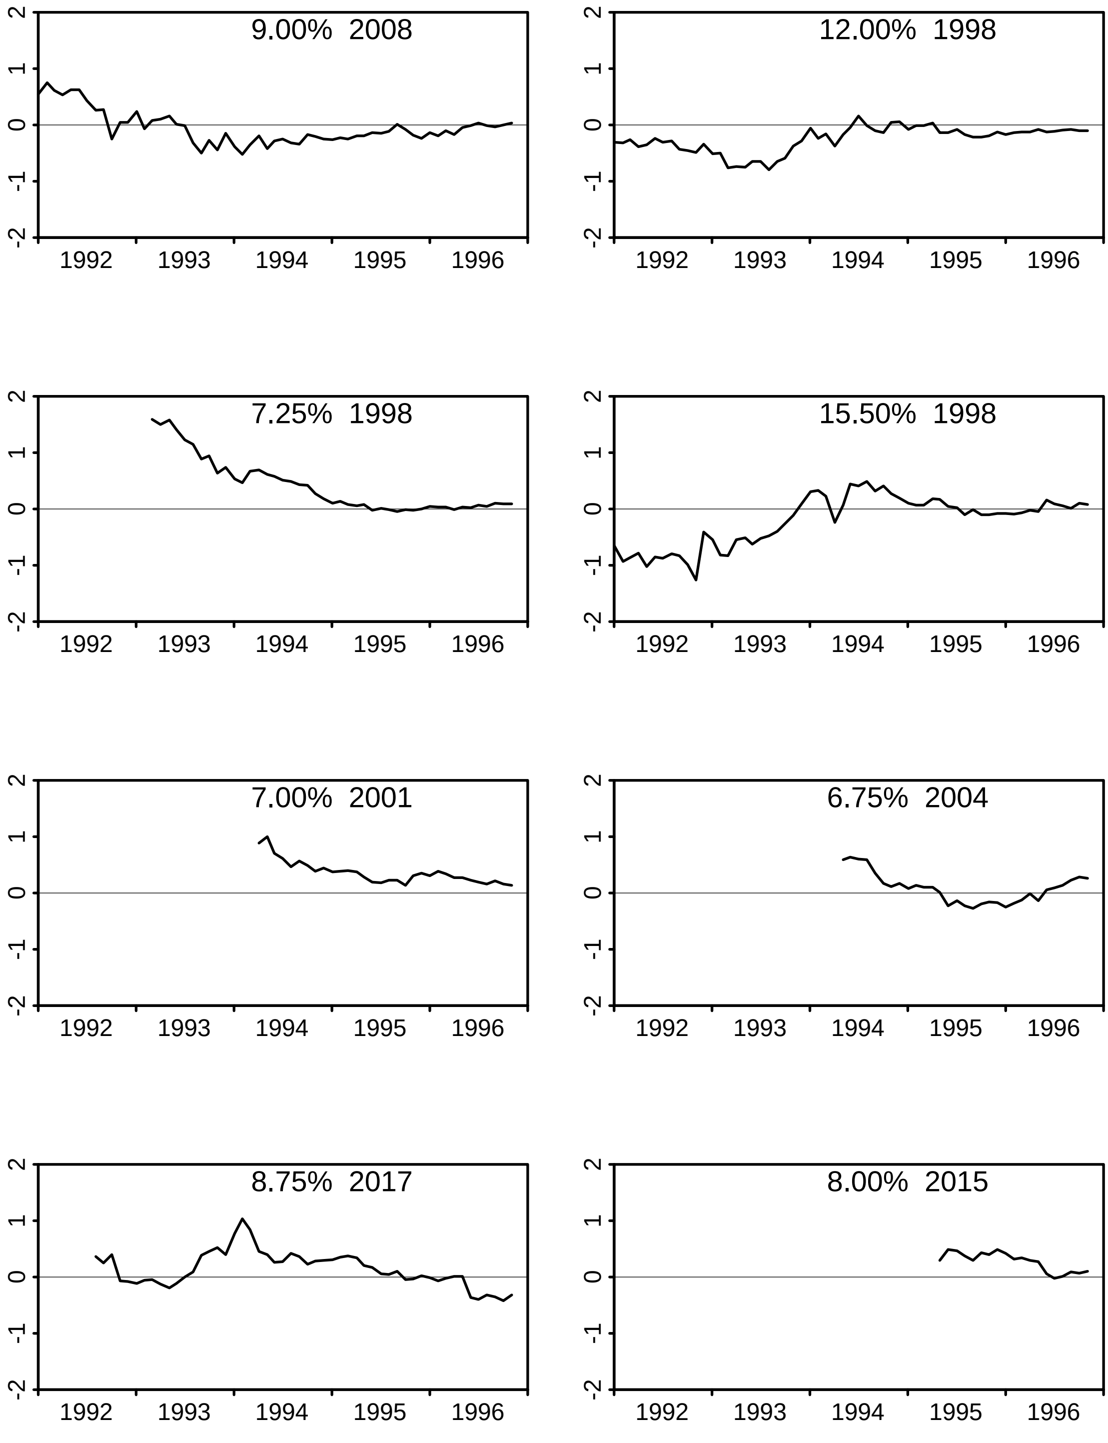

Figure 10: Evolution of price errors for specific stocks.

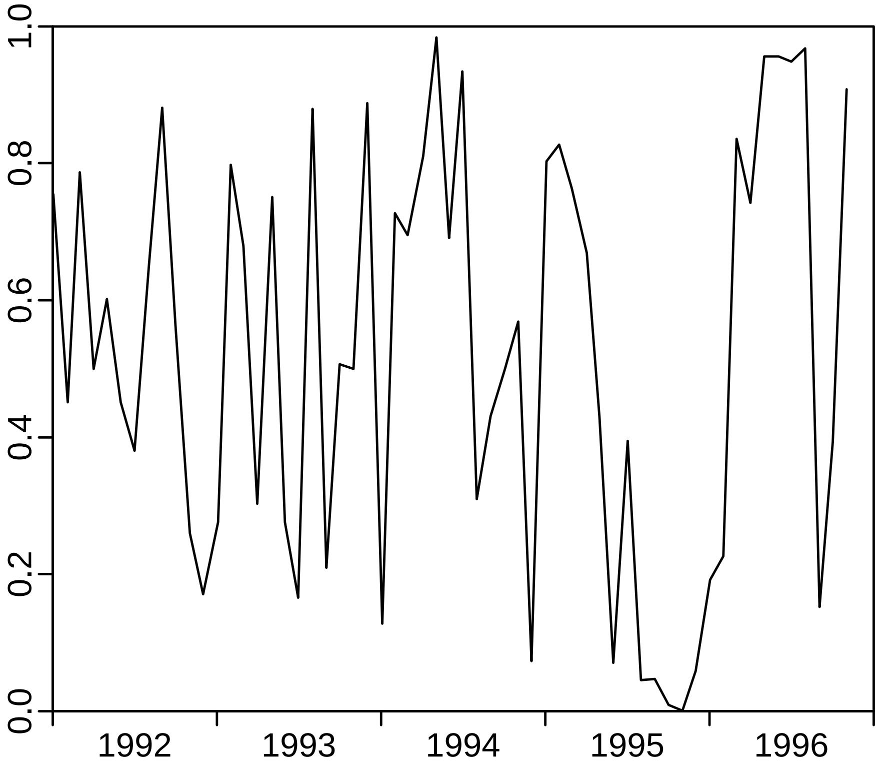

Figure 11: 1992 to 1996 data. Approximate p-values (calculated separately for each day) for the numbers of runs based upon the null-hypothesis that residuals z_i are independently distributed with $Pr(z_i > 0) = n_1/(n_1 + n_2)$.

1. The figures for 1992 to 1994 are somewhat misleading as, potentially, maturity correlations would be swamped by the size of the coupon effect (and coupons are, to a certain extent, randomly distributed through the maturity range).

2. There is a run of low p-values during 1995. This might be regarded as being evidence of non-independence. However, the number of runs will not be independent from one date to the next. A much longer run of data will be needed to establish that residuals on similar maturity dates are correlated.

3. Generally, however, the numbers of runs were not so low as to give serious grounds for concern. The outcome is reasonably acceptable without being too good.

4. It should, of course, be borne in mind that the test here is not entirely reliable. First, as mentioned above the pre-1995 coupon effect might distort the picture. Second, some stocks have benchmark status, for example, and trade at a premium. If these stocks are clustered together then there may again be a distorting effect.

5.5 *Confidence intervals*

5.5.1 Besides having a best estimate of the forward-rate curve it is useful also to know how reliable estimates of forward rates, spot rates, yields and prices are at various maturity dates.It is desirable, therefore, to be able to construct confidence intervals for these quantities.

5.5.2 We are given a set of price data for one date, x. Let $\hat{b}$ be the value of $b = (b_0, b_1, \ldots, b_4)^T$ which maximises the log-posterior distribution (thatis, the best estimate). The log-posterior density function can be written in the form

$$
\begin{aligned}
p(b|x) &= p(\hat{b}|x) - \frac{1}{2} db^T M db + o(db^2) \\
\Rightarrow db &\mathrel{\dot\sim} N(0, M^{-1}) \\
\Rightarrow db &\approx KZ \\
\text{where } KK^T &= M^{-1} \\
\text{and } Z &\sim N(0, I_n), \text{ and } n = 5 \text{ parameters}
\end{aligned}
$$

That is, the posterior distribution for b is approximately Normal with mean $\hat{b}$ and covariance matrix M^{-1}.

5.5.3 K and M are positive definite $n \times n$ matrices. K is not unique in the sense that there are infinitely many K such that $KK^T = M^{-1}$. The particular choice of K and similar matrices defined below (in particular A) does not affect the outcome of this analysis.

5.5.4 Now consider a quantity of interest y (for example, a 10-year instantaneous forward rate, a five-year spot rate or the price of a 17-year 8% coupon bond). Then

$$
y = y(\hat{b}) + h^T db + o(db)
$$

$$\text{where } db = b - \hat{b}$$
$$h = \frac{dy}{db}$$

5.5.5 Provided db is small, y is, therefore, approximately normally distributed with mean $y(\hat{b})$ and variance $h^T KK^T h = h^T M^{-1} h$.

5.5.6 Examples:

1. Forward rates: $y = f(t, t+s)$ and $h = (1, e^{-c_1 s}, e^{-c_2 s}, e^{-c_3 s}, e^{-c_4 s})^T$.

2. Spot rates:

$$y = R(t, t+s) = \frac{1}{s} \int_0^s f(t, t+u) du$$
$$\text{and } h = h_R(t, t+s) = \left(1, \frac{1-e^{-c_1 s}}{c_1 s}, \frac{1-e^{-c_2 s}}{c_2 s}, \frac{1-e^{-c_3 s}}{c_3 s}, \frac{1-e^{-c_4 s}}{c_4 s}\right)^T$$

3. Bond prices (maturity at time $t+s$ and coupon g per annum payable continuously):

$$y = P(t, t+s; g) = g \int_0^s e^{-uR(t,t+u)} du + e^{-sR(t,t+s)}$$
$$\text{and } h = -g \int_0^s u h_R(t, t+u) e^{-uR(t,t+u)} du - s h_R(t, t+s) e^{-sR(t,t+s)}$$

4. Bond yields corresponding to the bonds described above: the standard deviation of the estimated yield is equal to the standard deviation of the price per unit nominal divided by the duration of the bond.

5.5.7 Typical standard errors for these three quantities are plotted in Figure 12. The following comments can be made:

1. A 95% confidence interval for a given quantity is (using the Normal approximation) the best estimate plus or minus 2 times the standard error.

2. The errors clearly depend on maturity, s.

3. The errors have reduced substantially as the coupon effect disappeared.

4. The dips and humps in the curves correspond, roughly, to the clusters or otherwise of stocks with the same duration (forward and spot rates) or term (prices).

5. Standard errors get much larger outside the range of the stocks being use for estimation purposes. All quantities have finite limiting values as the term to maturity increases to infinity. This reflects the fact that the parameter b_0 has a finite variance.

6. Standard errors for spot prices are significantly lower than those for forward rates.

7. Spot rates and gross redemption yields give very similar results except at the longer end errors for gross redemption yields are lower.

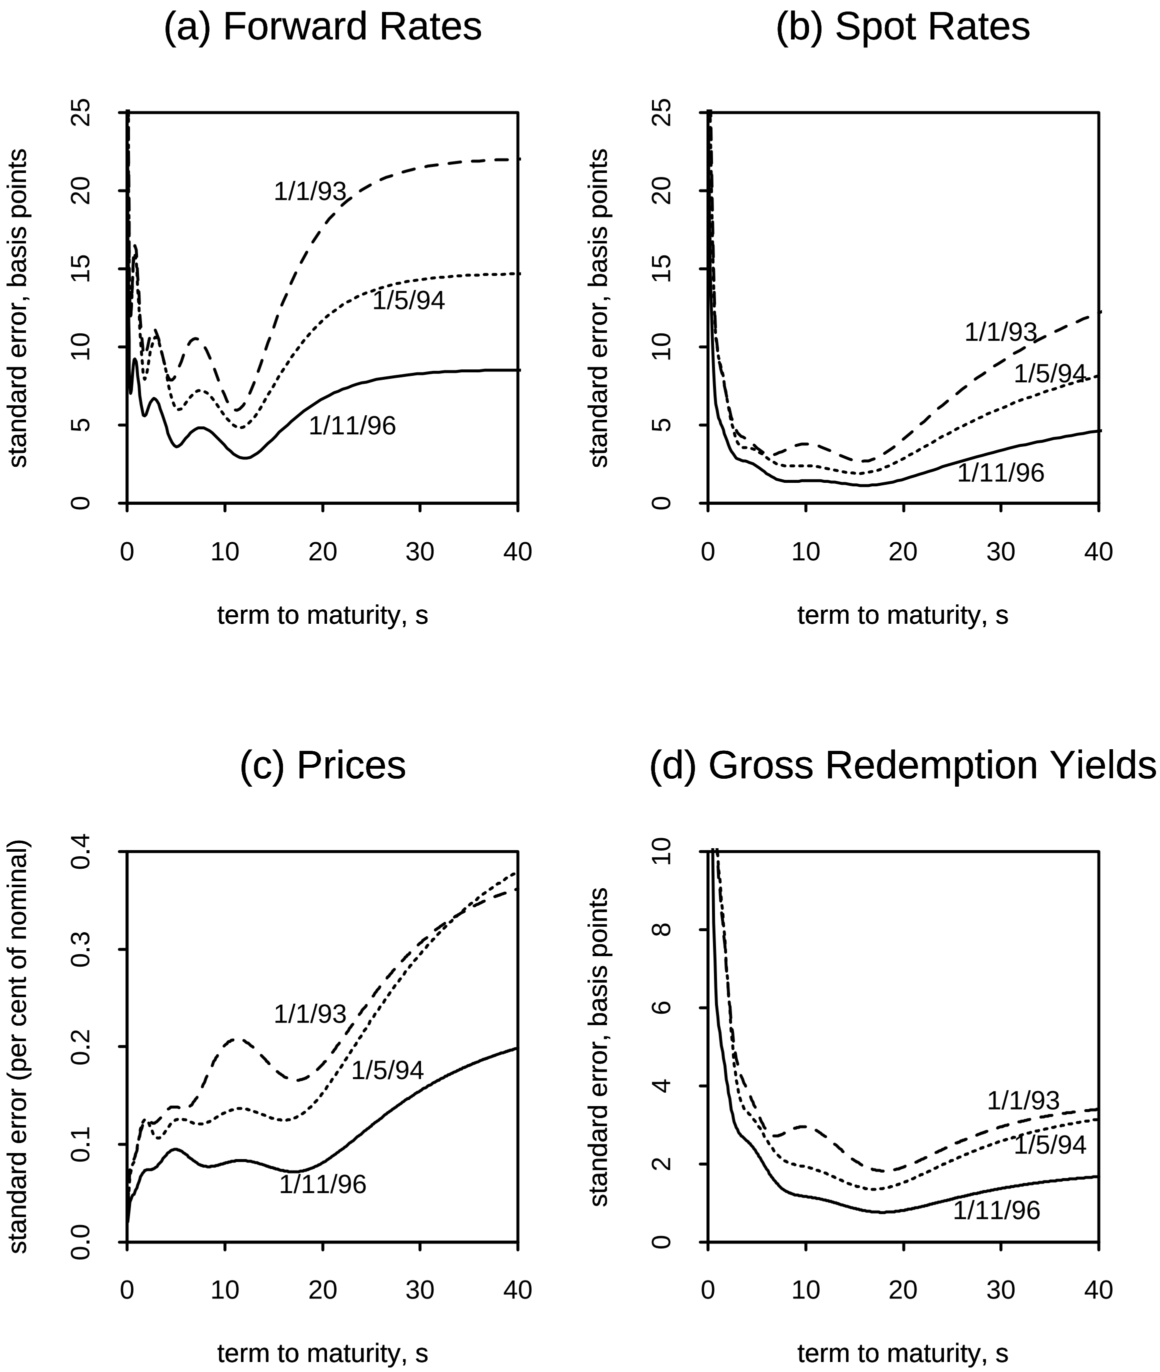

Figure 12: Standard errors for (a) forward rates, (b) spot rates, (c) prices of 8% coupon bonds at different maturities and (d) gross redemption yields. Errors are plotted for the three dates 1 November 1996 (solid line), 1 May 1994 (dotted) and 1 January 1993 (dashed).

5.6 *Correlations between quantities*

5.6.1 Besides the calculation of standard errors for specified quantities we can also calculate covariances and correlation coefficients. Let y_1 and y_2 be two quantities of interest (for example, spot rates at, say, 5 years and 10 years). Then as before, we have errors

$$\begin{aligned} dy_i &= h_i^T KZ \quad i = 1,2 \\ \Rightarrow Cov(dy_1, dy_2) &= E[(h_1^T KZ)(Z^T K^T h_2)] \\ &= h_1^T KK^T h_2 \\ &= h_1^T M^{-1} h_2 \\ \Rightarrow cor(dy_1, dy_2) &= \frac{h_1^T M^{-1} h_2}{\sqrt{(h_1^T m^{-1} h_1)(h_2^T M^{-1} h_2)}} \end{aligned}$$

5.6.2 These correlations should not be confused with changes in rates through time (as discussed in Feldman *et al.*, 1998, Appendix C).

5.6.3 Correlations between the errors in forward rates at different durations are plotted in Figure 13. It can be seen that there is a very complex structure.

5.6.4 To give an example, compare the rates at durations 3 and 10 years. The correlation is about -0.75. Therefore if the true 10-year rate is above its estimate then it is most likely that the 3-year rate will be below its estimate. Such negative correlations are not unexpected since an increase in one rate should be compensated by a decrease in another rate in order to maintain the fit of the model. The existence of negative correlations perhaps goes some way also to explaining why standard errors for spot rates (Figure 12) are lower than those for forward rates: spot rates are a weighted average of forward rates. Perhaps more surprising is the correlation of about 0.5 between rates at durations 2 and 25 years.

6 CATASTROPHIC JUMPS

6.1.1 The original motivation for this paper arose because of the existence of multiple minima in the Dobbie & Wilkie (1978) objective function which leads to catastrophic jumps from time to time. It was also shown in Section 3.5 that the same problem arises in in the Svensson (1994) model.

6.1.2 Here we discuss whether or not the same problem exists in the model described in Section 3.

6.2 *The zero-coupon bond market*

6.2.1 It can be seen that for a zero-coupon bond $\hat{P}_i(b)$ is linear in b (unlike coupon bonds). In Appendix D this linearity is utilised to prove that both the log-likelihood (which is quadratic in b) and log-posterior density functions have unique maxima.

6.2.2 Therefore if we use only zero-coupon bond prices there will be no catastrophic jumps. In the case of the Dobbie & Wilkie (1978), Nelson & Siegel (1987) and Svensson (1994) models

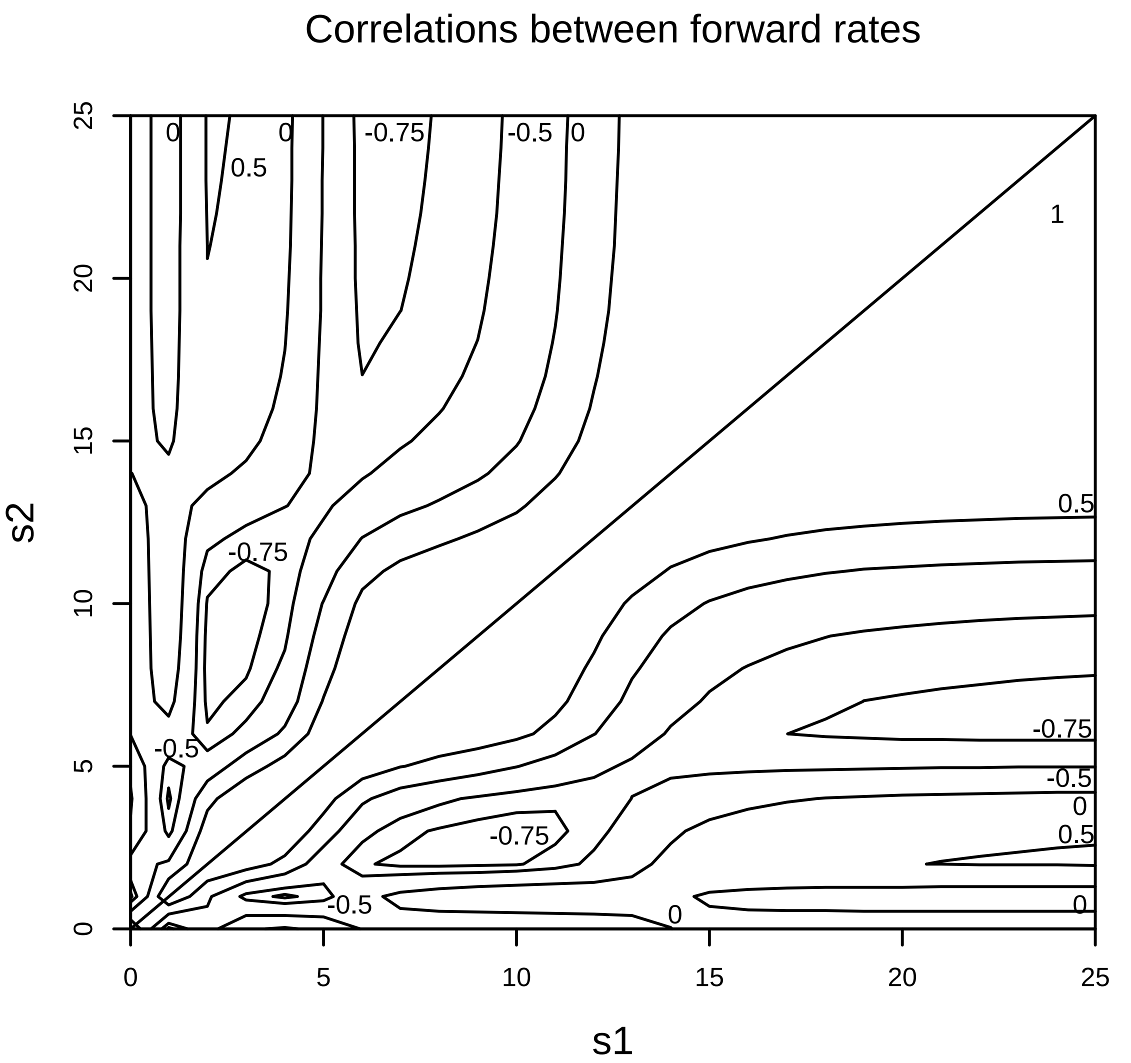

Figure 13: Contour plot of the correlation function between forward rates at different maturities s_1 and s_2.

the relevant functions to be optimised are linear in some parameters but non-linear in others. These non-linearities frequently cause multiple optima.

6.3 *Coupon bond markets*

6.3.1 In other circumstances it is necessary to use coupon bonds to estimate the parameters of the forward-rate curve.

6.3.2 Since $\hat{P}_i(b)$ is non-linear in b it is difficult to say with certainty that there is no possibility that more than one maximum exists. However, it is strongly suspected that multiple maxima do not exist in practice. There are two arguments for this view.

6.3.3 First, a number of specific dates were considered on an individual basis. On each date 100 starting points for the optimisation were chosen at random. In each case the algorithm converged to the same maxmimum. This, of course, does not provide conclusive proof of its uniqueness. However, in the case, for example, of the Svensson (1994) model it is easy to find multiple maxima. We can therefore say that if they do exist in the new model then they are much less frequent.

6.3.4 Second, one can argue that multiple maxima are less likely to exist on the following grounds. It was shown in Section 6.1 that there is a unique maximum if we have a zero-coupon bond market. As we move (in a continuous sense) from this ideal position to one which includes small coupons, then, at first, continuity tells us that within a finite range of values for the b_i the likelihood surface will have more-or-less the same shape as before. In particular, it will still only have one maximum. As we increase the significance of the coupons further this may cease to be true. But we can say that if all coupons are small enough (but non-zero) then there will be a unique maximum. It is not clear at what level of coupons (if any) would multiple maxima turn up and if this level would cover reasonable or only unrealistically high coupons.

7 A MODEL FOR INDEX-LINKED GILTS

7.1 *Description*

7.1.1 In this section we describe the development of a model for the forward-inflation curve implied by the index-linked gilts market. This curve, $r(t,t'+s)$, is parametrised in the same way as the forward-rate curve for fixed-interest stocks described in Section 3: that is,

$$\begin{aligned} r(t,t'+s) &= r_0(t) + \sum_{i=1}^{4} r_i(t)\exp(-c_i s) \\ &= r_0 + \sum_{i=1}^{4} r_i \exp(-c_i s) \end{aligned}$$

where t' is the time of the last known RPI. Thus $r(t,t'+s)$ is, in effect, measured from time t', as opposed to $f(t,t+s$ which is measured from today. A precise description of the pricing model is given in Appendix E. Paragraph E.2.2.5 discusses the reasons for not introducing a seasonal element into the forward-inflation curve.

7.1.2 The model allows us to project the future cashflows by starting from the most recent value of the RPI and projecting the index forward using the implied forward-inflation curve. Allowance is made for the 8-month time lag, so that the first coupon payment is usually known with certainty. The model can be fitted in the same way as the forward-rate curve was originally and without any prior distribution. This includes the use of a similar error structure to that for fixed-interest prices with zero weight given to stocks which mature in under 8 months, and low weight given to other short-dated stocks.

7.1.3 It was found initiallythat the short end of the forward inflation curve was very unreliable: volatile and not at all like recent observed price inflation (see Figure 14, dotted line). Taken at face value, this implies that the market view of the current instantaneous force of inflation was very volatile. It did not seem plausible that this could vary by up to 10% in one year when the annual rate of inflation was relatively stable.

7.1.4 This problem was addressed by restricting the short end of the forward inflation curve in the following way. The current annual rate of inflation is a known quantity and it seems sensible to make some sort of use of this when the curve is being fitted. We use the log of the change in the index over the last 12 months:

$$Q = \log \frac{RPI(\tilde{m})}{RPI(\tilde{m}-12)}$$

where $\tilde{m}$ is the month of the last known RPI. We use a period of 12 months for two reasons: first, it avoids the complications of seasonality; second, it is much more stable than the annualised month-on-month rate. This is then to be compared with

$$\hat{Q}(r) = \int_{-1}^{0} r(t, t' + s)ds$$

which projects the implied inflation curve backwards.

7.1.5 The quantity Q is put to use by specifying a prior distribution for $\hat{Q}(r)$: that is, a Normal distribution with mean Q and standard deviation $\sigma_Q = 0.005$ (50 basis points). If the standard deviation was much smaller then the quality of the fit to the price data deteriorated. If the standard deviation was much higher then the predicted short-term inflation would often be very far from recent, observed figures and also rather volatile. (Even with the use of Q it is felt that short-term implied inflation should be treated with caution.)

7.2 *Results for the 1992 to 1996 data*

7.2.1 In Figure 14 we plot results covering the period January 1992 to November 1996 allowing for standard deviations in $\hat{Q}(r)$ of 0.005 (solid line) and 0.1 (dotted line). In the top-right graph we compare Q (dots) with $\hat{Q}(r)$. In the second-row-left graph we compare Q (dots) with predicted inflation over the next year. In the bottom-right graph we plot the root-mean-squared-error (RMSE) for each set of results. In this case the RMSE is expressed relative to the error structure used in the estimation process. Thus we have (see Appendix E):

$$\log P_i \sim N\Big(\log \hat{P}_i, \nu^2(\hat{P}_i, d_i')\Big)$$
$$RMSE = \frac{1}{N}\sum_{i=1}^{N}\frac{(\log P_i - \log \hat{P}_i)^2}{\nu^2(\hat{P}_i, d_i')}.$$

(d_i' is the inflation duration of stock i.)

7.2.2 Thus *RMSE* should be of the order of 1 if the suggested error structure is reasonable.

7.2.3 There are a number of comments to be made on the results illustrated in Figure 14.

1. By comparing the solid and dotted lines we can see how volatile the implied inflation curve can be if it is not restrained at the short end.
2. A weaker short-term RPI constraint seems to have a less dramatic effect when one looks at implied price inflation over the next year (Figure 14, second row, left). Perhaps more curious is the fact that the one-year implied inflation figures do not appear to be particularly good predictors of inflation. (Note, for example, that the one-year inflation predicted on 1 January 1992 should be compared with observed inflation on 1 January 1993.) However, here there are only, in effect, 5 independent predictions and observations making it difficult to draw any firm conclusions.
3. Residuals are generally very low, but these shot up in the aftermath of the UK's exit from the ERM in September 1992. This reduced quickly at first but then took about 2 more years to settle back down to the previous level.
4. In times of relative calm the errors in real yields are of the order of 1 to 2 basis points which is rather lower than the lowest errors in the fixed-interest stocks.
5. The UK's exit from the ERM resulted in a jump of about 2% in long-term inflation expectations.
6. The peak RMSE in November 1995 was caused by a poor fit for a short-dated stock which was about to change to fixed interest status.

7.2.4 Overall the model seems to work well for long term inflation expectations: it is relatively stable and reacted in a sensible way to the UK's exit from the ERM. The model also fits price data very well.

7.2.5 In Figure 15 we plot results covering the same period with $\sigma_Q = 0.005$ but this time investigating the effect of excluding stocks with a duration of less the 8 months (solid line, and the same as in Figure 14) or excluding stocks with a duration of less than 3 years (dotted line, with weights tending to zero at duration 3 years).

7.2.6 It can be seen in Figure 15 that the RMSE is much more stable if stocks with a duration of less than 3 years are excluded from the analysis. On-the-other-hand short-term rates, which interpolate between observed inflation and the prices of stocks with durations over three years only, are more volatile (Figure 15, top right). Neither situation is entirely satisfactory

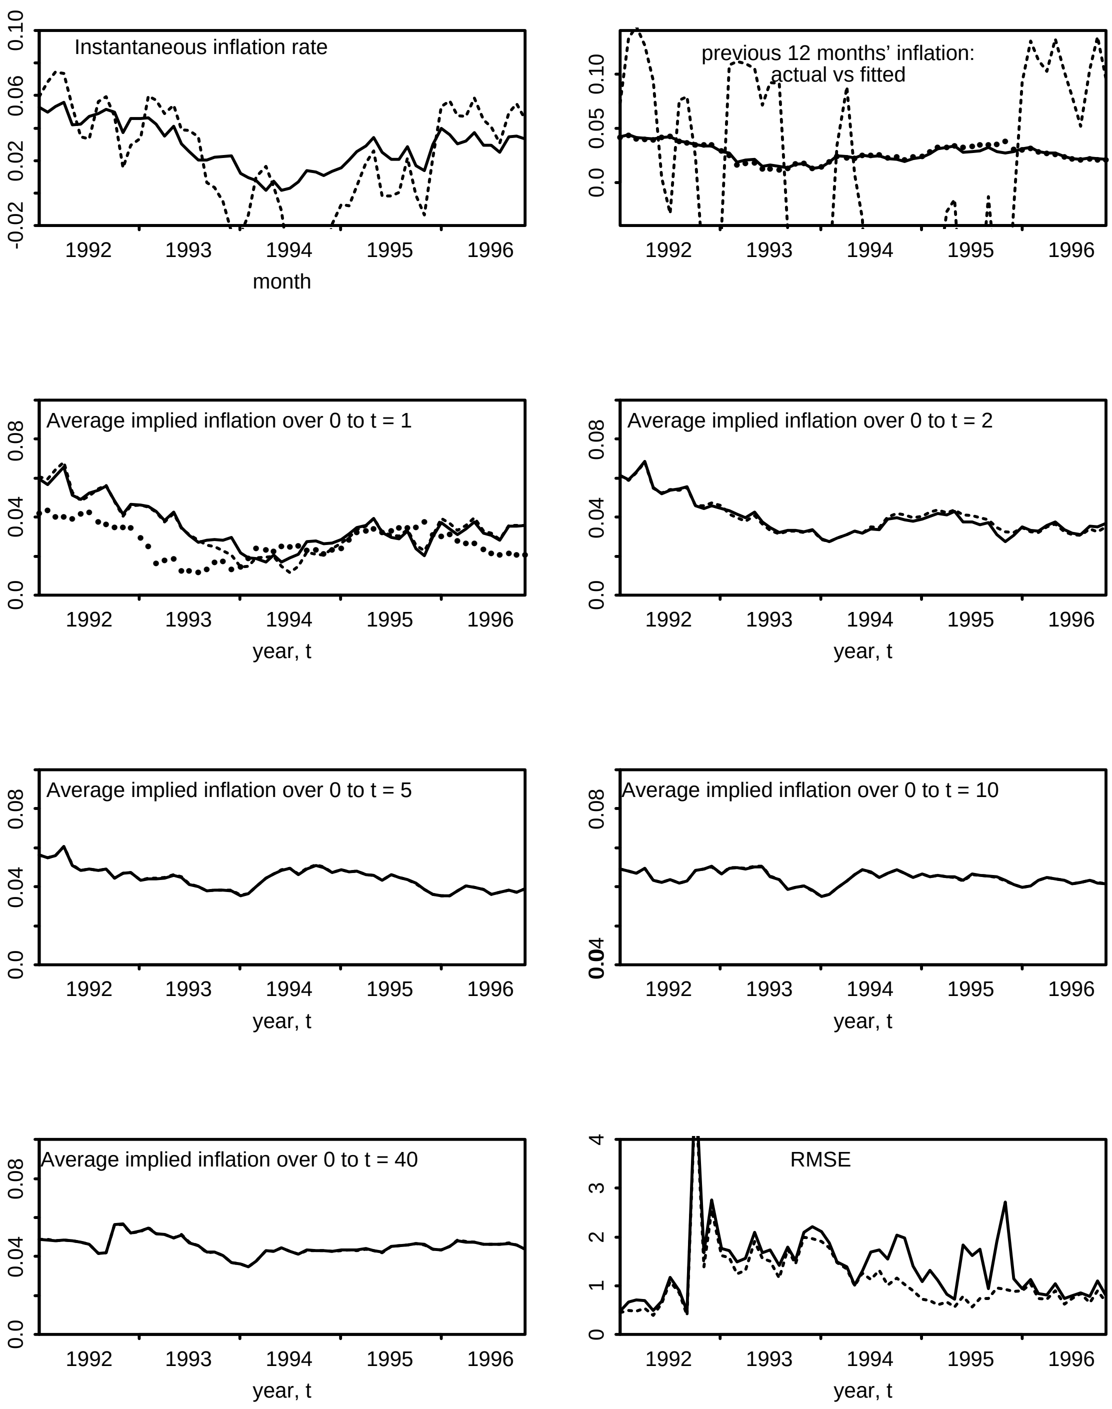

Figure 14: Results of fitting the forward-inflation curve. Solid lines: $\sigma_Q = 0.005$. Dotted lines: $\sigma_Q = 0.1$. Points: price inflation over the preceeding 12 months, Q.

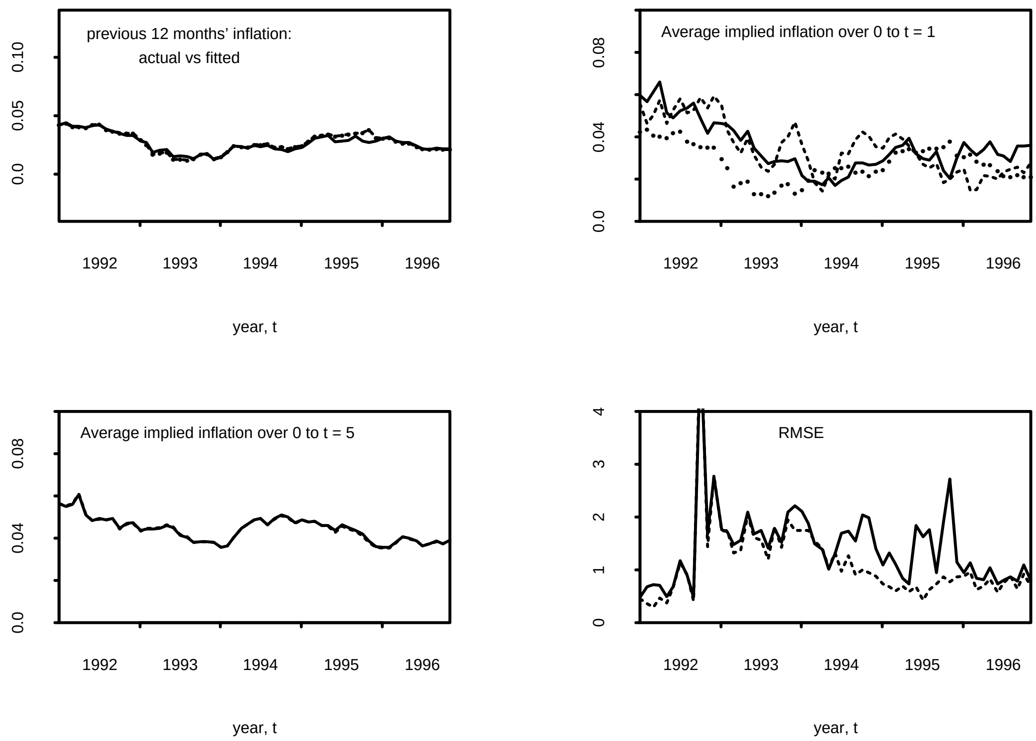

Figure 15: Results of fitting the forward-inflation curve with $\sigma_Q = 0.005$. Solid lines: stocks under 8 months excluded. Dotted lines: stocks under 3 years excluded. Points: price inflation over the preceeding 12 months, Q.

and future work could focus on a revision of the error structure in order that some use can be made of short-dated stocks to stabilise predicted short-term inflation but with a reduced risk that the RMSE will be disrupted by these short-dated stocks. In this way we should be able to make use of the limited information contained in these stocks.

7.2.7 In Figure 16, as an example, we plot fitted curves for 1 January 1994 ($\sigma_Q = 0.005$). In the top graph we plot nominal forward-rate curve (dotted line) estimated from fixed-interest price data, and the forward-inflation curve estimated using the nominal forward-rate curve and index-linked price data. The fitted real-forward-rate curve is plotted in the middle graph. (Note that the various time-lags involved mean that it is not possible to estimate this curve accurately from index-linked data alone.) In the bottom graph we plot implied spot inflation rates using the fitted curve (dotted line) and a simple 'bootstrapping' technique. (The bootstrapping technique assumes that the implied forward-inflation curve is piecewise constant with jumps at the times of maturity of each of the index-linked stocks. The number of degrees of freedom is equal to the number of stocks and the forward-inflation values are chosen to ensure that the theoretical prices precisely match the observed prices.) We can see that the bootstrap line is relatively stable and is closely approximated by the fitted, parametric curve. This is reflected in the low RMSE on that date.

7.2.8 In contrast to 1 January 1994, 1 October 1992 (just after Sterling's exit from the ERM) exhibits much more instability in the bootstrapped forward-inflation curve (particularly for the shorter-dated index-linked stocks). This resulted in a high RMSE on that date.

7.2.9 In Figure 17 we plot the development of real forward rates at various durations. It is clear that the real forward rate is relatively volatile at shorter durations. The long-term rate is very stable with only one dip towards the end of 1993. (Just as significantly nothing unusual happened in September 1992.) This can be compared with the same period plotted in Figure 6. It would appear that there was excess demand during this period for fixed interest stocks (worldwide) which did not carry over to index-linked stocks. As has been commented upon earlier (paragraph 4.2.1) this inflation of fixed interest prices was reversed during the first half of 1994 following the 0.25% interest rate rise in the United States. We see in Figure 17 that this corresponds to the return of the 40-year real forward rate to its more normal level of around 4%.

8 SUMMARY

8.1.10 This paper has reviewed the construction of the FTSE Actuaries Government Securities Yield Indices. The known risk of catastrophic jumps has been highlighted and has been shown to apply also to the Svensson (1994) forward-rate model.

8.1.11 A number of alternative descriptive models for the term structure were considered before the restricted exponential class

$$f(t,t+s) = b_0(t) + b_1(t)e^{-c_1 s} + b_2(t)e^{-c_2 s} + b_3(t)e^{-c_3 s} + b_4(t)e^{-c_4 s}$$

was settled upon. The restriction that the exponential parameters remain fixed is central to the analysis. First, this reduces (and possibly eliminates) the risk that there will be catastrophic jumps. Second, the restriction is known to be a requirement for consistency in the Hull & White

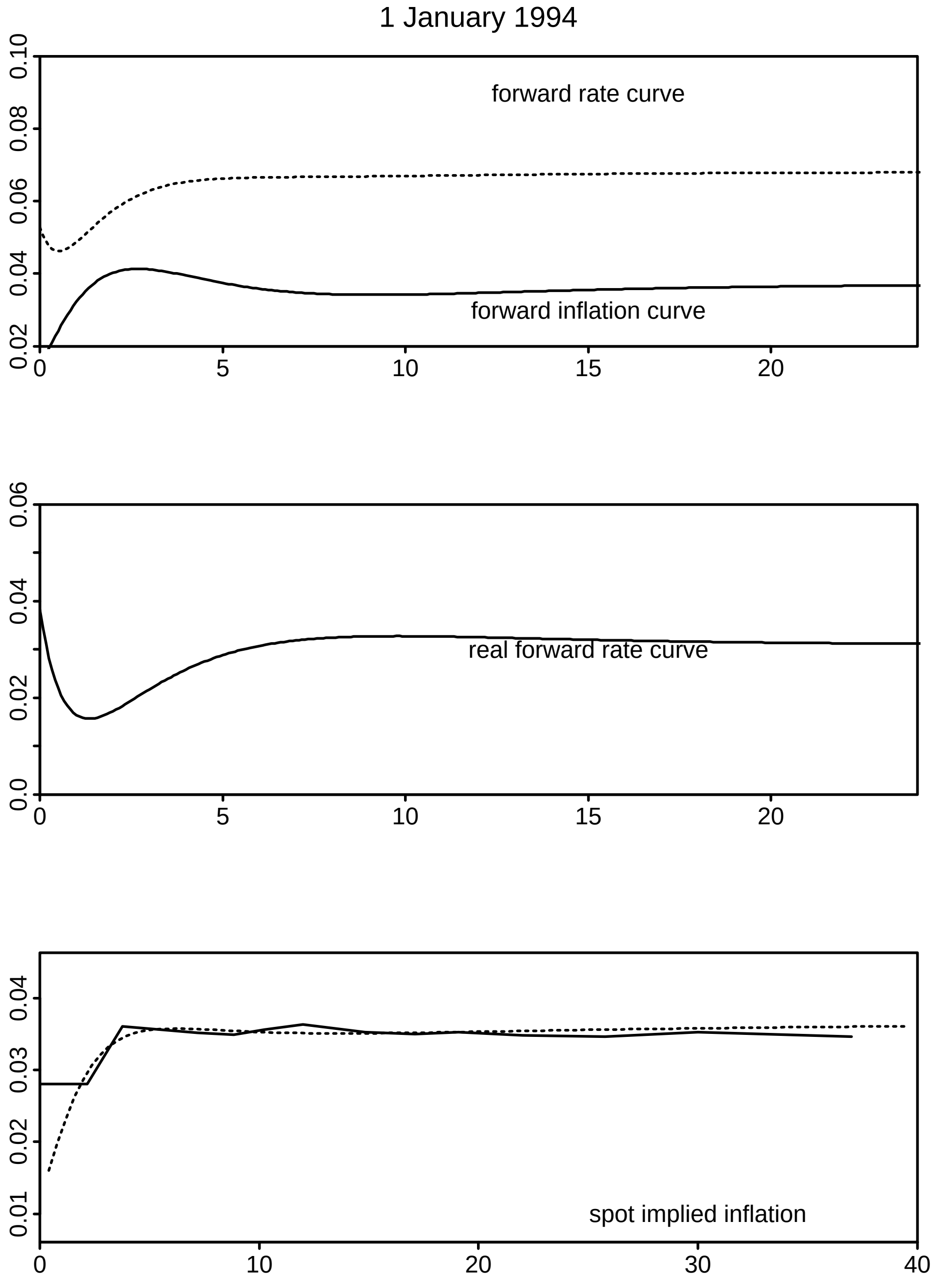

Figure 16: 1 January 1994. Top: forward rate curve for fixed interest stocks and forward inflation curve. Middle: real forward rate curve. Bottom: bootstrapped spot inflation (solid line) versus fitted spot inflation (dotted line).

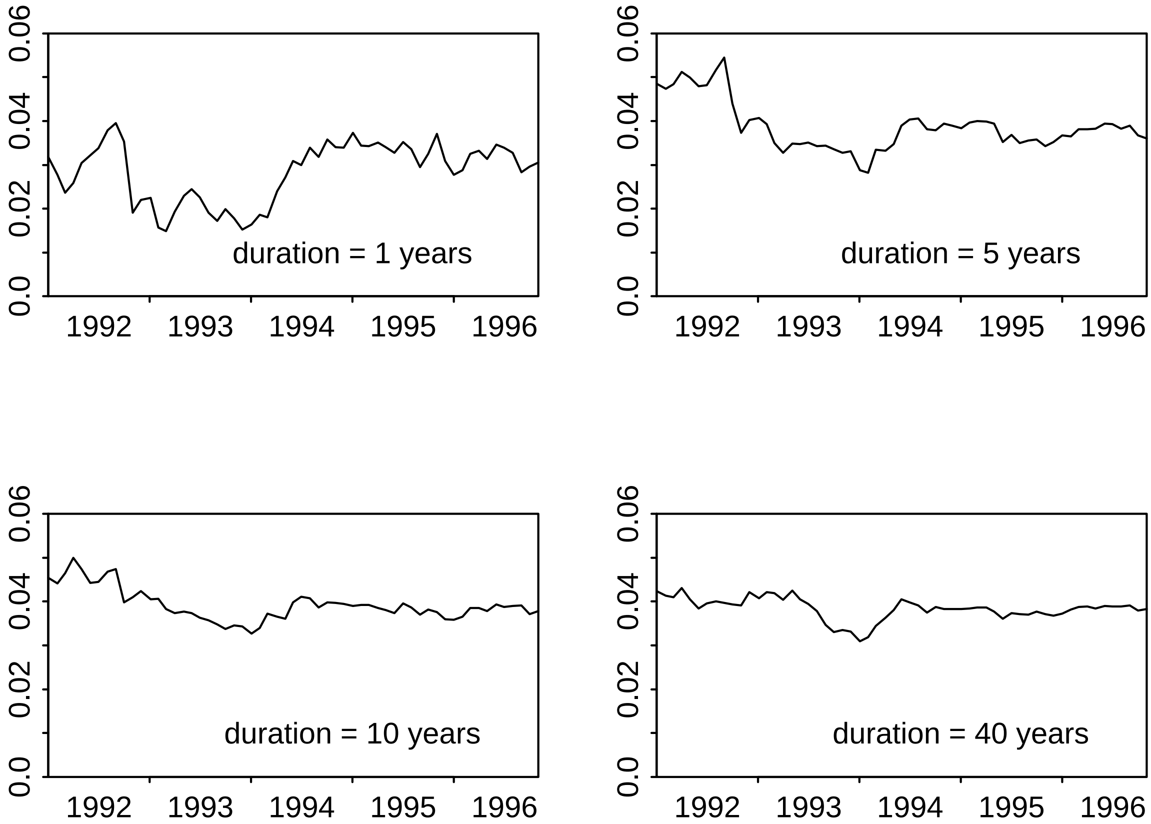

Figure 17: Real forward rates over the period 1992 to 1996 with $\sigma_Q = 0.005$.

(1990) and Ho & Lee (1986) evolutionary or no-arbitrage models for the term structure (Bjørk & Christensen, 1997).

8.1.12 It is felt that, for use in the construction of the yield indices, the 5-parameter model described above be used. It has been shown that 5 parameters have been required on a number of dates to ensure a good fit. However, the recommendation is principally made on the basis that the yield indices should not be susceptible to catastrophic jumps. Allowing the number of parameters to change from time to time would needlessly introduce catastrophic jumps at each point in time when the number of parameters changes.

8.1.13 A similar model for the forward-inflation curve implied by index-linked prices has also been proposed. This is found to be most effective (in a subjective sense) when recent, observed RPI data is added to price data. Even with this additional data, however, it was felt that less reliance could be placed on short-term implied inflation (which is subject to the whims of private, tax-paying investors) than on medium and long-term implied inflation. In particular, medium and long-term real rates of interest were found to be robust and stable.

ACKNOWLEDGEMENTS

The author would like to thank the various members of the Fixed Interest Indices Working Group. In particular, George Gwilt first suggested the use of the exponential class of forward-rate curves as an alternative to yield curves, while Geoff Chaplin, Keith Feldman and David Wilkie provided useful discussion of the model specification and the results. Other stimulating discussions took place with Andrew Smith, Julian Wiseman, Tomas Bjørk and Gerry Kennedy and with a scrutineer of the paper.

Most of the data used in the paper was supplied by Patrick Phillips for which the author is very grateful. Without this data this work would not have been possible.

Much of this research was carried out while the author was in receipt of a joint research grant from the Faculty of Actuaries and the Institute of Actuaries.

REFERENCES

ANDERSON, N., BREEDON, F., DEACON, M., DERRY, A. AND MURPHY, G. (1996) *Estimating and interpreting the yield curve.* Wiley, Chichester.

BJØRK, T. AND CHRISTENSEN, B. J. (1997) Forward rate models and invariant manifolds. Preprint, Stockholm School of Economics.

CAIRNS, A.J.G. (1995) Uncertainty in the modelling process. *Transactions of the 25th International Congress of Actuaries, Brussels* **1**, 67-94. (Correct version at `http://www.ma.hw.ac.uk/~andrewc/papers/ajgc3.ps`).

CHAPLIN, G.B. (1998) A review of term structure models and their applications. To appear in BAJ.

CLARKSON, R.S. (1979) A mathematical model for the gilt-edged market. *JIA* **106**, 85-132.

COX, J.C., INGERSOLL, J.E., AND ROSS, S.A. (1985) A theory of the term-structure of interest rates. *Econometrica* **53**, 385-407.

DEACON, M., AND DERRY, A. (1994) Estimating the term structure of interest rates. Bank of England Working Paper Series Number 24.

DOBBIE, G.M. AND WILKIE, A.D. (1978) The FT-Actuaries Fixed Interest Indices. *JIA* **105**, 15-27.

and *TFA* **36**:203-213.

FELDMAN, K.S. (1977) The gilt-edged market reformulated. *JIA* **104**, 227-240.

FELDMAN, K.S., BERGMAN, B., CAIRNS, A.J.G., CHAPLIN, G.B., GWILT, G.D., LOCKYER, P.R., AND TURLEY, F.B. (1998) Report of the fixed interest working group. To appear in BAJ.

GWILT, G.D. (1982) Presidential address. *TFA* **38**, 1-17.

HEATH, D., JARROW, R. AND MORTON, A. (1992) Bond pricing and the term structure of interest rates: a new methodology for contingent claims valuation. *Econometrica* **60**, 77-105.

HO, S.Y., AND LEE S-B. (1986) Term structure movements and pricing interest rate contingent claims. *Journal of Finance* **41**, 1011-1029.

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

MASTRONIKOLA, K. (1991) Yield curves for gilt-edged stocks: a new model.. Discussion paper number 49. Bank of England..

MCCULLOCH, J.H. (1971) Measuring the term structure of interest rates. *Journal of Business* **44**, 19-31.

MCCULLOCH, J.H. (1975) The tax-adjusted yield curve. *Journal of Finance* **30**, 811-830.

McCutcheon, J.J. and Scott, W.F. (1986) *An Introduction to the Mathematics of Finance.* Heinemann, Oxford.

Neave, H.R. (1978) *Statistics Tables.* George Allen and Unwin, London.

Nelson, C.R., and Siegel, A.F. (1987) Parsimonious modeling of yield curves. *Journal of Business* **60**, 473-489.

Poston, T. and Stewart, I.M. (1978) *Catastrophe theory and its applications.* Pitman, London.

Svensson, L.E.O. (1994) Estimating and interpreting forward interest rates: Sweden 1992-1994. *Working paper of the International Monetary Fund* **94.114**, 33 pages.

Vasicek, O. (1977) An equilibrium characterisation of the term structure. *Journal of Financial Economics* **5**, 177-188.

Wei, W.W.S. (1990) *Time Series Analysis.* Addison-Wesley, Redwood City.

Wiseman, J. (1994) The exponential yield curve model. J.P.Morgan, 16 December.

9 APPENDICES

APPENDIX A: STRIPPABLE BONDS

A.1 There are seven strippable gilts:

Name	Redemption Date
8% 2000	7/12/2000
7% 2002	7/6/2002
6.5% 2003	7/12/2003
8.5% 2005	7/12/2005
7.5% 2006	7/12/2006
7.25% 2007	7/12/2007
8% 2015	7/12/2015
8% 2021	7/6/2021

Each of the bonds has the same two coupon payment dates: 7 June and 7 December.

A.2 Each coupon payment will be strippable and sold as a zero-coupon bond. Similarly the capital payment on redemption can be sold as a zero-coupon bond.

APPENDIX B: DESCRIPTION OF THE FIXED INTEREST ERROR STRUCTURE

B.1 The first source of error (rounding errors in prices) gives us

$$\log P_i \sim N(\log \hat{P}_i, \sigma_0^2)$$

B.2 The second type of error described in the text focusses on errors in yields. Each stock has a theoretical yield $\hat{y}_i$ and an observed yield y_i. The differences between the two are assumed to be independent and identically distributed: that is,

$$\begin{aligned} y_i &\sim N(\hat{y}_i, \sigma^2) \\ \Rightarrow \log P_i &\sim N\left(\log \hat{P}_i\,,\ \sigma^2 d_i^2\right) \qquad \text{(approximately)} \\ \text{where } d_i &= \text{duration of stock } i \end{aligned}$$

B.3 Third, it was felt that errors in the prices of long-dated stocks should not be allowed to become too great. Thus we have, for high duration stocks:

$$\log P_i \sim N(\log \hat{P}_i, \sigma_\infty^2)$$

B.4 Blending these three sources of error together we get the following model for a given parameter set b:

$$
\begin{aligned}
\log P_i &\sim N\left(\log \hat{P}_i\ ,\ \nu^2(\hat{P}_i, d_i)\right) \\
\text{where } d_i &= \text{MacAulay duration of stock } i \\
&= \frac{\partial}{\partial y_i} \log P_i(y_i) \\
&= \frac{\partial P_i(y_i)/\partial y_i}{P_i(y_i)} \\
\text{and } y_i &= \text{the (continuously-compounding) gross} \\
&\qquad\qquad \text{redemption yield on stock } i \\
\nu^2(p,d) &= \frac{\sigma_0^2(p)\left[\sigma_\infty^2 d^2 b(p) + 1\right]}{\sigma_0^2(p) d^2 b(p) + 1} \\
b(p) &= \frac{\sigma_d^2}{\sigma_0^2(p)[\sigma_\infty^2 - \sigma_0^2(p)]}
\end{aligned}
$$

B.5 Some justification for the choice of formula for $\nu^2(p,d)$ is required! Note that

$$
\begin{aligned}
\nu^2(p,0) &= \sigma_0^2(p) \\
\nu^2(p,\infty) &= \sigma_\infty^2 \\
\frac{\partial(\nu^2)}{\partial(d^2)} &> 0 \\
\left.\frac{\partial(\nu^2)}{\partial(d^2)}\right|_{d=0} &= \sigma_d^2.
\end{aligned}
$$

B.6 Now

$$
\begin{aligned}
P_k &= \hat{P}_k \exp\left(\nu(\hat{P}_k, d_k) Z\right) \\
\text{where } Z &\sim N(0,1) \\
\Rightarrow P_k &\approx \hat{P}_k\left(1 + \nu(\hat{P}_k, d_k) Z\right) \\
\Rightarrow S.D.(P_k) &\approx \hat{P}_k \nu(\hat{P}_k, d_k)
\end{aligned}
$$

B.7 If $d_k = 0$ we use $S.D.(P_k) = 1/32$ to reflect rounding errors in prices and the bid/offer spread. Hence

$$\sigma_0^2(\hat{P}_k) = (32\hat{P}_k)^{-2} \approx (32P_k)^{-2}$$

B.8 One effect of having $\sigma_0^2 > 0$ is that the effective weight attached to short-dated stocks tends to 0 as they approach maturity.

B.9 For large d_k and fixed interest stocks we use $S.D.(P_k) = 0.001P_k$: that is, $\sigma_\infty^2 = 0.001^2$. Thus for a long-dated fixed-interest stock with a typical duration of 10 years this formula gives a standard deviation of about 1 basis point in the yield.

B.10 It is appropriate to include this limit (particularly for the index-linked data) since the duration of index-linked stocks can be relatively high. Before this limit was introduced, errors in the prices of long-dated stocks could be relatively large.

B.11 Temporarily, now assume that $\sigma_0^2(p) = 0$ and $\sigma_\infty^2 = \infty$. Then

$$\nu^2(p, d_i) \to \sigma_d^2 d_i^2 \quad \text{as } \sigma_0^2 \to 0 \text{ and } \sigma_\infty^2 \to \infty$$

B.12 This is consistent with real gross redemption yields y_k having mean $\hat{y}_k$ (consistent with $\hat{P}_k$) and variance σ_d^2: that is, the variances of the y_k do not depend upon the duration of the stock.

B.13 When $\sigma_0^2(p) > 0$ and $\sigma_\infty^2 < \infty$ this argument provides the logic behind the value

$$\left.\frac{\partial(\nu^2)}{\partial(d^2)}\right|_{d=0} = \sigma_d^2$$

B.14 In addition, the standard deviation of the errors on consols was scaled up by a factor of 10. This reflects the relative illiquidity of the irredeemable stocks.

B.15 It was felt desirable to keep two of the irredeemable stocks in to provide some stability in fitted log-term interest rates.

B.16 Finally, we can rewrite $\nu^2(p, d)$ in the following way:

$$\begin{aligned} \nu^2(p,d) &= \frac{\sigma_d^2 d^2}{w} \\ \text{where } w &= \frac{\left(\sigma_0^2(p) d^2 b(p) + 1\right) \sigma_d^2 d^2}{\sigma_0^2(p)\left(\sigma_\infty^2 d^2 b(p) + 1\right)} \end{aligned}$$

B.17 It can easily be shown that: $w \to 0$ as $d \to 0$; and as $d \to \infty$, w is asymptotically like $d^2\sigma_d^2/\sigma_\infty^2$ if $\sigma_\infty^2 < \infty$ or $w \to 1$ if $\sigma_\infty^2 = \infty$.

B.18 The same arguments can be applied to the error structure of index-linked bonds with the duration d_k on stock k replaced by the 'inflation' duration $d'_k = d_k - t_k^f$ (where t_k^f is the time to maturity on the day the stock ceases to have any element of indexation left). (See Appendix E.)

Appendix C: Equivalent least squares approach

C.1 It is well known that $\log x = (x-1) - \frac{1}{2}(x-1)^2 + \frac{1}{3}(x-1)^3 + \ldots$ and that if x is close to 1 then $\log x \approx x - 1$. Therefore, since P_i and $\hat{P}_i$ (per unit nominal) are close to 1 (that is, in the range 0.5 to 1.5) the log-likelihood function is approximately equal to

$$\tilde{g}_1(P|b) = -\frac{N}{2}\log(\sigma^2 d_i^2) + \frac{1}{2}\sum_{i=1}^{N}\log w_i - \frac{1}{2\sigma^2}\sum_{i=1}^{N} w_i \frac{(P_i - \hat{P}_i)^2}{d_i^2}.$$

C.2 Thus maximisation of the likelihood function with respect to b is equivalent to minimisation of the weighted least-squares function

$$\begin{aligned} S^2 &= \sum_{i=1}^{N} w_i \frac{(P_i - \hat{P}_i)^2}{d_i^2} \\ \text{where } w_i &= \frac{\sigma^2 d_i^2}{\sigma^2 d_i^2 + 1/3200^2} \\ &= \text{the weight attached to stock } i \\ &\rightarrow \begin{cases} \infty & \text{as } d_i \rightarrow \infty \\ 0 & \text{as } d_i \rightarrow 0 \end{cases} \end{aligned}$$

Appendix D: Uniqueness of maximum

D.1 For a zero-coupon bond with term τ_i to redemption, the theoretical price is

$$\hat{P}_i(b) = \exp\left[-b_0\tau_i - \sum_{k=1}^{4}\frac{b_k}{c_k}\left(1 - e^{-c_k\tau_i}\right)\right]$$

D.2 First consider the likelihood function. This is of the form

$$g_1(P|b) = -\sum_{i=1}^{N}\lambda_i[\log P_i - \log\hat{P}_i(b)]^2 + \text{constant}$$

where the $\lambda_i = w_i/(2\sigma_d^2\tau_i^2) > 0$ are constants depending on the τ_i.

D.3 Hence

$$\begin{aligned} g_1(P|b) &= -\sum_{i=1}^{N}\left[\log P_i + b_0\tau_i + \sum_{k=1}^{4}\frac{b_k}{c_k}\left(1 - e^{-c_k\tau_i}\right)\right]^2 + \text{constant} \\ &= -\sum_{i=1}^{N}\lambda_i\left(\log P_i + \sum_{k=0}^{4} b_k d_{ki}\right)^2 + \text{constant} \\ \text{where } d_{0i} &= \tau_i \quad \text{for all } i \\ d_{ki} &= \left(1 - e^{-c_k\tau_i}\right)/c_k \quad \text{for } k = 1,2,3,4 \end{aligned}$$

D.4 It is tidier to use vector and matrix notation. Thus $d_i = (d_{0i}, d_{1i}, \ldots, d_{4i})^T$ and we have

$$\begin{aligned} g_1(P|b) &= -\sum_{i=1}^{N} \lambda_i (\log P_i + d_i^T b)^2 \\ \Rightarrow \frac{dg_1}{db}(P|b) &= -2\sum_{i=1}^{N} \lambda_i d_i (\log P_i + d_i^T b) \\ \text{and } \frac{d^2 g_1}{db^2}(P|b) &= -2\sum_{i=1}^{N} \lambda_i d_i d_i^T \end{aligned}$$

D.5 Setting $dg_1(P|b)/db = 0$ gives

$$\tilde{b} = -\left(\sum_{i=1}^{N} \lambda_i d_i d_i^T\right)^{-1} \left(\sum_{i=1}^{N} \lambda_i d_i \log P_i\right)$$

D.6 For any vector h, $h^T(d_i d_i^T)h > 0$. Hence the matrix of second derivatives $d^2 g_1(P|b)/db^2$ is negative definite and $\tilde{b}$ is therefore the unique maximum of $g_1(P|b)$ (and $g_1(P|b)$ is a strictly concave function of b).

D.7 Second consider the log-prior density function, $g_2[f(t,t)] + g_2[f(t,t+10)] + g_2[f(t,t+\infty)]$, where $g_2(\delta) = (\alpha-1)\log\delta - \beta\delta + \text{constant}$ and $\delta = 0,\ 10,\ \infty$. Now $f(t,t+s)$ is linear in b so $d^2 f(t,t+s)/db^2 \equiv 0$. Hence

$$\begin{aligned} \frac{d^2 g_2(\delta)}{db^2} &= \frac{d^2 g_2(\delta)}{d\delta^2}\left(\frac{d\delta}{db}\right)\left(\frac{d\delta}{db}\right)^T \\ \text{where } \frac{d^2 g_2(\delta)}{d\delta^2} &= -\frac{(\alpha-1)}{\delta^2} < 0 \\ \left(\frac{d\delta}{db}\right)^T &= (1, e^{-c_1 s}, e^{-c_2 s}, e^{-c_3 s}, e^{-c_4 s}) \end{aligned}$$

Therefore $d^2 g_2(\delta)/db^2$ must be negative semi-definite.

D.8 The log-posterior density is

$$g_3(b|P) = g_1(P|b) + g_2[f(t,t)] + g_2[f(t,t+10)] + g_2[f(t,t+\infty)] + \text{constant}$$

D.9 Since $g_1(P|b)$ and $g_2(\delta)$ are (strictly) concave, $g_3(b|P)$ must also be strictly concave. Therefore any maximum $\hat{b}$ must be unique.

Appendix E: The inflation model

E.1.1 In this appendix we consider the pricing of index-linked gilts. The method takes the forward-rate curve for nominal gilts as given and estimates a forward inflation curve from index-linked prices.

E.1.2 We use the following notation:

$$
\begin{aligned}
f(t,t+s) &= \text{nominal forward-rate curve} \\
&= b_0(t)+\sum_{i=1}^{4} b_i(t)e^{-c_i s} = b_0+\sum_{i=1}^{4} b_i e^{-c_i s} \\
\text{For stock } k\text{:} & \\
m_{k0} &= \text{base month for RPI for stock } k \\
m_{k1} &= \text{month of next coupon payment} \\
m_{k2} &= \text{month of second coupon payment} \\
&\vdots \\
m_{kn_k} &= \text{month of final coupon and redemption} \\
m_{k,j+1} &= m_{kj}+6 \\
\tilde{m} &= \text{month of last known RPI} \\
t_{k1} &= \text{time of first coupon payment (years)} \\
&\vdots \\
t_{kn_k} &= \text{time of final coupon and redemption} \\
t_{k,j+1} &\approx t_{kj}+0.5 \\
RPI(m) &= \text{RPI for month } m,\ m \leq \tilde{m} \\
R\hat{P}I(m) &= \text{estimated RPI for month } m > \tilde{m}
\end{aligned}
$$

E.1.3 Let C_{kj} be the (estimated) amount of the cashflow at time t_{kj}. Then

$$
\begin{aligned}
C_{k1} &= \frac{g_k}{2}\frac{RPI(m_{k1}-8)}{RPI(m_{k0})} \\
C_{k2} &= \begin{cases} \frac{g_k}{2}\frac{RPI(m_{k2}-8)}{RPI(m_{k0})} & \text{if } m_{k2}-8 \leq \tilde{m} \\ \frac{g_k}{2}\frac{R\hat{P}I(m_{k2}-8)}{RPI(m_{k0})} & \text{if } m_{k2}-8 > \tilde{m} \end{cases} \\
C_{kj} &= \frac{g_k}{2}\frac{R\hat{P}I(m_{kj}-8)}{RPI(m_{k0})} \quad \text{for } j=3,\ldots,n_k-1 \\
C_{kn_k} &= \left(\frac{g_k}{2}+100\right)\frac{R\hat{P}I(m_{kn_k}-8)}{RPI(m_{k0})}
\end{aligned}
$$

E.1.4 The estimated dirty price of the stock for parameter set ϕ is thus

$$
\begin{aligned}
P_k(\phi) &= \sum_{j=1}^{n_k} C_{kj}(\phi) P(t, t+t_{kj}) \\
\text{where } P(t,t+s) &= \exp\left[-\int_0^s f(t,t+s)ds\right] \text{ is given}
\end{aligned}
$$

E.2 *The model for* $\hat{RPI}(m)$

E.2.1 Suppose that it is now time t. Let $s(m) = (m - \tilde{m})/12$ be the time (in years) from the last known RPI to the desired time for projection of RPI.

$$
\begin{aligned}
\hat{RPI}(m) &= RPI(\tilde{m}) \exp\left[r_0(t)s(m) + \sum_{i=1}^{4} r_i(t) \frac{1 - \exp(-c_i' s(m))}{c_i'}\right] \\
&= RPI(\tilde{m}) \exp\left[r_0 s(m) + \sum_{i=1}^{4} r_i \frac{1 - \exp(-c_i' s(m))}{c_i'}\right]
\end{aligned}
$$

E.2.2 The parameters r_0, r_1, r_2, r_3 and r_4 are estimated daily, while the c_i' are fixed, as in the nominal forward-rate curve. We note::

1. $r(t,t'+s) = r_0 + \sum_{i=1}^{4} r_i \exp(-c_i' s)$ is the implied, underlying force of inflation as to be estimated at time t. t' is the date of application of the most recently announced RPI before t and corresponds to $\tilde{m}$. s is the time from the last known RPI.

2. It was found from experimentation that $c_i' = c_i$ $(i = 1,2,3,4)$ gives the best results.

3. $c = (c_1, c_2, c_3, c_4) = (0.2, 0.4, 0.8, 1.6)$ gives significantly better results than $c = (0.1, 0.2, 0.4, 0.8)$: the residuals are smaller and the long end of the forward inflation curve is more stable.

4. Other forms for $\hat{RPI}(t)$ have been considered. One model allowed for an additional, initial shock term, δ, which adjusted for a supposed difference between the published RPI and its true 'underlying' value. This model did achieve a good fit to the data. Moreover, δ would generally progress smoothly from one day to the next except on the days on which RPI was announced when it would jump to a new level (presumably if the new RPI had removed the temporary deviation from its true level or vice versa). However, the pattern of δ's turned out to be difficult to explain with little evidence of seasonality in its values. Instead it was thought that it was just helping to keep the predicted price of the shortest-dated stock in line.

5. There has been no attempt to put the seasonal effect back into the estimated cashflows: that is, by using $\tilde{RPI}(m) = \hat{RPI}(m)S(m)$, where $S(m)$ is some seasonal function with period 12 months and average value 1. For one reason, historical data indicates that the seasonal nature of inflation has changed substantially over the years and will continue to change in the future. It would be difficult to predict such changes in the future. Coupon

and redemption payments are also roughly spread thoughout the calendar year so that errors may, in any event, balance out. Never-the-less, the absence of seasonality in $R\hat{P}I(m)$ could be significant for shorter-dated stocks.

E.3 *A statistical model for errors in prices*

E.3.1 Let the true parameter set be ϕ. P_k equals the actual price of stock k while $\hat{P}_k = \hat{P}_k(\phi)$ denotes the theoretical price given ϕ. The statistical model is this analysis given ϕ is:

1. the P_k are independent;
2. $\log P_k \sim N\left(\log \hat{P}_k, \nu^2(\hat{P}_k, d'_k)\right)$;
3. d'_k is the 'inflation' duration of stock k and is equal to $d_k - t_k^f$ where d_k is the duration of stock k and t_k^f is the time to maturity on the day on which the stock ceases to have any element of indextion;
4.

$$\begin{aligned}
\nu^2(p,d') &= \frac{\sigma_0^2(p)\left[\sigma_\infty^2 d'^2 b(p)+1\right]}{\sigma_0^2(p) d'^2 b(p)+1} \\
b(p) &= \frac{\sigma_d^2}{\sigma_0^2(p)[\sigma_\infty^2 - \sigma_0^2(p)]}
\end{aligned}$$

E.3.2 Note that

$$\begin{aligned}
\nu^2(p,0) &= \sigma^2(p) \\
\nu^2(p,\infty) &= \sigma_\infty^2 \\
\frac{\partial(\nu^2)}{\partial(d'^2)} &> 0 \\
\left.\frac{\partial(\nu^2)}{\partial(d'^2)}\right|_{d'=0} &= \sigma_d^2.
\end{aligned}$$

E.3.3 Now

$$\begin{aligned}
P_k &= \hat{P}_k \exp\left(\nu(\hat{P}_k, d'_k) Z\right) \\
\text{where } Z &\sim N(0,1) \\
\Rightarrow P_k &\approx \hat{P}_k\left(1+\nu(\hat{P}_k, d'_k) Z\right) \\
\Rightarrow S.D.(P_k) &\approx \hat{P}_k \nu(\hat{P}_k, d'_k)
\end{aligned}$$

E.3.4 If $d'_k = 0$ we use $S.D.(P_k) = 4/32$ to reflect rounding errors in prices and the bid/offer spread (and also to decrease reliance on the shortest-dated stocks which are marginally favoured by tax-payers). Hence

$$\sigma_0^2 = (8\hat{P}_k)^{-2} \approx (8P_k)^{-2}$$

E.3.5 One effect of having $\sigma_0^2 > 0$ is that the effective weight attached to short-dated stocks tends to 0 as they approach fixed interest status.

E.3.6 For large d'_k we use $S.D.(P_k) = 0.002P_k$: that is, $\sigma_\infty^2 = 0.002^2$.

E.3.7 It is appropriate to include this limit (more so than the fixed-interest data) since the duration of index-linked stocks can be relatively high. Before this limit was introduced, errors in the prices of long-dated stocks could be relatively large.

E.3.8 Temporarily, now assume that $\sigma_0^2 = 0$ and $\sigma_\infty^2 = \infty$. Then

$$\nu^2(p,d') \to \sigma_d^2 d_i'^2 \quad \text{as } \sigma_0^2 \to 0 \text{ and } \sigma_\infty^2 \to \infty$$

E.3.9 This is consistent with real gross redemption yields y_k having mean $\hat{y}_k$ (consistent with $\hat{P}_k$) and variance σ_d^2: that is, the variances of the y_k do not depend upon the duration of the stock.

E.3.10 When $\sigma_0^2 > 0$ and $\sigma_\infty^2 < \infty$ this argument provides the logic behind the value

$$\left.\frac{\partial(\nu^2)}{\partial(d^2)}\right|_{d=0} = \sigma_d^2$$

E.3.11 In this analysis σ_d was set to 0.002 (20 basis points). This reflects the problems in getting a reliable model for index-linked prices and also the desire to reduce the weight of the shortest-dated stocks.