

Testing term structure estimation methods: Evidence from the UK STRIPs market

James M. Steeley

Aston Business School
Aston University
Birmingham, B4 7ET, UK.

Abstract

Prices and yields of UK government zero-coupon bonds are used to test alternative yield curve estimation models. Zero-coupon bonds permit a more pure comparison, as the models are providing only the interpolation service and not also making estimation feasible. It is found that better yield curve estimates are obtained by fitting to the yield curve directly rather than by fitting first to the discount function. A simple procedure to set the smoothness of the fitted curves is developed, and a positive relationship between over-smoothness and fitting error is identified. A cubic spline function fitted directly to the yield curve provides the best overall balance of fitting error and smoothness, both along the yield curve and within local maturity regions.

This research received financial support from the Leverhulme Trust, Personal Research Fellowship, no.10505. I am very grateful for this funding and also for the support of H.M. Treasury's Debt Management Office. This revision has benefited from the generous and wise counsel of the anonymous referee, and from discussions with Mark Deacon and colleagues at the D.M.O., Mike Joyce and colleagues at the Bank of England, and Clint Cummins at TSP International. I am very grateful to them all. Address for correspondence: Professor James M. Steeley, Aston Business School, Aston University, Birmingham, B4 7ET, UK. Tel: 0121-204-3248. E-mail: j.m.steeley@aston.ac.uk.

Testing term structure estimation methods: Evidence from the UK STRIPs market

Abstract

Prices and yields of UK government zero-coupon bonds are used to test alternative yield curve estimation models. Zero-coupon bonds permit a more pure comparison, as the models are providing only the interpolation service and not also making estimation feasible. It is found that better yield curve estimates are obtained by fitting to the yield curve directly rather than by fitting first to the discount function. A simple procedure to set the smoothness of the fitted curves is developed, and a positive relationship between over-smoothness and fitting error is identified. A cubic spline function fitted directly to the yield curve provides the best overall balance of fitting error and smoothness, both along the yield curve and within local maturity regions.

1 Introduction

The term structure of interest rates defines the array of discount factors on a collection of default-free pure (zero-coupon) discount bonds that differ only in their time to maturity. It is used for a number of purposes. For example, it can be used to set the price of new issues of default-free debt, to set the interest rates on intra-governmental loans, or to test the theories of the stochastic evolution of the term structure, (e.g., Cox, Ingersoll and Ross, 1985).

The prices of government bonds, known as gilts in the UK, are generally used in the estimation of the term structure of interest rates. Following the introduction of the gilt STRIP facility in December 1997, which permitted the separate trading of the cash flows of a selection of conventional (coupon-bearing) gilts, direct observation of the term structure of interest rates became possible. The yield to maturity on a stripped cash flow, which is its per-period average rate of return until maturity, defines the unique default-free interest rate for that maturity. Unless there exists a continuum of zero-coupon bonds across the maturity range, however, there will be a need to use some form of interpolation method whenever interest rate estimates are required for dates that do not coincide with existing cash flows. The aim of this study is to compare a selection of interpolation methods in the case of the UK gilt STRIPs market.

Interpolation methods were introduced into the term structure literature long before zero-coupon bonds were first traded. At that time, it was necessary to use coupon-bearing bonds to estimate the term structure, and the use of interpolation methods met twin objectives. In addition to the desirability for interpolation, the use of such methods would, in many cases, make estimation feasible. In the absence of arbitrage, the prices of all coupon-bearing bonds, P_i , are related to the zero-coupon yield (spot interest rate) curve, $y(t)$, by

$$P_i = \sum_{j=1}^{n_i} c_i F_i d(t_{i,j}) + F_i d(t_{i,n_i}) \quad i = 1, 2, \dots, m \quad (1)$$

where the discount function, $d(t) = \exp(-ty(t))$, is the price of a unit payoff at date t and where bond i pays a periodic coupon at rate c_i and repays its face value, F_i , at date n_i . If the number of bonds equals the number of payment dates, the term structure can be exactly identified by solving equation (1) simultaneously for all bonds. In the case where there are more bonds than payment dates, equation (1) is over-identified, but an approximate fit can be obtained by using, for example, least squares. Carleton and Cooper (1976) successfully employed this method by selecting a sample of around thirty bonds that only incurred four payment dates per year and had a maximum maturity of seven years. In most bond markets, however, the number of payment dates greatly exceeds the number of bonds available, and so some means of reducing the dimensionality of the term structure (to less than the number of bonds) is needed. The solution, first suggested by McCulloch (1971), is to approximate the term structure with a low dimensional function, which depends on a small number of parameters. This function, being defined across the maturity spectrum, simultaneously undertakes the interpolation task.

Since the pioneering work of McCulloch (1971,1975), a huge literature has grown up devoted to both innovating and comparing alternative approximation functions.¹ With few exceptions, this literature has used coupon-bearing bonds to conduct the empirical analysis. With coupon-bearing bonds the comparison of alternative approximation functions is made a more complex task as the functions are both making the estimation feasible and providing the interpolation service. With

1. The pre-history of yield curve fitting is described in Anderson et al (1997). During this time, yield curves were either hand-drawn or fitted using ordinary least squares methods, and were applied to the yields on coupon-paying bonds rather than those on zero-coupon bonds. By focussing on zero-coupon bonds, the work of McCulloch marked a significant step forward.

zero-coupon bonds, estimation of at least some points along the term structure is always feasible, and so the approximation function is only providing the interpolation service. The feasibility and interpolation benefits from the use of approximation functions do not arise without a cost. Implicit in selecting the form of approximating function, is the selection of the amount of the fine shape of the term structure that will be lost through approximation. This cost, however, is incurred regardless of the coupon type of the bonds used in estimation. But, by using zero-coupon bonds, it is possible to directly see the trade-off between the benefits arising from estimating the entire term structure and the costs of applying a particular approximating function. In other words, a more pure comparison of alternative approximation functions is possible. This is the prime motivation for this study that will undertake such a comparison using data on UK gilt STRIPs.

Previous studies of alternative yield curve estimation models have compared the models as originally developed, which may not necessarily be making a fair comparison. In this study, all the models will be compared on as similar a basis as is possible, so that it is the pure shape representation ability of the alternative models that is being compared, rather than, perhaps, differences in degrees of freedom, application of constraints, or algorithms and accuracy in non-linear estimation. In addition, a method to assess whether the fitted curves have been over-smoothed is applied in this paper, and the relationship between goodness-of-fit and smoothness is explored. The main findings are that relatively simple, low parameterized, cubic spline functions fitted directly to the yield curve, provide extremely accurate fits to the sequence of yield curves, both on an in- and out-of-sample basis. Even with a small number of parameters the curves are not statistically significantly too smooth and are shown to fit well along the entire maturity spectrum of stylized yield curve shapes. Furthermore, the relation between smoothness and goodness of fit is found

to be negative and convex. This convexity reveals the losses to over-smoothing, as the goodness of fit will deteriorate at a faster rate than the statistical significance of the smoothing decreases.

The rest of the paper is structured as follows. Section 2 describes the models that will be compared, including a new model for the discount function, and sets them all within the broader term structure estimation literature. In Section 3, the many issues that arise in the estimation of these models are described and their resolution within the context of this study explained. These include the use of theoretically-motivated constraints, the dimensionality of the functions, the setting of certain parameters and a procedure to identify the desired level of smoothness. Section 4 begins with an explanation of the dataset used and then presents the results of the estimation exercises. These commence with the results of the procedures used to choose the dimensionality of each model to be estimated, and continue with both in- and out-of-sample measures of the overall fitting ability of the alternative models. The final analysis considers the ability of the various models to fit to local maturity regions of the yield curve. The conclusions are presented in Section 5.

2 Term Structure Estimation Methods

When term structure estimation methods are developed they are usually framed within one of the three alternative representations: the discount function (zero-coupon bond prices), $d(t)$, the zero-coupon (spot) yield curve, $y(t)$, or the implied forward rate function, $f(t) = -\partial \log(d(t)) / \partial t$, which measures the marginal return at maturity t of extending one's investment. Given that observations on zero-coupon bond prices and yields are available to this study, and because

calculation of the implicit forward rate curves requires at least an assumption of piecewise constant rates,² this study will compare models that either fit directly to the discount function, or directly to the yield curve.

The Weierstrass Theorem has been used to justify polynomial approximation of functions, and this is the first model that will be examined. For the discount function, this is,

$$d(t) = \sum_{i=0}^L a_i t^i \quad (2)$$

while for the yield curve, it is, similarly,

$$y(t) = \sum_{i=0}^L a_i t^i \quad (3)$$

This polynomial approximation to the yield curve is equivalent to the exponential polynomial approximation to the discount function suggested by Chambers, Carleton and Waldman (1984). In the application of this model to coupon-paying bonds, non-linear estimation is needed, whereas for data on zero-coupon yields to maturity, the model is linear in the parameters as the exponential transformation has occurred in the data and does not need to be introduced within the estimation.

The use of polynomial functions requires care that computed numerical values do not diverge rapidly to unbounded values. The solution, first proposed by McCulloch (1971) for term structure applications, is to use spline approximation that fits a separate polynomial within successive segments of the interval. Intuitively, this reduces the chances of any one polynomial diverging to unbounded values. The

2. As used, for example, by Fama and Bliss (1987) and Coleman et al (1992).

divisions between the segments of the interval and called "knots", and the polynomials are constrained to be continuous and smooth around each knot point. McCulloch (1971) applied quadratic polynomial splines to estimate the discount functions of US railway bonds from 1920-1938, and US government bonds from 1946-1966. As forward rates are first differences of discount factors, the estimated forward curves from the quadratic discount function were discontinuous. McCulloch (1975) overcame the "knuckle" shape of the forward rate function by applying a cubic spline to the discount function, and this has become the standard degree for discount function estimation. A cubic spline to the discount function can be represented by

$$d(t) = \sum_{i=0}^3 a_i t^i + \sum_{p=1}^{n-1} b_p (t - t_p)^3 D_p \quad (4)$$

where n is the number of segments of the maturity span, and where $D_p = 1$ if $t \geq t_p$, and $D_p = 0$ otherwise. Since that time, cubic splines in this form have been applied, for example, by Shea (1984) and Kikugawa and Singleton (1994) to estimating the discount function for Japanese bonds, by Mastronikola (1991), to the UK government bond market, and by Rumsey (1994) to Canadian government bonds.

While straightforward and widely adopted, numerical inaccuracies may arise using polynomial splines that consist of power and truncated power functions, see Powell (1981, p.227-8). Alternative component, or basis, functions can be used to construct a numerically stable cubic spline function. Shea (1984) and Steeley (1991) recommended the use of a cubic B-spline basis, and subsequent applications include Eom *et al* (1998), Lin and Paxson (1993,1995), Subramanian (2001) and Lin (2002). The increase in numerical precision that is available within programming software means that the potential instability of a basis of power functions and truncated power

functions may be much reduced now. To examine this, all the analysis was repeated using both representations of the cubic spline. Using double precision and an efficient programming structure and language produced identical results using either spline representation.³

A cubic spline model applied directly to the yield curve is equivalent to the exponential spline of the discount function proposed by Langetieg and Smoot (1989). This model approximates the yield curve as

$$y(t) = \sum_{i=0}^3 a_i t^i + \sum_{p=1}^{n-1} b_p (t - t_p)^3 D_p \quad (5)$$

where all terms have been defined previously, and, like the polynomial approximation to the yield curve, produces an exponential decay for the discount function. Again, like the exponential polynomial, this model is non-linear in the parameters when applied to the discount function, but linear in the parameters, when applied to the observable yields to maturity on zero-coupon bonds. McCulloch and Kochin (2000) have employed a similar functional form to estimate US real term structures, using a quadratic spline on the yield curve. By further requiring that the forward rate is constant beyond the maximum observed maturity and that both the

3. The programming was done using the algebraic and matrix capabilities of TSP4.5 (Hall and Cummins, 1999) along with the use of some in-built routines, running under double precision. For the B-spline representation, and because only cubic splines are used, the recurrence relation was used to expand the B-spline functions mathematically prior to the estimation, rather than use the recurrence relation within the estimation programme. In an earlier draft of this paper, the computation of the values of the basis functions using the recurrence relation led to some small numerical discrepancies between the alternative spline representations, and also to a much slower execution time. It is possible that, for higher degree splines or where there are many more parameters than are estimated here, differences between the alternative spline bases may reappear.

yield and forward curve are linear out to the first knot, an exact fit with a small dataset can be achieved.⁴ Lin (2002) has also applied a cubic spline approximation to a spot rate curve, but used the B-spline parameterization.

Vasicek and Fong (1982) recommended fitting a polynomial spline to an exponential transformation of maturity. However, Shea (1985) argued that the difficulties of fitting exponential decays by polynomial functions do not extend to local approximations (i.e., splines) to exponential functions, such as equation (4), and thus this method was unnecessarily complex. Responding to this complexity, Ferguson and Raymar (1998) proposed a simplified version of the Vasicek and Fong (1982) method that maintained the exponential functions, but now globally, and could (by setting the non-linear parameter, \bar{y} , equal to the observed long yield) be estimated by linear methods. Specifically, they specified the discount function to be of the form

$$d(t) = \sum_{i=0}^L a_i \exp(-i\bar{y}t) \quad (6)$$

and used both four and six parameter versions of this model. The application of this model in this study will also set the non-linear parameter equal to the observed long yield.

Functional forms involving exponential functions of maturity have also been proposed for the yield curve. Nelson and Seigel (1987) proposed the following model for zero-coupon yields to maturity

4. A linear spline is equivalent to linear interpolation when the data points coincide with the knot points. An unrestricted cubic spline will have 3 more parameters than data points if the knot points and data points coincide. Using quadratic splines and two such additional restrictions in this way, therefore, aligns the dimensions to achieve a similar, but higher degree, correspondence.

$$y(t) = a_0 + a_1 \left[\frac{1 - \exp(-t/t_1)}{t/t_1} \right] + a_2 \left[\frac{1 - \exp(-t/t_1)}{t/t_1} - \exp(-t/t_1) \right] \quad (7)$$

where a_0, a_1, a_2, t_1 are the parameters to be estimated, and which also produces an exponential decay for the discount function. This form for the yield curve actually comes from integrating the proposed functional form for the forward curve, that is,

$$f(t) = a_0 + a_1 \exp\left(-\frac{t}{t_1}\right) + a_2 \left[\frac{t}{t_1} \exp\left(-\frac{t}{t_1}\right) \right] \quad (8)$$

By focussing on the forward curve, Nelson and Seigel (1987) were responding to the difficulties in obtaining well-behaved forward curves from estimates of either the discount function or the yield curve. Shea (1985) had pointed out that even estimation methods that deliver smooth discount functions or smooth yield curves may not necessarily lead to smooth forward rate curves. So that if a smooth forward rate curve is desired, then it may be necessary to apply an approximation technique directly to the forward curve. McCulloch (1977, 2000) has shown that when the forward curve is extracted from bond prices, it will have a tendency to become increasingly poorly defined as maturity increases, unless some stable structure is imposed upon it.

In applications of this model, to coupon-paying bonds, to estimate a forward curve, or both, non-linear estimation is required. However, with data on zero-coupon bonds, the yield curve form of the model can be estimated by linear methods if the non-linear parameter, t_1 , is set to a pre-determined maturity point. This parameter influences the maturity of any hump in the estimated forward curve, and Svensson (1994) pointed out that additional flexibility could be introduced by adding additional "hump" terms to the estimated forward curve. Svensson recommended one further term, giving a yield curve of the form,

$$y(t) = a_0 + a_1 \left[\frac{1 - \exp(-t/t_1)}{t/t_1} \right] + a_2 \left[\frac{1 - \exp(-t/t_1)}{t/t_1} - \exp(-t/t_1) \right] + a_3 \left[\frac{1 - \exp(-t/t_2)}{t/t_2} - \exp(-t/t_2) \right] \quad (9)$$

where there are now two non-linear parameters, t_1, t_2 . Although some applications of this model, such as Deacon and Derry (1994), have estimated the non-linear parameters freely, the potential collinearity if the two parameter values coincide suggests that pre-setting them or using a rough grid spacing may be more robust. Restricting the non-linear parameters will reduce the dimensionality of this model, but this can be restored by adding more "Svensson" terms but all with fixed non-linear parameters. Therefore, the extended form of the Svensson model, which is tested in this study, can be refilled to the dimensions of the original model, yet estimated using linear methods.⁵

While the Nelson and Seigel and Svensson models have been found to generate smooth forward curves,⁶ Anderson and Sleath (1999) have shown that there can be unforeseen consequences of "over-smoothing" the forward rate curve. Specifically, they observe that the long maturity horizontal asymptotes can dominate the shape of the curve along its entire length.⁷ This problem is likely to be worse, however, when these models are estimated with free non-linear parameters. Whenever these parameters are drawn to the short end of the curve, the long end of the curve will be

5. In a sense, this treatment of the non-linear parameters makes this extended Svensson model much more like, intuitively at least, a spline function, where the knots fulfil a similar role of distinguishing the component functions. Of course, unlike the spline function, all the component functions in these models are non-zero across the entire maturity spectrum.

6. Applications and comparative studies include, Anderson (1994), Deacon and Derry (1994), Anderson and Sleath (2001), Bliss (1997) and Ioannides (2003).

7. While Dybvig et al (1996) and McCulloch (1999) have shown that there are no-arbitrage restrictions on the yield on infinite maturity zero-coupon bonds, the issue here is that certain functional forms used in yield curve estimation applications may enforce this restriction at unrealistically short maturities.

predominantly influenced by the constant term. By presetting the non-linear parameters, it will be seen in this study whether the relative influence of the constant term, which still appears in the yield curve form, is reduced.

The application of spline functions directly to the forward rate curve was first proposed by Fisher et al (1995), who also used a cubic B-spline basis. They introduced a large number of knot points into the maturity range and then constrained the roughness by introducing a penalty function, which is related to the curvature of the forward rate curve, into the objective function. The control of smoothness in this model has subsequently been refined in the variable roughness penalty (VRP) model of Waggoner (1997) and also by Anderson and Sleath (2001). The results in Anderson and Sleath (2001) and the comparative studies of Ioannides (2003) and Waggoner (1997), however, question the value of the extra effort expended to estimate yield curves through an approximation to the forward curve. While these studies find little differences between the fit of these models and simpler models that fit directly to the discount function or yields, the forward curve models require large numbers of non-linear parameters to be estimated and the forward rates themselves are not directly observable but must be constructed using an assumption of, at least, piecewise constant rates.

The mathematical parameterizations of the yield curve described so far have been developed without specific reference to economic theories that have been developed to explain the shape of the yield curve, yet there are many equilibrium and no-arbitrage models of the evolution of the yield curve. While some of these take a fitted yield curve as an input,⁸ the models of, for example, Vasicek (1977)

8. For example, Ho and Lee (1986), Hull and White (1990) and Heath, Jarrow and Morton (1992).

and Cox, Ingersoll and Ross (1985) (CIR) have explicit functional forms for the yield curve. These can, therefore, be used as alternatives to the purely mathematical parameterizations of the yield curve.

Specifically, the Vasicek model implies that the yield curve is given by

$$y(t) = -(b(t)(r_\infty - r_0) - r_\infty t - 0.25b^2(t)\sigma^2/\kappa)/t \quad (10)$$

where $b(t) = (1 - \exp(-\kappa t))/\kappa$ and $\kappa, r_0, r_\infty, \sigma^2$ are the parameters to be estimated.

Although these parameters are interpretable within the Vasicek model as, respectively, the rate of mean reversion of the short interest rate, the current value of the short interest rate, the current value of the longest maturity interest rate, and interest rate variance per unit time, these have no special status in the application of this model as a yield curve approximation method. They do, however, provide a guide to useful starting values in the non-linear estimation that is required for this model.

The CIR model is also a function of four parameters, and can be written as

$$d(t) = a(t) \exp(-b(t)r_0) \quad (11)$$

$$a(t) = \left[\frac{\phi_1 \exp(\phi_2 t)}{\phi_2 (\exp(\phi_1 t) - 1) + \phi_1} \right]^{\phi_3}$$

$$b(t) = \left[\frac{\exp(\phi_1 t) - 1}{\phi_2 (\exp(\phi_1 t) - 1) + \phi_1} \right]$$

where $r_0, \phi_1, \phi_2, \phi_3$ are parameters to be estimated. As in the Vasicek model, r_0 represents the current value of the short interest rate, and in certain combinations, all the estimated parameters can be used to recover interpretable parameters, see Brown and Dybvig (1986).

In addition to the preceding models that have been suggested in earlier studies, a new model for the discount function is proposed here. This model includes features of both the simplified Vasicek and Fong model for the discount function and the extended Svensson model for the yield curve, and is written as

$$d(t) = \sum_{i=0}^L a_i \exp(t_i t) \quad (12)$$

As in the Vasicek and Fong model, the function is a sum of exponential functions of maturity. But, instead of maturity being scaled by multiples of the long yield, the scaling is by a set of constants, t_i , $i = 0, 1, \dots, L$, that can be set at fixed maturity points. As in the extended Svensson model, these can be preset to maturity quantiles or optimized over some maturity grid. Wiseman (1998) has proposed a similar functional form for the forward rate curve rather than the discount function, and proposes that the estimation should include a second stage that attempts to improve the fit by varying the non-linear parameters given the values of linear parameters estimated at the first stage. However, his results suggest that the incremental benefit from this second stage appears very small if maturity quantiles were used in the first stage as pre-set values for the non-linear parameters. So, in this study, equation (12) will be applied using fixed values for the non-linear parameters.

3 Issues in Estimation

In addition to the selection of both the functional form and the curve that is to be estimated, there are other choices to be made prior to estimation. These include the imposition of constraints at the ends of the curves, the desired level of smoothness

of the function and the dimensionality of the chosen functional form, that is, the number of free parameters. This section outlines the decisions that need to be made and the procedures adopted to make them.

Constraints. During estimation of a discount function, it is natural to impose the no-arbitrage constraint that $d(0) = 1$. While the Cox et al (1985) model imposes this constraint endogenously the other discount function models need this constraint imposed on the parameters during estimation. The polynomial and cubic spline functions can be constrained by setting $a_0 = 1.0$, while the new exponential model and simplified Vasicek and Fong model require the constraint $\sum_i a_i = 1$.

There are, however, other constraints that one may wish to impose. For example, Schaefer (1981) introduced constraints to ensure that forward rates were always positive, another no-arbitrage condition. Shea (1984) suggested that constraining the slope of the discount function would achieve the same end. McCulloch and Kochin (2000) add curvature constraints at both ends of the yield curve. In many ways, these constraints are responding to the desire to control excess roughness. A known hazard in applying approximation functions is the difficulty in controlling the extreme regions of the function. For the discount function, the ability to force the function through unity at time zero mitigates this problem. There is, however, no equivalent no-arbitrage condition for the spot or forward rate curves at zero maturity. The closest that can be achieved is to set the short end of the yield curve equal to an estimate of a short term interest rate from another source. Anderson and Sleath (2001) explore the differences in fit that can be obtained using a constraint on the short interest rate. Specifically, they add the information contained in six general collateral gilt repo rates taking maturities of between 1 week and 6 months. Although it appears that the yield curve fit is much improved by constraining the

short end of the curve, the comparison is not truly like-for-like, as they also increase the number of free parameters in the estimated curves. In fact it is possible that the constraint is necessary to control excess flexibility induced by the increased dimensionality rather than to improve the fit at the short end of the yield curve itself. In this study, only the short end of the discount function will be constrained.

Dimensionality. Except for the Vasicek (1977) and CIR models, it is necessary to choose either the degree of or the number of terms in the functional form. This dimensionality choice in each of the alternative functional forms will determine both the fit of the discount function or yield curve and its smoothness. As the dimensionality increases, so the fit will improve, but the function also becomes less smooth (rougher between the observed values).⁹ For the polynomial model, the spline model, the simplified Vasicek and Fong model, the exponential model and the extended Svensson model, it is necessary to set the number of free parameters. For the polynomial models, this means adding higher degree terms. The simplified Vasicek and Fong model and the exponential model expand with similar ease, by the addition of further matching exponential functions and, in the latter case, also with finer maturity quantiles. The Nelson-Seigel model, as extended by Svensson, can be further extended by adding additional Svensson-style terms. So, in equation (9), the non-linear parameter t_1 , which is the only non-linear parameter in the Nelson and Seigel version of this model, is set at the median maturity for the term involving the linear free parameter a_1 . For the Svensson-style terms involving linear free parameters a_2 and higher, non-linear parameters t_1 and higher are set at even quantile

⁹ If smoothness is the overriding consideration, then Adams and van Deventer (1994), using an unpublished theorem of Vasicek, show that a fourth order polynomial spline, but without the cubic term, applied to the discount function generates the smoothest forward rate curve.

divisions of the maturity space.¹⁰ The spline functions can be expanded either through an increase in the degree of the spline or through an increase in the number of knot points, with the latter option being the standard choice, as cubic splines offer the most parsimonious choice that delivers continuous forward curves.

While all of the alternative specifications, except for the Vasicek and CIR models, can be expanded to any desired degree, the degree chosen is key to the success of the estimated curves. It is perhaps surprising then that there has been relatively little examination of this issue in prior studies. The earliest suggestion, from McCulloch (1975), who used cubic splines to approximate the discount function, is to set the number of knots equal to the closest integer to the square root of the number of bonds. By contrast, Litzenberger and Rolfo (1984) divided the bond market into "short", "medium" and "long" maturities, as generally understood by the market. Some of the more parsimonious exponential models, such as Nelson and Seigel (1987) have emphasized smoothness and so are deliberately low dimension. In this study, we will implement a new procedure suggested by McCulloch and Kochin (2000) to select the dimensionality of the curves.

McCulloch and Kochin (2000) pointed out that the autocorrelation of the fitting errors along a discount function can provide guidance as to whether or not the function is sufficiently smooth. If the autocorrelation is significantly positive, then the curve may be too stiff, while if the autocorrelation is significantly negative, then the curve may be too flexible, such that neighbouring errors alternate in sign too frequently to be consistent with chance. To implement this criterion, we estimate

10. So, in the basic Nelson and Seigel form of the model, t_1 , is set to the median maturity for the terms involving both a_1 and a_2 . In the Svensson form, t_1 , is again set to the median maturity for the term involving a_1 , but is set to the lower tercile on the term involving a_2 , while t_2 is set to the upper tercile on the term involving a_3 . When the model has a further "Svensson" style term added, quartiles are used, and so on.

each model for a range of different dimensions and compute the fitting errors and their autocorrelations. We use the combination of statistical significance of autocorrelation coefficients and the size of absolute errors, relative to estimates of dealing costs, to select the best dimension. This dimension is then set equal across all the models that are being compared.

Knots and non-linear parameters. In the case of the spline models, the Nelson-Seigel-Svensson model and the new exponential model, certain maturity-based parameters must also be selected. For the spline model, this means the selection of knot positions, while in the other two models, the non-linear parameters are set to certain maturity levels to generate a linear model. For both kinds of model, two forms of search exercise will be explored. In the first exercise, a grid search is used across non-overlapping whole-year maturity points. For example, for the cubic spline function applied to the discount function, this is a three-dimensional search in the range 0 to 40 years. This grid search maintains the same spacing for the maturity-based parameters throughout the sample period.¹¹ In the second, exercise, evenly spaced quantiles of maturity are used. Although the actual maturity spacings will change from day to day throughout the sample period, the quantile division is maintained.

Estimation. By pre-specifying the non-linear parameters in the simplified Vasicek and Fong model, the new exponential model and the extended Svensson model, it is possible to estimate all the models, except the Vasicek and CIR models,

¹¹ In practical applications, the re-optimization of the maturity-based parameters each day may well be too time consuming, meaning that a fixed selection based upon an historical exploration such as this, is likely to be the expedient option. Moreover, if the dual considerations of smoothness and fit are used, the ultimate parameter selection will be an additional subjective decision each time the yield curve is estimated.

using linear methods.¹² In all cases, a least squares criterion is used, so the estimated parameters are those that minimize the squared errors between the observed discount factors [spot yields] and fitted discount factors [spot yields], that is,

$$\text{Min}_{\hat{\theta}}(d(t) - d(t:\hat{\theta}))'(d(t) - d(t:\hat{\theta})) \quad (13)$$

for the discount function models, and

$$\text{Min}_{\hat{\theta}}(y(t) - y(t:\hat{\theta}))'(y(t) - y(t:\hat{\theta})) \quad (14)$$

for the yield curve models, and where θ is the vector of parameters to be estimated.

The accuracy of non-linear estimation relies on its ability to find the global minimum rather than stalling in a local minimum or failing to converge at all. In addition to judicious setting of the various control parameters in the minimization algorithm it is also important to use good starting values for the estimated parameters. The approach taken in this paper was to use the previous day's estimated parameters as the starting values for the following day's estimated parameters. On the first day of the sample, searches across various starting values were undertaken to ensure that the first day's starting values were themselves robust. This procedure was then repeated for each day in the first two weeks of the sample and the parameter estimates compared to those found using the previous day's starting values. There was no difference.

Error Statistics. In order to compare the models, three error measures are reported. These are: mean absolute pricing error, mean absolute yield error, and

12. Parameter instability and convergence problems in non-linear estimation are well known problems in the Svensson model, see, for example, Anderson and Sleath (2001), and so extending the flexibility of the Nelson and Seigel model through the addition of terms having unknown linear parameters could prove a worthwhile alternative.

weighted mean absolute pricing error. The discount function models are naturally assessed by a pricing error, while the yield curve models are naturally assessed by a yield error. In order to compare across models - to be able to say whether curves should be fitted to the discount function or to the yield curve - pricing errors are translated into yield errors and vice versa. In addition, and following many other studies, e.g., Bliss (1997) and Waggoner (1997), pricing errors weighted by the inverse of duration (maturity in the case of strips) are also reported. Weighting by inverse-duration controls the heteroscedasticity often found in unweighted pricing errors.

As one of the core functions of an estimated yield curve is the pricing of new issues, it is equally important to compare price and yield errors on an out-of-sample basis. To do this, a hold-out sample is identified from within the set of bonds available and the curves are estimated from the remaining (in-sample) zero-coupon bonds. These curves are then used to obtain the price and yield errors for the hold-out sample. This process is then repeated with the in-sample and hold-out sample reversed, and the errors from both hold-out samples are combined to provide an overall out-of-sample equivalent of the price and yield error measures. Of course, the separation of the sample into in-sample and hold-out sample is arbitrary but, as also found in Bliss (1997), the results regarding the relative errors of the different models are qualitatively similar for alternative sub-sample selection methods. Specifically, we explored methods based on choosing every other bond along the maturity spectrum, which is similar to a stratified random sample using annual strata, stratified random sampling based on short, medium and long maturity strata and random sampling across the entire maturity spectrum.

To measure smoothness, we first compute the Durbin-Watson autocorrelation statistic for the errors along a discount function or a yield curve on a given day. This

provides a measure of the excess roughness along the curve, the higher the statistic the rougher the curve. Adopting a parsimony criterion, we define a curve to be no longer too smooth, when the Durbin-Watson statistic is no longer statistically significant, that is, below the 5% d_L value. We then compute P-values for the statistic using a non-linear approximation to make a small sample adjustment to the asymptotic 5% d_L table as shown in Cummins and Hall (1999, p.241), and these are reported in the results tables. As with the mean error statistics, we compute the P-values for both the actual price (or yield) errors and those translated into yield (or price) errors.

4 Data and Results

The data used in this study are daily closing prices (GEMMA reference prices) of all UK government bond coupon strips over the period from the beginning of strips trading on December 8, 1997 to May 15, 2002. This provides a total of 1119 trading days of data. During the sample period the number of coupon strips fell mainly in the range 57 to 62; the changes being caused by coupons reaching their payment date. For the first 24 weeks of the sample, before the introduction of very long term strippable gilts, the number of coupon strips was around 46. As UK government bonds pay coupons semi-annually, and as the strippable bonds had common payment dates of June 7 and December 7, these coupon strips spanned maturities out to around 23 years. With the issue of the long term strippable gilts in 1998, the maximum maturity rose above 30 years. In April 2002, strippable bonds with a new set of coupon payment dates (March 7 and September 7) were introduced across the maturity spectrum, raising the number of coupon strips to 107.¹³ Each of

13. Further detail on the data set is available in the appendix to this paper.

the models of the discount function, equations (2), (4), (6), (11) and (12), and yield curve, equations (3), (5), (9) and (10) are fitted every day using all available coupon strips on that day.

Dimensionality. Prior to the main comparison between yield curve and discount function models, it is necessary to select the dimensionality of the functions. Table 1 summarises results of the selection exercise for the yield curve models. As the number of free parameters in each function is increased, so the goodness-of-fit increases and the over-smoothness, measured by significant error autocorrelation, decreases. Once each model has six free parameters, the change in yield error is less than one basis point and the level is less than three basis points. This level is well within empirical estimates of dealing costs in UK government bonds, see, for example, Proudman (1995) and Vitale (1998). Moreover, at this dimensionality, the positive autocorrelation in the errors is no longer statistically significant. As each model is estimated each day across the 1119-day sample period, the reported statistics are averages of each days' actual statistic. So, on average, across the sample period, yield curve functions with six free parameters cannot be regarded as too smooth. The results presented for the cubic spline model are further averaged across all possible combinations of (whole year maturity point) knot spacings consistent with that number of free parameters, and so understate the potential possible outcomes for that model.

The same smoothness and fit analysis was undertaken for the discount function models, using the fitting errors and autocorrelations along the discount function and, again six free parameters seemed to be sufficient to provide for insignificant positive autocorrelation in the discount function. As the discount function models are estimated subject to the restriction that $d(0) = 1$, the unrestricted dimension of the

discount function models is one higher than for the yield curve functions. This means that the discount function spline function models contain an additional knot relative to the yield curve spline.

By setting the number of free parameters equal to six, in all the models that permit this flexibility, provides for a comparison of the different models that will be based on their ability to capture the shape of the curves rather than any intrinsic advantage due to increased parameterization.¹⁴ The two models that permit only four free parameters, Vasicek and CIR, will be compared to the best performing of the six parameter models pruned back to their equivalent four parameter forms.

Parameter Selection. The columns of figures in Table 1 also suggest that improving the fit of the curve and reducing the autocorrelation of the errors are achieved together. As the spline function models allow further flexibility in that the spacing of the knot points can be varied for any given number of parameters, additional sampling of the relation between fit and smoothness can be obtained. The figures reported in Table 1 are averages across all non-redundent whole number year maturity combinations of the knot points with the maturity space out to 20 years, which is sufficient to identify the global minimum. The fitting error and autocorrelation P-value pairs for all the combinations, in the case of 6 free parameters for both the yield curve and the discount function are displayed in figures in the appendix to this paper. The relationship between fitting error and smoothness is broadly convex. This highlights the dangers of over-smoothing the curve, as the fit

¹⁴ Prior comparative studies, such as Anderson (1994), Deacon and Derry (1994), Anderson and Sleath (2001) and Ioannides (2003), have tended to examine the different models as each was originally proposed rather than try to equalize the number of free parameters and so it is difficult to draw robust conclusions from them. By contrast, Ferguson and Raymar (1998) compare several models, but do so by restricting them to have only four free parameters, which is likely to make the curves too smooth.

deteriorates at a faster rate than the statistical significance of the smoothing decreases, and indicates that this hazard is of more concern in the case of fitting to the discount function.

Table 1: Smoothness and Fitting Errors

For the three yield curve functions (polynomial, cubic spline and extended Svensson), fitted to the cross section of yields to maturity of UK coupon strips for each day (separately) between December 8, 1997 and May 15, 2002, this table contains the average (across the sample) mean (across the curve) absolute yield error (MAYE). P-value is the probability value of the Durbin-Watson statistic measuring the correlation of adjacent yield errors along the fitted curve. The P-values are calculated using a non-linear approximation to make a small sample adjustment to the asymptotic 5% d_L table, see Cummins and Hall (1999, p.241). The reported value is the average of this statistic for each day's curve, across the sample period. Each row contains these statistics for the specified number of free parameters in the function. The polynomial model with 6 free parameters is a 5th degree polynomial, while the extended Svensson model has six free linear parameters, fixed non-linear parameters using maturity quintiles, and contains two extra "Svensson" terms. The results for the cubic spline model are further averaged across all possible combinations of (whole year maturity point) knot spacings consistent with that number of free parameters, and under represent the best possible outcome for that model. The cubic spline model reduces to the cubic polynomial for four or fewer parameters.

Number of free parameters	Polynomial		Cubic-spline		Extended Svensson	
	MAYE	P-value	MAYE	P-value	MAYE	P-value
2	0.1190	0.0000	0.1190	0.0000		
3	0.0889	0.0001	0.0889	0.0001	0.0829	0.0004
4	0.0604	0.0018	0.0604	0.0018	0.0435	0.0107
5	0.0385	0.0295	0.0357	0.0265	0.0305	0.0497
6	0.0295	0.0654	0.0258	0.0847	0.0239	0.1173

The grid search over combinations of knot points, used to gauge the overall relation between fit and smoothness, also identifies the optimum spacing for the knot points, and the sensitivity of the fitting error and the smoothness to the positioning of the knot points. Once again, the detailed results of this exercise are contained in the appendix, and only a summary is provided here. For the yield curve, six free parameters requires two internal knot points in a cubic spline. As the discount function has one constraint imposed ($d(0) = 1$), there will be 3 within sample knot points to give 6 free parameters. In each case, the mean absolute error surface (or

hyper-surface) for the combination of knot spacings appears to have a single optimum. For the discount function, this is where the knots are set at maturities of 3, 10 and 16 years, whereas for the yield curve, the knot spacings are at 1 and 8 years.

The fitting error for the discount function varies from the optimum by less than 2 pence per £100 of nominal bonds in all cases, and in most cases is less than half that amount. For the yield curve, the fitting error is in most cases less than one half of one basis point different. As the optimum knot spacing for the discount function (yield curve) was set using price error (yield error) measures rather than yield error (price error) measures, an examination of whether the fit and smoothness of the curves is sensitive to this choice is made. For the cubic spline model fitted to the discount function, the minimum mean absolute price error using a price error criterion is less than 38 pence per £100,000 nominal bond value smaller than the minimum mean absolute price error using a yield error minimization criterion. Similarly, the yield curve fitting errors are little affected by the choice of error criterion. However, using the price error minimization criterion leads to much less autocorrelation in the errors along the fitted yield curves, than using the yield error minimization. For the discount function, the error criterion choice has little effect on smoothness.

Model Comparisons. The in-sample error statistics for the different models are contained in Table 2. Among the discount function models, the cubic spline function, with optimized knot positioning, and the exponential model provide the best fit and do not over smooth either the discount function or the derived yield curve according to the error autocorrelation criterion. Although the polynomial function, with the same number of free parameters, fits the discount function to an accuracy of only two pence per £100 less than the other two models, the errors in the derived

yield curve are around 50 percent larger. This is most likely due to the significant over-smoothing by this model, particularly in the case of the derived yield curve.¹⁵ The simplified Vasicek and Fong model produces errors, for both the discount function and the derived yield curve, that are around double those obtained from the spline and exponential models, and also appears to over-smooth both curves given the same number of free parameters. The CIR model, which has two fewer free parameters, not only produces errors that are up to ten times larger than those obtained with the spline and exponential models, but also performs less well than an equally parameterized spline function. This indicates that the variety of shapes permissible within the CIR model's yield curves are too restrictive as valuable models of the UK nominal yield curve.¹⁶

Among the yield curve models, the cubic spline and the extended Svensson model provide the best fit and also a smooth yield curve. By contrast with the discount function models, however, the relative smoothing of the yield curve and derived discount function by these best fitting models is quite different, using the same number of free parameters. For example, when the spline model is applied to the discount function, the derived yield curve is smoothed more than the discount function but neither is over-smoothed significantly. When the spline function is applied to the yield curve, the derived discount function is significantly over-smoothed, while the yield curve itself is allowed relatively more flexibility. Moreover, the extended Svensson model permits a more flexible yield curve than the spline function. The Svensson model, and the Nelson and Seigel model from which it was developed, has been criticized in the past for over-smoothing. It is clear

¹⁵. The cubic spline using quantile divisions of maturity to set the knot points did slightly worse than the polynomial function in terms of both error and smoothness. Results available on request.

¹⁶. By contrast, Brown and Schaefer (1994), have found that the model performs adequately for a real interest rate yield curve derived from the prices of UK indexed linked bonds.

now that this property was likely a function of the low parameterization rather than necessarily any inappropriate yield curve shape characteristics.¹⁷ The polynomial function performs less well than the spline function and extended Svensson model, but only by around one half of one basis point in terms of yield errors and does not over-smooth the yield curve. The Vasicek model, which has two fewer parameters, does less well than a polynomial function with the same number of free parameters.

Using the in-sample error statistics, the best two discount function models provide the best overall fits to the discount function, while the best two yield curve models provide the best overall fits to the yield curve. Out-of-sample, however, the discount function models provide the best fit to both the discount function and the yield curve. The results for each of the models, using the hold-out samples based on selecting every other bond along the yield curve, are displayed in Table 3. As with the in-sample results, there is little distance between the two best fitting models for either the discount function or the yield curve. The same rankings are preserved for the alternative hold-out sample selection procedures discussed in Section 3.

Visual shape representation. The error measures reported in the tables above have been both aggregated across the sample and, before that, across the maturity cross section of bonds. While certain models may appear to fit well and to provide adequate smoothing when considering the yield curve as a whole and all yield curves within a given sample, it is possible that some models may fit certain maturity ranges of the discount function or certain typical shape characteristics (effectively

17. While this extended Svensson model is different to the model as originally proposed by Svensson, it maintains the shape representing properties of the component exponential functions. No significant advantage was found by using fixed maturity points to set the non-linear parameters as opposed to using maturity quantiles.

sub-samples) better than others. To examine this, the fitted yield curves are graphed against the observed yields on days within the sample period that contain stylized shapes.

Figure 1 shows some fitted spot rate curves, minimizing yield errors, for two of the fitting method, the cubic spline and the extended Svensson method. Sample curves for the other methods and for all methods minimizing price errors are contained in the appendix to this paper. The two dates displayed here are at the extremes of the sample period and are representative of the two general yield curve shapes observed over the sample: downward sloping and humped. The estimated curves here and in the appendix demonstrate a number of results that apply more generally to the sample. First, fitting to the spot yield curve tends to give a better fit than fitting to the discount function and then deriving the spot yield curve. Second, the difference in fitting ability (particularly for those models that fit most accurately) is most noticeable at the short end of the curve. Models that fit to the yield curve directly, seem better at fitting the short end of the curve. Third, the fitting ability of simple polynomials is perhaps not as poor as the potential for numerical errors could initially suggest. Interestingly, the difference between fitting the same kind of model to either the discount function or the spot rate function is most noticeable for the polynomial model, but the choice is not clear cut. When the yield curve is humped, the polynomial fitted to the yield curve is better, whereas when the yield curve is downward sloping, there is little difference between the yield curves whether the model is fitted to the discount function or to the yield curve is better. Of course, the polynomial model is less accurate than both the spline models and the exponential models. Fourth, for the better fitting models, where also the dimensionality has removed excess smoothness according to the autocorrelation criterion, there is no obvious visual evidence that the curves remain too smooth, or indeed have become too rough.

Table 2: Error Statistics (in-sample)

Panel A contains errors statistics from fitting a polynomial model, a cubic spline model, a simplified Vasicek and Fong model, an exponential model and a CIR function to the prices of UK coupon strips between December 8, 1997 and May 15, 2002. Panel B contains error statistics from fitting a polynomial model, a cubic spline model, an extended Svensson model and a Vasicek model to the yields to maturity of the same UK coupon strips. P-value is the probability value of the Durbin-Watson statistic measuring the correlation of adjacent error terms along the discount function or the yield curve. All reported values are an average of the time series of statistics obtained daily across the sample.

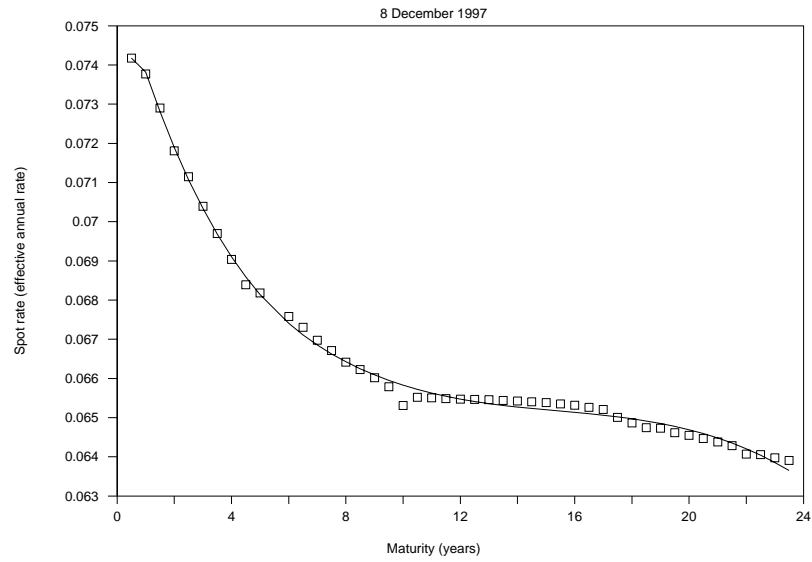
Model	Mean Absolute Price Error £ per £100 par	Mean Absolute Yield Error %	Weighted M.A.P.E £ per £100 par	Price Error autocorrelation p-value	Yield Error autocorrelation p-value	Free Parameters
Panel A: Discount function models						
Cubic spline	0.06808	0.02037	0.00070	0.08708	0.05913	6
Exponential	0.07072	0.02069	0.00071	0.07232	0.05679	6
Polynomial	0.08863	0.03219	0.00123	0.03384	0.01112	6
Vasicek & Fong	0.13455	0.05169	0.00215	0.00685	0.00037	6
Cubic spline	0.14673	0.05689	0.00236	0.00216	0.00022	4
CIR	0.67291	0.15599	0.00594	0.00258	0.00232	4
Panel B: Spot yield models						
Cubic spline	0.09137	0.01738	0.00055	0.01979	0.05379	6
Extended Svensson	0.10759	0.02392	0.00066	0.03121	0.11731	6
Polynomial	0.13491	0.02951	0.00081	0.02009	0.06541	6
Polynomial	0.29141	0.06039	0.00183	0.00004	0.00184	4
Vasicek	0.30949	0.06076	0.00185	0.00113	0.01168	4

Table 3: Error Statistics (out-of-sample)

Panel A of this table contains the out-of-sample errors statistics from fitting a polynomial model, a cubic spline model, a simplified Vasicek and Fong model, an exponential model and a CIR model to the prices of UK coupon strips between December 8, 1997 and May 15, 2002. Panel B of this table contains the error statistics from fitting a polynomial model, a cubic spline model, an extended Svensson model and a Vasicek model to the yields to maturity of the same UK coupon strips.

Model	Mean Absolute Price Error	Mean Absolute Yield Error	Weighted M.A.P.E	Free Parameters
	£ per £100 par	%	£ per £100 par	
Panel A: Discount function models				
Cubic spline	0.07395	0.02156	0.00074	6
Exponential	0.07640	0.02188	0.00075	6
Vasicek Fong	0.14421	0.05261	0.00215	6
Polynomial	0.09293	0.03300	0.00125	6
CIR	0.69467	0.15910	0.00600	4
Cubic spline	0.14966	0.05716	0.00236	4
Panel B: Spot yield models				
Cubic spline	0.09436	0.02487	0.00074	6
Svensson	0.10553	0.02665	0.00077	6
Polynomial	0.13571	0.03220	0.00091	6
Vasicek	0.26594	0.05392	0.00161	4
Polynomial	0.29291	0.06244	0.00190	4

Figure 1: Estimated Yield Curves
Cubic spline



Extended Svensson






Figure 1: Estimated Yield Curves (cont.)
Cubic spline



Extended Svensson



More specifically, the spline based models appear to fit to the typical shapes of the market data better than either the exponential model or the extended Svensson model, particularly so at the short end, although the extended Svensson model provides a very good fit to the long end of a hump-shaped yield curve. This suggests that while, on average, the exponential model and extended Svensson model may provide as good a fit as the spline model, they are less well able to represent localized shape characteristics, particularly at the short end. The overall recommendation of this study would be the use of a spline model fitted directly to the yield curve, with knots either optimized over a prior dataset or set at market-based or quantile maturity points.¹⁸

5 Conclusions

This paper has used a new dataset of zero-coupon bond prices to compare the ability of alternative models for estimating the yield curve. The use of zero-coupon bonds prices permits a purer comparison of models than using coupon bonds, since the models are not required to make estimation feasible but only to interpolate the curves. Several popular models are compared, using measures of price and yield errors and smoothness, and models are fitted to both the discount function and also directly to the zero-coupon yield curve.

By contrast to many previous studies, all the models are compared in a form that provides them with the same number of free parameters, so that each model has equal potential to fit to the data. In this way, the comparison is focussing on the ability of the different functional forms to capture the shapes of the yield curves

18. In the appendix, the results of further examination of local fitting, which consider the sensitivity to parameter selection choices, are reported. These show that optimizing the knot points of the spline models tends to improve the fit to the short end of the curve.

present in the data. The dimensionality is set at the smallest number for which the estimated yield curves are no longer statistically too smooth. Smoothness is measured by the correlation of adjacent errors along the yield curve and, when increasing the dimensions from five to six free parameters, there is a marked change in the flexibility of all the yield curve models. So, the models are compared using six free parameters.

The main findings of this study, and the implications for market practitioners and policy makers, are as follows. It appears that fitting directly to the yield curve, rather than fitting to the discount function and then deriving the yield curve, is likely to provide a better fit to the yield curve. It is also argued that directly fitting to the forward rate curve may be unnecessarily complex. For the models fitted here, relatively simple, low parameterized, spline functions provided extremely accurate fits to the sequence of yield curves. Although, other models using combinations of exponential-style functions, performed similarly well within the sample, out-of-sample tests suggested that spline functions would be more reliable to price new issues. The spline functions were also more able to fit to local maturity regions, particularly the short end of the yield curve, and were surprisingly robust to alternative knot spacings. It was also found that even if the knot positions were set extremely haphazardly, spline functions could outperform some other models, such as simple polynomials.

6 References

- Adams, K., and D. Van Deventer (1994) Fitting yield curves and forward rate curves with maximum smoothness, *Journal of Fixed Income*, June, 52-62.
- Anderson, N., (1994) Testing models of the yield curve, Bank of England mimeo.
- Anderson, N., F. Breedon, M. Deacon, A. Derry and G. Murphy (1996) *Yield Curve Estimation and Interpretation*, (Chichester: John Wiley).
- Anderson, N., and J. Sleath (2001) New estimates of the UK real and nominal yield curves, Bank of England working paper.
- Bjork, T., and B. Christensen (1999) Interest rate dynamics and consistent forward rate curves, *Mathematical Finance*, 9, 323-348.

- Bliss, R., (1997) Testing term structure estimation methods, *Advances in Futures and Options research*, 9, 197-231.
- Brown, R., and S. Schaefer (1994) Term structure of real interest rates and the Cox, Ingersoll and Ross model, *Journal of Financial Economics*, 35, 3-42.
- Brown, S., and P. Dybvig (1986) The empirical implications of the Cox, Ingersoll and Ross theory of the term structure of interest rates, *Journal of Finance*, 41, 617-630.
- Carleton, W., and I. Cooper (1976) Estimation and uses of the term structure of interest rates, *Journal of Finance*, 31, 1067-1083.
- Chambers, D., W. Carleton and D. Waldman (1984) A new approach to estimation of the term structure of interest rates, *Journal of Financial and Quantitative Analysis*, 19, 233-252.
- Coleman T., L. Fisher and R. Ibbotson (1992) Estimating the term structure of interest rates from data that include the prices of coupon bonds, *Journal of Fixed Income*, September, 85-116.
- Cox, J., Ingersoll, J., and S. Ross (1985) A Theory of the Term Structure of Interest Rates, *Econometrica*, 53, 531-552.
- Cummins, C., and B. Hall (1999) TSP4.5 Reference Manual, (TSP International, Stanford).
- Deacon, M. and A. Derry (1994) Estimating the term structure of interest rates, Bank of England working paper series, no. 23.
- Diaz, A., and F. Skinner (2001) Estimating corporate yield curves, *Journal of Fixed Income*, Sept., 95-103.
- Eom, Y., M. Subrahmanyam and J. Uno (1998) Coupon effects and the pricing of Japanese government bonds: an empirical analysis, *Journal of Fixed Income*, Sept., 69-86.
- Fama, E., and R. Bliss (1987) The information in log-maturity forward rates, *American Economic Review*, 77, 680-692.
- Ferguson, R., and S. Raymar (1998) A comparative analysis of several popular term structure estimation models, *Journal of Fixed Income*, March, 17-33.
- Fisher, M., D. Nychka and D. Zervos (1995) Fitting the term structure of interest rates with smoothing splines, Federal Reserve Board, Finance and Economics Discussion Paper, 95-1.
- Heath, D., R. Jarrow and A. Morton (1992) Bond pricing and the term structure of interest rates: A new methodology, *Econometrica*, 60, 77-105.
- Ho, T. and S. Lee (1986) Term structure movements and pricing interest rate contingent claims, *Journal of Finance*, 41, 1011-29.
- Hull, J., and A. White (1990) Pricing interest rate derivative securities, *Review of Financial Studies*, 3, 573-92.
- Ioannides, M., (2003), A comparison of yield curve estimation techniques using UK data, *Journal of Banking and Finance*, 27, 1-26.
- Kikugawa, T., and K. Singleton, (1994) Modeling the term structure of interest rates in Japan, *Journal of Fixed Income*, Sept., 6-16.
- Langtieg, T., and C. Smoot (1989) Estimation of the term structure of interest rates, *Research in Financial Services*, 1, 181-222.
- Lin, B-H., (2002) Fitting term structure of interest rates using B-splines: the case of Taiwanese Government bonds, *Applied Financial Economics*, 12, 57-75.
- Lin, B-H., and D. Paxson (1993) Valuing the "new issue" quality option in Bund futures, *Review of Futures Markets*, 12, 349-88.

- Lin, B-H., and D. Paxson (1995) Term structure volatility and Bund futures embedded options, *Journal of Business Finance and Accounting*, 22, 101-27.
- Litzenberger R., and L. Rolfo (1984) An international study of tax effects on governments bonds, *Journal of Finance*, , 1-22.
- Mastronikola, K., (1991) Yield curves for gilt-edged stocks: a new model, Bank of England Discussion Paper (Technical Series) 49.
- 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, 811-29.
- McCulloch, J.H., (1977) The cumulative unanticipated change in interest rates: Evidence on the misintermediation hypothesis, NBER Working Paper #222, December 1977.
- McCulloch, J.H., (1999) Long forward and zero-coupon rates indeed can never fall, but are indeterminate: A comment on Dybvig, Ingersoll and Ross, unpublished manuscript.
- McCulloch, J.H., (2000) The "Bellwether" 30-year Treasury Bond is an exceptionally bad investment, February 2000, unpublished manuscript.
- McCulloch, J.H., and L. Kochin (2000) The inflation premium in the US real and nominal term structures of interest rates, Ohio State University Economics Department Working Paper #98-12.
- Nelson and Siegel (1987) Parsimonious modelling of yield curves, *Journal of Business*, 60, 473-489.
- Proudman, J., 1995, The microstructure of the UK gilt market, Bank of England Working Paper, no.38.
- Rumsey, J., (1996) Comparison of tax rates inferred from zero-coupon yield curves, *Journal of Fixed Income*, March, 75-81.
- Shea, G. (1984) Pitfalls in smoothing interest rate term structure data: equilibrium models and spline approximations, *Journal of Financial and Quantitative Analysis*, 19, 253-269.
- Shea, G. (1985) Interest rate term structure estimation with exponential splines: a note, *Journal of Finance*, 40, 319-325.
- Schaefer, S. (1981) Measuring a tax-specific term structure of interest rates in the market for British government securities, *Economic Journal*, 91, 415-438.
- Steeley, J., (1991) Estimating the gilt-edged term structure: basis splines and confidence intervals, *Journal of Business Finance and Accounting*, 18, 513-529.
- Subramanian, K., (2001) Term structure estimation in illiquid markets, *Journal of Fixed Income*, June, 77-86.
- Svensson, L., (1994) Estimating and interpreting forward interest rates: Sweden 1992-1994, IMF working paper no.114.
- Vasicek O., (1977) An Equilibrium Characterisation of the Term Structure, *Journal of Financial Economics*, 5, 177-188.
- Vasicek O., and H. Fong (1982) Term structure modelling using exponential splines, *Journal of Finance*, 37, 339-356.
- Vitale, P., 1998, Two months in the life of several gilt-edged market makers on the London Stock Exchange, *Journal of International Financial Markets, Institutions and Money*, 8, 299-324.
- Waggoner, D., Spline methods for extracting interest rate curves from coupon bond prices, Federal Reserve Bank of Atlanta, working paper 97-10.
- Wiseman, (1994) European Fixed Income Research: Second Edition Technical Specification - The Exponential Yield Curve Model, mimeo, JP Morgan.