Long-Horizon Yield Curve Forecasts: Comparison of Semi-Parametric and Parametric Approaches

Ken Nyholm* and Riccardo Rebonato+
*European Central Bank,
Risk Management Division,
Strategy Unit
+Royal Bank of Scotland, Quantitative Research Group
Oxford University (OCIAM)
Imperial College, Tanaka Business School, London

August 13, 2007

Abstract

Two methods for evolving forward the yield curve are evaluated and contrasted within a Monte Carlo experiment: one is originally presented by Rebonato, Mahal, Joshi, Bucholz, and Nyholm (2005) and the other by Bernadell, Coche, and Nyholm (2005). A detailed account of how to implement the models is also presented. Results suggest that the two techniques are complementary and able to capture important cross sectional and time series properties of observed yield curve data. A joint examination of the results of the two models also gives interesting indications about the nature of the transition from one regime to the next. Our results are of interest to practitioners in the financial markets who need accountable and history consistent procedures for generating long-term yield curve forecasts.

1 Introduction

1.1 Motivation and Goals of the Paper

Many financial problems require that explicit views are formed on the shape and level of the whole yield curve over the medium-to-long term. Some examples are:

- asset and liability management in corporations, banks and pension funds require a coherent model to evaluate the cash-flows from assets and claims over long time horizons in a consistent manner;
- strategic asset allocation decisions, which by definition stretch beyond the short-term horizon, generally require long-term expected returns from

the portfolio assets as inputs. Key elements in generating such expected return forecasts are the shape and location of the yield curve and the term structure of risk premia for the relevant investment universe (e.g. spread products and equity);

- counterparty credit risk assessment requires the evaluation of the future values of derivatives transactions, conditional on yield curves that will prevail a long time in the future;
- testing the effectiveness of derivatives pricing models require a realistic long-term evolution at least of the yield curve.

This paper explores and contrasts two models that are relevant for the simultaneous evolution of yields of several maturities: one is the model by Rebonato, Mahal, Joshi, Bucholz, and Nyholm (2005), (for brevity referred to as the RMJBN model in the following), and the other is the model by Bernadell, Coche, and Nyholm (2005), (referred to as the BCN model). These approaches are among the very few that have been specifically designed to produce a truly multi-factor evolution the whole yield curve in the real-world measure over a forecasting horizon relevant for long-term financial decision making (ie, several years).

These two models lie, in a way, at opposite ends of the modelling spectrum, in that one assumes as little as possible about the underlying dynamics of the yield curve, while the other is structurally very 'rich'. The RMJBN model, in fact, is a semi-parametric approach built on a window-sampling technique of observed differences in yields. Its only 'structural' feature is the presence of a series of spring functions designed to replicate the behaviour of arbitrageurs in the market. The BCN approach, on the other hand, starts from a strong *Ansatz* about the nature of the yield curve dynamics, embodied by the specification of a three-state Markov chain process for the yield curve driver(s) and it relies on a parametric regime-shifting factor-description for the yield curve. Looked at in this manner, the two approaches are therefore almost 'orthogonal' to each other, and any common conclusion about complex features of the yield curve dynamics arrived at from such different starting points deserves careful examination. Indeed, without pre-empting future results, we will show, for instance, that the RMJBN model, that 'knows nothing' about discrete-state Markov chains 'discovers' state transitions and persistence patterns similar to those predicted by the BCN approach.

In the light of these considerations, the purpose of this paper is

- to give a unified presentation of the two models in terms of how they can be applied;
- to compare their strengths and weaknesses, to establish a common footing for the two models and to establish to what extent, despite their different starting point, they actually concur in their predictions; and

- to conduct a Monte Carlo experiment in order to assess the robustness of the parameter estimation and the parameter stability in the BCN approach by using the RMJBN model as the data-generating engine.

The second and third bullet points deserve some further elaboration. Starting from the second, since the models are designed to generate medium-to-long-term yield curve scenarios it is relevant to investigate whether the models statistically (and 'optically'[1]) match observed data. To this effect the paper explores, *inter alia*, whether, and to what extent, the models are capable of producing truly different shapes of the yield curve and whether they adequately capture several important statistical properties of observed yield data.

As for the third bullet point, models that require parameter estimation are usually 'trained' on historical data sets as long as available. It is always difficult, however, to test whether the available historical record is of sufficient length to determine the model parameters in a reliable manner. Since the observed history is but one realization of a population of possible 'universes' that might have occurred, when estimating the parameters (at least of stationary processes) one is implictly making an ergodic assumption, and replacing the unavailbale ensemble averages with the available time averages. The question then naturally arises: "How much time-serial data is necessary in order to unlock truly representative statistical features of the underlying population (ie, for the ergodic theorem to apply)?". This question is particularly pertinent when a model, such as the BCN, incorporates regime switches.

An appealing way to answer this question is by decreeing *by fiat* that a second model represents 'reality', and by 'playing God', ie, by creating those simultaneous, equi-probable parallel realizations of the trading universe that are unavailable in reality. Clearly, the exercise is only interesting if 'God's model' is sufficiently rich to be a reasonable representative of the real world: if the 'God's model' were, say, a joint diffusion based on PCA, there would be little point in testing whether the parameters of a regime-switching model can be reliably estimated. The results reported in Rebonato, Mahal, Joshi, Bucholz, and Nyholm (2005), however, make us believe that the RMJBN model is sufficiently rich to be a plausible candidate. We therefore use it as 'God's model' in our entry-level test.

A second, more interesting, question, that at the same time can provide some degree of cross-corroboration of 'God's model', is whether the parameters estimated over the population of simulated realizations are reasonably similar to those estimted (time-serially) using *the* one realization of the real world. This is clearly a much taller order, and we address this question as well.

The rest of the paper is organized in the following way. Section 2 present a survey of the literature. The following section introduces the two models and describes how they can be applied to evolve the yield curve under the empirical

[1]Rebonato et al (2005) show that some statistically plausible term structure models produce future yield curves that fail to resemble anything ever observed in the real world. This 'optical' analysis is a useful pointer to uncover deeper statistical inadequacies of the naive modelling approach.

measure. Section 4 elaborates on the empirical experiments conducted, section 5 describes the results and section 6 concludes the paper.

2 Relevant Results from the Literature

Beside the approaches investigated in this paper, there exist other types of models that can be used evolve the yield curve over long horizons. The field is vast, and therefore we do not attempt a systematic review. Instead we mention those strand of research that can put the models examined in this paper in sharper focus. The first is the class of models formulated under the risk-neutral measure ("term-structure" models); the second is the modelling approach based on PCA (or, more recently, on ICA); the third is the class of Vector-Auto-Regression-based (VAR-based) approaches; the last deals with explicitly regime-switching[2] models.

2.1 Modelling the Yield Curve in the Risk-Neutral Measure

Starting from the first class of models, there exist a number of relative-pricing models that evolve the yield curve (in the risk-neutral measure) over as long a time horizon as desired. Two sub-classes should be distinguished: models that make use of a joint specification of the risk premia and of the real-world dynamics ('fundamental' models), and models that directly describe the evolution of the yield curve in the risk-neutral measure ('reduced-form' models[3]). See Hughston and Brody (2000a) for an excellent unified treatment of the two.

2.1.1 Fundamental Models

Models belonging to the first ('fundamental') sub-class of pricing models separately posit a process for the driving factor(s) of the yield curve and a functional form for the market price of risk, which is ultimately linked to the risk aversion of the agent. The Vasicek (1977), Cox, Ingersoll, and Ross (1985), Longstaff and Schwartz (1992), Brennan and Schwartz (1982) are early examples of models that fall in this category. Market prices of options and bonds do not provide direct information about the parameters of the real-world yield curve dynamics (so, for instance, the reversion speeds that enter the Vasicek and CIR models are risk-neutral and not real-world speeds). The translation from the risk-neutral

[2]The term 'explcitly regime-switching' refers to those models where regime switching is explictly built into the model itself. As we shall see, however, also the RMJBN model displays features consitent with regime-switching, even if this feature is not imposed *a priori* from the outset.

[3]The term 'reduced form' is often used in the credit-derivatives literature to denote models that dispense with the specification of the evolution of the form assets, and deal directly with the probability of defualt. We do not imply such a technical meaning in this paper, but we still refer to models that by-pass a higher-level description (in our case, of the real-world dynamics *and* of the investors' preferences) and directly model the risk-neutral measure.

to the real-world drift could in principle be effected by taking into account the contribution to drift-related quantities stemming from risk aversion. This route is explored from the theoretical point of view in Hughston (2000), Hughston (2003) and Hughston and Brody (2000b). The estimation of the utility function that underlies these approaches is, however, notoriously difficult and the results tend to be more of qualitative than quantitative interest.

There is a more fundamentally problem with these approaches, that makes unsuitable for our purposes (see Section 1): since these models have been designed to price derivatives securities, the modelling choices of the real-world dynamics of the yield curve tend to be dictated by computational considerations: one-factor models, which imply approximately parallel moves of the yield curve, for instance, are common and the diffusive assumptions for the driver(s) is almost universally made. For these reasons, 'fundamental' pricing models are not appropriate for the applications highlighted in Section 1.

2.1.2 Reduced-Form Models

The specification of reduced-form pricing models is quite different, but also often dictated by considerations of computational speed rather than econometric realism (computationaol spped is a must because these models are often used by financial institution for the computationally intensive tasks of pricing and hedging entire books of complex transactions). Their realism of these models is therefore rather limited (while second-generation models are often multi-factor, they rarely move beyond the diffusive set-up). See Rebonato (2003) for a recent survey of diffusion-based reduced-form pricing models, Jamshidian (1997), Glasserman and Kou (2003) for extensions beyond continuous semi-martingales.

Whatever their strengths and weaknesses, the reduced-form pricing models are constructed for relative-pricing purposes, and inhabit the pricing rather than the objective measure. They therefore contain drift terms purely dictated by no-arbitrage considerations. These drift terms, perfectly appropriate for the relative-pricing application for which these models have been devised, drive a wedge between the real-world and the risk-adjusted evolution of the yield curve. Over long horizons, the evolutions produced by these approaches becomes totally dominated by the no-arbitrage drift term, and bear virtually no resemblance with the real-world evolution.

For the purposes stated in Section 1, reduced-form models may therefore be sometimes richer and more complex than fundamental models. The difficulty to disentangle and specify exogenously the risk-adjustments needed to recover the real-world evolution of the yield curve makes them, however, also unsuitable for our purposes.

2.2 Modelling the Yield Curve in the Objective Measure

2.2.1 PCA (ICA) Modelling

Moving to direct modelling in the objective measure, one the the most common modelling strands has been the use of Principal Component Analysis (PCA), or, more recently, of Independent Component Analysis. In both cases, a linear transformation of the 'original' correlated variables describing the yield curve (eg, yields or forward rates) is carried out to obtain new variables which have the property of being uncorrelated (PCA), or independent (ICA). The number of works based on PCA is vast (see, eg, Martellini and Priaulet (2001) for a recent review) and the salient results are the decomposition of the mode of deformation of the yield curve in terms of shifts in its level, slope, curvature, etc, and the fact that a small number of principal components appears to have a high explanatory power in accounting for changes in the yield curve.

Because of this latter property, when it comes to the real-world modelling of the yield curve the approach has often been taken to estimate the first few eigenvectors and the associated eigenvalues of the covariance matrix of yield changes; to assume that the processes of the 'original' variables are joint diffusions with deterministic volatility and identical and independent increments; and to simulate the evolution of the yield curve by assigning independent shocks to the first few principal components.

It is widely recognized that such a procedure only produces an approximation to the real-world evolution of the yield curve. However, the greatest shortcoming usually levelled at this approach is the fact that only a few principal components are typically retained. Rebonato, Mahal, Joshi, Bucholz, and Nyholm (2005) show that, *even if all the principal component were retained*, the resulting synthetic yield curves soon assume shapes not observed in reality. In particular, the distributions of yield curve curvatures obtained using PCA are markedly different from what observed in reality. We are not aware of similar investigations carried out using ICA, but this approach, which ensure *cross-sectional* (not serial) independence of the increments, is unlikely to change this feature. The RMJBN approach described below, by construction, not only asymptotically recovers exactly all the eigenvectors and eigenvalues from the original data; but also produces a very satisfactory recovery of the distribution of curvatures across the yield curve. This can be important for the applications mentioned in the opening section.

2.2.2 VAR Modelling

A different approach to the modelling of the evolution of the yield curve rests on a Vector Auto-Regressive (VAR) specification of changes in the short rate and in its spreads to longer rates. See, for example, Campbell and Shiller (1987), Hall, Anderson, and Granger (1992), Kugler (1990) and Lanne (2000). In this modelling framework rates are assumed to be cointegrated since changes in the short rate and the spreads to longer rates are treated as stationary variables. Naturally, nominal rates cannot be pure I(1) processes because they are bounded at zero and cannot go to infinity. The reason for modelling rates within a cointegrated system of variables is that the persistence pattern in the time-series evolution of the rates implied by the I(1) processess is better approximated than when rates are modelled as stationary variables. Applications of this modelling framework to interest rates has often been used to test the expectation hypothesis, but they are in principle equally well applicable to the prediction of rates. In this literature, however, applications are frequently limited to bivariate systems, ie, to the simultaneous evolution of two rates. While this can be adequate for a qualitative description of the evolution of yield curves, it is probably too restrictive for the applications mentioned in the introductory section. Although some studies falling in the VAR category of models do account for regime-switching (see e.g. Driffill and Sola (1994) and Engsted and Nyholm (2000)) the majority of the VAR papers model just one regime.

As a final comment on VAR modelling, we would like to point out that, of the VAR-based approaches referred to above, those which have been used to test the expectation hyothesis explictly assume that the yield curve embeds information about the future, ie, that agents trade today on the basis of their expectations of the future (and that agents have similar beliefs). While this assumption is reasonable and consistent with the general assumptions of asset pricing theory, its validity is not strictly required by the BCN and RMJBN approaches. Instead, BCN assume a contemporaneous relation between the state of the macroeconomy and the yield curve through a mechanism similar to the Taylor rule, and, because of its semi-parametric nature, the RMJBN methodology does not require any such assumption (at most, the RMJBN model is compatible with, but does not require, a particular microstructuraal interaction between noise traders and pseudo-arbitrageurs - see the discussion in Section 3.1). Naturally, the RMJBN and BCN approaches do not preclude the existence of a relationship between the yield curve and the state of the macroeconomy - it is just not a necessary assumption.

2.2.3 Regime-Switching Models

Numerous papers have shown that it is important to include regime shifting in the modelling of interest rates. By now substantial empirical evidence exists in favour of interest rates being modelled by non-linear techniques (see, among others, Stanton (1997), Aït-Sahalia (1996), Conley, Hansen, Luttmer, and Scheinkman (1997)). One argument against regime switches still remains, though; it can be argued that the econometrician resorts to this modelling approach in lack of a properly specified (non-regime-switching) model. This line of reasoning suggests that regime switching models are reduced-form representations of fuller specified economic models. However, with a reduced-form model comes parsimony, which in terms of model stability is advantageous. Hence, for practical modeling purposes 'regime switching' seems to constitute a relevant building block when it comes to capturing yield curve dynamics in the maturity and time-series dimensions. In this vein, for example, Hamilton (1988) shows that a univariate autoregressive model without regime switches is inconsistent with the time-series behaviour of the short rate in the US, while a model including regime-switches much better describes the time-series evolution of the analysed data series. Furthermore, analysing jointly *ex post* real interest rates and inflation, Garcia and Perron (1996) show the importance of regime shifts in the mean and variance of the modelled series: allowing for regime switches makes the real rate behave like a random walk (which is an important property, otherwise shocks to the real rate would only have temporary effects and the real rate would exhibit mean reversion). More recently a series of papers that fall within the affine class of models (see Duffie and Kan (1996)) also facilitate the incorporation of regime-switches, see for example, Ang and Bekaert (2002), Ang and Bekaert (2004), Banzal and Zhou (2002) ,Banzal, Tauchen, and Zhou (2003), Dai, Singleton, and Yang (2003) and Evans (2003). Ang and Bekaert (2002) use a model with two latent factors and one observable variable (inflation) to seperate nominal yields into the a real rate, the expected inflation and the inflation risk premium. Regime switches are incorporated in the means and variances of one of the latent factor, in the process for inflation and in the risk premium. The regimes are identified on the basis of the level and volatility of the real rate and inflation. On the dataset used they find that 77% of the observations fall in the regime characterised by a low and stable real rate combined with a high and stable inflation; 10% of the observations fall in the regime characterised by a high and stable real rate in combination with a low and stable inflation rate; 11% of the observations fall in the regime characterised by a low and volatile real rate and high and volatile inflation; while the remaining 2% fall in the regime of a high and volatile real rate combined with a low and volatile inflation rate.

Banzal and Zhou (2002) develop a term structure model where the short rate and the market price of risk are regime switching. In particular, the reversion speed, the level, and the volatility of the short rate process are allowed to exhibit regime switching behaviour. They show that the conditional joint dynamics of short and long rates are captured much better when regimes are allowed for in the modelling framework than when they are not. Banzal, Tauchen, and Zhou (2003) use a model setup that is identical to the one used by Banzal and Zhou (2002), but demonstrate further the merit of a modelling framework that incorporates regime switching. For US nominal yields they show that the estimated regimes are closely related to the business cycle as well as to the risk premium.

Evans (2003) use a model that is similar in spirit to Banzal and Zhou (2002) in that the short rate process incorporates regime switches in the reversion level and speed and that inflation and the risk premium are regime switching. The paper shows that three interest rate regimes can be identified on the basis of the shape of the yield curve (both real and nominal). The identified yield curve shapes are sharply upward sloping, upward sloping and inverse. Evans also shows that time-varying term premia and inflation risk premia contribute significantly to the changes in the dynamics of the spot rate and expected inflation. However, Evans (2003) does not link the identified yield curve shapes to the state of the macro economy. Finally, Dai, Singleton, and Yang (2003) use an arbitrage-free dynamic term structure model with regime switches in the variance term to illustrate the importance of letting the transition probabilities be regime-dependant.

The regime-switching model presented in Bernadell, Coche, and Nyholm (2005), and which is analysed in the current paper, does not build on the arbitrage-free modelling framework as the above mentioned literature does. Instead, yields are modelled directly under the empirical (objective) measure following a regime-switching expansion of Diebold and Li (2006). However, many of the empirical results cited above are also found by Bernadell, Coche, and Nyholm (2005). Diebold and Li (2006) build an intuitive model that captures both the time-series evolution of yields as well as their cross sectional relation (i.e. the relation between yield and maturity). The framework rests on a parametric factor model capturing the shape and location of the yield curve at a given point in time. To this end the three-factor model of Nelson and Siegel (1987) is used. The yield curve factors are then modelled in the time-series dimension using a one-lag VAR model. To estimate the model Diebold and Li (2006) use a two step approach where, in a first step, the yield curve factors for each date in the data sample are estimated. In a second step the time-series evolution of the yield factors are captured through a VAR(1) model. In the concluding section of the paper the authors suggest to estimate the factors and their time-series dynamics jointly using a state space model. This approach is followed in Diebold, Rudebusch, and Aruoba (2006) and Bernadell, Coche, and Nyholm (2005).

As mentioned above, the regime-switching model implemented in the present paper follows closely Bernadell, Coche, and Nyholm (2005); however, here we restrict the transition probabilities to be constant over time, as opposed to being regime-dependant as in Bernadell, Coche, and Nyholm (2005). BCN denote the three estimated yield curve regimes according to the size of the slope factor: the curve is very steep in one state (this corresponds to a highly negative slope estimate); in another state, the curve is inversely shaped (or flat) which corresponds to slightly positive (or close to zero) slope estimate; in the third state the curve is upward sloping but not as steep as in the first state. Following an argument similar in spirit to the Taylor rule (Taylor (1993)) the economic interpretation of the three estimated yield curve regimes are as follows:

(A) in periods where the activity of the economy is too low the central bank tends to lower the policy rate to ease the capital costs of firms, in an attempt to stimulate economic growth. A reaction pattern similar to this is likely to translate into a yield curve, which is steeper than normal, even if the long end of the curve falls too;

(B) in periods where the central bank feels that too strong inflationary pressures are present it will react by increasing the policy rate. Even if the long end of the curve reacts the movement in the short rate initiated by the central bank will produce a yield curve, which is likely to be flat or even inverse;

(C) when the central bank does not intervene in the financial markets by controlling the short rate, the yield curve will be upward sloping with a yield spread significantly lower than in case (A).

3 The Models Under Investigation

3.1 The RMJBN Approach

The RMJBN approach is presented in detail by Rebonato, Mahal, Joshi, Bucholz, and Nyholm (2005) and briefly summarized below to give a self-contained presentation. In order to generate history-consistent yield-curve projections, the RMJBN methodology relies on a re-sampling approach coupled with a behavioural assumption on how arbitrageurs exploit mispricings. To be more specific, the underlying data consist of over ten years' worth of yield curves observed daily at a discrete set of maturities[4]. To the starting (current) yield curve 'deterministic' and 'stochastic' innovations are added. The stochastic innovations are determined by

1. pre-choosing a fixed number of days (the 'window length', typically a number between 20 and 50);
2. randomly choosing a past date in the data set;
3. adding sequentially to the current yield curve, maturity by maturity, the (absolute or percentage) changes in the yields from the chosen random day to the end of the window;
4. choosing a new random past day as a new start for the resampling and
5. repeat from 3.

Since each daily vector of synchronous yield-curve changes for different maturities is treated as one entity, this way of adding the stochastic innovation ensures that a full cross section of yield curve changes is sampled at each resampling interval. This ensure the preservation of several important cross-sectional features in the resampled yield curves (see RMJBN for a discussion).

If the daily increments of the yields were iid, this sampling procedure with a window length of one would be enough to produce a statistically satisfactory evolution of the yield curve. However, in practice it is observed that, after a few 'months' of simulated evolution the resampled curves will not even vaguely resemble the shape of historically-observed curves. To produce plausible (and statistically desirable) future yield curves, a deterministic component to the innovation is added every day. This is done by modelling the behaviour of pseudo-arbitrageurs via a series of "spring constants"[5]. More precisely, if the yield at maturity τ_n at time t is significantly below (above) the yields (also at time t) for maturities τ_{n-1} and τ_{n+1}, then pseudo-arbitrageurs will act by receiving the high yields and paying the low yields. By so doing, they will trade

[4] The maturities observed in the RMJBN data set are 3m, 6m, 1y, 2y, 5y, 10y, 20y and 30y.

[5] It must be emphasized that the statistical validity of the approach does not hinge on the correctness of the interpretation of the springs as describing the action of the pseudo-arbitrageurs. Even if this 'picture' were not correct, Rebonato et al (2005) show that the springs correctly recover important features of the real-world yield curves.

away excessive kinks in the yield curve. To capture this intiution, springs are introduced, with varying stiffness across the maturity dimension, such that the yield at maturity τ_n will be pulled up (down) and the yields at maturities τ_{n-1} and τ_{n+1} will be pulled down (up). Springs have increasing stiffness across the maturity spectrum because pseudo-arbitrageurs are more confident in interpreting as arbitrageable 'mispricings' kinks that occur at longer than at shorter maturities (where expectations may play a stronger role).

The resampling methodology suggested by Rebonato, Mahal, Joshi, Bucholz, and Nyholm (2005) is applied as follows. Denote by Y_i the vector of yields observed at time i for a historical sample of $i = (1, 2, \dots, t)$ observations. In the RMJBN approach yields are evolved by using the following two equations.

$$y^{N+M} = y^N + \sum_{r=1,M} \Delta y^{U_r}, \tag{1}$$

$$\widetilde{y}^{N+M} = y^{N+M} + \sum_{r=1,M} k \odot \xi^{N+r}. \tag{2}$$

The yield curve forecast without the contribution from arbitrageurs is given by y^{N+M} where y^N is the initial yield curve. Δy^{U_r} represents the resampled yield curve changes for the U_r sampling length. Thus Δy^{U_r} is a matrix of yield curve changes spanning maturity in one dimension and time-series length in the other dimension, (r denotes the window length). The term $k \odot \xi^{N+r}$ then describes the actions of the arbitrageurs. The quantity k is a vector of spring constants, the symbol $\odot$ denotes the element-by-element multiplication operator and ξ^{N+r} is the vector of curvatures calculated for each curve covered by the forecasting horizon M. (Here M is the number of blocks necessary to cover the desired number of forward-projected yield curves.)

Rebonato, Mahal, Joshi, Bucholz, and Nyholm (2005) describe in detail the desirable statistical features of the approach. In the present context it is enough to recall

i) the exact asymptotic recovery of all the eigenvalues and eigenvectors of yield changes;

ii) the correct reproduction of the distribution of curvatures of the yield curve across maturities;

iii) the correct qualitative recovery of the transition from super- to sub-linearity as the yield maturity is increased in the variance of n-day changes, and

iv) a satisfactory account of the empirically-observed positive serial correlations in the yield changes.

3.2 The BCN Approach

As mentioned above, the BCN modelling approach can be seen as a regime-switching extension of Diebold and Li (2006), who model the yield curve (under the empirical measure) by applying a three-factor representation of observed yields relying on Nelson and Siegel (1987) and estimated via a two step approach:

in the first step yield curve factors are estimated and, in a second step, the time series properties of these factors are explored using a VAR-model. The BCN model is similar in spirit to this approach with two important modifications: regime-switches are incorporated at the level of the time-series dynamics of the yield curve factors and estimation of yield curve factors and factor dynamics are carried out using a state-space model.

The application of the regime-switching yield-curve model in the current paper is similar to Bernadell, Coche, and Nyholm (2005), however, we omit the functional dependence between macro economic variables and the transition probability matrix. This is because for comparing and contrasting the RMJBN and BCN approaches this link is of secondary importance.

3.2.1 Estimating the Model

In order to provide a self-contained treatment, a brief description of our version of the regime-switching yield-curve model and of the applied estimation technique is given below. As in Bernadell, Coche, and Nyholm (2005) we use the results from Kim and Nelson (1999) to combine Kalman and Hamilton filters.[6] The observation equation takes the following form:

$$Y_t = H\beta_t^j + e_t, \tag{3}$$

where $Y_t(\tau)$ is the observed yield curve at time t for maturities $\tau = (\tau_1, \tau_2, \ldots, \tau_T)$ months, β_t^j denotes the yield curve factors, superscript $j \in (State1, State2, \ldots, StateN)$ and $e \sim N(0, \Sigma)$ is the error term. The term $diag(\Sigma) = (\sigma_{e_1}^2, \sigma_{e2}^2, \ldots, \sigma_{eT}^2)$ contains the variances of the various yields. The H matrix displays the Nelson-Siegel functional form:

$$H = \begin{bmatrix} 1 & \frac{1-\exp(-\lambda\tau_1)}{\lambda\tau_1} & \frac{1-\exp(-\lambda\tau_1)}{\lambda\tau_1} - \exp(-\lambda\tau_1) \\ 1 & \frac{1-\exp(-\lambda\tau_2)}{\lambda\tau_2} & \frac{1-\exp(-\lambda\tau_2)}{\lambda\tau_2} - \exp(-\lambda\tau_2) \\ \vdots & \vdots & \vdots \\ 1 & \frac{1-\exp(-\lambda\tau_T)}{\lambda\tau_T} & \frac{1-\exp(-\lambda\tau_T)}{\lambda\tau_T} - \exp(-\lambda\tau_T) \end{bmatrix} \tag{4}$$

The time-series evolution of the yield curve factors is governed by the state equation, which in the current application takes the form of a regime-switching VAR(1) model. More precisely, the state equation is:

$$\beta_{t|t-1}^j = \begin{bmatrix} \beta_1^j \\ \beta_2^j \\ \beta_3^j \end{bmatrix}_{t|t-1} = \begin{bmatrix} c_1^j \\ c_2^j \\ c_3^j \end{bmatrix} + \begin{bmatrix} a_1 & 0 & 0 \\ 0 & a_2 & 0 \\ 0 & 0 & a_3 \end{bmatrix} \begin{bmatrix} \beta_1 \\ \beta_2 \\ \beta_3 \end{bmatrix}_{t-1|t-1} + \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix}_t, \tag{5}$$

where the error term is $v \sim N(0, \Psi)$ and $diag(\Psi) = (\sigma_{v_1}^2, \sigma_{v_2}^2, \sigma_{v_3}^2)$. We denote by F the matrix of autoregressive parameters, ie, $diag(F) = (a_1, a_2, a_3)$, and by $\beta_{t|t} = \left(\beta_{1,t|t}, \beta_{2,t|t}, \beta_{3,t|t}\right)'$.

[6] The implementation of the regime-switching model relies on Hamilton (1994, Ch. 22).

One problem that emerges in the estimation of the above model is how to integrate regime-switches in the iterative Kalman filer algorithm. A naive application would lead the state equation for the regime-switching β^j to grow by a factor N, equal to the number of regimes modelled, at each observation t. This would quickly become unmanageable. Kim and Nelson (1999) show how this 'curse of dimensionality' can be avoided by calculating a weighted average of the N number of β^js at each time t using the regime-switching probabilities as weights. This is outlined below; for further details see Kim and Nelson (1999).

As a consequence of the model set-up, the prediction errors in the Kalman filter are regime-dependent:

$$\eta^j_{t|t-1} = Y_t - H\beta^j_{t|t-1}. \tag{6}$$

However, since the mean specification of the state equation is assumed to capture all regime-switching behaviour, the conditional prediction-error variance can be calculated in the following way:

$$f_{t|t-1} = E\left[\eta^j_{t|t-1}\right] = HP_{t|t-1}H' + R, \tag{7}$$

where $P_{t|t-1}$ is the one-step-ahead predictor for the variance of β, which is:

$$P_{t|t-1} = FP_{t-1|t-1}F' + Q. \tag{8}$$

To complete the Kalman filter procedure, the beta and the variance terms need to be updated. While updating the variance proceeds according to 'canonical' Kalman filter technique, the betas need to take into account the modelled regime switches. As mentioned above, Kim and Nelson (1999) show how to integrate regime-switching into the Kalman filter by calculating weighted averages of β^j at each time t. To this end Hamilton's regime-switching methodology is applied (see, Hamilton (1994, Ch. 22)), and the density for each of the N regimes can be calculated by:

$$l^j_t(\theta) \propto \left|f_{t|t-1}\right|^{-0.5} \exp\left(-0.5\,\eta^j_{t|t-1} f^{-1}_{t|t-1} \eta^j_{t|t-1}\right). \tag{9}$$

Regime-witching probabilities can then be calculated as:

$$\pi_{t|t} = \frac{\pi_{t|t-1} \odot D_t}{I'_1\left(\pi_{t|t-1} \odot D_t\right)}, \tag{10}$$

where, as above, $\odot$ denotes the element-by-element multiplication, D collects the densities $l^j_t(\theta)$ in a column vector of dimension $N \times 1$, and I_1 is a $N \times 1$ unity vector. According to the Hamilton filter these probabilities can be predicted one-step ahead in the following way:

$$\pi_{t|t} = \widetilde{P}\pi_{t|t-1}, \tag{11}$$

with $\widetilde{P}$ being the transition probability matrix. Using (11) the Kalman filter updating steps take the following form:

$$P_{t|t} = P_{t|t-1} - P_{t|t-1} H' f_{t|t-1} H P_{t|t-1} \,, \tag{12}$$

$$\beta^j_{t|t} = \beta^j_{t|t-1} + P_{t|t-1} H' f^{-1}_{t|t-1} \eta^j_{t|t-1} \,, \tag{13}$$

To handle the curse of dimensionality (ie, the fact that the dimension of β^j grows by N at each t), following Kim and Nelson (1999), the 'reduction' in (14) is imposed:

$$\beta_{t|t} = B_{t|t} \pi_{t|t} \,, \tag{14}$$

where $B_{t|t} = \left[\beta^{j=state1}_{t|t} \middle| \beta^{j=state2}_{t|t} |\dots \middle| \beta^{j=stateN}_{t|t}\right]$, i.e. the regime-dependant betas are collapsed into a vector after each iteration. This is the vector that is used in (5). For the empirical application we choose $J = 3$ and we allow regime-switching in β_2 alone, ie in the slope factor.[7]

3.2.2 Model Projections

Yield curve projections can be formed once the parameters of the BCN model have been identified. An additional central ingredient is the matrix of projected regime probabilities, ie, $\pi_{t+s|t}$ for $s = (1, 2, \dots, S)$. The following steps outline how yield curve projections in the BCN model are obtained:

A) Obtain projections for $\pi_{t+s|t}$ for the desired projection horizon $s = (1, 2, \dots, S)$.

B) Use the result of (A) together with $\widehat{B}$ to obtain $\widehat{\beta}_{t+s|t}$ for $s = (1, 2, \dots, S)$ i.e. the projected regime dependant betas for the projection horizon S.

C) Insert the obtained betas into the observation equation using $\widehat{H}$, possibly perturbed by $\widehat{e}$, which could be bootstrapping from the empirical distribution.

Following these steps produces yield curve projection according to the BCN model.

4 The Monte Carlo Experiment

4.1 The Goals of the Analsysis

As mentioned in the introductory section, we intend to use a Monte Carlo experiment to answer several related questions. First, we want to explore more deeply than in Rebonato, Mahal, Joshi, Bucholz, and Nyholm (2005) whether the RMJBN methodology concurs in identifying important features of the yield curve surfaces discovered by other modelling approaches. In particular, we

[7] These choices are based on results from Bernadell, Coche, and Nyholm (2005) and Evans (2003).

want to see whether the RMJBN approach 'knows' about regime switches in the yield curve slope, and what it implies about the auto-regressive parameters in the state equation of the BCN model.

Second, we want to explore the stability of the BCN methodology by investigating whether parameter estimates are stable across simulations. If the parameter estimates of the BCN model obtained with yield curve surfaces generated by the RMJBN technique are sufficiently stable, we take it as an indication that the BCN model, 'right' or 'wrong' as it may be, can at least be robustly estimated.

Finally, we investigate whether the application of the BCN model to the simulated histories produced by the RJMBN method yields stable parameter estimates that conform with the parameter estimates obtained from real-world data (as estimated in Bernadell, Coche, and Nyholm (2005)). If this is the case, and given the very different nature of the two modelling approaches, we would take this as a strong indication that both thechiques are able to generate yield curve surfaces equivalent in important ways to the observed ones. Naturally, our test *by itself* can only provide evidence of the consistency of the two methodologies. However, when the radically different starting points are taken into account, one is naturally led from cross-corroboration to at least partial validation of both models with respect to real-world data.

4.2 The Methodology

In order to carry out the Monte Carlo experiment used to compare the RMJBN and the BCN modelling approaches, we proceed in two steps. The semi-parametric RMJBN methodology is first used to create, via re-sampling, history-consistent synthetic future yield-curve surfaces. Then, in a second step, following the principles outlined by BCN, the estimation is carried out. To be more precise: nominal US yield curve data are used, covering the period from 1-Jan-1986 to 7-May-2004 and observed at a daily frequency. Yields are observed at maturities (in months) $\tau = \{3,\ 6,\ 12,\ 24,\ 36,\ 48,\ 60,\ 84,\ 120\}$. The data therefore spans 4,788 time series observations for 8 maturities. In accordance with RMJBN the following steps are performed:

- Percentage (relative) yield curve changes are calculated using the observed data.
- A starting yield curve is chosen randomly from the data set of historical yield curves.
- Samples are drawn from relative yield curve changes by:
    - drawing blocks with a minimum length of 5 and a maximum length of 50 observations. For any given draw there is a 5% probability of exiting the block before the end of the current re-sampling[8];

[8]Here we deviate slightly from RMJBN by introducing a minimum block length and by using a slightly higher maximum number for the length of the block.

  - draws comprise all changes across the maturity spectrum i.e. one draw consists of relative yield curve changes for all eight maturities at a given time t;
  - 4,700 daily yield curve changes are drawn for each generated yield curve surface.

- The resampled yield changes are added successively to the starting yield curve.

- The 'springs' are applied to the resampled yield curve data as described in section 3.1, using the spring constants suggested in Rebonato, Mahal, Joshi, Bucholz, and Nyholm (2005): $\{0.10; 0.125; 0.125; 0.20; 0.20; 0.20\}$.

Performing the steps above generates a yield curve surface with a daily sampling frequency for the entire projection horizon. Once these future synthetic yield curves have been created, we move to the second (estimation) step. This is carried out as follows.

- To limit the estimation time, the daily resampled yield curve surfaces are converted to a monthly observation frequency by collecting every twentietth observation.[9]

- Some heuristics are applied to limit estimation time by trying to ensure a broad correspondence between the resampled yield curve surface and the regime-switching model. These heuristics are:

  - The average of the long rate should be between 2% and 15%.
  - There is a minimum of 100 basis points between the average of the slope and its minimum and maximum values.
  - The two most extreme yield curves produced by the resampled data cannot be inverse by more than 400 basis points and steeper by more than 700 basis points.

- The regime-switching state-space yield curve model is estimated on the regenerated data consisting of 235 monthly observations.

To complete the Monte Carlo experiment the above mentioned resampling-estimation steps are repeated 1,000 times. The parameter estimates, regime classifications and residuals from each of the 1,000 repetitions are analysed in the next section.

[9] Experiments show that there are no qualitative difference between results derived from daily and monthly data. However, a significant difference in calculation time is observed. On a PC (3.2 GhZ, 1 MB ram) it takes approximately ten minutes to obtain convergence when 235 time-series of monthly observations are used and more than three hours when the 4,700 daily observations are used.

5 Empirical Results

A number of stylised facts emerge from Bernadell, Coche, and Nyholm (2005) who analyse US yield curve data covering the period 1953 to 2004. In perticular, these authors find distinct regime classifications, with an almost binomial behaviour for the regime probabilities $\pi_{t|t}$. As can be seen from (11) (i.e. our parametrization of the model, which is similar to the one applied by Bernadell, Coche, and Nyholm (2005)), $\pi_{z|z}$ is a N-by-1 vector containing the probability that the zth observation is generated by either of the N hypothesised regimes. A binomial behaviour for these probabilities means that one entry in π tends to dominate the others (by being close to unity) at each 'time' z in the sample. Naturally, this does not imply that one regime dominates the other regimes at all times. In addition, regimes are found to exhibit a high degree of persistence.

The actual degree of persistence can be seen from the diagonal elements of the estimated transition probability matrix. To fix ideas, Figure 1 shows regime probabilities obtained from a randomly chosen sample from the Monte Carlo experiment we conduct.

[Figure 1 about here]

In this particular example, which uses resampled data, three regimes are identified and the behavior of π seems to fit the well with the patterns found by Bernadell, Coche, and Nyholm (2005). A clear persistence in the regime classifications is detected after observation 26: first Regime 2 dominates until observation 87, then Regime 1 dominates until observation 201, and finally Regime 3 dominates for the remaining part of the simulated data.

The observed persistence in the regime classifications corresponds well with macroeconomic data, which is believed to evolve in a similar way.

5.1 General Results from the Monte Carlo Experiment

5.1.1 Estimated Regimes

Four possible outcomes exist for the number of identified regimes in each simulation run. Any given resampled yield curve surface can result in three, two, one or zero regimes being identified; the latter corresponds to instances where the estimation algorithm fails to converge to an optimal value for the likelihood function. A given regime is classified as being effective if the corresponding probability (ie, π from [11]) is higher than 70% . This definition is applied for practical purposes and the results are not sensitive to the precise choice of the cut-off level. Table 1 shows the results across the 1,000 simulations: three regimes are effective in approximately 51% of the cases, two regimes are effective in approximately 46% of the cases and the remaining 2% of the cases only one regime is found. The used maximum likelihood estimation technique converges for all 1,000 simulated paths.

[Table 1: Summary statistics, about here]

We find these results to lend cross-corroboration for both the RMJBN and the BCN modes. It is somewhat surpising that the RMJBN technique should imply, via the BCN parameter estimation, this feature of clear regime identification. Recall that the resampling scheme applies to blocks of daily yield-curve time-series differences while the yield curve regimes are estimated and identified at the level of the yields themselves using a monthly sampling frequency. The results above indicate that the relative short (expected) sampling window of slightly over fifteen daily observations does not imply a similar duration for the yield regimes. The average duration of resampled yield regimes will depend entirely on how yield curve changes are mapped into yields and thus to yield regimes. This mapping process is unknown but seems to be replicated well by the RMJBN technique.

The surprising observation that the RMJBN model (with its short sampling window) can recover regimes that persist for periods much longer than the sampling windows can be explained if the transitions from one regime to another are both rare and sudden: if this were the case, in fact, even a relatively short data window of resampling can 'catch' the transition phase. Once the transition has occurred, the low probability of a regime switch would then ensure that the same regime perists over a long period of time. This possible explanation finds indirect corroboration in the study of drawdowns in the US$ yield curve. Rebonato and Gaspari (2006) find, in fact, that, in order to explain the observed statistical features of the drawdowns, one is naturally lead to posit two regimes, one made up of small, frequent ('normal') price changes and the other made up of rare, large and positively correlated ('exceptional') time changes. Related evidence of sudden transitions between regimes is provided in Rebonato and Kainth (2004), Rebonato (2006), who examine the US$ swaption market, and Rebonato and Chen (2007) who analyze the US$ yield curves.

Table 2, which shows the percentage share of the generated regimes substantiates further the applicability of the RMJBN approach. The methodology also seems to generate occurrence frequencies for the regimes that fit well with economic intuition[10]. Only 2% of the yield curves could not be classified according to the applied algorithm; 50%, 35% and 13% are generated/estimated to fit with classifications of main, steep and flat scenarios, respectively.

[**Table 2: Mean of the estimated regime probabilities, about here**]

The accuracy of the estimated regime classifications is explored in Table 3, by showing the mean of π conditional upon regime classifications. A pattern similar to that found by Bernadell, Coche, and Nyholm (2005) is detected: the average conditional regime probability is close to unity when the corresponding regime is effective and otherwise it is close to zero. Furthermore, little dispersion around these averages are seen, as indicated by the low standard deviations. This means that the algorithm assigns regime classifications with high certainty

[10]However, these numbers seems to be biased when compared to estimates based on the full sample, ie the original data covering 1986 to 2004. Here the occurence frequency of the regimes are approximately 32%, 40%, 28% for main, steep and flat, respectively.

within each of the 1,000 simulations, and thus that the hypothesised mixture density seems to fit data well.

[Table 3: Average number of consecutive observations in each state, about here]

5.1.2 Parameter Values and Stability

The mean values and associated sample standard deviations for the estimated parameters of the regime switching model across the 1,000 simulations are shown in Table 4. Again, the results are encouraging. The estimated regimes exhibit a degree of persistence comparable to that of observed data. All entries on the diagonal of the transition probability matrix $\widetilde{P}$, which measure the probability of staying in a regime once there, are in the area of 0.90. Using the mean values reported in the table suggests that the average duration is 18 months for the main regime, 14 months for the steep regime and 7 months for the flat regime. These average monthly durations clearly exemplifies that the average length of the sampling window applied to daily relative yield curve changes, as described above, does not map directly into a particular pattern of persistence for the estimated regimes.

The obtained average parameter values for the Nelson-Siegel slope factors lend support to the economic interpretation of the estimated regimes used through out the paper. In the regime labelled 'Main' the slope (i.e. the difference between the yield at infinite and zero maturities) is positive and on average takes a value of 227 bp[11]. The steep regime is characterised by upward sloping yield curves, which on average is 436bp; and in the flat regime the yield is on average inverse of approximately 47bp. Some variation is observed around the referred generic slope values as indicated by the shown standard deviations. However, the identified regimes are distant enough to make sense economically.

[Table 4: Parameter stability, mean and standard deviation across simulations, about here]

The remaining three parameters shown in Table 4 concern the non-regime-switching part of the model. First, λ is the time-decay parameter of the Nelson-Siegel functional form determining the effect of the yield curve factors at different maturities. In empirical applications, denoting τ in months, values for this parameter is reported in the region of 0.08 to 0.10. Our average estimate falls well within this range. Second, the level factor, which gives the yield for infinite maturity is quite low, avering 3.6%. However the dispersion around the mean is, as expected, quite large. Third, the curvature factor is on average -1.8 having a dispersion of 2.0.

[11] Note that the Nelson-Siegel form reverts the sign of the slope factor; hence a negative slope factor represents an upward sloping yield curve and, conversely, a positive slope factor represents an inversely shaped yield curve.

5.2 Possible Extensions

In the present paper we have presented two models that fulfil the purpose of evolving forward the whole yield curve for long horizons using two very different modelling techniques. Some extensions to both these frameworks are possible.

The RMJBN method lends itself to other interesting empirical analyses. For example, in addition to producing a single future path for the yield curve it could be envisaged to use the method to produce a full distribution for the future yield curve paths. Such a distribution could be calculated conditional on a specific yield curve environment. For example, if the analyst is of the opinion that the current yield environment is similar to one or more periods observed historically, then a yield distribution can be calculated by applying the resampling technique to selected historical samples rather than to the full data history. Constrained and unconstrained forecasts will thus differ if the distribution of yield curve changes for the selected periods differ significantly from the distribution of yield curve changes over the full data history.

As regards the BCN approach it seems necessary to investigate further the properties of the chosen econometric specification. Table 5 reports in Panel A the properties of residuals from the observation equation (3) and in Panel B the properties of the squared residuals. These results suggest that the econometric properties of the applied model setup might be improved.

[Table 5: Properties of the Estimation errors (Y_hat - Y for all maturities), about here]

The residuals might have a non-zero mean and seem to exhibit quite

strong first order serial autocorrelation. Taken together, this suggests that the three factors implied by the Nelson-Siegel functional form is not flexible enough to fully capture the variability of observed yield curves. In particular, problems about getting the yield level right seems to be present. The observed serial autocorrelation in the residuals is not a problem *per se*. It might suggest misspecification of the econometric model, as also indicated by the non-zero means of the residuals, and has an effect on the standard errors of the parameter estimates of the model. It is, however, not too surprising that serially autocorrelated errors are found. Yield curve movements are themselves serially autocorrelated and it is thus questionable if the AR(1) term incorporated in the state equation (5) is of high enough dimensionality to provide a proper fit to the data. Also, the squared residuals exhibit serial autocorrelation, which indicates that GARCH effects should be modelled explicitly. Taking this into account, one possible expansion of the BCN method could be to use a four or five factor model in the observation equation e.g. following Söderlind and Svensson (1997) and Björk and Christensen (1999), respectively and adding GARCH terms. It is unclear however, if these modelling step would add much to the economic significance of the results.

6 Conclusions

Two methods that can be used to evolve in time the yield curve for several maturities simultaneously are presented and compared in this paper. One is the semi-parametric approach of Rebonato, Mahal, Joshi, Bucholz, and Nyholm (2005) (denoted by RMJBN) and the other is Bernadell, Coche, and Nyholm (2005) (denoted by BCN). We give a detailed account of how to implement the techniques and test them empirically. In a Monte Carlo experiment we show that the two approaches broadly agree. Yield curve surfaces generated by the RMJBN technique exhibit regime-switching behaviour similar to what is found in the observed data. In general, three regimes are detected and reflect the yield curve shape in generic states of the world: one state is equivalent to a recession state where the yield curve is quite steep, another is equivalent to an inflationary state where the curve is flat or even inverted, and the last state is the residual, which can then be called the "normal" state.

In addition to matching the economic intuition yield curve states generated by the RMJBN technique also exhibit an occurrence frequency and persistence as found in observed data.

Our results also lends support to the BCN technique by demonstrating that estimated parameters are very stable across different realisations of yield curve surfaces.

In summary, our results demonstrate that the two techniques are complementary and able to capture important cross sectional and time series properties of observed yield curve data. Hence, our results are of interest to practitioners in the financial markets who need accountable and history-consistent procedures for generating long-term yield curve forecasts.

References

Aït-Sahalia, Y. (1996): "Testing Continuous-Time Models of the Spot Interest Rate," *Review of Financial Studies*, 9, 385–426.

Ang, A., and G. Bekaert (2002): "Short Rate Nonlinearities and Regime Switches," *Journal of Economic Dynamics and Control*, 26, 1243–1274.

Ang, A., and G. Bekaert (2004): "The Term Structure of Real Rates and Expected Inflation," *Columbia University Working Paper*.

Banzal, R., G. Tauchen, and H. Zhou (2003): "Regime-Shifts, Risk Premiums in the Term Structure, and the Business Cycle," *Journal of Business and Economics Statistics*, October, 396–409.

Banzal, R., and H. Zhou (2002): "Term Structure of Interest Rates with Regime Shifts," *Journal of Finance*, 57, 1997–2043.

Bernadell, C., J. Coche, and K. Nyholm (2005): "Yield Curve Prediction for the Strategic Investor," *Working Paper no 472, Risk Management Division, European Central Bank*.

Björk, T., and B. Christensen (1999): "Interest Rate Dynamics and Consistent Forward Rate Curves," *Mathematical Finance*, 9, 323–348.

Brennan, M. J., and E. S. Schwartz (1982): "An Equilibrium Model of Bond Pricing and a Test of Market Efficiency," *Journal of Financial Quantitative Analysis*, 17, 301–329.

Campbell, J. Y., and R. J. Shiller (1987): "Cointegration and Tests of Present Value Models," *Journal of Political Economy*, 95, 1062–1088.

Conley, T. G., L. P. Hansen, E. G. J. Luttmer, and J. A. Scheinkman (1997): "Short-Term Interest Rates as Subordinated Diffusions," *Review of Financial Studies*, 10, 525–577.

Cox, J., J. E. Ingersoll, and S. A. Ross (1985): "A Theory of the Term Structure of Interest Rates," *Econometrica*, 53, 385–407.

Dai, Q., K. J. Singleton, and W. Yang (2003): "Regime Shifts in a Dynamic Term Structure Model of U.S. Treasury Bond Yields," *Working paper, Stern School of Business*.

Diebold, F. X., and C. Li (2006): "Forecasting the Term Structure of Government Bond Yields," *Journal of Econometrics*, 130, 337–364.

Diebold, F. X., G. D. Rudebusch, and S. B. Aruoba (2006): "The Macroeconomy and the Yield Curve: A Dynamic Latent Factor Approach," *Journal of Econometrics*, 131, 309–338.

Driffill, J., and M. Sola (1994): "Testing the Term Structure of Interest Rates Using a Stationary Vector Autoregression with Regime Switching," *Journal of Economic Dynamics and Control*, 18, 601–628.

Duffie, D., and R. Kan (1996): "A Yield-Factor Model of Interest Rates," *Mathematical Finance*, 6, 379–406.

Engsted, T., and K. Nyholm (2000): "Regime Shifts in the Danish Term Structure of Interest Rates," *Empirical Economics*, 25, 1–13.

Evans, D. D. (2003): "Real Risk, Inflation Risk, and the Term Structure," *The Economic Journal*, 113, 345–389.

Garcia, R., and P. Perron (1996): "An Analysis of the Real Interest Rate under Regime Shifts," *Review of Economics and Statistics*, 78, 111–125.

Glasserman, P., and S. G. Kou (2003): "The Term Structure of Simple Forward Rates with Jump Risk," *Mathematical Finance*, 13, 383–410.

Hall, A. D., H. M. Anderson, and W. J. Granger (1992): "A Cointegration Analysis of Treasury Bill Yields," *Review of Economics and Statistics*, 74, 116–126.

Hamilton, J. D. (1988): "Rational-Expectations Econometric Analysis of Changes in Regime: An Investigation of the Term Structure of Interest Rates," *Journal of Economic Dynamics and Control*, 12, 385–423.

——— (1994): *Time Series Analysis*. Princeton University Press, Princeton, New Jersey.

Hughston, L. (2000): *The New Interest Rate Models*. Risk Books, London.

——— (2003): "The Past, Present and Future of Term Structure Modelling," in *Modern Risk Management: A History*. Risk Books, London.

Hughston, L., and D. C. Brody (2000a): "Modern Theory of Interest Rates - With Extensions to Foreign Exchange, Inflation and Credit," *Working Paper, King's College*.

——— (2000b): "Modern Theory of Interest Rates - With Extensions to Foreign Exchange, Inflation and Credit," *Imperial College and King's College London lecture notes.*

Jamshidian, F. (1997): "LIBOR and Swap Market Models and Measures," *Finance and Stochastics*, 1, 293–330.

Kim, C. J., and C. R. Nelson (1999): *State Space Models with Regime Switching*. MIT Press, Cambridge, Massachussetts.

Kugler, P. (1990): "The Term Structure of Euro Interest Rates and Rational Expectations," *Journal of International Money and Finance*, 9, 234–244.

LANNE, M. (2000): "Near Unit Roots, Cointegration, and the Term Structure of Interest Rates," *Journal of Applied Econometrics*, 15, 513–529.

LONGSTAFF, F. A., AND E. S. SCHWARTZ (1992): "Interest Rate Volatility and the Term Structure: A Two-Factor General Equilibrium Model," *Journal of Finance*, 47, 1259–1282.

MARTELLINI, L., AND P. PRIAULET (2001): *Fixed-Income Securities.* John Wiley, Chichester.

NELSON, C. R., AND A. F. SIEGEL (1987): "Parsimonious Modeling of Yield Curves," *Journal of Business*, 60, 473–489.

REBONATO, R. (2003): "Term-Structure Models: A Review," *Working Paper, Royal Bank of Scotland.*

——— (2006): "Forward-Rate Volatilities and the Swaption Matrix: Why Neither Time-Homogeneity Nor Time Dependence Will Do," *International Journal of Theoretical and Applied Finance*, 9, 705–746.

REBONATO, R., AND J. CHEN (2007): "Evidence for State Transition and Altered Serial Co-Dependence in US*InterestRates*,″ *submitted to Quantitative Finance.*

REBONATO, R., AND V. GASPARI (2006): "Analysis of Drawdowns and Drawups in the US Interest-Rate Market," *Quantitative Finance*, 6, 297–326.

REBONATO, R., AND D. KAINTH (2004): "A Two-Regime Stochastic-Volatility Extension of the LIBOR Market Model," *International Journal of Theoretical and Applied Finance*, 7, 1–21.

REBONATO, R., S. MAHAL, M. JOSHI, L.-D. BUCHOLZ, AND K. NYHOLM (2005): "Evolving Yield Curves in the Real-World Measure: A Semi-Parametric Approach," *forthcoming Journal of Risk.*

SÖDERLIND, P., AND L. SVENSSON (1997): "New Techniques to Extract Market Expectations from Financial Instruments," *Journal of Monetary Economics*, 40, 383–429.

STANTON, R. (1997): "A Nonparametric Model of Term Structure Dynamics and the Market Price of Interest Rate Risk," *Journal of Finance*, 52, 1973–2002.

TAYLOR, J. B. (1993): "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.

VASICEK, O. (1977): "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics*, 5, 177–188.

7 Tables

	Active Regimes		
	3	2	1
Observations	513	463	24

Table 1: Summary statistics of regime classificatrions

Table1: This table reports the number of active regimes across the 1,000 simulations. "3" indicates that a given resampled yield curve surface resulted in the estimation of all three possible regimes i.e. (main, steep, inverse). "2" indicates that estimations resulted in identification of two of the three possible regimes, and "1" indicates that just one of the three possible regimes were identified i.e. that no regime-shifts could be detected. A regimes is identified if the corresponding state probability ($\pi_{t|t}$ from equation [11]) exceeds 0.70. Statistics are reported across the 1,000 replications.

	Estimated Regimes			
	Main	Steep	Flat	NA
Percentage share	49.54	35.36	12.94	2.16

Table 2: Distribution of estimated states

Table2: A state classification is defined when the corresponding state probability exceeds 0.70 (i.e. $\pi_{t|t} > 0.70$). "Main" refers to the state where the yield curve is upward sloping and the slope in not as steep as it is in the "steep" state. "Flat" refers to the state where the yield curve is flat or inverted. "NA" refers to observations that could not be classified according to the applied classification rule. Statistics are reported across the 1,000 replications.

	State probability		
	π(Main)	π(Steep)	π(Flat)
Mean	0.9918	0.9924	0.9910
Std.dev	0.0326	0.0318	0.0348

Table 3: Statistical properties of the estimated regime probabilities

Table3: Statistics conditional on the state classification are reported for the state probability $\pi_{t|t}$ across the 1,000 replications. A state classification is defined when the corresponding probability exceeds 0.70, hence the table reports the mean, standard deviation, skewness and kurtosis of $\pi_{t|t}(j) \left| \pi_{t|t}(j) > 0.70 \right.$.

Table4: The "Mean" and the "Std.dev" (standard deviation) are reported for the distribution of parameter estimates across the 1,000 replications. "λ" refers to the time decay parameter in the Nelson-Siegel equation. "Level" and "curvature" are the first and third

	λ	Level	Curvature	Slope Main	Steep	Flat	P11	P22	P33
Mean	0.093	3.624	-1.882	-2.269	-4.355	0.467	0.945	0.931	0.847
Std.dev	0.103	2.389	2.041	1.044	1.405	1.006	0.137	0.145	0.223

Table 4: Parameter Stability

Nelson-Siegel yield curve factor. P_{11}, P_{22}, P_{33} refer to the diagonal of the transition probability matrix. Summary statistics for the parameter estimates of the "slope" are calculated conditional on the state-classification $\left|\pi_{t|t}(j) > 0.70\right.$ for j={Main, Steep, Flat}.

	Maturity 1	2	3	4	5	6	7	8
	Panel A: Residuals							
Mean	-0.140	-0.120	-0.117	-0.131	-0.197	-0.210	-0.195	-0.156
Std.dev	0.566	0.629	0.806	1.010	1.026	0.946	0.841	0.729
	Panel B: Squared residuals							
Mean	0.340	0.410	0.664	1.038	1.092	0.939	0.745	0.556
Std.dev	1.946	1.990	2.759	3.879	3.163	1.950	2.195	1.912

Table 5: Residuals

Table5: This table reports summary statistics for the estimation errors obtained from the observation equation. For each of the eight maturities the "Mean" and "Std.dev" (standard deviation) are reported in panel A. Panel B reports the same summary statistics for the squared residuals .

8 Figures

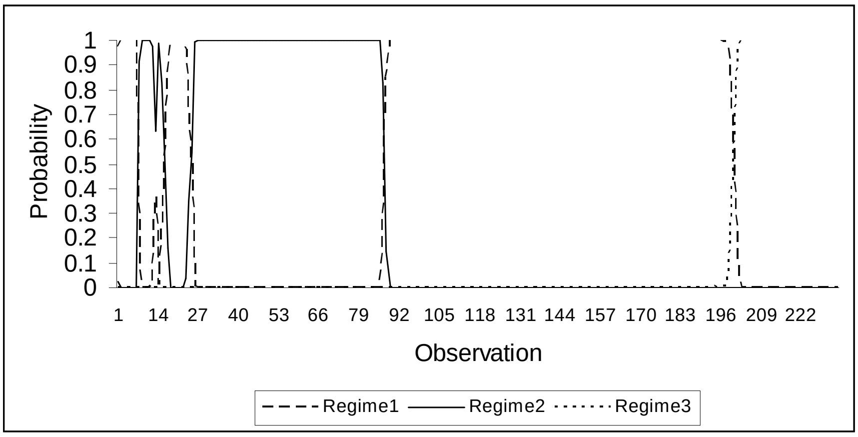

Regime classifications from a randomly chosen sample