BANK OF ENGLAND

Working Paper No. 518

Evaluating the robustness of UK term structure decompositions using linear regression methods

Sheheryar Malik and Andrew Meldrum

December 2014

BANK OF ENGLAND

Working Paper No. 518

Evaluating the robustness of UK term structure decompositions using linear regression methods

Sheheryar Malik(1) and Andrew Meldrum(2)

Abstract

This paper evaluates the robustness of UK bond term premia from affine term structure models. We show that this approach is able to match standard specification tests. In addition, term premia display countercyclical behaviour and are positively related to uncertainty about future inflation, consistent with previous findings for the United States. Premia are robust to correction for small sample bias and the inclusion of macro variables as unspanned factors. Including survey information about short rate expectations, which is a common way of improving identification of affine term structure models, however, results in inferior performance using UK data, as measured by standard specification tests and the economic plausibility of the estimated premia. Finally, we show that imposing the zero lower bound within a shadow rate term structure model does not have a large impact on estimates of long-maturity term premia.

Key words: Affine term structure model, term premia, bias correction, interest rate surveys, unspanned macro risks, shadow rate model.

JEL classification: E43, G10, G12.

(1) International Monetary Fund. Email: smalik2@imf.org
(2) Bank of England. Email: andrew.meldrum@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. The authors are grateful to Andy Blake, Iryna Kaminska, Chris Young and an anonymous referee for helpful comments. Contributions to this paper by Sheheryar Malik were made while he was working at the Bank of England. This paper was finalised on 19 November 2014.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email publications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

Summary

The term structure of interest rates on government bonds, commonly known as the 'yield curve', relates time to maturity to the average return, or yield, of a bond over its life. If investors were unconcerned about the risk of changes to interest rates over the investment period, the yield would equal the average expectation of the UK Monetary Policy Committee's short-term policy rate, Bank Rate, over the lifetime of the bond. This is because investors can choose between buying a long-term bond or investing in a series of short-term bonds and the expected returns on the two strategies must be equal to rule out the possibility of risk-free, or arbitrage, profits. This is sometimes known as the 'pure expectations hypothesis' of the yield curve. In practice, however, investors are risk-averse and demand a premium in the form of a higher expected return for investing in long-term bonds. This additional expected return is often known as the 'term premium'.

Both components of yields - expectations of future policy rates and term premia - contain useful information for policymakers. Expectations of future policy rates reflect investors' views about how the economy will evolve and how monetary policy will respond. Term premia can provide a guide to how uncertain investors are about the future and their attitudes towards that risk. Unfortunately, the two components cannot be observed separately and we need models of the yield curve in order to obtain a decomposition. That introduces uncertainty because we cannot be sure that we have a good model or that we have estimated the true parameter values for any given model.

The most commonly used vehicle for decomposing government bond yields over the last decade has been the 'Gaussian no-arbitrage affine term structure model' (ATSM). These make assumptions about the variables that affect yields, known as 'pricing factors', and how those factors behave over time. These pricing factors are often assumed to be principal components of bond yields. The models further impose theoretical restrictions on the yield curve to rule out the possibility that investors can make risk-free, or arbitrage, profits from investing in different maturity bonds. We first demonstrate that four (rather than the usual three) pricing factors are required to achieve a good fit to the yield curve and to match standard specification tests. This is consistent with recent studies of US data. Term premia turn out to be countercyclical - i.e. they are higher in relatively bad economic times - which is intuitive, since the compensation investors require for risk is likely to rise during those bad times. They are also increasing in the amount of uncertainty about future inflation. Again, this is intuitive, since investors are likely to require more compensation for risk if there is greater uncertainty about the real value of the returns they will receive on bonds.

One problem is that we typically only have quite short samples of data with which to estimate how the factors in the model behave. This increases the possibility that we obtain biased or imprecise parameter estimates. We therefore explore the robustness of our decompositions. For example, our benchmark sample period starts in May 1997 to reduce the possibility of changes in the model parameters (associated with the granting of operational independence to the Bank of England at that time) causing biases in the estimates of term premia. Extending the sample back to October 1992, when inflation targeting was introduced, makes only a small difference to

the estimated premia. But in longer samples the term premium estimates are typically lower, which is a result of the much higher average levels of policy interest rates over that sample: for a given bond yield today, if the model forecasts that the policy rate will revert back to a higher average level in future, then this must equate to a lower term premium.

We show that long-maturity term premium estimates are not materially affected if we apply standard statistical techniques for correcting for small sample bias in the parameter estimates, which suggests that such biases are not a substantial concern for our data set. Moreover, including additional macroeconomic variables as factors does not have a large impact on long-maturity premia. On the other hand, introducing information on expectations of future policy rates from surveys of professional economists (a common way of providing these models with more information about how interest rates behave over time) results in long-maturity term premium estimates that are significantly different from our benchmark model at times. However, there is evidence from standard tests that a model that includes surveys is misspecified; and the resulting estimates of term premia are no longer countercyclical or significantly related to the uncertainty about future inflation.

One drawback with ATSMs is that they do not impose non-zero nominal interest rates. When bond yields are substantially above zero, this does not matter, because the probability of negative rates implied by these models is very small. But when yields have been low, as in recent years, this may lead to misleading results. One popular recent approach has been to modify ATSMs so that the short rate is non-negative. Using one of the possible techniques for estimating such a model, we show that the impact of allowing for the zero lower bound on long-maturity term premia is likely to be fairly small, which is consistent with previous results for the United States.

1 Introduction

Understanding movements in the term structure of government bond yields is of considerable interest to financial market participants and policy-makers. Long-term bond yields can be decomposed into expectations of future short-term interest rates and a term premium that reflects the additional expected return required by investors in long-term bonds relative to rolling over a series of short-term bonds. Both components provide useful information for policy makers. Expected short rates reflect investors' views about the outlook for monetary policy, while the term premium reflects, among other things, the uncertainty around future short-term interest rates and the compensation that investors demand for bearing risk. Unfortunately, however, the decomposition of yields into expectations and term premia cannot be observed purely from market prices; rather, we need to use a model of bond yields to estimate the decomposition.

The workhorse model used by central banks for decomposing bond yields over the last decade has been the Gaussian no-arbitrage essentially affine term structure model (ATSM) of Duffee (2002): the long list of studies published by central bankers using these models includes Kim and Wright (2005) and Adrian et al. (2013) for the US and Joyce et al. (2010) and Guimarães (2014) for the UK.

One practical problem with using these models is that we typically only have short samples of highly persistent bond yield data with which to estimate the time series dynamics of the pricing factors. This can result in small sample biases and poor identification of some of the large number of parameters in these models (a point that is discussed in depth by Bauer et al. (2012) among others). This in turn will be reflected in greater uncertainty and / or biases in estimates of term premia. The main aim of this paper is therefore to study the robustness of UK term structure decompositions using the Gaussian affine class of term structure models.

Previous studies that have estimated ATSMs using UK data have typically relied on likelihood-based estimation techniques (e.g. Lildholdt et al. (2007); Joyce et al. (2010); Kaminska (2013); D'Amico et al. (2014); and Guimarães (2014)). These are numerically complex, involving high-dimensional non-linear optimisations over a likelihood surface that has many local optima and undefined regions. In this paper, we instead apply the estimation approach proposed by Adrian et al. (2013) (ACM henceforth) to UK data. They split the estimation into a number of steps, each of which involves a linear regression. Even though the technique does not impose no-arbitrage restrictions on the cross-section of yields during the estimation, we show that this has almost no impact on estimated premia, so is not likely to be a problem in practice.

Much of the previous literature on nominal bond yields has used models with three pricing factors. We show that four factors are required to match key features of the UK data. This is in keeping with several recent studies using US, which find evidence that support the use of more than three factors (e.g. Cochrane and Piazessi (2005), Duffee (2011) and ACM). We show that UK term premia are countercyclical (consistent with findings by Bauer et al. (2012) for the US) and positively related to the uncertainty around future inflation (consistent with findings by Wright (2011) for a panel of countries, including the UK); and the model matches the 'linear projections of yields' (LPY) specification tests proposed by Dai and Singleton (2002).

The estimates of long-maturity term premia implied by the model are robust in a number of different respects. First, while our benchmark model is estimated using data starting in May 1997 - which is when the Bank of England gained operational independence for monetary policy, thereby minimising the possibility of structural breaks in the sample - estimated term premia are not affected substantially by extending the sample period back to October 1992, when the UK first adopted an inflation targeting framework for monetary policy. Estimating the model over a longer sample, starting in 1979, results in term premium estimates that have broadly the same dynamics but a lower level, reflecting the higher average level of interest rates over this sample.

Second, two of the approaches that have been suggested in the literature for improving estimates of the time series dynamics: (i) adjusting for small sample bias (as suggested by Bauer et al. (2012)); and (ii) allowing for unspanned macroeconomic risk factors (Joslin et al. (2014)) have little impact on the estimates of UK long-maturity term premia and on the performance of the models against the LPY tests. The model uncertainty associated with these different specifications is much smaller than the parameter uncertainty within a given model.

Third, the results are not robust to the inclusion of survey expectations of short-term interest rates to help identify term premia, as proposed by Kim and Wright (2005), Kim and Orphanides (2012), Joyce et al. (2010) and Guimarães (2014) among others. Term premia have the same *average* level in a model that incorporates survey expectations but - consistent with previous studies - are less volatile than in a model without surveys. But the model that includes surveys has markedly inferior performance against the LPY tests and term premia are not countercyclical, which suggests that a model including surveys is mis-specified for our data set.

Fourth, although ATSMs do not impose the zero lower bound on nominal interest rates, we show that a shadow rate term structure model (which does impose the lower bound), estimated using a variant of the technique proposed by Bauer and Rudebusch (2014), does not result in substantially different estimates of long-maturity term premia compared with our benchmark ATSM. This is consistent with the findings of Kim and Priebsch (2013) for the US.

The remainder of this paper proceeds as follows. Section 2 sets out the benchmark model and the estimation technique of ACM. Section 3 introduces the data set and explores issues of model specification: the choices of the appropriate number of factors and the sample period. Section 4 demonstrates that the benchmark results are robust to using the alternative estimation technique of Joslin et al. (2011). Section 5 examines the impact of small sample bias; including survey data on interest rate expectations; and including unspanned macroeconomic variables as factors. Section 6 considers the possible impact of imposing the zero lower bound. Section 7 concludes.

2 Benchmark model

2.1 Excess returns in an ATSM

This section derives an expression for excess bond returns in an ATSM, following the exposition from ACM. In Gaussian ATSMs (see Piazzesi (2003) for an overview), a $K \times 1$ vector of pricing factors, $\mathbf{x}_t$, evolves according to a Gaussian VAR (1):

$$\mathbf{x}_{t+1} = \boldsymbol{\mu} + \boldsymbol{\Phi}\mathbf{x}_t + \mathbf{v}_{t+1}, \tag{1}$$

where the shocks $\mathbf{v}_{t+1} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ are conditionally Gaussian, homoscedastic and independent across time. We denote the time-t price of a zero-coupon bond with a maturity of n by $P_t^{(n)}$. The the assumption of no-arbitrage implies the existence of a pricing kernel M_{t+1} such that

$$P_t^{(n)} = E_t[M_{t+1}P_{t+1}^{(n-1)}]. \tag{2}$$

The pricing kernel is assumed to be exponentially affine in the factors:

$$M_{t+1} = \exp(-r_t - \frac{1}{2}\boldsymbol{\lambda}_t'\boldsymbol{\lambda}_t - \boldsymbol{\lambda}_t'\boldsymbol{\Sigma}^{-1/2}\mathbf{v}_{t+1}), \tag{3}$$

where $r_t = \ln P_t^{(1)}$ denotes the continuously compounded one-period risk-free rate, which is assumed to be affine in the factors:

$$r_t = \delta_0 + \boldsymbol{\delta}_1'\mathbf{x}_t, \tag{4}$$

and the market prices of risk ($\boldsymbol{\lambda}_t$) are affine in the factors, as in Duffee (2002):

$$\boldsymbol{\lambda}_t = \boldsymbol{\Sigma}^{-1/2}(\boldsymbol{\lambda}_0 + \boldsymbol{\lambda}_1\mathbf{x}_t). \tag{5}$$

The log excess one-period holding return of a bond maturing in n periods is defined as

$$rx_{t+1}^{(n-1)} = \log P_{t+1}^{(n-1)} - \log P_t^{(n)} - r_t. \tag{6}$$

Using (3) and (6) in (2), ACM show that

$$E_t[\exp(rx_{t+1}^{(n-1)} - \frac{1}{2}\boldsymbol{\lambda}_t'\boldsymbol{\lambda}_t - \boldsymbol{\lambda}_t'\boldsymbol{\Sigma}^{-1/2}\mathbf{v}_{t+1})] = 1, \tag{7}$$

and, under the assumption of joint normality of $\{rx_{t+1}^{n-1}, \mathbf{v}_{t+1}\}$, they demonstrate that

$$E_t[rx_{t+1}^{(n-1)}] = Cov_t[rx_{t+1}^{(n-1)}, \mathbf{v}_{t+1}'\boldsymbol{\Sigma}^{-1/2}\boldsymbol{\lambda}_t] - \frac{1}{2}Var_t[rx_{t+1}^{(n-1)}]. \tag{8}$$

By denoting $\boldsymbol{\beta}_t^{(n-1)\prime} = Cov_t[rx_{t+1}^{(n-1)}, \mathbf{v}_{t+1}']\boldsymbol{\Sigma}^{-1}$ and using (5), (8) can be rewritten as

$$E_t[rx_{t+1}^{(n-1)}] = \boldsymbol{\beta}_t^{(n-1)\prime}[\boldsymbol{\lambda}_0 + \boldsymbol{\lambda}_1\mathbf{x}_t] - \frac{1}{2}Var_t[rx_{t+1}^{(n-1)}]. \tag{9}$$

This in turn can be used to decompose the unexpected excess return into a component that is correlated with $\mathbf{v}_{t+1}$ and another which is conditionally orthogonal:

$$rx_{t+1}^{(n-1)} - E_t[rx_{t+1}^{(n-1)}] = \boldsymbol{\beta}_t^{(n-1)\prime}\mathbf{v}_{t+1} + e_{t+1}^{(n-1)}, \tag{10}$$

where $e^{(n-1)}_{t+1} \backsim iid(0, \sigma^2)$. The return generating process for log excess returns can be written as

$$rx^{(n-1)}_{t+1} = \boldsymbol{\beta}^{(n-1)\prime}(\boldsymbol{\lambda}_0 + \boldsymbol{\lambda}_1 \mathbf{x}_t) - \frac{1}{2}(\boldsymbol{\beta}^{(n-1)\prime}\boldsymbol{\Sigma}\boldsymbol{\beta}^{(n-1)} + \sigma^2) + \boldsymbol{\beta}^{(n-1)\prime}\mathbf{v}_{t+1} + e^{(n-1)}_{t+1}. \tag{11}$$

The first to fourth terms on the right-hand side of (11) are the expected return, convexity adjustment, priced return innovation and return pricing error respectively. We observe returns for $t = 1, 2, ..., T$ and maturities $n = n_1, n_2, ..., n_N$. Stacking the system across n and t, ACM construct the expression

$$\mathbf{rx} = \boldsymbol{\beta}^{(n-1)\prime}(\boldsymbol{\lambda}_0 \boldsymbol{\iota}_T' + \boldsymbol{\lambda}_1 \mathbf{X}_{\_}) - \frac{1}{2}(\mathbf{B}^* vec(\boldsymbol{\Sigma}) + \sigma^2 \boldsymbol{\iota}_N)\boldsymbol{\iota}_T' + \boldsymbol{\beta}'\mathbf{V} + \mathbf{E}, \tag{12}$$

where $\mathbf{rx}$ is an $N \times T$ matrix of excess returns; $\boldsymbol{\beta} = [\boldsymbol{\beta}^{(1)}, \boldsymbol{\beta}^{(2)}, ..., \boldsymbol{\beta}^{(n)}]$ is a $K \times N$ matrix of factor loadings; $\boldsymbol{\iota}_N$ and $\boldsymbol{\iota}_T$ are $N \times 1$ and $T \times 1$ vectors of ones respectively; $\mathbf{X}_{\_}$ is a $K \times T$ matrix of lagged pricing factors; $\mathbf{B}^* = [vec(\boldsymbol{\beta}^{(1)}\boldsymbol{\beta}^{(1)\prime})....vec(\boldsymbol{\beta}^{(N)}\boldsymbol{\beta}^{(N)\prime})]'$ is an $N \times K^2$ matrix; and $\mathbf{V}$ and $\mathbf{E}$ are matrices of dimensions $K \times T$ and $N \times T$ respectively.

2.2 Estimation

In this section we provide a brief outline of the estimation technique, which involves four linear regressions. We refer the reader to ACM for further details.

1. In the first step, we estimate (1) by OLS, giving estimates $\widehat{\boldsymbol{\Phi}}$ and $\widehat{\boldsymbol{\Sigma}}$. Following ACM, we calibrate $\boldsymbol{\mu} = \mathbf{0}$, which ensures that the means of the factors (mean-zero principal components) are equal to their sample averages.

2. In the second step, we estimate the reduced-form of (12), $\mathbf{rx} = \mathbf{a}\boldsymbol{\iota}_T' + \boldsymbol{\beta}'\widehat{\mathbf{V}} + \mathbf{c}\mathbf{X}_{\_} + \mathbf{E}$, giving estimates $\widehat{\mathbf{a}}$, $\widehat{\boldsymbol{\beta}}$ and $\widehat{\mathbf{c}}$, as well as $\widehat{\sigma}^2 = \mathrm{tr}(\widehat{\mathbf{E}}\widehat{\mathbf{E}}')/NT$.

3. Using $\widehat{\mathbf{B}}^*$ constructed from $\widehat{\boldsymbol{\beta}}$ and noting that $\mathbf{a} = \boldsymbol{\beta}'\boldsymbol{\lambda}_0 - \frac{1}{2}(\mathbf{B}^* vec(\boldsymbol{\Sigma}) + \sigma^2\boldsymbol{\iota}_N)$ and $\mathbf{c} = \boldsymbol{\beta}'\boldsymbol{\lambda}_1$, we can estimate the parameters of the price of risk using cross-sectional regressions:

$$\widehat{\boldsymbol{\lambda}}_0 = (\widehat{\boldsymbol{\beta}}\widehat{\boldsymbol{\beta}}')^{-1}\widehat{\boldsymbol{\beta}}(\widehat{\mathbf{a}} + \frac{1}{2}(\widehat{\mathbf{B}}^* vec(\widehat{\boldsymbol{\Sigma}}) + \widehat{\sigma}^2\boldsymbol{\iota}_N)), \tag{13}$$

$$\widehat{\boldsymbol{\lambda}}_1 = (\widehat{\boldsymbol{\beta}}\widehat{\boldsymbol{\beta}}')^{-1}\widehat{\boldsymbol{\beta}}\widehat{\mathbf{c}}. \tag{14}$$

4. ACM assume the short-term interest rate is measured with error $\varepsilon_t \sim iid(0, \sigma_\varepsilon^2)$:

$$r_t = \delta_0 + \boldsymbol{\delta}_1' \mathbf{x}_t + \varepsilon_t. \tag{15}$$

We can estimate (15) by OLS to obtain estimates $\widehat{\delta}_0$ and $\widehat{\boldsymbol{\delta}}_1$.

2.3 Model-implied zero-coupon yields and term premia

Given the above assumptions, it is straightforward to show that zero-coupon bond yields are affine functions of the factors:

$$y_t^{(n)} = -\frac{1}{n}\left[a_n + \mathbf{b}_n' \mathbf{x}_t\right]. \tag{16}$$

where the coefficients a_n and $\mathbf{b}_n$ follow the recursive equations

$$a_n = a_{n-1} + \mathbf{b}_{n-1}'(\boldsymbol{\mu} - \boldsymbol{\lambda}_0) + \frac{1}{2}(\mathbf{b}_{n-1}' \boldsymbol{\Sigma} \mathbf{b}_{n-1} + \sigma^2) - \delta_0 \tag{17}$$

$$\mathbf{b}_n' = \mathbf{b}_{n-1}'(\boldsymbol{\Phi} - \boldsymbol{\lambda}_1) - \boldsymbol{\delta}_1' \tag{18}$$

and where we have the initial conditions $a_0 = 0$ and $\mathbf{b}_0 = \mathbf{0}$.

Following the definition of Dai and Singleton (2002), an n-period bond yield can be decomposed as

$$y_t^{(n)} = \frac{1}{n}\sum_{i=0}^{n} E_t r_{t+i} + TP_t^{(n)}, \tag{19}$$

where the first-term on the right-hand side of (19) is the average expected short rate over the next n periods; and $TP_t^{(n)}$ is a term premium that is the additional compensation investors require for investing in a long-term bond rather than rolling over a series of short-term investments. The expectations term can be computed as:

$$\frac{1}{n}\sum_{i=0}^{n} E_t r_{t+i} = -\frac{1}{n}\left[\widetilde{a}_n + \widetilde{\mathbf{b}}_n' \mathbf{x}_t\right], \tag{20}$$

where the coefficients $\widetilde{a}_n$ and $\widetilde{\mathbf{b}}_n$ follow the recursive equations

$$\widetilde{a}_n = \widetilde{a}_{n-1} + \widetilde{\mathbf{b}}'_{n-1}\boldsymbol{\mu} - \delta_0 \quad (21)$$

$$\widetilde{\mathbf{b}}'_n = \widetilde{\mathbf{b}}'_{n-1}\boldsymbol{\Phi} - \boldsymbol{\delta}'_1 \quad (22)$$

and we have the initial conditions $\widetilde{a}_0 = 0$ and $\widetilde{\mathbf{b}}_0 = \mathbf{0}$.

One way of assessing the effect of changing a model specification on term premia is to view it in the light of the uncertainty around term premia associated with *parameter* uncertainty within a benchmark model. In the following sections, we therefore report 80% confidence intervals for 10-year term premia in the benchmark four-factor model, computed using a bootstrap procedure. The algorithm we use to bootstrap the confidence interval has the following steps.

1. Estimate the model using the procedure explained above.

2. At iteration h of the bootstrap, generate a bootstrap sample of the factors for periods $t = 1, 2, ..., T$ (where the sample length T is same as for the data). For $t = 1$ randomly select $\mathbf{x}_1^{(h)}$ from the original sample. For $t > 1$ randomly resample (with replacement) $\mathbf{v}_t^{(h)}$ from the estimated residuals from (1) and compute $\mathbf{x}_t^{(h)} = \widehat{\boldsymbol{\Phi}}\mathbf{x}_{t-1}^{(h)} + \mathbf{v}_t^{(h)}$.

3. For periods $t = 1, 2, ..., T$ and maturities $n = n_1, ..., n_N$, randomly resample (with replacement) residuals $e_{t+1}^{(n-1)(h)}$ from (11) and compute $\mathbf{rx}^{(h)} = \widehat{\mathbf{a}}\boldsymbol{\iota}'_T + \widehat{\boldsymbol{\beta}}'\mathbf{V}^{(h)} + \widehat{\mathbf{c}}\mathbf{X}_{-}^{(h)} + \mathbf{E}^{(h)}$.

4. For periods $t = 1, 2, ..., T$, randomly resample (with replacement) residuals $\varepsilon_t^{(h)}$ from (15) and compute $r_t^{(h)} = \widehat{\delta}_0 + \widehat{\boldsymbol{\delta}}'_1\mathbf{x}_t^{(h)} + \varepsilon_t^{(h)}$.

5. Re-estimate the model using the h^{th} bootstrapped sample, saving the term premium estimates.

6. While $h < 10,000$ increase h by one and return to step 2.

7. Compute the percentiles of the bootstrapped term premium estimates.

3 Data

We estimate the model using month-end zero-coupon yields constructed using the smoothed cubic spline method of Anderson and Sleath (2001) (selected maturities from which are published by the Bank of England).[1] The pricing factors are principal components extracted from yields with maturities of $3, 4, ..., 120$ months; and we use excess returns for $n = 18, 24, 30, ..., 114, 120$ months, giving a cross-section $N = 18$ maturities.[2]

In our benchmark specification we use a four-factor model and a sample starting in May 1997 (when the Bank of England gained operational independence, in order to reduce the risk of biases associated with shifts in monetary policy regimes) and ending in December 2013. We discuss the choice of the number of factors in Section 3.1 and the sensitivity of term premium estimates to choosing longer sample periods in Section 3.2.

3.1 Number of factors

The standard approach when modelling the term structure of US nominal interest rates has been to assume three pricing factors, since three principal components are sufficient to explain the large majority of the cross-sectional variation in yields (Litterman and Scheinkman (1991)). Table 1 illustrates this for the UK. It shows the cumulative proportion of the sample variance of yields with maturities of $3, 4, ..., 120$ months for the period May 1997-December 2013 that are explained by the first five principal components of those yields. The first principal component explains more than 97% of the sample variation. Adding a second principal component takes the cumulative proportion explained to nearly 99.8% and a third to over 99.9%.

ACM propose that one way for testing for redundant factors is to use a Wald test of whether the ith column of $\boldsymbol{\beta}'$ is equal to zero, which amounts to testing whether the loadings on the i^{th} factor are zero. Under the null hypothesis that $\boldsymbol{\beta}_i = \mathbf{0}$, the test statistic

$$W_{\beta_i} = \widehat{\beta}_i' \widehat{V}_{\beta_i}^{-1} \widehat{\beta}_i \tag{23}$$

[1] http://www.bankofengland.co.uk/statistics/pages/yieldcurve/default.aspx.

[2] The published data set does not always include yields below one year, depending on the maturity of the shortest UK government bond available. We fill the gaps by linearly interpolating between Bank Rate (the UK monetary policy rate) - which we take to have zero maturity - and the shortest available maturity bond yield.

Table 1: Principal component analysis of UK bond yields

Principal component	Proportion of total variance explained	Cumulative proportion
1	97.452%	97.452%
2	2.325%	99.777%
3	0.200%	99.976%
4	0.019%	99.996%
5	0.003%	99.999%

Estimated using zero-coupon yields with maturities of 3, 4, ..., 120 months using a sample period of May 1997-December 2013.

is asymptotically distributed $\chi^2(N)$. Unreported results show that if we apply this to our data set, we obtain evidence in favour of a six-factor specification (i.e. more than the five factors ACM use to model the US term structure). While a six-factor model may be justified according to this test, however, the marginal improvement in the in-sample fit to bond yields gained from using such a large number of factors is economically trivial. Table 2 shows annualised percentage point root mean squared error statistics for bond yields. In a three-factor model, the RMSE of pricing errors is more than 4 basis points for all the considered maturities and is close to 9 basis points for short maturities. In a four-factor model, the RMSE is reduced below 3 basis points for all maturities. Adding further factors can actually result in *larger* pricing errors in some instances (recall that using the ACM estimation technique the model is not fitted directly to yields but to excess returns).

Table 2: Root mean squared fitting errors from models with different numbers of factors

Maturity (months)	Yield RMSE (percentage points)			
	3 factors	4 factors	5 factors	6 factors
6	0.088	0.016	0.030	0.022
18	0.078	0.023	0.012	0.011
24	0.058	0.019	0.007	0.008
60	0.044	0.008	0.025	0.008
120	0.049	0.026	0.076	0.013

In-sample root mean squared fitting errors (RMSE) for yields for models with different numbers of factors. All models are estimated using a sample period of May 1997-December 2013.

We next show that moving beyond four factors does not significantly affect the ability of the model to match the two 'linear projections of yields' (LPY) specification tests proposed by Dai and Singleton (2002). The first test, LPY(i), considers whether the model parameters imply the same pattern of slope coefficients as from the Campbell and Shiller (1991) regression

$$y_{t+1}^{(n-1)} - y_t^{(n)} = \alpha^{(n)} + \phi^{(n)}(y_t^{(n)} - r_t)/(n-1) + \zeta_t^{(n)} \tag{24}$$

that is observed in the data. This is a test of the capacity of the model to replicate the historical dynamics of yields. The second test, LPY(ii), considers a risk-adjusted version of (24):

$$y_{t+1}^{(n-1)} - y_t^{(n)} - e_t^{(n)}/(n-1) = \alpha^{*(n)} + \phi^{*(n)}(y_t^{(n)} - r_t)/(n-1) + \zeta_t^{*(n)}, \tag{25}$$

where $e_t^{(n)} \equiv E_t[\ln P_{t+1}^{(n-1)} - \ln P_t^{(n)} - r_t]$ is the expected excess one-period return on an n-period bond. If the model is able adequately to capture the dynamics of term premia, then we should recover a coefficient $\phi^{*(n)} = 1$.

Tables 4 and 5 report results from these two tests applied to models with different numbers of factors. The pattern of LPY(i) slope coefficients observed in the data (shown in the first column) is decreasing with maturity (as Dai and Singleton (2002) find for the US). To obtain values of $\phi^{(n)}$ for the models, we simulate time series of 100,000 periods conditional on the point estimates of the parameters. The estimated standard errors for the coefficients obtained using the data are large at long maturities, which means that it is difficult to distinguish between the models. Nevertheless, at short maturities, a three-factor model has coefficient estimates that are more than two estimated standard errors away from the coefficient obtained from the data, whereas the models with four or more factors do not.

Turning to the LPY(ii) test, a three-factor model is clearly mis-specified, with estimated slope coefficients that are much greater than one, and significantly so at shorter maturities. There is less to choose between models with four or more factors; and while all these models recover coefficients that are larger than one for all maturities, the differences from one are not significant. This again suggests that the advantages of moving beyond four factors are likely to be quite small.

We can shed further light on these results using model-implied estimates of term premia. Figure 1 shows that the 1-year term premium from the three-factor model behaves quite differently from the models with more factors. For 10-year premia (Figure 2), the estimates are more similar. Taking these results together, three factors are clearly insufficient to match key features of the short end of

Table 3: Performance against LPY(i) test by models with different numbers of factors

Maturity (months)	Data	3 factors	4 factors	5 factors	6 factors
6	1.964	0.426	2.081	1.676	1.761
	(0.308)				
18	1.037	0.169	0.601	0.745	0.796
	(0.587)				
24	0.658	0.052	0.235	0.403	0.446
	(0.667)				
60	−0.084	−0.361	−0.406	−0.328	−0.262
	(0.904)				
120	−0.429	−0.472	−0.869	−0.782	−0.693
	(1.188)				

Estimated slope coefficients ($\widehat{\phi}^{(n)}$) from the regression $y_{t+1}^{(n-1)} - y_t^{(n)} = \alpha^{(n)} + \phi^{(n)}(y_t^{(n)} - r_t)/(n-1) + \zeta_t^{(n)}$. Standard errors for the data are shown in parentheses. Model-implied slope coefficients are estimated using a data set with 100,000 periods simulated from the model, conditional on the point estimates of the parameters. All models are estimated using a sample period of May 1997-December 2013.

Figure 1: One-year term premium estimates from models with different numbers of factors

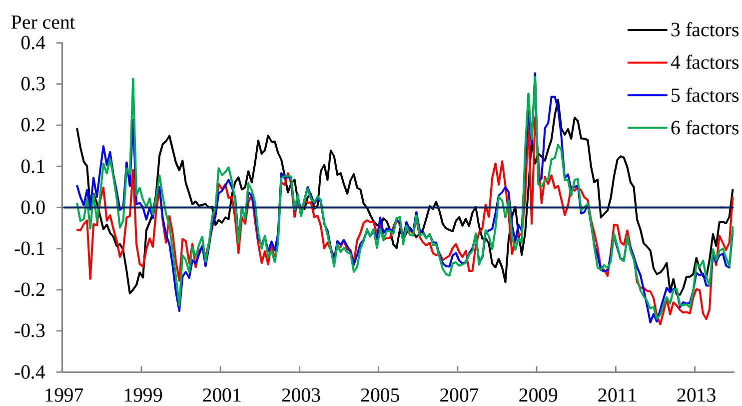

the yield curve. There is relatively little to distinguish between models with four or more factors, so we prefer a relatively parsimonious four-factor model as our benchmark model in the remainder of this paper.

3.2 Sample period

When estimating dynamic term structure models using UK data, the choice of sample period is particularly challenging given the high persistence of yields and the likelihood of several structural breaks. Figure 3 shows yields at selected maturities between January 1979 and December

Table 4: Performance against LPY(ii) test by models with different numbers of factors

Maturity (months)	Data	3 factors	4 factors	5 factors	6 factors
6	1	2.285	1.343	1.443	1.308
		(0.307)	(0.330)	(0.327)	(0.331)
18	1	1.972	1.437	1.294	1.319
		(0.560)	(0.644)	(0.664)	(0.668)
24	1	1.773	1.413	1.260	1.281
		(0.631)	(0.718)	(0.756)	(0.756)
60	1	1.488	1.234	1.178	1.166
		(0.885)	(0.942)	(0.961)	(0.961)
120	1	1.199	1.176	1.088	1.090
		(1.057)	(1.007)	(0.998)	(1.002)

Estimated slope coefficients ($\widehat{\phi}^{*(n)}$) from the regression $y_{t+1}^{(n-1)} - y_t^{(n)} - e_t^{(n)}/(n-1) = \alpha^{*(n)} + \phi^{*(n)}(y_t^{(n)} - r_t)/(n-1) + \zeta_t^{*(n)}$. Standard errors are shown in parentheses. All models are estimated using a sample period of May 1997-December 2013.

Figure 2: Ten-year term premium estimates from models with different numbers of factors

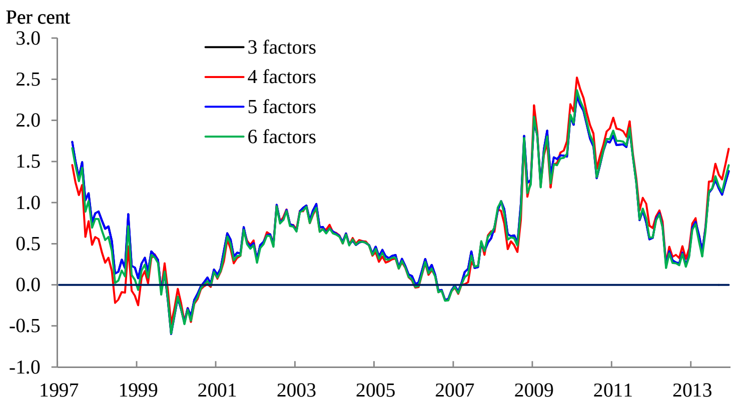

2013. Yields drifted downwards steadily during the 1980s. There were also sharper falls following the introduction of an inflation targeting regime in the UK in October 1992 and the granting of operational independence for monetary policy to the Bank of England in May 1997.

As explained above, in our benchmark model we have preferred to use a sample starting in May 1997 in order to reduce the possibility of structural breaks associated with the change in monetary policy regimes. Other studies have used different samples. For example, D'Amico et al. (2014) use a sample starting in January 1993, Joyce et al. (2010) in October 1992 and Guimarães (2014) in January 1972. In this section we therefore consider how changing the sample period affects estimates of term premia in bond yields.

Figure 4 illustrates how the term premium component of the ten-year yield compares when the

Figure 3: UK end-month bond yields, January 1979-December 2013

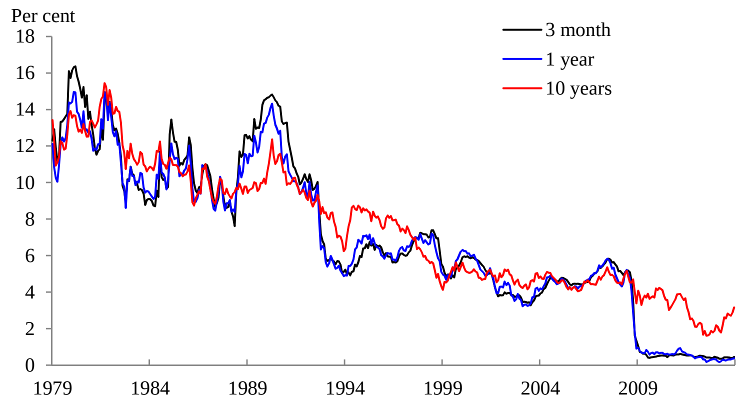

Figure 4: 10-year term premium estimates from four-factor models with different sample start dates

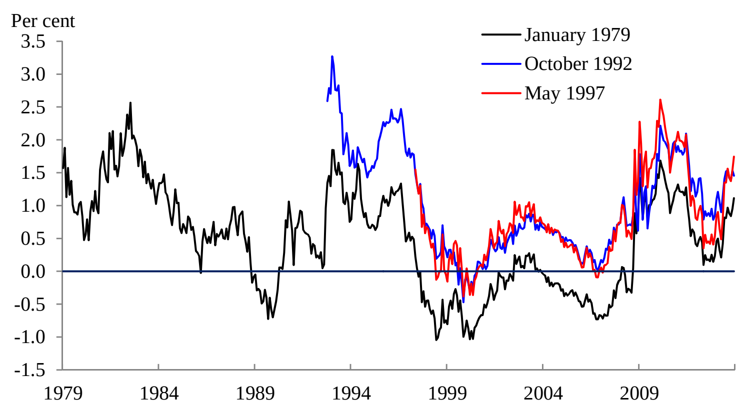

models are estimated using different sample periods. The broad trends and cycles across models estimated using samples starting in January 1979, October 1992 (the start of inflation targeting in the UK) and May 1997 are similar, with the timing of upturns and downturns broadly coinciding. While the levels of term premium estimates are similar from models using samples starting in October 1992 and May 1997, they are generally lower if we use a longer sample, which largely reflects the higher average level of bond yields over this period. The unconditional mean short rate estimated from the sample starting in January 1979 is 7.22%, compared with 4.18% from the sample starting in October 1992 and 3.68% from our benchmark sample starting in May 1997. If short rates are forecast to revert back to a higher mean, the term premium component of yields is generally lower. For instance, in December 2013, the expected paths of short-term interest rates from the three models - shown in Figure 5 - were relatively similar at short maturities but deviated significantly at longer horizons, reflecting the different long-run means.

Figure 5: Expected path of short-term risk-free interest rates in December 2013, taken from models with different sample start dates

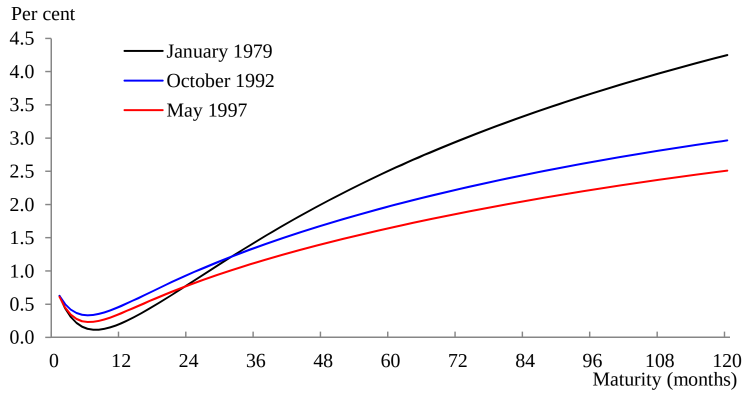

4 Robustness of the estimation technique

As mentioned previously, Gaussian ATSMs have more typically been estimated using maximum likelihood techniques. These exploit distributional assumptions and the linear state-space representation of the model to estimate (potentially unobserved) pricing factors and coefficients together. Maximum likelihood could be considered a natural way to estimate these models given they provide a complete characterisation of the joint distribution of yields. However, in practice the large number of parameters to be estimated combined with a likelihood function that has many global optima and undefined regions makes convergence to a global maximum computationally challenging.

Aside from ACM, a number of studies have applied multi-step methods to reduce the numerical challenges associated with the estimation of ATSMs. Examples include Moench (2008), Joslin et al. (2011) (JSZ henceforth), Hamilton and Wu (2012), Kaminska (2013) and Andreasen and Meldrum (2014). All of these methods involve some non-linear optimisation, which ACM avoid because they do not impose the bond pricing recursions (17) and (18) inside the estimation procedure. They show, however, that the factor loadings implied by their approach satisfy these restrictions to a high degree of precision. Unreported results show the same finding for UK data using our benchmark model.

To illustrate the robustness of this estimation technique for UK term premia further, we compare our term premium estimates with those obtained using maximum likelihood imposing the no-arbitrage restrictions, again using the first four principal components as pricing factors, assuming

Figure 6: 10-year term premium from benchmark model estimated using the ACM technique and using the JSZ normalisation and maximum likelihood (JSZ-ML)

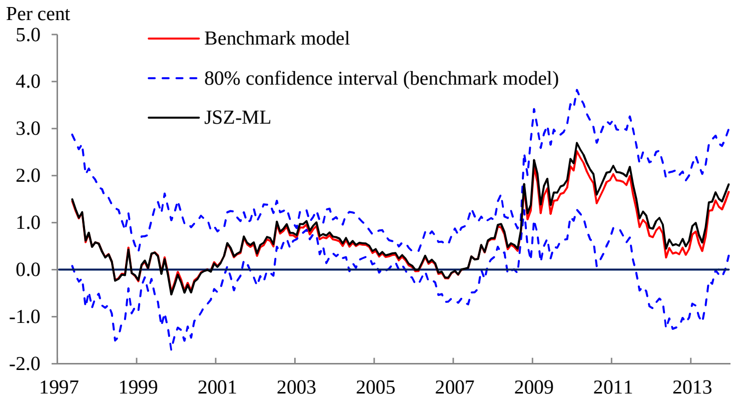

that they are measured without error. We apply the JSZ normalisation, which re-parameterises the model in terms of $\{\boldsymbol{\mu}, \boldsymbol{\Phi}, \boldsymbol{\Sigma}, r_{\infty}^{\mathbb{Q}}, \boldsymbol{\lambda}^{\mathbb{Q}}\}$, where $r_{\infty}^{\mathbb{Q}}$ is the unconditional mean of the short rate and $\boldsymbol{\lambda}^{\mathbb{Q}}$ is a vector of the eigenvalues of $\boldsymbol{\Phi} - \boldsymbol{\lambda}_1$. The estimation of the model involves the following steps:

1. Estimate (1) by OLS to obtain estimates $\widehat{\boldsymbol{\Phi}}$ and $\widehat{\boldsymbol{\Sigma}}$. We again set $\boldsymbol{\mu} = \mathbf{0}$.

2. Conditional on the parameter estimates at the first step, estimate $r_{\infty}^{\mathbb{Q}}$ and $\boldsymbol{\lambda}^{\mathbb{Q}}$ by maximum likelihood. We fit the model to yields with maturities of 1, 3, 6, 12, 24, 36, 60, 84 and 120 months, assuming that all yields are measured with error:

$$y_t^{(n)} = -\frac{1}{n}\left(A_n + \mathbf{B}_n'\mathbf{x}_t\right) + w_t^{(n)}, \tag{26}$$

where $w_t^{(n)} \sim \mathcal{N}\left(0, \sigma_w^2\right)$.

Figure 6 compares 10-year term premium estimates from our benchmark model with those obtained from the model estimated using this alternative technique. The two estimates are almost identical, which again suggests that not imposing the no-arbitrage restrictions in estimation has little impact on the results.

5 Robustness of the time series dynamics of the factors

5.1 Small sample bias

An issue recently highlighted by Bauer et al. (2012) (BRW henceforth) is that estimates of the time series dynamics of the factors in term structure models may suffer materially from bias due to the high degree of persistence of interest rates, particularly in cases where the sample period is short. Since the pricing factors are linear combinations of interest rates, such as principal components, this persistence will also be imparted to the factors. Biased parameter estimates for the VAR in (1) translate into biased estimates of expected future short-term interest rates and term premia. The bias typically goes in the direction of making the estimated system less persistent than the data-generating process: future short rates are forecast to revert to their unconditional mean too quickly and, as a consequence, the component of long-term yields reflecting expected future short rates is too stable and term premia are too volatile.

A multi-step estimation procedure such as the method of ACM makes bias-correction at the stage of estimating the time series dynamics of the factors straightforward. As BRW document, there are alternative bias correction methods that could be applied. We employ bootstrap (mean) bias correction in which data are simulated using a distribution-free residual bootstrap taking the OLS estimates as the data-generating parameters.[3] The algorithm is set out below.

1. Estimate (1) by OLS as before, storing the residuals and the OLS estimate $\widehat{\mathbf{\Phi}}$.
2. At iteration h of the bootstrap, generate a bootstrap sample for periods $t = 1, 2, ..., T$ (where the sample length T is same as for the data). For $t = 1$ randomly select $\mathbf{x}_1^{(h)}$ from the original sample. For $t > 1$ randomly resample (with replacement) residuals $\mathbf{v}_t^{(h)}$ and construct $\mathbf{x}_t^{(h)} = \widehat{\mathbf{\Phi}}\mathbf{x}_{t-1}^{(h)} + \mathbf{v}_t^{(h)}$.
3. Calculate the OLS estimate $\widehat{\mathbf{\Phi}}^{(h)}$ using the h^{th} bootstrapped sample.
4. While $h < 10,000$ increase h by one and return to step 2.
5. Calculate the average over all bootstrap samples: $\overline{\mathbf{\Phi}} = \frac{1}{H}\sum_{h=1}^{H}\widehat{\mathbf{\Phi}}^{(h)}$.

[3] Preliminary investigations showed that using inverse bias correction algorithm of BRW makes little difference to term premium estimates when compared with the bootstrap correction used here.

Figure 7: 10-year term premium estimates from benchmark model and a model with bootstrapped bias correction

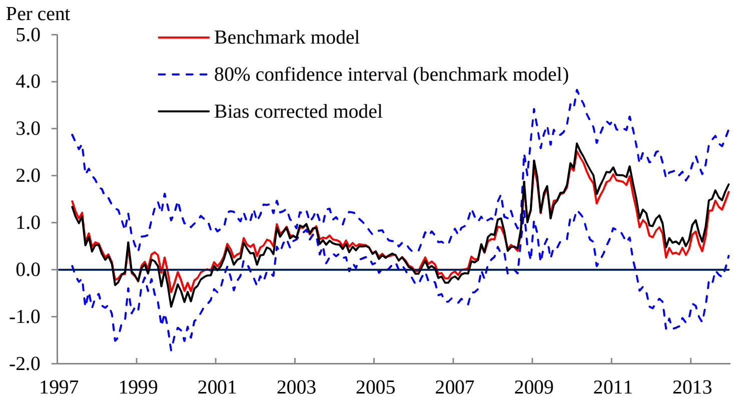

6. The bootstrap bias corrected estimate is: $\widetilde{\mathbf{\Phi}}^H = \widehat{\mathbf{\Phi}} - [\overline{\mathbf{\Phi}} - \widehat{\mathbf{\Phi}}] = 2\widehat{\mathbf{\Phi}} - \overline{\mathbf{\Phi}}$.

The remaining steps of the estimation (i.e. steps 2 to 4 in Section 2.2) are unaffected.[4] Figure 7 shows estimates of 10-year term premia with and without this bias correction applied and Tables 5 and 6 show the impact on the models' performance against the LPY tests. The bias correction has almost no effect for our data set, so we do not consider it further.

5.2 Inclusion of interest rate survey expectations

Aside from the issue of small sample bias, the short time series of persistent yield factors means that the parameters governing the time series dynamics of the factors are weakly identified. In this and the following sub-section, we consider the robustness of term premium estimates to two common approaches that have been suggested to sharpen inference about the time series properties of yields.

First, one common approach used to help identify the time series dynamics of the pricing factors has been to include survey data on short-term interest rate expectations in the model - for example, Kim and Wright (2005) and Kim and Orphanides (2012) for the US and Joyce et al. (2010) and Guimarães (2014) for the UK. In this section, we estimate a version of the model that incorporates information from the Bank of England's Survey of External Forecasters, which is

[4]In cases where the bias corrected estimate $\widetilde{\mathbf{\Phi}}^H$ has eigenvalues outside the unit circle, a standard approach is to shrink the bias corrected estimate towards the OLS estimate using the scheme proposed by Kilian (1998). This implies that the persistence of the pricing factors will depend entirely on the ad hoc shrinkage, which may have significant implications for the estimated persistence of short rates and hence term premium estimates. In our case, however, the bias-corrected estimate remains in the stationary region.

Table 5: Model performance against LPY(i) test

Maturity (months)	Data	Benchmark	Bias corrected	Survey	Unspanned macro
6	1.964	2.081	1.859	1.540	2.107
	(0.308)				
18	1.037	0.601	0.588	0.567	0.699
	(0.587)				
24	0.658	0.235	0.312	0.574	0.354
	(0.667)				
60	−0.084	−0.406	−0.285	1.193	−0.303
	(0.904)				
120	−0.429	−0.869	−0.728	0.968	−0.809
	(1.188)				

Estimated slope coefficients ($\widehat{\phi}^{(n)}$) from the regression
$y_{t+1}^{(n-1)} - y_t^{(n)} = \alpha^{(n)} + \phi^{(n)}(y_t^{(n)} - r_t)/(n-1) + \zeta_t^{(n)}$. Standard errors for the data are shown in parentheses. Model-implied slope coefficients are estimated using a data set with 100,000 periods simulated from the model, conditional on the point estimates of the parameters. The 'benchmark' model is a four-factor model estimated using data for May 1997-December 2013. The 'bias corrected' model applies the bootstrap bias correction, as described in Section 5.1. The 'survey' model includes survey expectations of future short-term interest rates, as described in Section 5.2. The 'unspanned macro' model incorporates macroeconomic variables as unspanned factors, as described in Section 5.3.

published quarterly from November 1999 and asks respondents (who are professional economists) for their expectations of Bank Rate at horizons one, two and three years ahead, which we take as a measure of the expected value of the one-month risk-free rate. This is the same set of Bank Rate surveys used by Guimarães (2014) (who also incorporates surveys of future inflation in a joint model of nominal and real bond yields).

Previous studies that have incorporated interest rate survey data do so within a maximum likelihood framework, jointly estimating the pricing factors and all parameters of the model together. We instead adapt ACM's approach to allow the inclusion of forecast data in the following way. First, we estimate the parameters δ_0 and $\boldsymbol{\delta}_1$ in (15). Second, we augment (1) to include the h-month ahead survey expectations (where as above we set $\boldsymbol{\mu} = \mathbf{0}$):

$$\begin{aligned} \begin{bmatrix} \mathbf{x}_{t+1} \\ E_t\left[r_{t+h}\right] \end{bmatrix} &= \begin{bmatrix} \mathbf{0} & \mathbf{\Phi} \\ \widehat{\delta}_0 & \widehat{\boldsymbol{\delta}}_1' \mathbf{\Phi}^h \end{bmatrix} \begin{bmatrix} 1 \\ \mathbf{x}_t \end{bmatrix} + \boldsymbol{\eta}_t \\ \boldsymbol{\eta}_t &\sim \mathcal{N}\left(\mathbf{0}, \mathbf{R}_\eta\right) \\ \mathbf{R}_\eta &= diag\left(\begin{bmatrix} \mathbf{\Sigma} & \sigma_s^2 \end{bmatrix}\right) \end{aligned} \tag{27}$$

Table 6: Model performance against LPY(ii) test

Maturity (months)	Data	Benchmark	Bias corrected	Survey	Unspanned macro
6	1	1.343	1.338	1.583	1.166
		(0.330)	(0.330)	(0.321)	(0.330)
18	1	1.437	1.369	1.473	1.238
		(0.644)	(0.664)	(0.563)	(0.660)
24	1	1.413	1.337	1.163	1.226
		(0.718)	(0.749)	(0.608)	(0.751)
60	1	1.234	1.184	−0.366	1.111
		(0.942)	(1.023)	(0.735)	(1.008)
120	1	1.176	1.271	−0.487	1.098
		(1.007)	(1.177)	(0.624)	(1.061)

Estimated slope coefficients ($\widehat{\phi}^{*(n)}$) from the regression $y_{t+1}^{(n-1)} - y_t^{(n)} - e_t^{(n)}/(n-1) = \alpha^{*(n)} + \phi^{*(n)}(y_t^{(n)} - r_t)/(n-1) + \zeta_t^{*(n)}$. Standard errors are shown in parentheses. The 'benchmark' model is a four-factor model estimated using data for May 1997-December 2013. The 'bias corrected' model applies the bootstrap bias correction, as described in Section 5.1. The 'survey' model includes survey expectations of future short-term interest rates, as described in Section 5.2. The 'unspanned macro' model incorporates macroeconomic variables as unspanned factors, as described in Section 5.3.

We include survey expectations 12, 24 and 36 months ahead, assuming that the mean of the individual responses is the expectation of Bank Rate at each horizon. We estimate this system (27) using maximum likelihood. This allows survey data to inform the path of expectations but the presence of measurement error on the surveys (with variance σ_s^2) means that the model is not constrained to fit the surveys exactly (as in previous studies that have incorporated surveys into ATSMs). Finally, we can proceed as before with steps 2 and 3 of the standard ACM technique set out in Section 2.2.

Figure 8 reports estimates of the ten-year term premium from a four-factor model that includes surveys alongside our benchmark model that does not. Including surveys has a marked effect on the estimates: although the average term premium is almost identical (at 0.75%), there is less variation in premia from the survey model. Although for the most part the ten-year term premium estimates from the survey model fall within the 80% confidence interval for the benchmark model, there are periods when this is not the case. For example, the survey model does not suggest that there was a sharp rise in term premia during the financial crisis in 2008/09. Moreover, the level of the term premium from the survey model ends the sample around 1.25 percentage points lower than from the benchmark model. Figure 9 illustrates why, by showing the expected path of Bank

Figure 8: 10-year term premium from benchmark model without surveys and an equivalent model using interest rate survey data

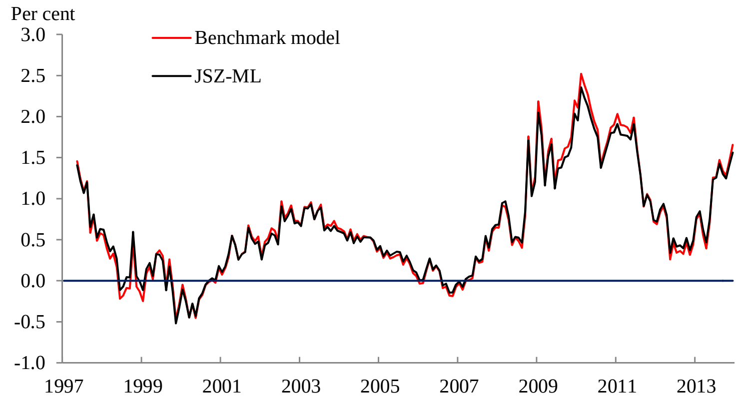

Figure 9: Expected path of short-term risk-free interest rates in October 2013, taken from benchmark model (without surveys) and a model estimated using survey data; with expectations of short rates from the Survey of External Forecasters and the observed forward curve

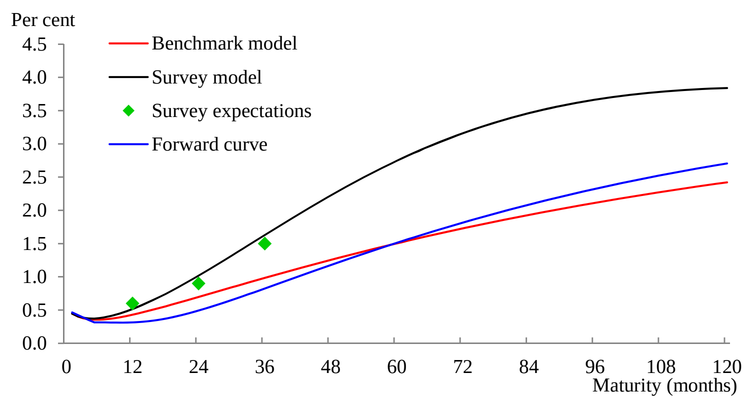

Rate from the two models at the end of October 2013 (the date of the final survey observation in the sample). The expected rate path from the survey model rises much more sharply, broadly in line with survey expectations at the time, whereas the expected path in the benchmark model rises more slowly and is more similar to market forward rates.

One way of assessing the weight we should place on models that include surveys is to consider how they affect performance according to the LPY tests, which is illustrated in Tables 5 and 6. These tests suggest that including the surveys has a detrimental effect on the model. The slope coefficients from the unadjusted Campbell and Shiller (1991) regressions (LPY(i)) are below those observed in the data at short maturities and above at long maturities, although the differences are not significant. But slope coefficients from the risk-adjusted regressions, (LPY(ii)) are negative and significantly below one at longer maturities.

One possible explanation for the poor performance of this particular specification for the UK is that we only incorporate surveys at horizons up to three years ahead: Kim and Wright (2005) argue that the inclusion of *long*-horizon surveys (six to eleven years ahead) plays an important role in pinning down US term premium estimates for long maturities. Unfortunately, we are not aware of any equivalent long-horizon surveys for UK short-term interest rates. On the other hand, our results appear consistent with findings by Piazzesi et al. (2013) that survey-based forecasts of interest rates are made as if respondents believed the level and slope of the term structure are more persistent than in the data. This results in more variation in long-horizon forecasts of interest rates.

Another, less formal, way of assessing whether models are reasonable is to consider whether the resulting term premium estimates are economically plausible. For example, previous studies have argued that term premia should be countercyclical (for example, Bauer et al. (2012)) and positively related to the degree of uncertainty around inflation (Wright (2011)). We explore this further by regressing 10-year term premia ($TP_t^{(120)}$) on four macroeconomic variables:

$$TP_t^{(120)} = \alpha_0 + \alpha_1 U_t + \alpha_2 INF_t + \alpha_3 UGDP_t + \alpha_4 UINF_t + u_t, \tag{28}$$

where U_t is the seasonally adjusted Labour Force Survey measure of UK unemployment among those aged 16 and over; INF_t is the Consensus Forecasts survey-based RPI inflation expectation for the next twelve months; and $UGDP_t$ and $UINF_t$ are Consensus Forecasts survey dispersion measures for next-year real GDP growth and inflation respectively. Consensus Forecasts reports RPI inflation and GDP growth (percentage change over the previous year) forecasts for a panel of professional forecasters. The inflation expectation measure is computed as the mean of the point forecasts across respondents; and the dispersion measures as the standard deviations of the point forecasts. Although cross-sectional dispersion in beliefs is not the same as aggregate uncertainty, Rich et al. (1992) show that measures of inflation dispersion are positively and significantly related to measures of inflation uncertainty for the US so (following Wright (2011)) it is reasonable to think of these measures as proxies for aggregate uncertainty.

Coefficient estimates and standard errors for the benchmark model and the model including surveys are reported in Table 7. In the benchmark model the term premium is a countercyclical variable: the coefficient on unemployment is positive and significant (particularly in a univariate regression). It rises during economic downturns (when unemployment rises) and falls during upswings (when unemployment falls). Figure 10 illustrates this graphically. Such cyclical behaviour is intuitively reasonable and is consistent with the findings of Bauer et al. (2012) for the US. In contrast, the sign on unemployment in the survey model is negative in both multivariate and univariate regressions. We can gain further insight into this result by considering a regression of ten-year term premia on the component of yields that reflects expectations of future short-term interest rates:

$$TP_t^{(120)} = \alpha_0 + \alpha_1 \left(y_t^{(120)} - TP_t^{(120)} \right) + u_t. \tag{29}$$

Results from these regressions are reported in Table 8. If the term premium is countercyclical and if interest rate expectations tend to fall in recessions and rise during periods of strong growth, we would expect a negative correlation between the term premium and expectations of future interest rates. This is indeed the case in the benchmark model but not in the survey model.

The results in Table 7 also show that in the benchmark model term premia are significantly higher when inflation expectations and inflation uncertainty, as proxied by the inflation dispersion measure, are high. Figure 11 illustrates this graphically for the inflation uncertainty measure. This is consistent with the results of Wright (2011), who finds that the largest declines in term premia since 1990 took place in countries that had changed their monetary policy frameworks to introduce inflation targeting and/or increased central bank independence; these countries witnessed larger declines in inflation uncertainty.[5] The survey model, in contrast, does not suggest there is a significant relationship between inflation uncertainty and ten-year term premia.

5.3 Unspanned macro factors

The fact that estimated term premia are countercyclical and related to the uncertainty about inflation motivates an exploration of whether the inclusion of macroeconomic variables as factors within the model affects term premium estimates. The inclusion of macroeconomic variables in ATSMs has become increasingly popular in recent years, following Ang and Piazzesi (2003) among others.

[5] In addition, Buraschi and Whelan (2012) find that an inflation dispersion measure can help explain excess bond returns.

Figure 10: 10-year term premium from benchmark and LFS seasonally adjusted unemployment rate among those aged 16 and over

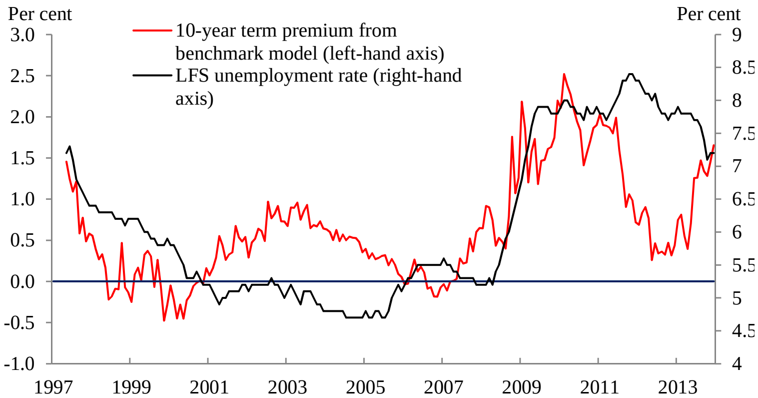

Figure 11: 10-year term premium from benchmark model and inflation dispersion measure based on Consensus surveys

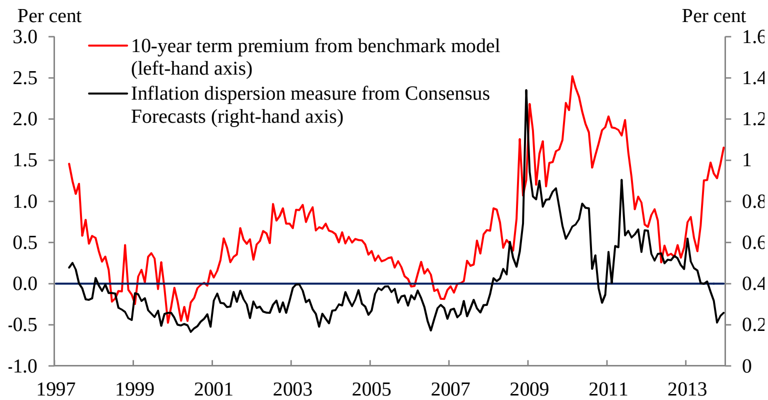

Table 7: Regressions of ten-year term premia on macroeconomic variables

Benchmark model					
Constant	-1.913^{***}	-1.371^{***}	-0.565	0.035	-0.281^{***}
	(0.288)	(0.192)	(0.305)	(0.160)	(0.082)
U_t	0.080^{*}	0.330^{***}	–	–	–
	(0.048)	(0.031)			
INF_t	0.423^{***}	–	0.486^{***}	–	–
	(0.109)		(0.120)		
$UGDP_t$	0.558^{**}	–	–	1.224^{***}	–
	(0.258)			(0.300)	
$UINF_t$	1.816^{***}	–	–	–	2.351^{***}
	(0.304)				(0.186)
Survey model					
Constant	0.381	1.138^{***}	0.188	0.652^{***}	0.815^{***}
	(0.252)	(0.153)	(0.203)	(0.108)	(0.071)
U_t	-0.186^{***}	-0.077^{***}	–	–	–
	(0.042)	(0.024)			
INF_t	0.479^{***}	–	0.187^{**}	–	–
	(0.096)		(0.080)		
$UGDP_t$	0.128	–	–	0.018	–
	(0.226)			(0.203)	
$UINF_t$	0.390	–	–	–	-0.384^{**}
	(0.266)				(0.160)

Coefficient estimates and estimated standard errors (in parentheses) from a regression of 10-year term premia on the variables listed in the table:
$TP_t^{(120)} = \alpha_0 + \alpha_1 U_t + \alpha_2 INF_t + \alpha_3 UGDP_t + \alpha_4 UINF_t + u_t$. * indicates significance at the 10% level, ** indicates significance at the 5% level and *** indicates significance at the 1% level. The 'benchmark model' is a four-factor model estimated using data for May 1997-December 2013. The 'survey' model includes survey expectations of future short-term interest rates, as described in Section 5.2.

More recently, Joslin et al. (2014) find evidence that macroeconomic variables are 'unspanned' by the yield curve, yet contain information about future excess returns on bonds beyond that which is contained in bond yields and are therefore important variables for the time series dynamics of yield curve factors and hence may be important for correctly measuring term premia. In this section, we therefore consider an alternative model with unspanned macroeconomic variables.

Evidence that the four macroeconomic variables considered in the previous section are unspanned by UK bond yields is provided by regressing those variables on the first four principal components of yields:

$$z_t = \alpha_0 + \alpha_1 x_{1,t} + \alpha_2 x_{2,t} + \alpha_3 x_{3,t} + \alpha_4 x_{4,t} + u_t \tag{30}$$

Table 8: Regressions of ten-year term premia on the expectations component of ten-year yields

	Benchmark model	Survey model
Constant	-1.798^{***}	-0.481^{***}
	(0.092)	(0.114)
Expected rates	-0.309^{***}	0.311^{***}
	(0.023)	(0.030)

Coefficient estimates and estimated standard errors (in parentheses) from a regression of 10-year term premia on the expected component of 10-year yields: $TP_t^{(120)} = \alpha_0 + \alpha_1 \left(y_t^{(120)} - TP_t^{(120)} \right) + u_t$. * indicates significance at the 10% level, ** indicates significance at the 5% level and *** indicates significance at the 1% level. The 'benchmark model' is a four-factor model estimated using data for May 1997-December 2013. The 'survey' model includes survey expectations of future short-term interest rates, as described in Section 5.2.

where z_t denotes one of U_t, INF_t, $UGDP_t$ or $UINF_t$. Results are reported in Table 9. The yield curve factors can explain around half of the variation in the unemployment rate, inflation expectations and the inflation dispersion measure; and much less for the GDP dispersion measure, with an R^2 of 0.16: consistent with the findings of Joslin et al. (2014) for the US, the macroeconomic variables are not (linearly) spanned by the yield curve.

Table 9: Regressions of macroeconomic variables on yield curve factors

	U_t	INF_t	$UGDP_t$	$UINF_t$
Constant	6.166^{***}	2.522^{***}	0.512^{***}	0.401^{***}
	(0.050)	(0.018)	(0.010)	(0.010)
$x_{1,t}$	-0.873^{***}	-0.103^{***}	-0.018^{*}	-0.111^{***}
	(0.050)	(0.018)	(0.010)	(0.010)
$x_{2,t}$	0.209^{***}	-0.030	0.033^{***}	0.068^{***}
	(0.050)	(0.018)	(0.010)	(0.010)
$x_{3,t}$	0.231^{***}	0.189^{***}	0.014	0.026^{***}
	(0.050)	(0.018)	(0.010)	(0.010)
$x_{4,t}$	0.311^{***}	0.167^{***}	-0.043^{***}	-0.024^{***}
	(0.050)	(0.018)	(0.010)	(0.010)
R^2	0.66	0.54	0.16	0.53

Table provides coefficient estimates and standard errors (in parentheses) from the regression $z_t = \alpha_0 + \alpha_1 x_{1,t} + \alpha_2 x_{2,t} + \alpha_3 x_{3,t} + \alpha_4 x_{4,t} + u_t$, where z_t denotes one of the variables in the column headings. * indicates significance at the 10% level, ** indicates significance at the 5% level and *** indicates significance at the 10% level.

At least some of these macroeconomic variables contain information that helps predict future excess returns on UK bonds. Table 10 reports results of regressions of annual excess returns on a 10-year bond relative to a 1-year bond ($rx_{t,t+12}^{(120)}$) on yield curve factors and the same set of macroeconomic variables:

$$\begin{aligned} rx_{t,t+12}^{(120)} &= \alpha_0 + \alpha_1 x_{1,t} + \alpha_2 x_{2,t} + \alpha_3 x_{3,t} + \alpha_4 x_{4,t} \\ &\quad + \alpha_5 U_t + \alpha_6 INF_t + \alpha_7 UGDP_t + \alpha_8 UINF_t + u_t. \end{aligned} \tag{31}$$

The results are broadly consistent with the regressions of term premia on macroeconomic variables reported in Table 7. The joint hypothesis that the coefficients on the four macroeconomic variables are equal to zero is clearly rejected, with an F-test statistic that is significant at the 1% significance level. The coefficients on unemployment, the inflation expectations and inflation dispersion measures are individually significant, although the coefficient on the GDP dispersion measure is insignificant even at the 10% significance level.

These results provide clear motivation for a no-arbitrage term structure model that includes three unspanned macro variables - unemployment and the inflation expectations and dispersion measures - in addition to the four spanned yield factors used in the benchmark model. The vector of pricing factors can be written as

$$\mathbf{x}_t = \begin{bmatrix} \mathbf{x}_t^{s\prime} & \mathbf{x}_t^{u\prime} \end{bmatrix}',$$

where $\mathbf{x}_t^s$ is a vector of the four spanned factors and $\mathbf{x}_t^u$ contains the three unspanned macro factors. The requirement that yields do not load on the unspanned factors contemporaneously requires that:

$$\begin{aligned} \boldsymbol{\delta}_1 &= \begin{bmatrix} \boldsymbol{\delta}_{1,s}' & \mathbf{0}_{(3\times 1)}' \end{bmatrix}', \\ \boldsymbol{\Phi}^{\mathbb{Q}} &= \boldsymbol{\Phi} - \boldsymbol{\lambda}_1 = \begin{bmatrix} \boldsymbol{\Phi}_{ss}^{\mathbb{Q}} & \mathbf{0}_{4\times 3} \\ \boldsymbol{\Phi}_{us}^{\mathbb{Q}} & \boldsymbol{\Phi}_{uu}^{\mathbb{Q}} \end{bmatrix}. \end{aligned}$$

We also follow the convention of setting $\boldsymbol{\Phi}_{us}^{\mathbb{Q}} = \boldsymbol{\Phi}_{us}$ and $\boldsymbol{\Phi}_{uu}^{\mathbb{Q}} = \boldsymbol{\Phi}_{uu}$, since these parameters are not separately identified. We refer the reader to ACM for details of how the estimation procedure is modified to allow for unspanned factors.

Figure 12 shows estimates of ten-year term premia from this model. These are close to those from the benchmark model for most of the sample. Since 2012, the estimates from the model with

Table 10: Regressions of annual excess returns on ten-year bonds on yield curve factors and macroeconomic variables

	Including macro variables	Yield factors only
Constant	-8.943^{**}	3.259^{***}
	(2.910)	(0.692)
$x_{1,t}$	1.626^{***}	-0.372
	(0.466)	(0.450)
$x_{2,t}$	3.196^{***}	3.846^{***}
	(0.374)	(0.861)
$x_{3,t}$	1.238^{***}	2.775^{***}
	(0.294)	(0.590)
$x_{4,t}$	2.401^{***}	2.155^{***}
	(0.280)	(0.820)
U_t	2.347^{***}	–
	(0.579)	
INF_t	-2.536^{**}	–
	(1.245)	
$UGDP_t$	3.373	–
	(2.412)	
$UINF_t$	5.257^{*}	–
	(3.231)	
R^2	0.693	0.567
$F(4,179)$	–	42.433^{***}

The first column reports coefficient estimates and estimated standard errors (in parentheses) from the regression
$rx^{(120)}_{t,t+12} = \alpha_0 + \alpha_1 x_{1,t} + \alpha_2 x_{2,t} + \alpha_3 x_{3,t} + \alpha_4 x_{4,t} + \alpha_5 U_t + \alpha_6 INF_t + \alpha_7 UGDP_t + \alpha_8 UINF_t + u_t$
where $rx^{(120)}_{t,t+12}$ is the excess holding period return from a 10-year bond relative to a one-year bond between months t and $t+12$. The second column reports the equivalent results from a regression including only the yield curve factors, i.e. with $\alpha_5 = \alpha_6 = \alpha_7 = \alpha_8 = 0$. The final row reports the F-test statistics of this joint hypothesis. Standard errors are estimated using the Newey-West method with a lag length of 12. * indicates significance at the 10% level, ** indicates significance at the 5% level and *** indicates significance at the 1% level.

macroeconomic factors has been lower than those from the benchmark model but remain well within the 80% confidence interval for the benchmark model. The unspanned macro-factor model delivers extremely similar results compared with the benchmark model in terms of performance against the LPY(i) and LPY(ii) tests (Tables 5 and 6), which suggests that while these macroeconomic variables may have explanatory power for annual excess returns, if the object of interest is the term premium in UK long-term bond yields there is little to choose in practical terms between a model that incorporates those variables as unspanned factors and one that does not.

Figure 12: 10-year term premium from benchmark model and unspanned macro factor model

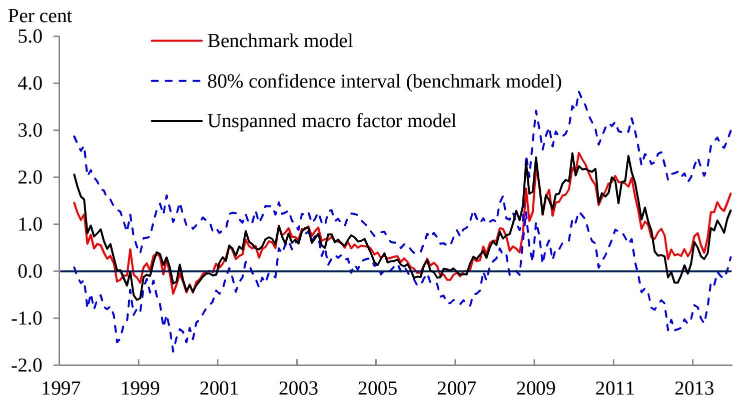

6 Impact of the zero lower bound

One drawback of Gaussian ATSMs is that they do not impose the zero lower bound on nominal rates. If bond yields remain well above zero, this is likely to be a trivial concern, since the model-implied probability of negative rates will be small. As a number of other studies have documented, however, as nominal interest rates fall close to zero, as in recent years, Gaussian ATSMs can imply a substantial probability of negative rates (e.g. Andreasen and Meldrum (2013) and Bauer and Rudebusch (2014)). A number of recent studies have therefore used alternative no-arbitrage models of yields, such as the shadow rate framework proposed by Black (1995). In a multi-factor shadow rate model, the law of motion for the factors (1) and specification of the price of risk (5) are the same as in an ATSM. However, rather than being affine in the factors, the short-term interest rate is assumed to be the maximum of zero and a 'shadow rate' (s_t) that is affine in the factors and can therefore be negative:

$$r_t = \max\{0, s_t\}, \tag{32}$$

$$s_t = \delta_0 + \boldsymbol{\delta}_1' \mathbf{x}_t. \tag{33}$$

Kim and Priebsch (2013) find that estimates of long-term term premia are not substantially affected by imposing the lower bound using a shadow rate model. To explore whether a similar result applies for the UK, we estimate a shadow rate model using our data set and a variation of a estimation technique proposed by Bauer and Rudebusch (2014). They first estimate an ATSM using a sub-sample that ends before the recent period of low bond yields. They then assume that the parameters of the model are stable across the full sample period and filter the unobserved pricing factors for the full sample.

In our case, we first estimate the parameters of an ATSM using the ACM method applied to a sample ending in December 2008, when Bank Rate was 2% (the UK Monetary Policy Committee subsequently lowered its policy rate to 1.5% in January 2009, to 1% in February 2009 and to 0.5% in March 2009). Over the period up to and including December 2008, we take the pricing factors to be observed principal components of yields as before, i.e. we assume that the mapping between yields and factors is linear and unaffected by the zero bound when yields are at historically normal levels. Given the non-linear mapping between yields and factors in the shadow rate model, this assumption is less likely to be reasonable when yields are low. We therefore use the Central Difference Kalman Filter of Norgaard et al. (2000) to filter the factors from January 2009 onwards.[6] Since it is not possible to solve for bond prices in the shadow rate model analytically, to obtain bond prices during this period we apply the technique proposed by Priebsch (2013), who approximates yields by taking a second-order approximation to the pricing equation (2) (we refer the reader to that paper for details). Term premia can then be computed as the difference between model-implied yields and the average expectation of future short rates, which are given by the conditional mean of a truncated Normal variable, i.e.:

$$TP_t^{(n)} = y_t^{(n)} - \frac{1}{n}\sum_{i=0}^{n-1} E_t \max\{0, s_{t+i}\}. \tag{34}$$

These conditional expectations can be computed straightforwardly using the formulae provided by Priebsch (2013).

Figure 13 shows estimates of 10-year term premia from the shadow rate model. In addition to the estimates from our benchmark ATSM reported above, the figure also shows estimates obtained from an alternative ATSM estimated over the sub-sample ending in December 2008, to ensure a fair comparison. The estimates from all three models are extremely similar, which suggests that the proximity of short-term yields to the zero lower bound has little impact on estimates of long-term term premia, at least if we use this method to estimate the shadow rate model.

[6] We initialise the filter by imposing that the posterior mean for the factors in December 2008 is equal to the observed principal components; and that the posterior covariance for the factors in this period is $10^{-16} \times \mathbf{I}$.

Figure 13: 10-year term premium from benchmark model and shadow rate model

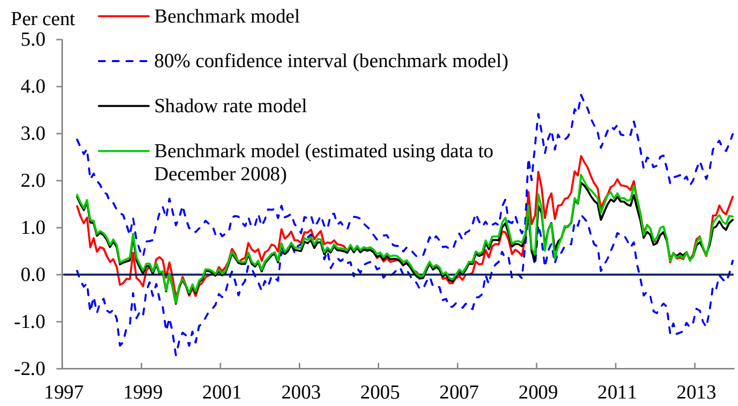

7 Conclusion

This paper applies the method for estimating affine term structure models recently proposed in Adrian et al. (2013) to UK nominal government bond yields. We find evidence that four pricing factors are required to achieve a good fit to the short end of the yield curve and to match the LPY tests of Dai and Singleton (2002). The resulting term premium estimates display countercyclical behaviour (consistent with Bauer et al. (2012) for the US) and are positively related to uncertainty about future inflation (consistent with Wright (2011) for a panel of countries, including the UK).

Term premium estimates are robust to correcting for small sample bias as proposed by Bauer et al. (2012), or including unspanned macroeconomic variables as factors, similar to the approach of Joslin et al. (2014). On the other hand, including short-term interest rate survey data at horizons up to three years ahead to help identify the time series dynamics of the pricing factors has a marked impact on premia, which are much less volatile compared with a model that does not include surveys. The survey model performs relatively poorly against the LPY tests and premia are not countercyclical or significantly associated with future inflation uncertainty. Finally, although the ATSM does not impose the zero lower bound on nominal interest rates, we show using a shadow rate term structure model estimated using a variant of the technique proposed by Bauer and Rudebusch (2014) that this does not have a substantial impact on estimates of long-maturity term premia, consistent with previous findings by Kim and Priebsch (2013) for the US.

References

Adrian, T., R. K. Crump, and E. Moench (2013). Pricing the term structure with linear regressions. *Journal of Financial Economics 110*, 110–138.

Anderson, N. and J. Sleath (2001). New estimates of the UK real and nominal yield curves. *Bank of England Working Paper 126.*

Andreasen, M. M. and A. C. Meldrum (2013). Likelihood inference in non-linear term structure models: the importance of the zero lower bound. *Bank of England Working Paper 481.*

Andreasen, M. M. and A. C. Meldrum (2014). Dynamic term structure models: the best way to enforce the zero lower bound. *Unpublished working paper*.

Ang, A. and M. Piazzesi (2003). A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. *Journal of Monetary Economics 50*, 745–787.

Bauer, M. D. and G. D. Rudebusch (2014). Monetary policy expectations at the zero lower bound. *Federal Reserve Bank of San Francisco Working Paper 2013-18.*

Bauer, M. D., G. D. Rudebusch, and J. C. Wu (2012). Correcting estimation bias in dynamic term structure models. *Journal of Business & Economic Statistics 30*, 454–467.

Black, F. (1995). Interest rates as options. *Journal of Finance 50*, 1371–1376.

Buraschi, A. and P. Whelan (2012). Term structure models with differences in beliefs. *Unpublished working paper*.

Campbell, J. Y. and R. J. Shiller (1991). Yield spreads and interest rate movements: a bird's eye view. *NBER Working Paper 3153.*

Cochrane, J. H. and M. Piazessi (2005). Bond risk premia. *American Economic Review 95*, 138–160.

Dai, Q. and K. J. Singleton (2002). Expectation puzzles, time-varying risk premia, and affine models of the term structure. *Journal of Financial Economics 63*, 415–441.

D'Amico, S., D. H. Kim, and M. Wei (2014). Tips from TIPS: the informational content of Treasury Inflation-Protected Security prices. *Federal Reserve Board Finance and Economics Discussion Series 2014-24*.

Duffee, G. R. (2002). Term premia and interest rate forecasts in affine models. *Journal of Finance 57*, 405–443.

Duffee, G. R. (2011). Information in (and not in) the term structure. *Review of Financial Studies 24*, 2895–2934.

Guimarães, R. (2014). Expectations, risk premia and information spanning in dynamic term structure model estimation. *Bank of England Working Paper 489*.

Hamilton, J. D. and J. C. Wu (2012). Identification and estimation of affine term structure models. *Journal of Econometrics 168*, 315–331.

Joslin, S., M. A. Priebsch, and K. J. Singleton (2014). Risk premiums in dynamic term structure models with unspanned macro risks. *Journal of Finance 69*, 1197–1233.

Joslin, S., K. J. Singleton, and H. Zhu (2011). A new perspective on Gaussian dynamic term structure models. *Review of Financial Studies 24*, 926–970.

Joyce, M. A. S., P. Lildholdt, and S. Sorensen (2010). Extracting inflation expectations and inflation risk premia from the term structure: a joint model of the UK nominal and real yield curves. *Journal of Banking and Finance 34*, 281–294.

Kaminska, I. (2013). A no-arbitrage structural vector autoregressive model of the UK yield curve. *Oxford Bulletin of Economics and Statistics 75*, 680–704.

Kilian, L. (1998). Small-sample confidence intervals for impulse response functions. *Review of Economics and Statistics 80*, 218–230.

Kim, D. H. and A. Orphanides (2012). Term structure estimation with survey data on interest rate forecasts. *Journal of Financial and Quantitative Analysis 47*, 241–272.

Kim, D. H. and M. A. Priebsch (2013). Estimation of multi-factor shadow-rate term structure models. *Unpublished working paper*.

Kim, D. H. and J. H. Wright (2005). An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates. *Federal Reserve Board Finance and Economics Discussion Series 2005-33*.

Lildholdt, P., C. Peacock, and N. Panigirtzoglou (2007). An affine macro-factor model of the UK yield curve. *Bank of England Working Paper 322*.

Litterman, R. and Scheinkman (1991). Common factors affecting bond returns. *Journal of Fixed Income 1*, 54–61.

Moench, E. (2008). Forecasting the yield curve in a data-rich environment: A no-arbitrage factor-augmented VAR approach. *Journal of Econometrics 146*, 26–43.

Norgaard, M., K. Poulsen, and O. Ravn (2000). Advances in derivative-free state estimation for non-linear systems. *Automatica 36*, 1627–1638.

Piazzesi, M. (2003). *Handbook of Financial Econometrics, Volume 1*, Chapter 12 Affine term structure models, pp. 691–758. Elsevier: Amsterdam, Netherlands.

Piazzesi, M., J. Salomao, and M. Schneider (2013). Trend and cycle in bond risk premia. *Unpublished working paper*.

Priebsch, M. A. (2013). Computing arbitrage-free yields in multi-factor Gaussian shadow-rate term structure models. *Federal Reserve Board Finance and Economics Discussion Series 2013-63*.

Rich, R. W., J. E. Raymond, and J. S. Butler (1992). The relationship between forecast dispersion and forecast uncertainty: evidence from a survey data - ARCH model. *Journal of Applied Econometrics 7*, 131–148.

Wright, J. H. (2011). Term premia and inflation uncertainty: empirical evidence from an international panel dataset. *American Economic Review 101*, 1514–1534.