Journal of Banking & Finance 67 (2016) 85–102

Contents lists available at ScienceDirect

Journal of Banking & Finance

journal homepage: www.elsevier.com/locate/jbf

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

Sheheryar Malik [1], Andrew Meldrum *

Macro Financial Analysis Division, Bank of England, Threadneedle Street, London EC2R 8AH, UK

ARTICLE INFO

Article history:
Received 24 February 2015
Accepted 17 February 2016
Available online 5 March 2016

JEL classification:
E43
G10
G12

Keywords:
Affine term structure model
Term premia
Bias correction
Interest rate surveys

ABSTRACT

Dynamic no-arbitrage affine term structure models (ATSMs) have become the standard framework for monetary policy-makers to decompose long-term bond yields into expectations of future short-term risk-free interest rates and the term premia that compensate investors in long-term bonds for risk. This paper presents estimates of ATSMs for the UK and explores how much weight users of these models can place on point estimates of term premia. Over much of the period since the early 1990s, broad movements in estimated premia are robust across a wide range of reasonable specifications. But there is substantial model and parameter uncertainty associated with these models and estimates of the time-series dynamics of yields may be biased in short samples. This model uncertainty is greater towards the end of our sample period, when bond yields have been well below historically normal levels.

1. Introduction

Understanding movements in the term structure of government bond yields is of considerable importance to financial market practitioners and public policy-makers. Long-term bond yields reflect both expectations of average future short-term interest rates and a term premium - 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: expected short rates reflect investors' views about the outlook for monetary policy, while the term premium reflects (among other things) uncertainty around future short-term interest rates and investors' aversion to bearing risk. Unfortunately, however, the two components cannot be observed separately.

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. In this paper, we find that point estimates of UK term premia from a four-factor ATSM estimated over the period October 1992 to December 2013 appear reasonable by a number of metrics. They 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). The model matches the 'linear projections of yields' (LPY) specification tests proposed by Dai and Singleton (2002). And broad movements in premia appear plausible: they fell in the late 1990s, which may reflect improvements in the credibility of monetary policy and an increased demand for the safety of government bonds following the Asian crisis; they were relatively low through much of the 2000s, including the 'Greenspan conundrum' period; and they rose sharply but temporarily during the financial crisis of 2008/09.

A general difficulty, however, is that estimation of ATSMs is fraught with problems associated with weak identification and computationally intensive optimization steps.[2] This has two conse-

☆ The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. The authors are grateful to Andy Blake, Iryna Kaminska, Chris Young and our anonymous referees for helpful comments.

* Corresponding author. Tel.: +44 (0)20 7601 5607.
E-mail addresses: smalik2@imf.org (Sheheryar Malik), andrew.meldrum@bankofengland.co.uk (A. Meldrum).

[1] Present address: International Monetary Fund, 700 19th Street N.W., Washington D.C. 20431, United States. Tel.: +1 202 623 8769.

[2] For example, previous studies that have estimated ATSMs using UK data have typically relied on maximum likelihood estimation (e.g. Lildholdt et al., 2007; Joyce et al., 2010; Kaminska, 2013; D'Amico et al., 2014; Guimarães, 2014). This involves a high-dimensional non-linear optimization over a likelihood surface that has many local optima and undefined regions. Hamilton and Wu (2012) discuss these issues in greater detail.

http://dx.doi.org/10.1016/j.jbankfin.2016.02.006
0378-4266/

quences. First, estimates of term premia from a single model have wide confidence bounds; and second, the challenges associated with estimating just a single model mean that establishing which properties of estimated term premia are robust across models has inevitably proved somewhat challenging. An obvious question is therefore just how much weight can policy-makers and other users place on estimates of term premia obtained from these models? To address this question, we apply the estimation approach proposed by Adrian et al. (2013) (ACM henceforth), who split the estimation into a series of linear regressions, which greatly reduces the time taken to estimate the model.[3]

In some respects, estimates of term premia for the UK are extremely robust. We highlight in particular that it makes little difference to estimates of term premia if we include additional macroeconomic variables as unspanned factors, as proposed by Joslin et al. (2014); if we extend the model to allow for a lower bound on nominal interest rates, as suggested by Black (1995); or if we vary the number of pricing factors between three and six. There is nevertheless substantial model and parameter uncertainty associated with estimating the time-series dynamics of term premia, which suggests that we should be particularly cautious about drawing strong inferences based on dynamic term structure models at times when interest rates are a long way from normal levels, as has been the case in recent years. We highlight two issues in particular. First, as discussed by Bauer et al. (2012), OLS estimates of the dynamics of the pricing factors driving yields will be biased in small samples; and correcting for this small-sample bias can have a substantial impact on estimated US term premia. Encouragingly, we find that applying a similar bias correction to the UK does not result in a materially different interpretation of past movements in bond yields for the majority of our sample period. But estimates of term premia towards the end of the sample from a bias-corrected version of the model are higher than from our benchmark model. The bias correction increases the estimated persistence of short-term interest rates, which means that it takes substantially longer for model-implied expected short rates to rise from the very low levels experienced since 2009; and with the expected path of short rates lower, the term premium is correspondingly higher.

Second, it is hard to estimate the persistence of yields precisely given the sample of yields available. A popular approach for improving the identification of dynamic term structure models is to include additional information in the form of survey expectations of future short-term interest rates (first proposed by Kim and Orphanides, 2012 and applied to UK data by Joyce et al., 2010; Guimarães, 2014). We adapt the ACM method to allow the inclusion of survey expectations of future short-term interest rates. Encouragingly, we again find that the broad pattern of movements in term premia from our 'survey model' is similar to that from our benchmark model. In contrast to the bias-corrected model, however, the main difference is that term premia from the survey model are slightly *lower* towards the end of the sample.

In the face of this model uncertainty, what practical advice can we offer policy-makers and others users of these models when seeking to estimate term premia for the UK? While it may be tempting to prefer a model in which the time-series dynamics have been bias-corrected, there are reasons to be cautious about this. In particular, the bias-corrected model performs relatively poorly against the Dai and Singleton (2002) LPY specification tests and the model-implied path of expected short rates towards the end of the sample is extremely low – perhaps implausibly so. There are also reasons to be cautious about the potential gains from including surveys in the model for the UK, since this also results in substantially inferior performance against the Dai and Singleton (2002) LPY specification tests. One possible explanation may be that UK interest rate survey expectations are only available for relatively short forecast horizons, so may be less informative about the persistence of interest rates than in the US.

The remainder of this paper is organized as follows. Section 2 sets out the standard no-arbitrage ATSM and the estimation technique of ACM. Section 3 introduces the data set and discusses the estimates of term premia from our benchmark model. Section 4 discusses the most important sources of uncertainty in estimates of UK term premia, related to the difficulty in specifying and precisely estimating the time-series dynamics of bond yields. Section 5 evaluates the robustness of term premium estimates to the inclusion of unspanned macroeconomic factors (as proposed by Joslin et al., 2014) and to the imposition of a zero lower bound on nominal interest rates. Section 6 concludes. The appendices to the paper provide a number of additional robustness checks.

2. Gaussian affine term structure models

2.1. Excess returns

This section sets out the key equations of a standard Gaussian ATSM, following the exposition from ACM. 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, homoskedastic 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 assumption of no-arbitrage implies the existence of a pricing kernel M_{t+1} such that

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

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

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

where $r_t = -\ln P_t^{(1)}$ denotes the continuously compounded one-period risk-free rate, which is 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

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

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

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

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

[3] A number of other 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), Kaminska (2013) and Andreasen and Meldrum (2015a). All of these methods involve some non-linear optimization, which ACM avoid entirely because they do not impose the no-arbitrage restrictions inside the estimation procedure. They show, however, that the factor loadings implied by their approach satisfy these restrictions to a high degree of precision. In Appendix A we show that our benchmark estimates of term premia are almost identical to those obtained using more standard maximum likelihood techniques to estimate the cross-sectional dynamics of the yield curve while also imposing no-arbitrage.