Journal of Econometrics 159 (2010) 166–182

Contents lists available at ScienceDirect

Journal of Econometrics

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

No-arbitrage macroeconomic determinants of the yield curve

Ruslan Bikbov[a], Mikhail Chernov[b,*]

[a] Federal Reserve Board, Washington, DC 20551, USA

[b] London Business School, London School of Economics, and CEPR, London, NW1 4SA, UK

ARTICLE INFO

Article history:
Received 29 October 2008
Received in revised form
8 October 2009
Accepted 24 May 2010
Available online 2 June 2010

JEL classification:
C13
C22
G12

Keywords:
Macro-finance models
Term structure
Variance decomposition
Kalman filter

ABSTRACT

No-arbitrage macro-finance models use variance decompositions to gauge the extent of association between the macro variables and yields. We show that results generated by this approach are sensitive to the order of variables in the recursive identification scheme. In a four-factor model, one may obtain 18 different sets of answers out of 24 possible. We propose an alternative measure that is based on levels of macro variables as opposed to shocks. We account for the correlation between the macro and latent factors via projection of the latter onto the former. As a result, the association between macro variables and yields can be computed uniquely via an R^2. Macro variables explain 80% of the variation in the short rate and 50% of the slope, and 54% to 68% of the term premia.

1. Introduction

The important work of Ang and Piazzesi (2003) extends the traditional no-arbitrage term structure models by modelling jointly the dynamics of yields and nonfinancial variables. The nonfinancial variables chosen by researchers are typically macro variables, so we refer to the models as no-arbitrage macro-finance models. One of the most important issues that the setting allows researchers to address is whether yields are associated with macro variables. This issue could be broken down further. First, is the association between yields and macro variables determined by the monetary policy? There is a strong intuition from the Taylor rule literature that suggests that such macro variables as inflation and real activity should matter for the interest rate, which is the monetary policy instrument.[1] Second, is the association between yields and macro variables determined by the yield risk premia? It could be that macro variables affect the yields exclusively through the spot interest rate. Finally, regardless of the channel, how large is this association quantitatively? A precise answer to these questions will provide a quantitative benchmark for equilibrium models. If structural implications of these models are to be taken seriously, they should be able to generate, at the very least, the same common variation in yields and macro variables as the benchmark.

The answers to these questions that are found in the literature are mixed. On the one hand, as is well-known from the term structure studies using financial variables only, such as Litterman and Scheinkman (1991) or Dai and Singleton (2000), two or three factors are sufficient for capturing 95–98% of the variation in the yield curve. This conclusion suggests a small role for macro variables in a empirically realistic term structure model. On the other hand, because the macro-finance models can be cast in a vector autoregression (VAR) framework, it is natural to quantify the common variation in yields and macro variables via variance decomposition. Studies that follow this route attribute a large fraction (40–80%) of the unconditional variance in yields to the joint contribution of real activity and inflation, which suggests that macro determinants are very important.[2]

The reason why two strands of the term structure literature produce seemingly conflicting results is that the VAR literature

---

* Corresponding address: London Business School, Sussex Place, London, NW1 4SA, UK. Tel.: +44 20 7000 8258.

E-mail addresses: ruslan.bikbov@frb.gov (R. Bikbov), mchernov@london.edu (M. Chernov).

[1] However, the degree of association is not obvious. In fact, interest rate rules and, therefore, their ingredients, do not matter at all in the absence of frictions, such as price stickiness (e.g., Bils and Klenow, 2004; Woodford, 2003).

[2] For example, Ang et al. (2004), Ang and Piazzesi (2003), Dai and Philippon (2004), Rudebusch and Wu (2005), among others.

0304-4076/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.jeconom.2010.05.004

uses *shocks* to macro variables to explain variation in yields, while the work based only on financial data focuses on the *levels* of variables to explain the same variation. It should be possible to reconcile the two approaches because levels are aggregates of shocks. In this paper, we propose a general procedure that allows us to obtain a correct measurement of the association between yields and macro variables.

A simple example illustrates how such a reconciliation can be achieved. Consider the Ang and Piazzesi (2003) model, where macro variables are independent of the latent variables. The infinite horizon variance decompositions from the VAR representation of the model characterize how much of the unconditional yield variance can be attributed to shocks in each of the macro variables. Our main interest is in disentangling the contribution of the macro and latent variables, rather than in disentangling every single variable. Therefore, we can exploit the independence of the two types of variables (macro and latent) and determine the proportion of the yield variance that can be explained by macro variables simply by computing the corresponding R^2.

The two approaches will yield the same result under two important conditions. First, the shocks have to be identified correctly. Researchers often employ the recursive identification scheme in the context of reduced-form VARs, such as macro-finance models. There is no guarantee that a particular order of variables will produce the correct identification; hence, it is possible for there to be multiple sets of measures. Second, the VAR has to provide a good fit to the data, in that the variance of yields in the data and in the model should be similar. This is because the computation of variance decomposition does not involve actual data. In contrast, an R^2 measures how much of the actual variance is explained by a model.

The coefficient of determination, R^2, appears to be a more robust measure because it relies on observed macro variables as opposed to unobserved shocks and, therefore, is not subject to assumptions about identification. We emphasize that the R^2 will not only be a unique measure of association between macro variables and yields but also be a correct one. However, this measure cannot be implemented if macro and latent variables are correlated with each other. This would be the case if, for example, there is a feedback from the financial sector into the real economy.

We resolve this problem by exploiting the special nature of latent variables. Latent variables do not have any macro interpretation. As Dai and Singleton (2000) point out, these factors can be re-parameterized, or rotated, in a number of different ways without changing the value of yields. Thus, latent factors serve no purpose other than to fit the data well. These observations lead us to propose a novel procedure that allows us to express the traditional latent factors via macro variables and new "residual" latent variables that are conditionally orthogonal to the macro ones. As a result, we disentangle the contribution of the macro and the new latent variables to the variation of yields. Now we can implement the R^2 that we advocated above.

The key to our procedure is a conditional model-based projection of the latent variables onto the macro ones. The projection effectively breaks the latent factors down into (i) a component that consists of the current and past macro variables, and (ii) new "projection residual" latent variables that are conditionally orthogonal to the entire history of macroeconomic variables. By construction, the variance of the new latent factors is minimized and the variance of the macro component is maximized. The spot interest rate becomes a linear function of macro variables, their lags, and a set of new latent factors.[3]

The empirical implementation of our procedure is based on a four-factor Gaussian model with two latent factors and inflation and real activity acting as observable factors. To estimate the model, we use a panel of eight yields ranging from 3 months to 10 years, with inflation and real activity observed at a monthly frequency from 1970 to 2002. We provide diagnostics that indicate that the model fits the data reasonably well.

We document that the variance decomposition results are extremely sensitive to the order of the model factors chosen for the recursive identification. For example, the fraction of the interest rate variance that is explained by the joint shocks to inflation and real activity can vary between 36% and 77% depending on the assumed order of the factors. Similarly, the fraction of the explained slope variance varies between 25% and 81%.[4] In particular, the standard order of the variables, that is, macro variables followed by latent variables, produces 52% for the short interest rate and 37% for the slope. Our procedure produces a unique set of explained variances precisely because it is based on observed macro variables. In particular, 80% of the variation in the interest rate and 52% of the variation in the slope can be explained by inflation and real activity. A similar set of figures (76% and 59%, respectively) is obtained via the variance decomposition if the order of the variables is real activity followed by latent variables, which are in turn followed by inflation. Clearly, there is no way of knowing this "correct" order of shocks in the context of a reduced-form model.

The point that variance decomposition results are sensitive to the order of variables in the recursive identification scheme is not new in itself. Nevertheless, given the popularity of the identification scheme in the macro-finance literature on the term structure, it is worth revisiting. The uncovered magnitudes of disagreement between the variance decompositions based on different orders of variables seem to warrant a consideration of an alternative approach. To be sure, our projection-based procedure is not a silver bullet. It cannot replace variance decompositions because it can only work in the presence of latent variables and has no causal implications. We view our procedure as a useful complement to the toolkit of a researcher studying VARs that are related to the yield curve.

We further illustrate the usefulness of our procedure by using it to address the issues raised at the beginning of the introduction.[5] First, we investigate how much of the short rate variation can be explained by the two macro variables, as opposed to by the latent factors. It seems natural to start with the null hypothesis that macro variables do not contribute to explaining the short rate, given the excellent explanatory power of latent factor models that we discussed earlier. As we show, this hypothesis places restrictions on various parameters that control the dynamics and cross-section of the yield curve. A number of previous studies have simply imposed these restrictions, thereby limiting the ability of yields to vary with the macroeconomy. We test these restrictions in terms of the size of the R^2 that can be explained by macro variables and confirm that macro variables are extremely important. Our projection-based approach can explain 80% of the short rate variation, using inflation, real activity, and their lags exclusively as the basis.

Second, we investigate whether macroeconomic variables are useful for explaining bond term premia and not just short-term interest rates.[6] Our approach enables us to determine that

---

[3] Our lag structure is not arbitrary: recursive projection formulas imply reliance on all lags and the loadings on these lags are optimal, due to the fact that they are selected to minimize the variance of the residuals.

[4] We measure the slope of the yield curve as the difference between the 10-year and 3-month yields.

[5] More recent works by Chernov and Mueller (2008) and Mueller (2008) provide further examples of applications of our methodology.

[6] In fact, there is currently a debate in the literature over this point. Ang et al. (2004) and Ludvigson and Ng (2009) find that inflation and real activity measures play an important role in explaining bond risk premia. In contrast, Duffee (2007) concludes that macro variables make a minimal contribution to the term premia.