International Journal of Forecasting 33 (2017) 373–389

Contents lists available at ScienceDirect

International Journal of Forecasting

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

Does realized volatility help bond yield density prediction?

Minchul Shin[a,*], Molin Zhong[b,1]

[a] University of Illinois, United States

[b] Federal Reserve Board, United States

ARTICLE INFO

Keywords:
Dynamic factor model
Forecasting
Stochastic volatility
Term structure of interest rates
Dynamic Nelson-Siegel model

ABSTRACT

We suggest using "realized volatility" as a volatility proxy to aid in model-based multivariate bond yield density forecasting. To do so, we develop a general estimation approach to incorporate volatility proxy information into dynamic factor models with stochastic volatility. The resulting model parameter estimates are highly efficient, which one hopes would translate into superior predictive performance. We explore this conjecture in the context of density prediction of U.S. bond yields by incorporating realized volatility into a dynamic Nelson-Siegel (DNS) model with stochastic volatility. The results clearly indicate that using realized volatility improves density forecasts relative to popular specifications in the DNS literature that neglect realized volatility.

1. Introduction

Time-varying volatility exists in US government bond yields.[2] In this paper, we introduce volatility proxy data in the hopes of better capturing this time-varying volatility for predictive purposes. To do so, we develop a general estimation approach to incorporate volatility proxy information into dynamic factor models with stochastic volatility. We apply it to the dynamic Nelson–Siegel (DNS) model of bond yields. We find that the higher frequency movements of the yields in the realized volatility data contain valuable information for the stochastic volatility and lead to significantly better density predictions, especially in the short term.

Our approach can be applied to the existing classes of dynamic factor models with stochastic volatility. Specifically, we can account for stochastic volatility on the latent factors or stochastic volatility on the measurement errors. We derive a measurement equation to link realized volatility to the model-implied conditional volatility of the original observables. Incorporating realized volatility improves estimation of the stochastic volatility by injecting precise volatility information into the model.

The DNS model is a dynamic factor model that uses latent level, slope, and curvature factors to drive the intertemporal movements of the yield curve. This reduces the high-dimensional yields to be driven by just three factors. The level of the yield curve has traditionally been linked to inflation expectations while the slope is linked to the real economy. Our preferred specification introduces stochastic volatility on these latent factors. This leads to a nice interpretation of the stochastic volatility as capturing the uncertainty surrounding well-understood aspects of the yield curve. It also reduces the dimension of modeling the time-varying volatility of the yield curve.

We then compare this specification to several others in the DNS framework, including random walk dynamics for the factors and stochastic volatilities, as well as stochastic volatility on the yield measurement equation. In

* Correspondence to: 214 David Kinley Hall, 1407 W. Gregory Dr., Urbana, Illinois, 61801.

E-mail addresses: mincshin@illinois.edu (M. Shin), Molin.Zhong@frb.gov (M. Zhong).

1 20th Street and Constitution Avenue N.W., Washington, DC, 20551, United States.

2 Early empirical work was done by Engle, Ng, and Rothschild (1990). In the years that followed, there have been many attempts to empirically validate and analyze time-varying volatility in US government bond yields/prices (See, for example, papers cited in Diebold & Rudebusch, 2012).

http://dx.doi.org/10.1016/j.ijforecast.2016.11.003
0169-2070/

a forecasting horserace on US bond yields, our preferred specification features slight improvements in the point forecast performance and significant gains in the density forecast performance. The realized volatility data injects accurate conditional second moment information into the model, which aids in both the extraction of the current volatility state and also the estimation of the volatility process parameters. Both are important for density prediction, with the accurate volatility state estimation effect dominating for short horizon forecasts and the accurate parameter estimation effect dominating at longer horizons. We also find that allowing for time-varying volatility is important for density prediction, especially in the short run. Unlike conditional mean dynamics, modeling volatility as first-order autoregressive processes rather than random walks leads to better predictive performance. Furthermore, having stochastic volatility on the factor equation better captures the time-varying volatility in the bond yield data when compared to stochastic volatility on the measurement equation.

Our paper relates to the literature in three main areas. First, our paper relates to work started by Barndorff-Nielsen and Shephard (2002) in incorporating realized volatility in models with time-varying volatility. Takahashi, Omori, and Watanabe (2009) use daily stock return data in combination with high-frequency realized volatility to more accurately estimate the stochastic volatility. Maheu and McCurdy (2011) show that adding realized volatility directly into a model of stock returns can improve density forecasts over a model that only uses level data, such as the EGARCH. Jin and Maheu (2013) propose a model of stock returns and realized covariance based on time-varying Wishart distributions and find that their model provides superior density forecasts for returns. There also exists work adding realized volatility in observation-driven volatility models (Hansen, Huang, & Shek, 2012; Shephard & Sheppard, 2010). As opposed to the other papers, we consider a dynamic factor model with stochastic volatility on the factor equation and use the realized volatility to help in the extraction of this stochastic volatility. In this sense, we bring the factor structure in the conditional mean to the conditional volatility as well. Cieslak and Povala (2016) have a similar framework in a no-arbitrage term structure model. Furthermore, we are the first paper to investigate the implications of realized volatility for bond yield density predictability.

Second, we contribute to a large literature on bond yield forecasting. Most of the work has been done on point prediction (see for example, Diebold & Rudebusch, 2012; Duffee, 2012, for excellent surveys). There has been, however, a growing interest in density forecasting. Egorov, Hong, and Li (2006) were the first to evaluate the joint density prediction performance of yield curve models. They overturn the point forecasting result of the superiority in random walk forecasts and find that affine term structure models perform better when forecasting the entire density, especially on the conditional variance and kurtosis. However, they do not consider time-varying conditional volatility dynamics in the bond yield predictive distribution. Hautsch and Ou (2012) and Hautsch and Yang (2012) add stochastic volatility to the DNS model by considering an independent AR(1) specification for the log volatilities of the latent factors. They do not do formal density prediction evaluation of the model, but give suggestive results of the possible improvements in allowing for time-varying volatility. Carriero, Clark, and Marcellino (2013) find that using priors from a Gaussian no-arbitrage model in the context of a VAR with stochastic volatility improves short-run density forecasting performance. Building on this previous work, we introduce potentially highly accurate volatility information into the model in the form of realized volatility and evaluate bond yield density predictions to see whether this extra information about the bond yield volatility can improve the quality of the predictive distribution.

Another related class of bond yield prediction papers in the literature uses external information to improve the quality of prediction. Altavilla, Giacomini, and Ragusa (2013) exploit information contained in survey expectations data and use it to restrict model-implied forecasts via a flexible informational projection method. van Dijk, Koopman, van der Wel, and Wright (2014) use various sources of external information, including survey expectations, to capture a shift in the endpoint of the yield curve. These papers attempt to improve the point prediction by incorporating external information. On the other hand, our paper exploits external information to improve the density prediction by accurately estimating the latent volatility states and their related parameters.

Finally, we also add to a growing literature on including realized volatility information in bond yield models. Andersen and Benzoni (2010) and Christensen, Lopez, and Rudebusch (2014) view realized volatility as a benchmark on which to compare the fits of affine term structure models. Cieslak and Povala (2016) are interested in using realized covariance to better extract stochastic volatility and linking the stochastic volatility to macroeconomic and liquidity factors. These papers focus on in-sample investigations of incorporating realized volatility in bond yield models. Another stream of research exploits information in high-frequency movements of bond prices to achieve better point prediction performance. For example, Wright and Zhou (2009) report that the realized jump mean measure constructed from Treasury bond futures improves excess bond return point prediction by 40%. Our paper, in contrast to these others, considers the improvement from using realized volatility in out-of-sample bond yield density prediction.

In Section 2, we introduce our methodology for incorporating volatility proxies into dynamic factor models in the context of the DNS model and other competitor specifications. We discuss the data in Section 3. We present our estimation and forecast evaluation methodology in Section 4. In Section 5, we present in-sample and out-of-sample results. We conclude in Section 6.

2. Model

We introduce the dynamic Nelson–Siegel model with stochastic volatility (DNS-SV) proposed by Bianchi, Mumtaz, and Surico (2009), Hautsch and Ou (2012), and Hautsch and Yang (2012). Then, we discuss the incorporation of realized volatility information into this framework. Finally, we consider alternatives to our main approach.