Journal of Macroeconomics 57 (2018) 248–265

Contents lists available at ScienceDirect

Journal of Macroeconomics

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

What does the yield curve imply about investor expectations?

Eric Gaus[a], Arunima Sinha[*,b]

[a] Department of Economics Haverford College, 370 Lancaster Ave, Haverford, PA 19041, United States
[b] Department of Economics, Fordham University, 113 West 60th street, New York, NY 10023, United States

ARTICLE INFO

JEL classification:
E43
E47
D83
C5

Keywords:
Adaptive learning
Investor beliefs
Inflation expectations
Excess returns

ABSTRACT

We use daily data to model investors' expectations of U.S. yields, at different maturities and forecast horizons. We consider two adaptive learning algorithms to characterize the conditional yield forecasts. Our framework yields the first empirical estimates of the pace of learning by investors. The superior performance of the endogenous learning mechanism suggests that investors account for structural change and respond to significant, persistent deviations by modifying the amount of information they use. Our results provide strong empirical motivation to use the class of adaptive learning models considered here for modeling and analyzing expectation formation by investors.

1. Introduction

> The Federal Reserve's ability to influence economic conditions today depends critically on its ability to shape expectations of the future, specifically by helping the public understand how it intends to conduct policy over time, and what the likely implications of those actions will be for economic conditions. (Vice-Chair Janet Yellen, At the Society of American Business Editors and Writers 50th Anniversary Conference, Washington, D.C., April 4, 2013).

Investor expectations about the term structure of yields are central to the conduct of monetary policy. Influencing these expectations through the different instruments available to the Federal Reserve, has been important during the Great Moderation. During the Great Recession and its aftermath, this strategy has been at the forefront of the central bank's policy. As the accommodative monetary policy stance of the Federal Reserve kept the federal funds rate at the zero-lower bound from December 2008 to November 2015, one of the main channels through which monetary policy affected longer yields (and the subsequent consumption and savings decisions of economic agents), was by affecting the formation of conditional expectations by market investors.

The contribution of the present analysis is to characterize the expectations formation process of market investors about the term structure of yields at different forecasting horizons. We further explore whether differences exist in the expectations formation process between the Great Moderation and the Great Recession periods, i.e., during periods of low and high macroeconomic volatility. To do this, we develop a novel methodology to model the evolution of investor beliefs using daily data on the U.S. nominal yield curve. Using the Great Moderation as the baseline period, we extend the results to include the Great Recession. Our analysis allows for the comparison of investor beliefs about the entire yield curve, across a cross-section of forecast horizons.

Our strategy is briefly described as follows: we use the daily yield curve factors estimated by Gürkaynak et al. (2007, henceforth, GSW) to construct yield forecasts. Following recent studies,[1] we first construct conditional expectations of yields (and associated

* Corresponding author.
E-mail addresses: egaus@haverford.edu (E. Gaus), asinha3@fordham.edu (A. Sinha).

[1] Examples include Diebold and Li (2006) and Aruoba et al. (2006).

https://doi.org/10.1016/j.jmacro.2018.06.003
Received 10 March 2017; Received in revised form 5 June 2018; Accepted 6 June 2018
Available online 14 June 2018
0164-0704/

latent factors) using a vector auto-regressive (VAR) model with constant coefficients. We evaluate the forecasting performance of the model, and a series of rationality tests of the implied forecasting errors confirm that these errors are biased, systematic, and correlated with revisions in yield forecasts. In addition to these findings on the forecast errors, the framework also imposes the restriction that investors must be placing identical weights on past information while forecasting the short and long yields over different forecasting horizons. Thus, it does not allow investors to endogenously adapt to any structural breaks that they might perceive in the evolution of the average yields, or the yield curve slope.

The above results motivate our hypothesis that market investors are using other models of expectations formation. Theoretical analyses, such as Piazzesi et al. (2015) and Sinha (2016), incorporate adaptive learning into the expectations formation of optimizing agents in models of the yield curve. The implied term structures are more successful at matching the properties of the empirical yield curve, relative to models with time-invariant beliefs. Therefore, we explore a class of adaptive learning models for the formation of conditional forecasts of the nominal and real yield curve factors, and subsequent yields: constant gain learning (CGL) and an endogenous learning (EGL) algorithm that we develop here. The main innovation is that investors are now allowed to vary the weights they place on past information about yields; they are also able to adapt to the size of large and persistent deviations observed in the yield curve factors data.

We find that there are significant improvements in forecasting performance of the model with the learning processes. Our results are based on the implied forecasts for the two sample periods, and the methodology characterizes the speed of learning by market investors using high frequency data, This, to our knowledge, is a first for the adaptive learning literature. The parameters of the learning models - the updating coefficient or the "gain", the adjustment factor in case of large deviations, and the time period used to compute deviations with respect to historical data, are all estimated from the daily yield curve data, for different forecasting horizons. The main result is that at different pairs of yield maturity and forecast horizon, the endogenous learning forecast improves upon the constant gain algorithm. For example, at the 1-month forecasting horizon, for the nominal 1-year yield, endogenous learning improves upon the constant gain mean square forecast error by 36%; at the 6-month horizon, the improvement is close to 18%. This improvement in the performance persists across yield maturities as well.

We investigate the forecasts obtained from the learning models on two separate dimensions. First, we subject the forecast errors to the same set of rationality tests, as done for the constant-coefficients model. The results show that forecast errors from both the learning models improve upon the findings for the time-invariant case: the errors are unbiased (or significantly less so), the predictive content of current forecasts for the errors is either insignificant or much reduced, and systematic nature of the forecast errors is also greatly diminished for the 1- and 10-year yields. Thus, we conclude that the forecast errors of the learning model improve upon those from the constant-coefficients model.

We then test whether the out-of-sample learning forecasts improve upon random walk forecasts.[2] Since yields are very persistent, the random walk model is difficult to out-perform in these out-of-sample forecasts. At the shortest forecasting horizon (1-month), we find that this is still the case. However, at the longer forecasting horizons, the learning models improve upon the random walk forecasts. For example, the EGL 6-month forecast for the 10-year yield improves upon the random walk model by approximately 13%.

In addition to the superior performance of the EGL mechanism on the above dimensions, the estimation of the endogenous gain parameters yields several insights into the expectations formation process of agents: (a) the implied conditional expectations of investors display substantial time-variation and adapt to large deviations in the data during periods of low and high macroeconomic volatility (we present estimates of the gains from the Great Moderation as well as the Great Recession to demonstrate this); (b) while constructing forecasts of the 10-year yield, the investors' expectations are largely invariant to large deviations in the observed data, and they account for structural change to a significantly smaller degree, relative to the 1- and 5-year yields, during the Great Moderation. This contrasts with the expectations formation during the Great Recession, where the 1- and 3-month forecasts of the 10-year nominal yield give more weight to the more recent level and slope factor data. This suggests that during periods of low macroeconomic volatility, investors do not expect structural change in the data for the long-term yield, but become much more attentive during highly volatile periods. Thus, monetary policy actions that target the long-end of the yield curve during a recession may be more successful at influencing the savings and investment decisions of agents. This result supports findings elsewhere in the literature: for example, Coibion and Gorodnichenko (2015) find that survey forecasters exhibit greater information rigidities during the Great Moderation, compared to the earlier recessions. Using observed data on the yield curve, we show that investors are, in fact, forming conditional expectations differently over the business cycle. Thus, the mechanism offers a tractable way of incorporating time-variation in expectations formation, and can aid in policy analysis: less attention to more recent data during recessions (for different yield maturities and forecast horizons) indicates that investors will less more rapidly to policy changes during recessions.

Our methodology allows for investors to allow the gains to vary across different yield curve factors and across forecast horizons; this provides a more intuitive way to allow for the investors to update their information. For example, while forming forecasts, the investors may place more or less weight on the history of the level of yields, than on the slope of the yield curve. If they believe that there were several structural breaks in the average level of the yield curve, they may prefer to place more weight on the recent past observations, instead of the longer history. If such breaks are not perceived to exists in the yield curve slope, the investors may place almost equal weight on past observations. These gain parameters are therefore central to the bounded rationality approach, since they determine the persistence in expectations formation, and how investors will react to permanent versus transitory shocks. In this analysis, we use fixed baseline time periods (for the Great Moderation and the Great Recession period) to find the optimal gains.

Given the success of the endogenous learning mechanism in modeling conditional forecasts of investors, we then consider the

[2] In the accompanying appendix, we also compare the learning forecasts with those of the Diebold and Li (2006) model.