International Journal of Forecasting 29 (2013) 469–478

Contents lists available at SciVerse ScienceDirect

International Journal of Forecasting

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

Are macroeconomic variables useful for forecasting the distribution of U.S. inflation?

Sebastiano Manzan [a], Dawit Zerom [b,*]

[a] Department of Economics & Finance, Baruch College, CUNY, United States
[b] California State University at Fullerton, United States

ARTICLE INFO

Keywords:
Distribution
Economic indicators
Inflation
Predictability
Quantiles

ABSTRACT

Much of the inflation forecasting literature examines the ability of macroeconomic indicators to predict the mean inflation accurately. For the period after 1984, the existing empirical evidence largely suggests that the likelihood of predicting inflation accurately using macroeconomic indicators is no better than a random walk model. We expand the scope of inflation predictability by exploring whether macroeconomic indicators are useful in predicting the distribution of inflation. We consider six commonly-used macro indicators and core/non-core versions of the Consumer Price Index (CPI) and the Personal Consumption Expenditure (PCE) deflator as measures of inflation. Based on monthly data and for the forecast period after 1984, we find that some of the macro indicators, such as the unemployment rate, housing starts and the term spread, provide significant out-of-sample predictability for the distribution of core inflation. An analysis of the quantiles of the predictive distribution reveals interesting patterns which would otherwise be ignored by existing inflation forecasting approaches that rely only on forecasting the mean. We also illustrate the importance of inflation distribution forecasting in evaluating some events which are of policy interest by focusing on predicting the likelihood of deflation.

1. Introduction

Forecasting the behavior of inflation plays a central role in the conduct of monetary policy, due to the lagged impact of the central bank's actions on economic activity. Thus, accuracy is important when predicting the effects of the many shocks that hit the economy on the future dynamics of inflation. The standard approach for forecasting inflation has been the Phillips curve (PC) model, which, in its expectation-augmented version, assumes a trade-off between unexpected inflation and unemployment, or, more generally, indicators of real economic activity. Despite its long-lasting success, recent empirical evidence on the effectiveness of the PC model has been far from unanimous. Stock and Watson (1999) provide a detailed study of the out-of-sample forecast accuracy of the PC by using an extensive set of macroeconomic variables. Using the forecast evaluation period January 1970–September 1996, their conclusion is that PC models have better forecasting performances (compared to univariate time series models) when using the unemployment rate as well as other leading indicators of economic activity (e.g., the output gap and capacity utilization). They also find that combining information or models might provide better results than simply relying on a few indicators. However, Atkenson and Ohanian (2001) provide empirical evidence suggesting the opposite, albeit for a different forecast evaluation period, January 1984–November 1999, where they report that PC models are no better than the naïve model, which assumes that the inflation expected over the next 12 months will be equal to the inflation experienced over the previous 12 months. Fisher, Liu, and Zhou (2002) conduct a systematic comparison of the forecasting accuracy at the one-year horizon in different sub-periods, and find that the

* Corresponding author.
E-mail address: dzerom@fullerton.edu (D. Zerom).

0169-2070/$ – see front matter
http://dx.doi.org/10.1016/j.ijforecast.2013.01.005

PC forecasts only outperformed the naïve forecasts in the 1977–1984 window. There is evidence that the declining predictive power of macroeconomic indicators is typical also of other variables (see D'Agostino, Giannone, & Surico, 2006; and Rossi & Sekhposyan, 2010). For a comprehensive survey and a discussion of the outstanding issues in inflation forecasting, see Stock and Watson (2008).

In this paper, departing from the existing focus on conditional mean forecasting, we explore whether leading indicators of economic activity are useful in predicting the distribution of future inflation. Despite the extensive body of literature on inflation forecasting which is available, little or no attention has been paid to the examination of whether indicators of economic activity carry useful information about the dynamics of higher moments, beyond the mean. For example, having some idea of the conditional second-order moment of future inflation can be vital to an assessment of the risk to inflation stability as a result of macroeconomic shocks. Greenspan (2004) discusses this issue in the following terms: "Given our inevitably incomplete knowledge about key structural aspects of an ever-changing economy and the sometimes asymmetric costs or benefits of particular outcomes, a central bank needs to consider not only the most likely future path for the economy, but also the distribution of possible outcomes about that path. The decision-makers then need to reach a judgment about the probabilities, costs, and benefits of the various possible outcomes under alternative choices for policy" (p. 37). While the average future inflation may signal the direction of the economy, it cannot help policy-makers to evaluate the risks of deviations from the most likely path, or the cost to the economy of such deviations. In a recent paper, Kilian and Manganelli (2008) introduced a model in which the monetary policy maker is viewed as a risk manager who is trying to balance the risks to inflation and output stability. In this framework, if the preferences of the policy maker are assumed to be quadratic and symmetric, then the only relevant moment (of the inflation and output distributions) is the conditional mean. However, they provide evidence of a departure of the preferences from such a benchmark. All of the elements above point to the suggestion that forecasting the distribution of inflation represents a relevant tool in the conduct of monetary policy. In fact, for many years the Bank of England has been publishing the so-called "fan charts", which represent the Bank's subjective forecasts about the future distribution of inflation.

We use quantile regression to incorporate macroeconomic variables into the prediction of the conditional distribution of future inflation. The approach considers several conditional quantiles of future inflation, and by doing so, offers more flexibility (than, for example, the conventional PC models) in capturing the possible roles of macroeconomic indicators in predicting the different parts of the inflation distribution. For instance, one may be able to investigate whether some periods of low or high inflation are driven by some macroeconomic indicators. Such information surely cannot be delivered by PC-type models that deal only with predicting the average inflation. We find strong empirical evidence of predictability for U.S. core monthly inflation using indicators of economic activity, particularly the unemployment rate, housing starts and the term spread. Importantly, the empirical findings apply to a forecast evaluation period that is intentionally chosen to be post-1984, when the existing literature shows that the macroeconomic indicators were not relevant for predicting future average inflation (see Stock & Watson, 2007). We attribute this result to the ability of our approach to account for the varying predictive effects of economic indicators on core inflation at different quantiles of its distribution. For example, when considering the unemployment rate, we find that it adds no predictive value to the average core inflation (consistent with the prevailing empirical evidence), while providing statistically significant improvements to the left tail (or lower quantiles) of the inflation forecast distribution. Reflecting this finding, when we use a forecast density that incorporates the unemployment rate to predict the likelihood of deflation (out of sample), the approach appropriately assigns a probability close to zero to the likelihood of deflation for the period 1996–2007. This clearly illustrates the way in which a macroeconomic variable may offer weak or no predictability of inflation in the conditional mean sense, but can still help improve the forecast density of inflation.

A possible explanation for these findings is that the public might form inflation expectations that are more responsive to macroeconomic news when inflation is in the tails of the distribution. This can happen if the monetary policymaker has a target (either implicit or explicit) which is located near the center of the distribution. When inflation fluctuates around the center of the distribution, expectations are well-anchored to the target, in the sense that the public believes that the central banks will be able to maintain price stability. However, when inflation deviates from the target significantly, the public might distrust the ability of monetary policy to bring inflation back to the target, and thus may form expectations that rely more on macroeconomic variables which represent more reliable indicators of future inflation.

The number of papers that deal with density forecasts of inflation is quite limited, and much of the existing literature focuses largely on evaluating the absolute accuracy of inflation density forecasts with respect to dynamical and distributional mis-specifications. For example, Clark (2011) considers a Bayesian VAR model with stochastic volatility, and finds that accounting for the time variation in the volatility is essential to the production of accurate density forecasts (in an absolute sense) of U.S. inflation, especially for the post-1985 period. A similar conclusion was also reached by Jore, Mitchell, and Vahey (2010) using AR and VAR models that allow for structural breaks. On the other hand, our approach is a model selection one, where we examine the ability of macroeconomic indicators to improve density forecasts of U.S. inflation. We do this by comparing the relative accuracy of the density forecasts (using scoring rules defined from density and quantile forecasts) implied by a quantile regression approach that incorporates a macroeconomic indicator with that of a benchmark model that does not include the macroeconomic indicator. Amisano and Giacomini (2007) also consider a model selection approach, but focus on the influence of nonlinearity on the accuracy of U.S. inflation density forecasts. In particular, they compare the relative accuracy of the out-of-sample density forecasts implied by the