North American Journal of Economics and Finance 22 (2011) 77–87

Contents lists available at ScienceDirect

North American Journal of Economics and Finance

Real-time inflation forecast densities from ensemble Phillips curves☆

Anthony Garratt[a,*], James Mitchell[b], Shaun P. Vahey[c], Elizabeth C. Wakerly[c]

[a] *Birkbeck College, University of London, Malet Street, Bloomsbury, London, WC1E 7HX, UK*
[b] *National Institute of Economic and Social Research, UK*
[c] *Centre for Applied Macroeconomic Analysis, Australian National University, Australia*

ARTICLE INFO

Article history:
Received 7 September 2009
Accepted 28 September 2010
Available online 23 December 2010

JEL classification:
C32
C53
E37

Keywords:
Density combination
Ensemble forecasting
Phillips curve

ABSTRACT

We examine the effectiveness of recursive-weight and equal-weight combination strategies for forecasting using many time-varying models of the relationship between inflation and the output gap. The forecast densities for inflation reflect the uncertainty across models using many statistical measures of the output gap, and allow for time-variation in the ensemble Phillips curves. Using real-time data for the US, Australia, New Zealand and Norway, we find that the recursive-weight strategy performs well, consistently giving well-calibrated forecast densities. The equal-weight strategy generates poorly-calibrated forecast densities for the US and Australian samples. There is little difference between the two strategies for our New Zealand and Norwegian data. We also find that the ensemble modelling approach performs more consistently with real-time data than with revised data in all four countries.

1. Introduction

A number of applied macro-econometric studies have found that forecast combination using recursive weights, based on historical forecast performance, is an ineffective strategy for improving point

☆ Financial support from the ESRC (grant no. RES-062-23-1753) and the ARC (grant no. LP0991098) is gratefully acknowledged. We benefitted greatly from comments by Dean Croushore, Domenico Giannone, James Hamilton, Kirdan Lees, Simon van Norden, Les Oxley, David Papell and two referees. Anne Sofie Jore and Tim Robinson kindly provided real-time data for Norway and Australia, respectively. We are also indebted to conference and seminar participants at the Reserve Bank of Australia and the CIRANO Data Revisions Workshop October 2008.

* Corresponding author. Tel.: +44 207 631 6410; fax: +44 631 6416.
E-mail address: a.garratt@bbk.ac.uk (A. Garratt).

1062-9408/$ – see front matter © 2010 Published by Elsevier Inc.
doi:10.1016/j.najef.2010.09.003

forecasts. Stock and Watson (2004), and Clark and McCracken (2010), among others, have found that an equal-weight strategy is more effective in terms of root mean squared forecast error. However, forecasters and policymakers are often interested in forecast densities rather than point forecasts. Considering the evidence for density forecasting performance, Jore, Mitchell, and Vahey (JMV, 2010) report strong density forecasting performance from the recursive-weight strategy, but not from the equal-weight strategy, with vector autoregressions using US data. Garratt, Mitchell, and Vahey (GMV, 2009) report similar findings for recursive weights for US inflation using a Phillips curve relationship based on the output gap, but do not consider the equal-weight strategy.

In this paper, we investigate the generality of the JMV finding by examining the recursive and the equal-weight strategies for inflation forecast densities in four countries. The model space is similar to that of GMV. That is, we consider many time-varying models of the relationship between inflation and the output gap. The forecast densities for inflation in each country reflect the uncertainty across models using many statistical measures of the output gap, and allow for time-variation in the ensemble Phillips curves.

To implement our recursive-weight strategy, we adopt the forecasting methodology proposed by JMV in which a real-time forecaster (or policymaker) recursively selects a combination of component forecasts from a set of models to produce an ensemble forecast density. Each component forecast density for inflation is produced by a single Phillips curve model for inflation based on lags of inflation and lags of the output gap. We utilise a “linear opinion pool” (LOP) to take out of sample density combinations (see Timmermann, 2006, p. 177) using the logarithmic score, as a measure of the Kullback-Leibler distance, to generate component weights. The resulting ensemble approximates the unknown potentially non-linear data generating process for inflation by using time-varying weights; and, the ensemble forecast densities are not restricted to be Gaussian. To implement our equal-weight strategy, we adopt an analogous ensemble methodology, again using LOP, to take equal-weighted forecast density combinations. We evaluate the recursive and equal-weight strategies by examining the probability integral transforms (*pits*) of the ensemble densities for inflation.

We consider real-time data for the US, Australia, New Zealand and Norway. For each data set, we compare and contrast the inflation forecasts from equal-weight and recursive-weight ensembles. The recursive-weight strategy performs well across the real-time data sets, consistently giving well-calibrated forecast densities. The equal-weight strategy performs less consistently, generating poorly-calibrated forecast densities for the US and Australian samples in particular. There is little difference between the two strategies for our New Zealand and Norwegian data. We also find that the ensemble modelling approach performs more reliably with real-time data than with revised data in all four countries.

The remainder of this paper is structured as follows. In Section 2, we outline the component models. In Section 3, we describe our methods for ensemble forecasting and density evaluation. In Section 4, we apply our methodology to US, Australian, New Zealand and Norwegian data and present the results. In the final section we conclude.

2. Component models

Following Orphanides and van Norden (2005) and GMV, we start with Phillips curve forecasting models of the Linear Gaussian form:

$$\pi_{t+h} = \alpha_1^j + \sum_{p=1}^{P} \beta_{1,p}^j \pi_{t-p+1} + \sum_{p=1}^{P} \gamma_{1,p}^j y_{t-p+1}^j + \sigma_1^j \varepsilon_{1,t+h}^j, \tag{1}$$

where inflation is defined as the log difference in the price level, and there are many output gap measures denoted y_t^j, where $j = 1, \ldots, J$; the number of lags of inflation and the output gap is denoted P, $\varepsilon_{1,t+h}^j \sim$i.i.d. $N(0, 1)$ and h is the forecast horizon.[1] The predictive densities for π_{t+h} denoted $g(\pi_{\tau,h}|I_{i,\tau})$,

[1] We set $h = 1$ in our applications that follow.