BANK OF CANADA
BANQUE DU CANADA

Staff Working Paper/Document de travail du personnel 2018-39

Sluggish Forecasts

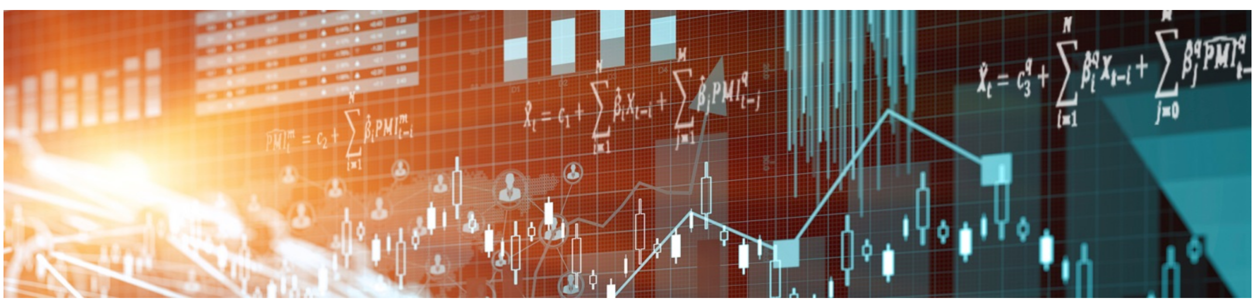

by Monica Jain

Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

www.bank-banque-canada.ca

Bank of Canada Staff Working Paper 2018-39

August 2018

Sluggish Forecasts

by

Monica Jain

Financial Markets Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
jain@bankofcanada.ca

ISSN 1701-9397 © 2018 Bank of Canada

Acknowledgements

Earlier versions of this paper were entitled "Conservatism in Inflation Forecasts" and "Sluggish Inflation Forecasts." I am thankful to Gregor Smith, Allan Gregory, Xuguang Sheng, Tara Sinclair, Kajal Lahiri, Peter Christoffersen, Nicolas Vincent, Marc-Andre Gosselin, Jonathan Witmer, Antonio Diez de los Rios, Bruno Feunou, Bo Young Chang, Yuriy Gorodnichenko, Jason Allen, seminar participants at the Bank of Canada Brown Bag Seminar (2013), the Society of Nonlinear Dynamics and Econometrics Conference (2013), the Canadian Economics Association Conference (2013), the International Symposium of Forecasting annual conference (2013), the Computing and Economics and Finance Conference (2013), the Bank of Canada Fellowship Learning Exchange (2014), the European Economics Association (2015), the Federal Reserve Bank of Cleveland Inflation: Drivers & Dynamics Conference (2018), and the North American Summer Meeting of the Econometrics Society (2018) for helpful comments.

The views and opinions expressed in this paper are those of the author and do not necessarily represent the views of the Bank of Canada.

Abstract

Given the influence that agents' expectations have on key macroeconomic variables, it is surprising that very few papers have tried to extrapolate agents' "true" expectations directly from the data. This paper presents one such approach, starting with the hypothesis that there is sluggishness in inflation and real GDP growth forecasts. Using individual-level data on 29 U.S. professional forecasters from the Survey of Professional Forecasters, I find that some degree of sluggishness is present in about 40% of inflation forecasts and in 60% of real GDP growth forecasts. The estimates of sluggishness are then used to recover a series of sluggishness-adjusted expectations that are more volatile and, at times, more accurate than the raw survey forecasts.

Bank topics: Econometric and statistical methods; Inflation and prices
JEL codes: E31, E37

Résumé

Vu l'influence que les anticipations des agents économiques exercent sur de grandes variables macroéconomiques, on peut s'étonner du très petit nombre d'études qui ont tenté d'extrapoler les « véritables » anticipations des agents directement à partir des données. Cette étude présente une telle approche, fondée sur l'hypothèse que les prévisions de l'inflation et du PIB réel comportent une certaine inertie. Se servant des données de 29 prévisionnistes professionnels américains tirées de l'enquête auprès des prévisionnistes professionnels, l'auteure constate qu'environ 40 % des prévisions de l'inflation et 60 % de celles du PIB réel présentent un certain degré d'inertie. Elle utilise ensuite les estimations de l'inertie pour obtenir des séries d'anticipations corrigées de l'inertie, lesquelles sont plus volatiles et, parfois, plus exactes que les prévisions brutes basées sur des enquêtes.

Sujets : Méthodes économétriques et statistiques; Inflation et prix
Codes JEL : E31, E37

Non-Technical Summary

Survey data are commonly used in research and policymaking to obtain a measure of agents' expectations. However, if agents do not report their true expectations or do not update their forecast annoucements as new information arrives, then these data are less relevant and the corresponding research and policy conclusions less accurate. This paper presents a method to measure this 'sluggishness' bias in the forecast announcements made by professional forecasters and uses these estimates to back-out forecasters' sluggishness-adusted forecasts. The focus on individual forecasters in this paper also allows us to consider forecasters' different objectives: some forecasters smooth their forecasts for reputational reasons, while others may not make revisions because of informational rigidities they face. Others may simply be interested in making good, accurate forecasts. Allowing for these subtle differences can deepen our understanding of how expectations are formed.

This paper finds that there is some sluggishness in inflation forecasts, and a larger degree of sluggishness in real GDP growth forecasts. In addition, there is heterogeneity in the estimates of sluggishness across individual forecasters, ranging from those who have a relatively high degree of sluggishness in both their inflation and real GDP growth forecasts, to those that do not exhibit sluggishness in either of these forecasts. Finally, when the estimates of sluggishness are used to back-out forecasters' underlying or sluggishness-adjusted expectations, it yields an expectations series that is more volatile, and, in terms of accuracy, performs at least as well as the survey data for inflation and real GDP growth forecasts for most horizons. The new adjusted expectations series presented in this paper thus provides a useful alternative to the observed survey data.

1 Introduction

Expectations are a key driver of the macroeconomy. Through their expectations, agents make micro-level decisions that collectively guide the direction of major macroeconomic variables. Indeed, the recent explosion of research on expectation formation speaks to its importance. While a useful and tractable approach to modelling expectations, the assumption of full information rational expectations (FIRE) has been unable to capture all features of expectations reported in data. This might be because agents have incentives to not report their true beliefs (Laster et al. (1999), Ottaviani and Sorensen (2006), Ehrbeck and Waldmann (1996)) or because they are subject to informational rigidities, such as sticky information (Mankiw, Reis and Wolfers (2003)) or noisy information (Woodford (2002), Sims (2003) and Mackowiak and Wiederholt (2009)). There is, however, limited empirical work in the former class of models, and the latter models are unable to explain the smoothness found in the data (Andrade and Le Bihan (2013)).

Developing a good understanding of the data-generating process underlying agents' expectations is vital to conducting accurate welfare analysis and policymaking. One approach towards this objective is to extrapolate agents' "true" expectations directly from the data. This paper takes a step in this direction. Specifically, I conduct a micro-level analysis that estimates the degree of sluggishness in inflation and real GDP growth forecasts from the U.S. *Survey of Professional Forecasters* (SPF). The degree of sluggishness is measured by the extent to which rational forecasters have an aversion to revising their previously announced forecasts, and as a result, do not necessarily announce their rational expectations forecast. A dual objective is captured in their loss function: forecasters would like to avoid having large forecast errors, but also would like to avoid making large revisions to their previous forecasts. Modelling individual agents and allowing for a forecaster-specific expectation formation process across agents are key to developing a better understanding of the distribution of macroeconomic expectations.

The main findings of this paper are as follows: i) there is some sluggishness in inflation forecasts, and a larger degree of sluggishness in real GDP growth forecasts; ii) there is heterogeneity in the estimates of sluggishness across individual forecasters, ranging from those who have a relatively high degree of sluggishness in both their inflation and real GDP growth forecasts, to those that do not exhibit sluggishness in either of these forecasts; and iii) when the estimates of sluggishness are used to back-out forecasters' underlying or sluggishness-adjusted expectations, it yields an expectations series that is more volatile and, at times, more accurate than the reported survey data.

Sluggishness, or forecast smoothing, may follow if the end-users of the forecasts (such as businesses or policymakers) distrust forecasters who make frequent forecast revisions.[1] As noted by Nordhaus (1987), forecasters may smooth their forecasts because more accurate but jumpy forecasts would "drive customers crazy," since forecast revisions that occur too frequently could lead to investment reversals. Nordhaus also adds that a sluggishness bias may be attributed to agents' tendency to hold on to old notions that are more familiar, rather than adjusting to surprises. Forecast smoothing may also occur as a result of informational rigidities. Coibion and Gorodnichenko (2015) (henceforth, CG) test the rational expectations hypothesis using an approach that suggests a theoretical link between *ex post* mean forecast errors and *ex ante* mean forecast revisions. They interpret their rejection of FIRE as driven by informational rigidities. Mertens and Nason (2015) jointly estimate an unobserved components model with a time-varying sticky information law of motion and find that survey forecasts have become more "sticky" over recent decades than they were prior to the mid-1980s. Irrespective of the source of sluggishness, however, the common outcome in many of these cases is a divergence between announced inflation forecasts and true expectations. The method used to recover the series of sluggishness-adjusted inflation expectations presented in this paper can counteract some of this divergence.

The focus in this paper is on individual forecasters. The majority of empirical work done on expectation formation uses a consensus forecast, given by either the mean or median forecast. However, recent research has found that focusing on a consensus forecast instead of individual-level forecasts can lead to different conclusions. Dovern et al. (2014), for example, find that the process of averaging induces additional stickiness into forecasts of real GDP growth. Pesaran and Weale (2006) find weak evidence against rational expectations at the individual level compared with what is observed when average forecasts are used.

[1] This view is supported by Batchelor and Dua (1992), Batchelor (2007) and Kirchgassner and Muller (2006).

This paper is organized as follows. Section 2 discusses the loss function and derives the first order conditions used throughout the paper. Section 3 discusses the inflation announcement data. Section 4 outlines the estimation strategy and results. Section 5 provides an application to CG and Lahiri and Sheng (2008). Section 6 discusses the measurement of sluggishness-adjusted expectations, and conducts an out-of-sample exercise. Section 7 concludes.

2 A Sluggish Loss Function

A typical forecast announcement is denoted by $x^a_{jt}(t+h)$, which is forecaster j's announcement of variable x, x^a, for time $t+h$, conditional on time t information. For example, $x^a_{jt}(t+1)$ and $x^a_{jt-1}(t+1)$ would be forecaster j's announcements of x one period ahead (based on time t information) and two periods ahead (based on time $t-1$ information), respectively. While used often, the assumption that forecasters announce their true conditional expectations usually does not hold in practice:

$$x^a_{jt}(t+h) = E_{jt}x(t+h), \quad \text{for } h = 1, ..., H, \tag{1}$$

where $E_{jt}x(t+h)$ is forecaster j's true expectation of x in period $t+h$. In this paper, this assumption will be relaxed.

Notice that $x^a_{jt}(t+1)$ and $x^a_{jt-1}(t+1)$ are both forecasts of the same thing: the realized value of x in period $t+1$. However, forecaster j may make revisions to the two-period-ahead forecast as her information set changes from one period to the next. Hence, a one-period-ahead, announced, forecast revision may be written as:

$$x^a_{jt}(t+1) - x^a_{jt-1}(t+1).$$

This revision is made by forecaster j at time t. In general, an h-steps-ahead announced forecast revision to the forecast for $x(t+1)$ may be written as:

$$x^a_{jt+1-h}(t+1) - x^a_{jt-h}(t+1).$$

I assume a forecasters' loss function follows from their dual objective: to avoid having large forecast errors, and to avoid making large revisions to their previous period's forecast. The former captures forecasters' primary objective, which is to make accurate forecasts, and the latter represents an objective that may follow if forecasters have some incentive to smooth their forecasts—perhaps if the end-users of the forecasts prefer that forecasters do not make frequent revisions to their previous forecasts, or if forecasters are subject to informational rigidities.

The loss function for forecaster j's one-step-ahead forecast is as follows:

$$\underset{x^a_{jt}(t+1)}{Min} \; E_{jt}\{[x(t+1) - x^a_{jt}(t+1)]^2 + \lambda_j[x^a_{jt}(t+1) - x^a_{jt-1}(t+1)]^2\}, \tag{2}$$

where λ_j is a weight that captures the extent to which the secondary objective influences j's forecast announcement. It is assumed to be time-invariant, and can take on any non-negative value. The first term penalizes inaccuracy while the second term penalizes revisions. Quadratic loss functions have the advantage that they penalize large deviations from the target more than small deviations.[2] In addition, the first derivative of the loss function is linear so that the forecasting rule may be represented by a linear function. In this case, the first-order conditions show that forecaster j should set her announcement so that:

$$x^a_{jt}(t+1) = \omega_{j1}E_{jt}x(t+1) + (1 - \omega_{j1})x^a_{jt-1}(t+1), \tag{3}$$

where $E_{jt}x(t+1)$ is forecaster j's true expectation of x in period $t+1$ and $\omega_{j1} = \frac{1}{1+\lambda_j}$. Notice this loss function implies that the forecaster's announcement is a combination of her true expectations and the previous period's two-steps-ahead announcement for x.

To solve for $x^a_{jt-1}(t+1)$, we can set up forecaster j's loss function for the two-steps-ahead

[2] It is also possible to consider an asymmetric loss function; however, Clements (2014) conducts individual-level analysis using data from the *Survey of Professional Forecasters* (SPF) and finds little evidence of asymmetric loss for inflation forecasters. Hence, I continue the analysis here under the assumption of symmetric loss.

announcement for time $t+1$:

$$\underset{x^a_{jt-1}(t+1)}{Min} \quad E_{jt-1}\{[x(t+1) - x^a_{jt-1}(t+1)]^2 + \lambda_j[x^a_{jt-1}(t+1) - x^a_{jt-2}(t+1)]^2$$
$$+\delta\lambda_j[x^a_{jt}(t+1) - x^a_{jt-1}(t+1)]^2\} \tag{4}$$
$$s.t. \quad E_{jt-1}x^a_{jt}(t+1) = \omega_{j1}E_{jt-1}x(t+1) + (1-\omega_{j1})x^a_{jt-1}(t+1),$$

where δ is a discount factor and $E_{jt-1}x^a_{jt}(t+1)$ is forecaster j's expected one-step-ahead inflation announcement.

The two-steps-ahead loss function is more complicated than the one-step-ahead case. This occurs as forecasters now are not only concerned with revisions to their past three-steps-ahead forecast, but also with future revisions that will be made to the two-steps-ahead forecast.[3] All revisions are made one period apart and therefore assigned the same weight, λ_j. In general, the problem of determining the optimal inflation announcement contains more revision terms at larger forecast horizons. Table 1 provides a detailed timeline of a typical forecaster's information set and loss function over time for the forecast of $x(t+1)$.

The first-order condition for the two-steps-ahead problem implies:

$$x^a_{jt-1}(t+1) = \omega_{j2}E_{jt-1}x(t+1) + (1-\omega_{j2})x^a_{jt-2}(t+1), \tag{5}$$

where $\omega_{j2} = \frac{1+\delta\lambda_j\eta_{j2}}{1+\lambda_j+\delta\lambda_j\eta_{j2}}$ and $\eta_{j2} = \omega^2_{j1}$. As in the solution to the one-step inflation announcement, notice the sum of the weights on true inflation expectations and the previous period's announcement add to one. In general, the first-order conditions show that any h-steps-ahead forecast announcement of $x(t+1)$ ought to be a weighted average of true expectations and the

[3] Note that it is possible to consider an extension in which forecasters choose a sequence of forecasts based on their current information set, and are concerned with both their future forecast revisions *and* their future forecast errors. This type of loss function is explored in CG:

$$Min \sum_{i=1}^{h} \gamma^{i-1}E_{jt+1-h}\{[x(t+1) - x^a_{jt+i-h}(t+1)]^2 + \lambda_j[x^a_{jt+i-h}(t+1) - x^a_{jt-1+i-h}(t+1)]^2\}.$$

The first-order conditions in this case still retain a similar relationship between forecast errors and forecast revisions as presented in this paper. Specifically, (3) can be rewritten as:

$$x(t+1) - x^a_{jt+i-h}(t+1) = \frac{1-\omega_{j1}}{\omega_{j1}}(x^a_{jt}(t+1) - x^a_{jt-1}(t+1)) + error.$$

Table 1: Timeline of Forecaster j's Loss Functions for Forecasts of $x(t+1)$[1]

Time (Forecast Horizon):	$t-2$ (4)	$t-1$ (3)	t (2)	$t+1$ (1)
Forecast for $x(t+1)$:	$x^a_{jt-3}(t+1)$	$x^a_{jt-2}(t+1)$	$x^a_{jt-1}(t+1)$	$x^a_{jt}(t+1)$
Current Revision:	$x^a_{jt-3}(t+1)-x^a_{jt-4}(t+1)$	$x^a_{jt-2}(t+1)-x^a_{jt-3}(t+1)$	$x^a_{jt-1}(t+1)-x^a_{jt-2}(t+1)$	$x^a_{jt}(t+1)-x^a_{jt-1}(t+1)$
Future Revisions:	$x^a_{jt-2}(t+1)-x^a_{jt-3}(t+1)$ $x^a_{jt-1}(t+1)-x^a_{jt-2}(t+1)$ $x^a_{jt}(t+1)-x^a_{jt-1}(t+1)$	$x^a_{jt-1}(t+1)-x^a_{jt-2}(t+1)$ $x^a_{jt}(t+1)-x^a_{jt-1}(t+1)$	$x^a_{jt}(t+1)-x^a_{jt-1}(t+1)$	—
Loss Function:	(iv)	(iii)	(ii)	(i)

$$\underset{x^a_{jt-3}(t+1)}{Min}\; E_{jt-3}\{[x(t+1)-x^a_{jt-3}(t+1)]^2+\lambda_j[x^a_{jt-3}(t+1)-x^a_{jt-4}(t+1)]^2+\delta\lambda_j[x^a_{jt-2}(t+1)-x^a_{jt-3}(t+1)]^2 \tag{iv}$$
$$+\delta\lambda_j[x^a_{jt-1}(t+1)-x^a_{jt-2}(t+1)]^2+\delta^2\lambda_j[x^a_{jt}(t+1)-x^a_{jt-1}(t+1)]^2\}$$
$$s.t.\quad E_{jt-3}x^a_{jt-2}(t+1)=\omega_{j3}E_{jt-3}x(t+1)+(1-\omega_{j3})x^a_{jt-3}(t+1)$$
$$s.t.\quad E_{jt-3}x^a_{jt-1}(t+1)=\omega_{j2}E_{jt-3}x(t+1)+(1-\omega_{j2})x^a_{jt-2}(t+1)$$
$$s.t.\quad E_{jt-3}x^a_{jt}(t+1)=\omega_{j1}E_{jt-3}x(t+1)+(1-\omega_{j1})E_{jt-3}x^a_{jt-1}(t+1)$$

$$\underset{x^a_{jt-2}(t+1)}{Min}\; E_{jt-2}\{[x(t+1)-x^a_{jt-2}(t+1)]^2+\lambda_j[x^a_{jt-2}(t+1)-x^a_{jt-3}(t+1)]^2+\delta\lambda_j[x^a_{jt-1}(t+1)-x^a_{jt-2}(t+1)]^2 \tag{iii}$$
$$+\delta^2\lambda_j[x^a_{jt}(t+1)-x^a_{jt-1}(t+1)]^2\}$$
$$s.t.\quad E_{jt-2}x^a_{jt-1}(t+1)=\omega_{j2}E_{jt-2}x(t+1)+(1-\omega_{j2})x^a_{jt-2}(t+1)$$
$$s.t.\quad E_{jt-2}x^a_{jt}(t+1)=\omega_{j1}E_{jt-2}x(t+1)+(1-\omega_{j1})E_{jt-2}x^a_{jt-1}(t+1)$$

$$\underset{x^a_{jt-1}(t+1)}{Min}\; E_{jt-1}\{[x(t+1)-x^a_{jt-1}(t+1)]^2+\lambda_j[x^a_{jt-1}(t+1)-x^a_{jt-2}(t+1)]^2+\delta\lambda_j[x^a_{jt}(t+1)-x^a_{jt-1}(t+1)]^2\} \tag{ii}$$
$$s.t.\quad E_{jt-1}x^a_{jt}(t+1)=\omega_{j1}E_{jt-1}x(t+1)+(1-\omega_{j1})x^a_{jt-1}(t+1)$$

$$\underset{x^a_{jt}(t+1)}{Min}\; E_{jt}\{[x(t+1)-x^a_{jt}(t+1)]^2+\lambda_j[x^a_{jt}(t+1)-x^a_{jt-1}(t+1)]^2\} \tag{i}$$

[1] This timeline depicts announcements made (conditional on previous period information) and respective loss functions discussed in section 2.

previous $h+1-$step-ahead forecast:

$$x^a_{jt+1-h}(t+1) = \omega_{jh}E_{jt+1-h}x(t+1) + (1-\omega_{jh})x^a_{jt-h}(t+1), \tag{6}$$
$$h = 1,...,H \quad \text{and} \quad j = 1,...,J,$$

where $\omega_{jh} = \frac{1+\delta\lambda_j\eta_{jh}}{1+\lambda_j+\delta\lambda_j\eta_{jh}}$, $\eta_{jh} = \omega^2_{jh-1} + \delta(1-\omega_{jh-1})^2\eta_{jh-1}$ and $\eta_{j1} = 0$. For a given forecast horizon, the trade-off between accuracy and continuity in forecasts may be pinned down by the value of λ_j. A value of λ_j close to zero, and hence a value of ω_{jh} close to one, would mean that the announced forecast is close to the forecaster's true expectation. On the other hand, a high value of λ_j would suggest a high degree of sluggishness in inflation announcements, as forecasters place a higher weight on the continuity of their announcement over its accuracy.[4]

3 Forecast Data

The announced forecast data come from the *Survey of Professional Forecasters* (SPF), the oldest quarterly survey of macroeconomic forecasts in the U.S. The survey began in 1968 and was conducted by the American Statistical Association and the National Bureau of Economic Research. The Federal Reserve Bank of Philadelphia assumed responsibility for the survey in June 1990. A survey participant's affiliation is kept confidential, but the individual responses are coded with an identification number for each forecaster that will be used to track their forecasts. The inflation forecasts used are quarterly forecasts for the CPI inflation rate (seasonally adjusted, at annual rates, in percentage points) and real GDP (seasonally adjusted, chain-weighted). These quarterly forecasts are annualized quarter-over-quarter percent changes and run from 1981:Q3 to 2017:Q4.

Estimation will be focused on the 29 forecasters in the sample that have submitted at least 50 inflation announcements. Of these 29 professional forecasters, 8 are financial service providers, 12 are nonfinancial service providers and the industry of 9 forecasters are either

[4] Note that the very first forecast that is made would not have any sluggishness (since there is no previous announcement that forecasters would be revising), so it would not be possible until the second period and there would be no revision ever made to the first one-step ahead forecast. Subsequent one-step ahead forecasts would be subject to revisions and therefore would exhibit a larger bias than forecasts made at longer horizons. In other words, as the forecast horizon, h, gets larger, past announcements would play less of a role.

missing or unknown.

Data for actual inflation and real GDP growth come from FRED® (Federal Reserve Economic Data), a database of over 14,000 U.S. economic time series maintained by the Federal Reserve Bank of St. Louis. Monthly data running from 1981 to 2017 are averaged to quarterly data and are then converted to an annualized quarterly rate to be consistent with the survey data.

The advantage of using CPI inflation is that the monthly U.S. CPI releases by the Bureau of Labour Statistics are typically not subject to large revisions. In contrast, GDP data are usually subject to large revisions, which could result in forecasts that differ systematically from announced values solely because of data revisions. Hence, the results that use GDP data presented in the following section focus on the second GDP release, but have also been repeated using the first and final releases.[5]

4 Estimation Strategy and Results

A natural question to consider is the value of λ_j in practice, along with its variation across forecasters. Estimating (6) using OLS is not be feasible since $E_{jt+1-h}x(t+1)$ is not observed. Instrumental variable (IV) estimation would be an alternative if a valid instrument could be obtained. If we impose FIRE, then actual inflation data could be used. In this case, $x(t+1) = E_{jt+1-h}x(t+1) + \varepsilon_{jt+1-h}(t+1)$ and we could write (6) as:

$$x^a_{jt+1-h}(t+1) = \omega_{jh}x(t+1) + (1-\omega_{jh})x^a_{jt-h}(t+1) - \omega_{jh}\varepsilon_{jt+1-h}(t+1), \tag{7}$$

where $\varepsilon_{jt+1-h}(t+1)$ is forecaster j's h-period rational expectations error. Under rational expectations, directly estimating (7) is problematic as the forecast error, $\varepsilon_{jt+1-h}(t+1)$, is correlated with actual inflation. Ang, Bekaert and Wei (2007) find that the SPF inflation forecast data are one of the leading forecasts of inflation. By definition, however, we cannot use this as an

[5] These various GDP releases are obtained from the Federal Reserve Bank of Philidelphia's Real-Time Data Set for Macroeconomists. The results in the following section hold regardless of whether the first, second, or final release of real GDP data is considered.

instrument since the SPF data will be used to capture inflation announcements.[6]

I opt to take an Anderson-Rubin (1949) (AR) approach, which allows a test of different values of λ_j (and hence $1-\omega_{jh}$). This procedure has frequently been used to perform testing in the presence of weak instruments. For example, Dufour, Khalaf and Kichian (2006) test the empirical relevance of Gali and Gertler's (1999) new-Keynsian Phillips curve (NKPC) equations using AR tests for both U.S. and Canadian data. Nason and Smith (2008) use the AR approach to provide a new set of tests of the forward-looking inflation model within the hybrid NKPC. The main advantage of the AR approach is that while the more conventional IV tests could be affected by weak identification, test statistics that result from the AR approach apply whether identification is weak or not.

We start by taking the future value of $x(t+1)$, to the left-hand side in (7):

$$x^a_{jt+1-h}(t+1) - \omega_{jh}x(t+1) = [1-\omega_{jh}-\alpha]x^a_{jt-h}(t+1) - \omega_{jh}\varepsilon_{jt+1-h}(t+1). \qquad (8)$$

Then, I choose a particular value of λ_j, labelled λ^0_j, from a grid of possible values, which would then yield a corresponding value of ω^0_{jh}.[7] While we cannot use this regression to estimate a value, it can be used to test any value of this weight on the expected future realization of x. To test whether $\lambda_j = \lambda^0_j$, the AR procedure involves a test of the null hypothesis that $\alpha = 0$. The estimate of λ_j is obtained by jointly estimating the first-order conditions after FIRE is imposed on the loss functions $(i)-(iv)$ given in Table 1. This yields an estimate of λ_j for when x represents inflation, and another estimate of λ_j when x depicts real GDP growth.[8] In this estimation, the AR procedure yields an F-statistic known as the AR-statistic, which follows a

[6] Stock and Yogo (2005) propose a test for weak instruments based on the Cragg-Donald (1993) statistic. For this test, I consider three possible instruments for $x(t+1)$: 3-month U.S. Treasury Bill interest rates, a lag of the log change in real WTI oil prices, and the previous quarter's one-quarter-ahead forecast error, all of which have been linked to future inflation rates. When this test is conducted using different combinations of these instruments, the resulting test statistics do not allow us to reject the null hypothesis that the instruments are weak. The consequence of the lack of explanatory power of these instruments is a larger bias in the estimated IV coefficients. This challenge of obtaining instruments for $x(t+1)$ in a similar context has been documented in past research (for examples, see Stock and Watson (1999) and Hansen, Lunde and Nason (2011)).

[7] The grid used for λ_j ranges from 0 to 10 and increases in increments of 0.1.

[8] The discount factor δ is set to 0.95 in both sets of estimation. To account for any correlation between the error term and time t variables, the following instruments are also included in the estimation: the corresponding median inflation or real GDP growth forecasts from the SPF in inflation and real GDP growth estimation, respectively, as well as the previous period's 3-month U.S. Treasury Bill interest rate, averaged into a quarterly series, and a lag of the log change in real WTI oil prices.

Fisher distribution under the null hypothesis. AR statistics are pivotal in finite samples, provide exact tests and are also robust to weak and omitted instruments. The interested reader may refer to Dufour (2003) for further details on the AR statistic as an approach to the weak instruments problem.

Table 2 presents the minimum-F-statistic (or maximum p-value) estimates of λ_j for both sets of estimation, along with the p-values of the test that $1 - \omega_{jh} = \alpha$ for all forecast horizons considered ($h = 1, .., 4$). Some interesting findings emerge. First, less than half of the sample of 29 forecasters have some degree of sluggishness in their inflation announcements, while over 60% of forecasters exhibit sluggishness in their real GDP growth forecasts. Only a minority of forecasters (about 25% of the sample) do not exhibit any sluggishness in either their inflation or real GDP growth forecasts. These forecasters simply announce their conditional expectations. Second, while sluggishness exists in both inflation and real GDP announcements, the degree of sluggishness is relatively more subdued in the former.[9] Third, the p-values associated with the minimum-F estimates of λ_j indicate that the model would not be rejected at the 1% significance level for both inflation and real GDP growth for the vast majority of forecasters.[10]

5 The Predictability of Forecast Errors and the Efficiency of Using Public Information

The relationship between forecast errors and forecast revisions has received attention in different contexts. CG, under the assumption of FIRE, derive the following relationship between *ex-post* forecast errors and forecast revisions:

$$x(t+1) - x_t^a(t+1) = -\theta^1(x_t^a(t+1) - x_{t-1}^a(t+1)) + \theta^2(x_{t-1}^a(t+1) - x_{t-2}^a(t+1)) + error. \quad (9)$$

CG find that forecast smoothing yields the predictability of *ex-post* forecast errors. They pool mean inflation forecasts and find no evidence of predictive power for the lagged forecast revisions

[9] The distinction in the findings between inflation and real GDP growth has been documented in past research (Faust and Wright (2009), Patton and Timmermann (2012)).

[10] Note, at the associated estimate of λ_j, the p-values on a joint test of the instruments also indicate that their coefficients are not significanty different from zero in most cases.

Table 2: Minimum-F Values of λ_j

Forecaster	**Inflation**	**RGDPG**
1 (U, ID: 20)	0.0	0.9
	(0.08)	(0.10)
2 (U, ID: 40)	0.0	0.0
	(0.07)	(0.00)
3 (U, ID: 65)	0.2	0.0
	(0.33)	(0.50)
4 (U, ID: 84)	0.0	6.4
	(0.05)	(0.27)
5 (U, ID: 94)	0.0	0.0
	(0.93)	(0.09)
6 (U, ID: 99)	0.0	0.0
	(0.81)	(0.32)
7 (F, ID: 407)	0.2	0.1
	(0.07)	(0.02)
8 (NF, ID: 411)	0.3	2.8
	(0.55)	(1.00)
9 (F, ID: 420)	2.6	0.0
	(0.00)	(0.73)
10 (F, ID: 421)	0.0	0.5
	(0.04)	(0.32)
11 (F, ID: 422)	0.0	0.1
	(0.02)	(0.18)
12 (F, ID: 426)	0.3	0.4
	(0.00)	(0.01)
13 (NF, ID: 428)	0.1	6.3
	(0.48)	(0.09)
14 (NF, ID: 429)	0.1	0.5
	(0.97)	(0.52)

Notes: (1) p-values of the estimation of $(i) - (iv)$ in Table 1 (described in section 2) are in parentheses. (2) Instruments included are median inflation or RGDPG forecasts from the previous period, 3-month T-bill rates, and a lag of the log change in WTI oil prices. (3) $\delta = 0.95$ by assumption. (4) Acronyms used are: F: financial industry forecaster; NF: nonfinancial industry forecaster; U: industry unknown; RGDPG: real GDP growth.

Table 2 cont'd: Minimum-F Values of λ_j

Forecaster	**Inflation**	**RGDPG**
15 (F, ID: 431)	0.0	2.7
	(0.40)	(0.82)
16 (NF, ID: 433)	0.0	0.3
	(0.52)	(0.51)
17 (NF, ID: 446)	0.1	2.7
	(0.39)	(0.70)
18 (F, ID: 456)	0.0	1.0
	(0.00)	(0.29)
19 (U, ID: 463)	0.0	0.0
	(0.70)	(0.99)
20 (NF, ID: 472)	4.2	0.3
	(0.88)	(0.78)
21 (U, ID: 483)	1.8	1.1
	(0.98)	(0.27)
22 (NF, ID: 484)	0.9	0.0
	(0.79)	(0.96)
23 (F, ID: 504)	0.0	0.0
	(0.90)	(0.73)
24 (NF, ID: 506)	0.0	0.0
	(0.18)	(0.11)
25 (U, ID: 507)	0.0	0.0
	(0.10)	(0.87)
26 (NF, ID: 510)	0.0	0.2
	(0.21)	(0.04)
27 (NF, ID: 512)	0.1	0.0
	(0.59)	(0.05)
28 (NF, ID: 518)	0.0	4.6
	(0.47)	(0.24)
29 (NF, ID: 520)	7.6	4.8
	(0.32)	(0.97)

Notes: (1) p-values of the estimation of $(i) - (iv)$ in Table 1 (described in section 2) are in parentheses. (2) Instruments included are median inflation or RGDPG forecasts from the previous period, 3-month T-bill rates, and a lag of the log change in WTI oil prices. (3) $\delta = 0.95$ by assumption. (4) Acronyms used are: F: financial industry forecaster; NF: nonfinancial industry forecaster; U: industry unknown; RGDPG: real GDP growth.

and the coefficient on the contemporaneous forecast revision is positive, consistent with the existence of informational rigidities.

Lahiri and Sheng (2008) (henceforth, LS) investigate the efficiency of using public information in a Bayesian learning framework by studying a panel of countries' aggregate forecasts. They derive a testable relationship between forecast errors and forecast revisions that may be used to perform an efficiency test on the forecasts made by professional forecasters as follows:

$$x(t+1) - x_t^a(t+1) = \theta(x_t^a(t+1) - x_{jt-1}^a(t+1)) + u_t, \tag{10}$$

where under the null hypothesis $\hat{\theta}_{jt+1-h}$ should be zero, suggesting that forecasters attach an efficient weight to their prior beliefs. Conversely, a positive (negative) estimate of θ would be observed if the forecaster overweights (underweights) her prior belief (or underweights [overweights] public information). They find that forecasters tend to underweight new public information in the medium term in their real GDP forecasts. In the shorter and longer horizons, though, they tend to overweight public information.

In the context of sluggish forecasts and the model presented in this paper, it is possible to rewrite the first-order condition with FIRE imposed in a similar form as (9) and (10) (see footnote 3) and test for the predictability of forecasters and efficiency of using public information at the individual forecaster level. Table 3 presents the results.

Panel A shows the coefficients on the contemporaneous and lagged forecast revisions. To be consistent with CG, the regression also includes two lags of log changes in oil prices. There are some interesting distinctions from CG when this relationship is estimated for individual forecasters, notably that for the majority of forecasters, the coefficient on contemporaneous revisions is negative, suggestive of forecast-smoothing behaviour. The coefficient on lagged forecast revisions is also primarily negative for most forecasters (though only statistically significant in a handful of cases), whereas CG find no evidence of predictive power in lagged forecast revisions.

Panel B shows the estimates of (10), which also contain negative coefficients for the majority of forecasters; however, these estimates are statistically significant in only a few cases, suggesting that most forecasters efficiently use new public information when $h = 1$. For the few forecasters

Table 3: Estimates of Forecast Error Predictability and the Efficiency of Using Public Information

Forecaster	Panel A: CG(2015)		Panel B: LS(2008)
	$\hat{\theta}_j^1$	$\hat{\theta}_j^2$	$\hat{\theta}_{j1}$
1	-0.52**	-0.65**	-0.27**
	(0.18)	(0.18)	(0.09)
2	-0.46**	-0.75**	-0.20
	(0.14)	(0.24)	(0.22)
3	-0.52**	-0.41*	-0.38
	(0.15)	(0.21)	(0.21)
4	-0.14	-0.24	0.11
	(0.13)	(0.24)	(0.14)
5	-0.71*	-0.19	-0.23
	(0.16)	(0.22)	(0.22)
6	-0.42	0.83	0.08
	(0.24)	(0.54)	(0.07)
7	0.12	0.02	0.04
	(0.20)	(0.37)	(0.21)
8	0.21*	-0.07	-0.16
	(0.12)	(0.20)	(0.18)
9	0.08	0.15	-0.28
	(0.26)	(0.29)	(0.18)
10	0.24	0.08	0.00
	(0.19)	(0.36)	(0.19)
11	-0.18	-0.87**	-0.52
	(0.24)	(0.34)	(0.28)
12	-0.35**	-0.43*	-0.07
	(0.17)	(0.26)	(0.13)
13	0.71**	-0.24	-0.10
	(0.29)	(0.39)	(0.16)
14	0.01	0.20	-0.45**
	(0.12)	(0.24)	(0.19)

Notes: (1) Estimation of (9) for inflation and (10) for RGDPG are presented in panel A and B, respectively. (2) Standard errors are in parentheses. * denotes significance at the 10% level, ** denotes significance at the 5% level.

Table 3 cont'd: Estimates of Forecast Error Predictability and the Efficiency of Using Public Information

Forecaster	**Panel A: CG(2015)**		**Panel B: LS(2008)**
	$\hat{\theta}_j^1$	$\hat{\theta}_j^2$	$\hat{\theta}_{j1}$
15	0.15	-0.30	-0.04
	(0.19)	(0.31)	(0.18)
16	-0.18	-0.22	-0.11
	(0.13)	(0.26)	(0.16)
17	0.26	0.57	0.26
	(0.19)	(0.55)	(0.19)
18	-0.35*	-0.66*	-0.46**
	(0.19)	(0.35)	(0.15)
19	0.15	-0.09	0.00
	(0.10)	(0.26)	(0.11)
20	-0.35**	0.26	0.19
	(0.13)	(0.22)	(0.17)
21	0.36	0.53	-0.32
	(0.50)	(0.85)	(0.13)
22	-1.04**	-0.85	-0.24
	(0.32)	(0.65)	(0.19)
23	-0.10	-0.19	-0.34**
	(0.11)	(0.18)	(0.14)
24	-0.11	0.08	-0.08
	(0.11)	(0.24)	(0.21)
25	-0.07	-0.30	-0.25
	(0.17)	(0.33)	(0.14)
26	-0.47*	-0.74**	-0.35
	(0.26)	(0.27)	(0.22)
27	-0.43	-0.77*	-0.14
	(0.30)	(0.44)	(0.22)
28	-0.23	-0.22	-0.03
	(0.22)	(0.25)	(0.20)
29	0.06	-1.93**	0.21
	(0.48)	(0.65)	(0.28)

Notes: (1) Estimation of (9) for inflation and (10) for RGDPG are presented in panel A and B, respectively. (2) Standard errors are in parentheses. * denotes significance at the 10% level, ** denotes significance at the 5% level.

that overweight new public information, this is consistent with the findings of LS at shorter horizons. The intuition is that as the target date of the forecast approaches, base-year GDP growth numbers become available with more certainty; as a result, forecasters place a higher weight on newly arrived public information.

These results provide an interesting complement to the aggregated data used by CG and LS, suggesting that the use of mean forecasts can yield different conclusions than when forecasters are analyzed at the individual level: there appears to be more evidence of forecast smoothing in inflation and forecasters place a higher weight on new public information in their current period real GDP growth forecasts.

6 Sluggishness-Adjusted Expectations

The volatility in inflation forecasts that come from survey data has fallen in recent years (Clark and Davig (2011)), more so than the volatility of actual inflation (Leduc, Sill and Stark (2007)). Survey forecasts also react less to macroeconomic shocks than the actual data (Jain (2017)). Despite this, surveys have still been shown to outperform other methods of forecasting (Ang, Bekaert and Wei (2007)). It may still be possible to improve on survey forecasts, however, if we could adjust them for any sluggishness that may be present. This section explores this possibility.

Notice that from the first-order conditions derived in section 2, once we have an estimate for the degree of sluggishness, $\hat{\lambda}_j$, we can solve for forecaster j's implicit or sluggishness-adjusted expectations since we have data on inflation announcements. That is, rearranging (6) and denoting implicit expectations of x by $\hat{E}_{jt+1-h}x(t+1)$ yields:

$$\hat{E}_{jt+1-h}x(t+1) = [x^a_{jt+1-h}(t+1) - (1-\hat{\omega}_{jh})x^a_{jt-h}(t+1)]/\hat{\omega}_{jh}. \tag{11}$$

This is computed for each forecaster j. For the sake of compactness, Figures 1A and 1B show a comparison between median announcements reported in the SPF (captured by the blue solid line), and the corresponding median sluggishness-adjusted series (captured by the red dashed line). For reference, the height of the bar shows the corresponding level of realized inflation and

Figure 1A: Median Announced (Blue) and Median Implied (Red) Series for Inflation

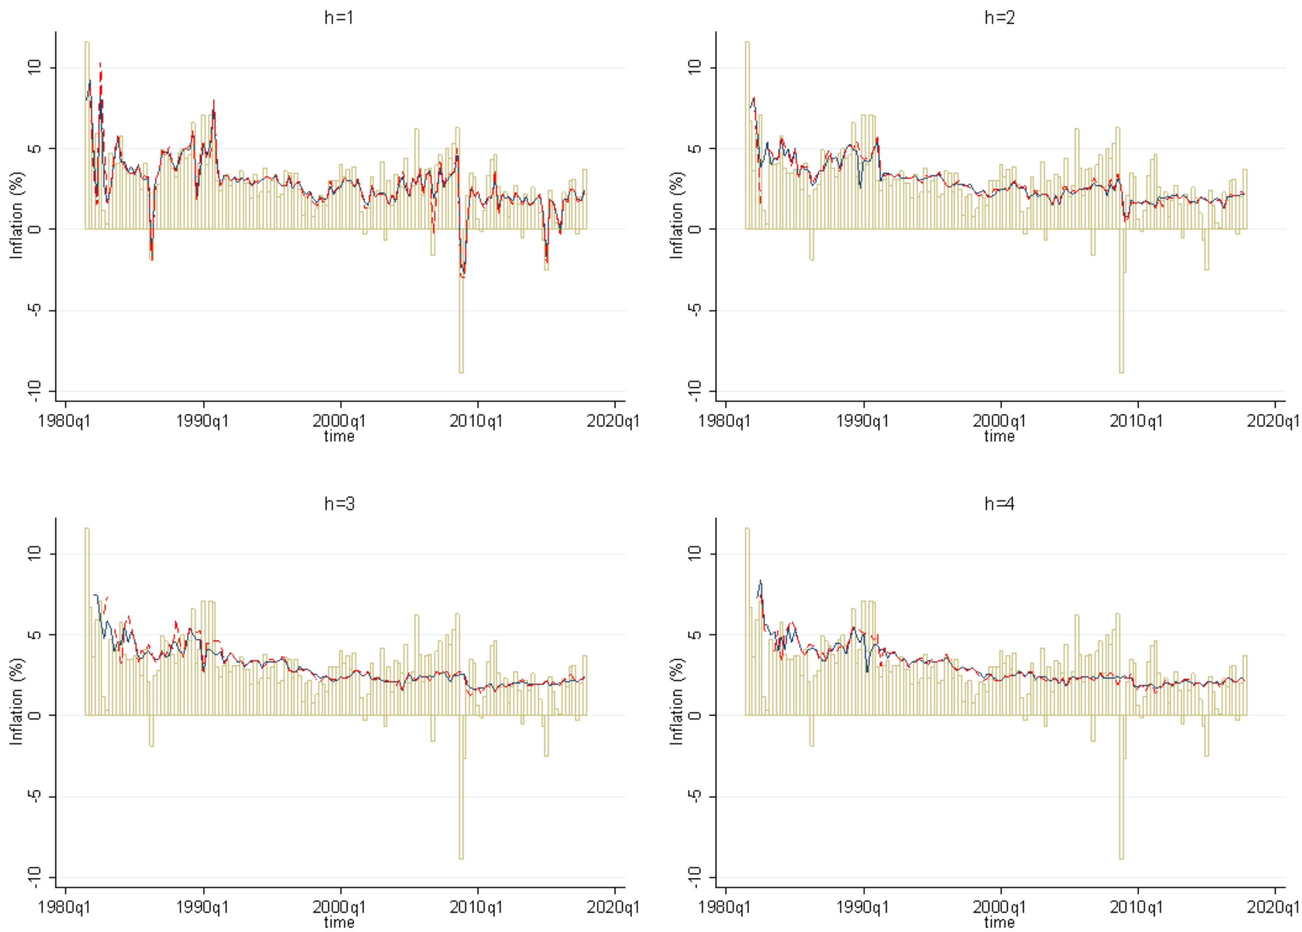

Notes: The implied or sluggishness-adjusted expectations series (red, dashed) is computed using (11) for each forecaster j. The median of the series is depicted above. The median of the announced forecasts (blue, solid) from the SPF is computed for the 29 forecasters in the sample. The bars indicate the realized value of the series corresponding to the forecasts.

Figure 1B: Median Announced (Blue) and Median Implied (Red) Series for Real GDP Growth

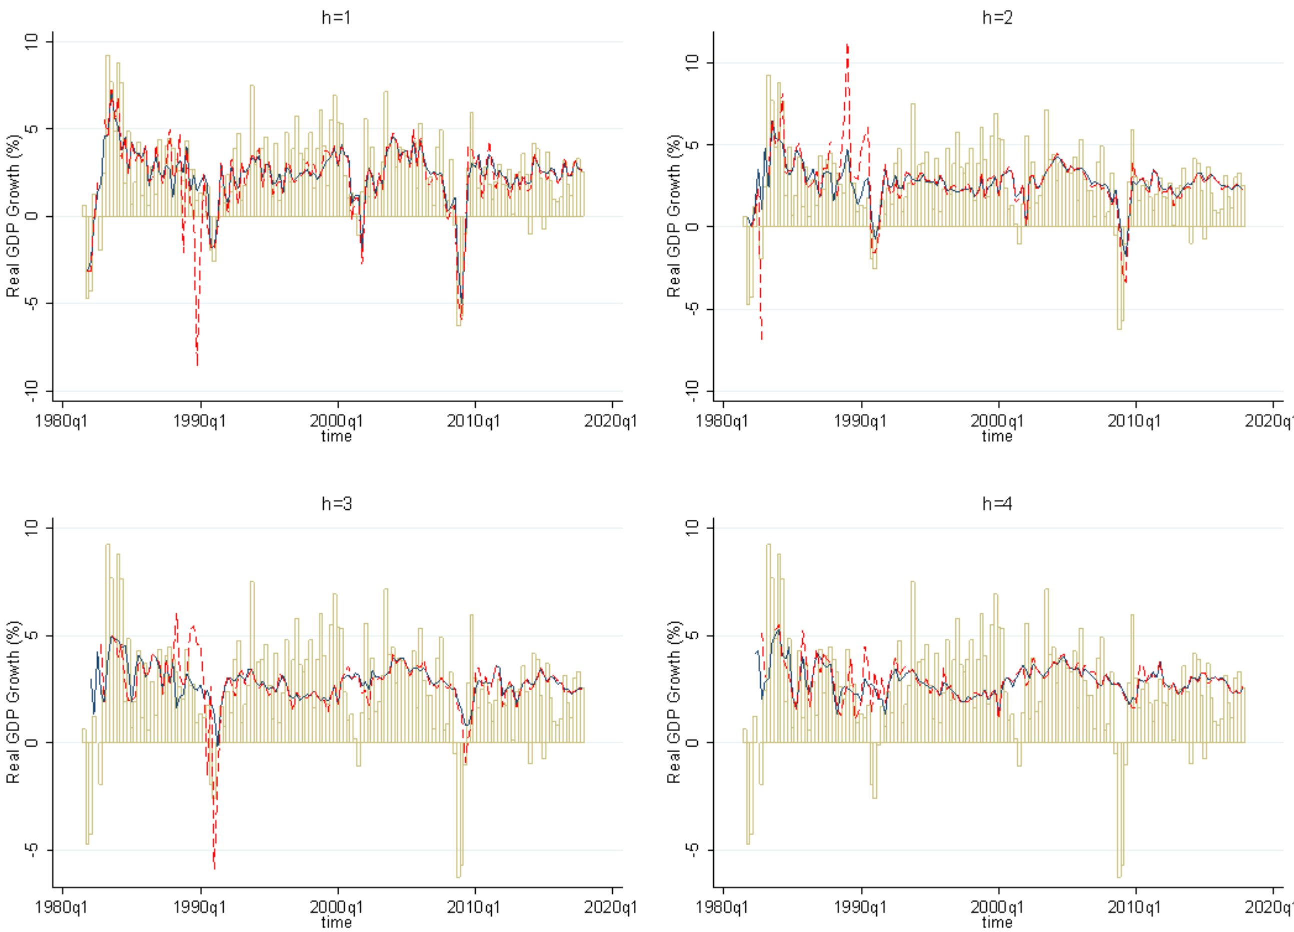

Notes: The implied or sluggishness-adjusted expectations series (red, dashed) is computed using (11) for each forecaster j. The median of the series is depicted above. The median of the announced forecasts (blue, solid) from the SPF is computed for the 29 forecasters in the sample. The bars indicate the realized value of the series corresponding to the forecasts.

real GDP growth, respectively, throughout the sample period. Note that in most cases, implicit expectations are a more volatile series than announcements, as the model would suggest. In some cases, implicit expectations are very close to the corresponding announced value, and in some cases they are not. In some situations, notice the implicit series is closer to the realized series than announced values, though in some cases it is not.

To obtain a sense of whether the sluggishness-adjusted expectations series provides any additional accuracy over the announcements found in the data, we compute the relative RMSEs using the median of the sluggishness-adjusted series to the RMSEs of the median of the announcements. Table 4 summarizes these calculations. When the full sample is considered in Panel A, the sluggishness-adjusted expectations series does not perform better than the survey data. This is mostly due to a median based on very few participants in the survey at the beginning of the sample (hence the extreme volatility in the real GDP growth series early in the series, which has a larger adjustment because of higher sluggishness in GDP). Panel B considers a comparison from 1995Q1 onwards, when more forecasters have joined the survey: the sluggishness-adjusted series now does at least as well as the announced series for both inflation and real GDP growth for almost all horizons. In Figure 1B we note that the sluggishness-adjusted series does a better job at capturing major falls in real GDP growth, so consider Panel C, which includes only the last decade of data—here the sluggishness-adjusted series improves further in terms of accuracy over the announced data.

Table 4: Relative RMSEs

	Panel A: 1981Q3–2017Q4		Panel B: 1995Q1–2017Q4		Panel C: 2008Q1–2017Q4	
h	**Inflation**	**RGDPG**	**Inflation**	**RGDPG**	**Inflation**	**RGDPG**
1	1.00	1.15	0.91	0.99	0.93	0.95
2	1.02	1.06	0.99	0.95	0.99	0.90
3	1.05	1.07	1.00	1.00	1.00	0.99
4	1.00	1.06	0.99	1.02	1.00	1.01

Notes: Ratio of RMSE of the median of sluggishness-adjusted expectation series to the RMSE of the median SPF forecast is presented.

The exceptions are in the four-quarters-ahead horizon, where the advantage of using the sluggishness-adjusted series appears to diminish, particularly for real GDP growth. This occurs as the sluggishness-adjusted series becomes less volatile as h increases, and closer to the announced data.

While many papers have looked at the presence of deviations from FIRE in survey data, perhaps as a result of reputational concerns or informational rigidities faced by forecasters, very few have proposed making specific adjustments to improve these data. Ang et al. (2007) consider a linear and non-linear bias adjustment to survey data, but find that the raw survey forecasts tend to outperform the adjusted forecasts. While this may be true during certain periods of the survey (for example, when participation is low), the adjusted series presented here can provide a very useful alternative series that contains more volatility than its raw data counterpart.

A natural question to consider is whether sluggishness-adjusted expectations can be informative when x has not yet been realized. To address this, I re-estimate the degree of sluggishness, $\hat{\lambda}_j$, using data from 1981Q3–2014Q4, and recompute sluggishness-adjusted expectations for each forecaster using (11).[11] This new series is somewhat more volatile than both the announced data and the adjusted series calculated using the full sample. A similar median relative RMSE analysis is presented in Table 5, which suggests that during the out-of-sample portion,

Table 5: Relative RMSEs for Out-of-Sample Exercise

	Period: 2015Q1–2017Q4	
h	**Inflation**	**RGDPG**
1	0.98	1.06
2	1.00	0.95
3	0.99	1.02
4	0.97	1.00

Notes: (1) Ratio of RMSE of the median of sluggishness-adjusted expectation series to the RMSE of the median SPF forecast is presented. (2) Out-of-sample period begins in 2015Q1.

the accuracy of inflation between the sluggishness-adjusted series still remains at least as accurate

[11] While it would be ideal to have a larger adjusted expectations series out-of-sample, I choose a shorter period since it is also important to ensure we have enough observations for each forecaster j to perform the estimation of λ_j.

as the survey data. For real GDP growth the out-of-sample results are mixed: the relative accuracy deteriorates when $h = 1$, improves when $h = 2$, is slightly worse when $h = 3$, and has similar accuracy as the survey data when $h = 4$. Hence, even if the degree of sluggishness is not estimated over the full sample period, it can still provide an informative, more volatile, and at times more accurate alternative over its survey data counterpart.

7 Conclusions

Given the influence that expectations have on the level of realized macroeconomic variables, it is vital to have a deeper understanding of the process of forecasting. This paper considers the situation in which professional forecasters have rational expectations that may not be reflected in their forecast announcements. In the sluggishness model, this discrepancy between true expectations and forecast announcements is fuelled by forecasters' tendency to maintain continuity in their forecasts, though it compromises their accuracy.

From the point of view of a policymaker, for example, a central bank interested in understanding whether inflation expectations are well-anchored could find that since inflation expectations as measured by survey data have not changed much in response to inflationary or deflationary pressures, agents are more confident that the central bank would not let inflation persistently deviate from the target. Suppose, however, that inflation announcements instead simply adjust too slowly or by too little to provide a useful guide for monetary policy.[12] From the perspective of a forecaster, they may be subject to informational rigidities or be required to give a compelling and well-supported explanation any time a previous forecast is revised, causing them to eliminate certain shocks from their forecasts.

Starting with a relatively simple loss function, this is indeed what is found in many of the inflation and real GDP growth forecasts. The focus on individual forecasters in this analysis also reminds us that forecasters are different. Some forecasters smooth their forecasts for reputational reasons, while others may not make revisions because of informational rigidities they face. Others are simply interested in making good, accurate forecasts. These subtle differences can deepen

[12] Research has also found other reasons why survey data may be an imperfect guide for research and monetary policy. For example, Ghysels and Wright (2009) emphasize that survey forecasts often become stale because they are released fairly infrequently.

our understanding of how expectations are formed, but are lost when analysis is focused entirely on mean or median forecasts.

Finally, if sluggishness occurs at the cost of accuracy in forecasts, then we should observe at least some improvement in forecast accuracy with the sluggishness-adjusted expectations series when it is compared with the raw survey data. The adjusted series presented in this paper is found to perform at least as well as the survey data for inflation and real GDP growth forecasts for most horizons.

There still remains significant potential for further research in this area. Survey data have been shown to do quite well at forecasting, outperforming many traditional time-series models, and proposing different types of adjusted expectations series based on these survey data is an area of research that still remains underdeveloped. With more research, this can lead to superior forecasting.

References

[1] Anderson, Theodore W., and Herman Rubin. 1949. Estimation of the Parameters of a Single Equation in a Complex System of Stochastic Equations. *Annals of Mathematical Statistics*, 20: 46-63.

[2] Andrade, Philippe, and Herve Le Bihan. 2013. Inattentive Professional Forecasters. *Journal of Monetary Economics*, 60: 967-982.

[3] Ang, Andrew, Geert Bekaert, and Min Wei. 2007. Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better? *Journal of Monetary Economics*, 54: 1163–1212.

[4] Batchelor, Roy. 2007. Bias in Macroeconomic Forecasts. *International Journal of Forecasting*, 23: 189–203

[5] Batchelor, Roy, and Pami Dua. 1992. Conservatism and Consensus-seeking among Economic Forecasters. *Journal of Forecasting*, 11: 169-181.

[6] Clark, Todd E., and Troy Davig. 2011. Decomposing the Declining Volatility of Long-term Inflation Expectations. *Journal of Economic Dynamics and Control*, 35(7): 981-999.

[7] Clements, Michael P. 2014. US Inflation Expectations and Heterogeneous Loss Functions, 1968-2010. *Journal of Forecasting,* 33: 1-14.

[8] Coibion, Olivier, and Yuriy Gorodnichenko. 2015. Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts. *American Economic Review,* 105(8): 2644–2678.

[9] Dovern, Jonas, Ulrich Fritsche, Prakash Loungani, and Natalia Tamirisa. 2014. Informational Rigidities: Comparing Average and Individual Forecasts For a Large International Panel. *International Journal of Forecasting*, 21: 144-154.

[10] Dufour, Jean-Marie. 2003. Weak Instruments and Statistical Inference in Econometrics. *Canadian Journal of Economics*, 36(4): 767-808.

[11] Dufour, Jean-Marie, Lynda Khalaf, and Maral Kichian. 2006. Inflation Dynamics and the New Keynesian Phillips Curve: An Identification-Robust Econometric Analysis. *Journal of Economic Dynamics and Control*, 30: 1707-1727.

[12] Ehrbeck, Tilman, and Robert Waldmann. 1996. Why Are Professional Forecasters Biased? Agency versus Behavioral Explanations. *Quarterly Journal of Economics*, 111(1): 21-40.

[13] Gali, Jordi, and Mark Gertler. 1999. Inflation Dynamics: A Structural Econometric Analysis. *Journal of Monetary Economics*, 44: 195-222.

[14] Ghysels, Eric, and Jonathan H. Wright. 2009. Forecasting Professional Forecasters. *Journal of Business and Economic Statistics*, 27(4): 504-516.

[15] Faust, Jon, and Jonathan H. Wright. 2009. Comparing Greenbook and Reduced Form Forecasts Using a Large Realtime Dataset, *Journal of Business and Economic Statistics*, 27(4): 468-479.

[16] Hansen, Lars Peter R., Asger Lunde, and James M. Nason. 2011. Model Confidence Sets for Forecasting Models. *Econometrica*, 79: 453-497

[17] Jain, Monica. 2017. Perceived Inflation Persistence. *Journal of Business and Economic Statistics,* forthcoming.

[18] Kirchgassner, Gebhard, and Ulrich K. Muller. 2006. Are Forecasters Reluctant to Revise Their Predictions? Some German Evidence. *Journal of Forecasting,* 25: 401-413.

[19] Lahiri, Kajal and Xuguang Sheng. 2008. Evolution of Forecast Disagreement in a Bayesian Learning Model. *Journal of Econometrics*, 144: 325-340.

[20] Laster, David, Paul Bennett, and In Sun Geoum. 1999. Rational Bias in Macroeconomic Forecasts. *Quarterly Journal of Economics*, 114(1): 293-318.

[21] Leduc, Sylvain, Keith Sill, and Tom Stark. 2007. Self-fulfilling Expectations and the Inflation of the 1970s: Evidence from the Livingston Survey. *Journal of Monetary Economics*, 54: 433-459.

[22] Maćkowiak, Bartosz, and Mirko Wiederholt. 2009. Optimal Sticky Prices under Rational Inattention. *American Economic Review*, 99: 769-803.

[23] Mankiw, Gregory N., Ricardo Reis, and Justin Wolfers. 2003. Disagreement about Inflation Expectations. *NBER Macroeconomics Annual*, 209-248.

[24] Mertens, Elmar, and James M. Nason. 2015. Inflation and Professional Forecast Dynamics: An Evaluation of Stickiness, Persistence, and Volatility, mimeo.

[25] Nason, James M., and Gregor W. Smith. 2008. Identifying the New Keynesian Phillips Curve. *Journal of Applied Econometrics*, 23: 525-551.

[26] Nordhaus, William D. 1987. Forecasting Efficiency: Concepts and Applications. *Review of Economics and Statistics*, 69 (4): 667-674.

[27] Ottaviani, Marco, and Peter N. Sorensen. 2006. Reputational Cheap Talk, *RAND Journal of Economics*, 37(1): 155-175.

[28] Patton, Andrew J., and Allan Timmermann. 2012. Predictability of Output Growth and Inflation: A Multi-horizon Survey Approach, *Journal of Business and Economic Statistics*, 29(3): 397-410.

[29] Pesaran, Hashem M., and Martin Weale. 2006. Survey Expectations. In *Handbook of Economic Forecasting, Volume 1*, edited by Graham Elliott, Clive W.J.Granger, and Allan Timmermann, 715-776.

[30] Sims, Christopher A. 2003. Implications of Rational Inattention. *Journal of Monetary Economics*, 50(3): 665-690.

[31] Stock, James H., and Mark W. Watson. 1999. Forecasting Inflation. *Journal of Monetary Economics*, 44: 293-335.

[32] Stock, James H., and Motohiro Yogo. 2005. Testing for Weak Instruments in Linear IV Regression. In *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, edited by D. W. K. Andrews and J. H. Stock, 80-108. Cambridge: Cambridge University Press.

[33] Woodford, Michael. 2002. Imperfect Common Knowledge and the Effect of Monetary Policy. In *Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, edited by P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, Princeton University Press.