

Conservatism in Inflation Forecasts

Monica Jain¹

October 2014

Abstract:

This paper examines the extent to which professional forecasters exhibit conservatism in their announced inflation forecasts. Conservatism occurs as a result of a secondary objective captured in the forecaster's loss function: they would like to avoid having large forecast errors, but also would like to avoid making large revisions to previously announced forecasts. As a result of forecasters' reluctance to make large revisions, announced inflation forecasts change less than true inflation expectations. The Anderson-Rubin approach is used to determine the degree of conservatism in U.S. inflation forecasts, and evidence suggests it may be quite substantial. The findings in this paper have implications for both research and monetary policy that uses inflation forecast data. In the case of a central bank that is interested in understanding the extent to which inflation expectations are well-anchored, the presence of conservatism could lead to a series of inflation announcements that simply adjust too slowly or by too little to provide a useful guide for monetary policy. This paper presents an estimate of the degree of conservatism in inflation announcements as well as a method to undo the conservatism bias and therefore recover a series of forecaster's "true" or conservatism-adjusted inflation expectations. This true inflation expectations series is more volatile than the announced inflation forecasts observed in the data for most forecasters, suggesting that inflation expectations may not be as well-anchored as inflation announcements may suggest.

JEL classification: E31, E37

Keywords: conservatism, forecast smoothing, rational bias, loss function, Survey of Professional Forecasters

¹Monica Jain: Canadian Economic Analysis Department, Bank of Canada, jain[at]bankofcanada.ca. I am thankful to Gregor Smith and Allan Gregory for their guidance and suggestions. I would also like to thank Peter Christoffersen, Nicolas Vincent, Marc-Andre Gosselin, 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), and the Bank of Canada Fellowship Learning Exchange (2014) 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.

1 Introduction

The formation of inflation expectations and inflation forecasts is a process that is generally not straightforward to understand. This is because forecasts are usually not directly derived from the outcome of a model, but through a combination of both model-based forecasts and judgement. Keane and Runkle (1990) investigated the rationality behind reported forecasts and found strong evidence that professional forecasters are indeed rational. On the contrary, a large, subsequent literature documents departures from rationality (Batchelor, 2007), (Bennett, Geoum and Laster (1999)) along with substantial disagreement and biases across forecasters (Capistran and Timmermann, 2009). Models of inattentive agents such as sticky information models (Mankiw, Reis and Wolfers, 2003) and noisy information models (Woodford, 2002, Sims, 2003 and Mackowiak and Wiederholt, 2009) have been proposed. Papers such as Andrade and Le Bihan, 2013 have tested these models with expectations data to help reconcile the observed departure from full rationality and heterogeneity across forecasters, and have managed to capture some features of the data. Based on the predictions of these models of informational rigidity, Coibion and Gorodnichenko (2012) 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 full information rational expectations as a deviation from full information rather than a deviation from rational expectations, driven by informational rigidities.

The goal of this paper is to consider a departure from purely rational forecasts that reflects forecasters' dual objective: to avoid having large forecast errors, but also to avoid making large revisions to their previous forecasts. Like Bennett, Geoum and Laster (1999), the framework presented in this paper is one of *rational bias*. Bennett, *et al.* present a model in which forecasters have common information and identical expectations, though their inflation forecasts differ to the extent that their wages depend on the publicity that their forecasts are able to generate. However, in this paper, forecasters have an incentive to compromise the accuracy of their forecasts to make conservative forecasts. In the context presented here, the term *conservative* does not reflect one's political views, principles or habits; it can be viewed as a form of rational sluggishness in setting inflation forecasts, measured relative to the adjustment cost of making revisions to previously announced inflation forecasts. Conservatism may follow if the end-users of the forecasts (such as businesses or policy makers), distrust the forecasters who make frequent forecast revisions. As noted by Nordhaus (1987), forecasters may smooth their forecasts because more accurate but jumpy forecasts would "drive customers crazy," since forecast revisions which occur too frequently could lead to the reversal of investment plans too often. Nordhaus also adds that a conservative bias may

be attributed to agents tendency to hold on to old notions that are more familiar, rather than adjusting to surprises. It may also arise if forecasters have an informational advantage that they would not like to disclose via making a forecast revision. In any of these cases, this latter objective introduces a potential source of sluggishness into the inflation forecasting process which can create a divergence between announced inflation forecasts and true inflation expectations.

In this paper, forecasters have a very specific alternative to making accurate forecasts; that is, to retain continuity in their forecasts which may be captured by their previous forecast announcements. Forecasters' true inflation expectations are rational, but they may not necessarily announce this true expectation because they are reluctant to make large revisions to previous forecasts. Kirchgassner and Muller (2006) test a similar hypothesis using semi-annual forecasts of German macroeconomic time series for national income, private consumption, private investment, exports and imports made by the Association of German Economic Research Institutes. Using seemingly unrelated regression (SUR) estimation for these variables, they find evidence that forecasters are indeed reluctant to revise their predictions arising from an effort to protect the reputation of the forecasting institution.¹ Messina, Sinclair, and Stekler (2014) examine the informational content of forecast revisions to the Greenbook forecasts and find no evidence of forecast smoothing. When a similar exercise is performed using the revisions made by U.S. professional forecasters, however, there did appear to be some evidence of forecast smoothing.

The primary contribution of this paper is that it not only provides estimates of the degree of sluggishness in inflation announcements for a group of individual forecasters, but also a method to undo their conservative bias by recovering an alternative series of forecasters' "true" expected inflation. Using the Anderson-Rubin (1949) approach which allows for reliable inference in the presence of weak instruments, findings suggest there may be a substantial degree of sluggishness in inflation announcements as a result of forecasters' reluctance to make large revisions to their forecasts, which becomes more prevalent at forecast horizons beyond the current quarter. As a result, for most forecasters, "true" inflation expectations are more volatile than their announced inflation forecasts.

The presence of conservatism does not only have implications for research that uses this data, but could be misleading from the perspective of monetary policy. A central bank interested in understanding whether inflation expectations are well-anchored or in obtaining

¹Batchelor and Dua (1992) also examine the extent of conservatism and consensus-seeking in a panel from the Blue Chip Economic Indicators. Using a method-of-moments estimator and monthly data running from August 1976 to December 1986, they find that most forecasters are too conservative and give too much weight to their own past forecasts. They find this to be true for forecasts for real GNP growth, inflation in the GNP deflator, the unemployment rate and short-term interest rates.

measures of credibility, could find that since inflation expectations have not changed much in response to inflationary or deflationary pressures, that 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. A recent IMF report indicates that this may be true for data on long-term inflation announcements, and therefore, shorter term forecasts ought to be used to inform monetary policy questions (Modhadam, Teja and Berkmen, 2014). This paper focuses on understanding the conservatism bias in forecasts made for shorter horizons, and finds that these forecasts may still be an imperfect guide for monetary policy.

2 A Conservative Loss Function

A typical forecast announcement is denoted by $\pi_{jt}^a(t+h)$ which is forecaster j 's inflation announcement, π^a , for time $t+h$, h periods ahead. For example, $\pi_{jt}^a(t+1)$ and $\pi_{jt-1}^a(t+1)$ would be forecaster j 's inflation announcements one-period ahead (hence, based on time t information) and two-periods ahead (based on time $t-1$ information), respectively. It is fairly standard to assume that the following is true:

$$\pi_{jt}^a(t+h) = E_{jt}\pi(t+h), \quad \text{for } h = 1, \dots, H, \quad (1)$$

where $E_{jt}\pi(t+h)$ is forecaster j 's true expectation of inflation in period $t+h$. This states that for any forecast horizon, forecasters announce their true expectations of inflation. In this paper, this assumption will be relaxed.

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

$$\pi_{jt}^a(t+1) - \pi_{jt-1}^a(t+1).$$

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

$$\pi_{jt+1-h}^a(t+1) - \pi_{jt-h}^a(t+1).$$

A forecaster's 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 a secondary objective which may follow if the end-users of the forecasts prefer that forecasters do not make frequent revisions to their previous forecasts.

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

$$\pi_{jt}^a(t+1) \underset{Min}{E_{jt}} \{ [\pi(t+1) - \pi_{jt}^a(t+1)]^2 + \lambda_j [\pi_{jt}^a(t+1) - \pi_{jt-1}^a(t+1)]^2 \}, \quad (2)$$

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.² The first term penalizes inaccuracy while the second term penalizes revisions. Quadratic loss functions are frequently used in monetary policy problems. They have the advantage that they penalize large deviations from the target more than small deviations. 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 his or her announcement so that:

$$\pi_{jt}^a(t+1) = \frac{1}{1+\lambda_j} E_{jt} \pi(t+1) + \frac{\lambda_j}{1+\lambda_j} \pi_{jt-1}^a(t+1), \quad (3)$$

where $E_{jt} \pi(t+1)$ is forecaster j 's true expectation of inflation in period $t+1$. Notice our conservative loss function implies that forecasters may not necessarily announce their true inflation expectations; instead their announcement is a combination of their true inflation expectations and the previous period's two-step ahead inflation announcement.

To solve for $\pi_{jt-1}^a(t+1)$, we can setup forecaster j 's loss function for the two-step ahead inflation announcement for time $t+1$:

$$\pi_{jt-1}^a(t+1) \underset{Min}{E_{jt-1}} \{ [\pi(t+1) - \pi_{jt-1}^a(t+1)]^2 + \lambda_j [\pi_{jt-1}^a(t+1) - \pi_{jt-2}^a(t+1)]^2 + \lambda_j [\pi_{jt}^a(t+1) - \pi_{jt-1}^a(t+1)]^2 \} \quad (4)$$

$$s.t. \quad E_{jt-1} \pi_{jt}^a(t+1) = \frac{1}{1+\lambda_j} E_{jt-1} \pi(t+1) + \frac{\lambda_j}{1+\lambda_j} \pi_{jt-1}^a(t+1), \quad (5)$$

where $E_{jt-1} \pi_{jt}^a(t+1)$ is forecaster j 's expected one-step ahead inflation announcement.

The two-step 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-step ahead forecast, but also with future revisions that will be made to the two-step ahead forecast. All revisions are made one period apart and therefore assigned the same weight, λ_j .

²Shortly, I will introduce a reparameterization such that any λ_j will return a value between $[0, 1]$.

In general, the problem of determining the optimal inflation announcement becomes more complicated at larger forecast horizons.

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

$$\pi_{jt-1}^a(t+1) = \frac{1 + \frac{\lambda_j}{(1+\lambda_j)^2}}{1 + \lambda_j + \frac{\lambda_j}{(1+\lambda_j)^2}} E_{jt-1} \pi(t+1) + \frac{\lambda_j}{1 + \lambda_j + \frac{\lambda_j}{(1+\lambda_j)^2}} \pi_{jt-2}^a(t+1). \quad (6)$$


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 -step ahead forecast announcement of $\pi(t+1)$ ought to be a weighted average of true inflation expectations and the previous $h+1$ step ahead forecast:

$$\begin{aligned} \pi_{jt+1-h}^a(t+1) &= \omega(\lambda_j, h) E_{jt+1-h} \pi(t+1) + (1 - \omega(\lambda_j, h)) \pi_{jt-h}^a(t+1), \\ h &= 1, \dots, H \quad \text{and} \quad j = 1, \dots, J \end{aligned} \quad (7)$$

where $\omega(\lambda_j, h) \in [0, 1]$ is a function of λ_j and h . For example, $\omega(\lambda_j, 1)$ is given by $\frac{1}{1+\lambda_j}$ as in (3) and $\omega(\lambda_j, 2)$ is given by $\frac{1 + \frac{\lambda_j}{(1+\lambda_j)^2}}{1 + \lambda_j + \frac{\lambda_j}{(1+\lambda_j)^2}}$ as in (6). 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 $\omega(\lambda_j, h)$ close to one, would mean that the announced forecast is close to the forecaster's true inflation 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. Figure 1 depicts the different values of λ_j along with the corresponding weights on past inflation announcements $(1 - \omega(\lambda_j, h))$ for $h = 1, 2, 3$. Note that the very first forecast that is made would not have any conservatism (since there is no previous announcement that forecasters would be revising), so conservatism 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 (as figure 1 suggests). In other words, as the forecast horizon, h , gets larger, past announcements would play less of a role.

It is also possible to think of conservatism in a multi-period model, though unlike the static case, it cannot be estimated without further assumptions on forecast announcements. Section 8 derives the optimal forecasts in this case, and shows that they have a similar structure as in the static case with different weights for horizons beyond $h = 1$.

Figure 1: Weights on Past Announcements



Notes: Each curve represents the weights on past announcements as λ increases for a given forecast horizon. To see the corresponding weights for higher values of λ , please refer to the static model section of Table 2 in Section 8.

3 Alternative Interpretations of Conservatism

Inflation announcements that differ from the full information, rational expectation outcome could be a result of a deviation from full information or a deviation from rational expectations. Coibion and Gorodnichenko (2012), henceforth CG, explore the former set of arguments while this paper aims to explore the latter set of arguments.

CG find that sticky information and imperfect information models both yield a similar result as the first-order conditions from the conservative loss function: the current average forecast is a weighted average of the previous period's forecast and the true rational expectation of inflation for the current period. They argue that reputational concerns are not the source of conservatism, instead it is a manifestation of informational rigidities faced by professional forecasters. To disentangle these alternative interpretations, both the data and the analysis used in this paper differs from CG in several ways. For example, CG suggest also that reputational concerns are not at play as they do not find that SPF forecasts are worse than market driven expectations. This result, however, is based on a mean forecast across SPF forecasters. Zarnowitz and Braun (1993) as well as Bennett, *et al.* (1999) document that consensus forecasts tend to be more accurate than virtually all individual forecasts. In fact, Manski (2010) shows that this empirical regularity follows from Jensen's inequality. Hence, it is possible that reputational concerns are still present among individual forecasters, which may not be detected when analysis is performed with a consensus forecast. The focus here will be on analyzing a small group of forecasters to detect whether conservatism

is present at the individual level.³

Similarly, CG suggest that academic forecasters are more likely to have reputational concerns than industry forecasters, whose main focus is on profit-generating activities. They find more rigidity in forecasts from private industry forecasters than academics and interpret this evidence as more suggestive of informational rigidities than reputational concerns. On the other hand, Bennett, *et al.* (1999) propose and estimate a model in which part of forecasters' objective is to gain publicity for their firms and find evidence of a substantial strategic component in professional forecasting. Hence, it is not entirely clear whether academic forecasters would focus less on minimizing forecast errors and have greater reputational concerns than private industry forecasters, as both would likely defend their forecast. The end users of these forecasts, whether for policy-making purposes or business or financial decisions, would likely have a preference toward stability in forecasts, particularly for longer term decision making, therefore reputational concerns cannot be ruled out. To address this, I allow for a mix of both financial and nonfinancial industry forecasters at the individual level.

The work of CG is complementary here as they do not find evidence of reputational concerns or strategic incentives in the consensus forecast, a result which they suggest may not hold at the individual level. By focusing the analysis here on individual forecasters, the objective of this paper is to highlight the possibility that reputational concerns are likely to be a bigger part of the picture.

Batchelor and Dua (1992) use an alternative measure of conservatism. Their measure is based on the assumption that each forecaster prepares his or her forecast as a weighted average of his or her previous month's forecast, the previous month's consensus forecast, and a forecast based on new economic news arriving in the current month. Unlike Batchelor and Dua (1992), the setup in this paper starts with a loss function for each forecaster which implies an optimal forecast as determined by the first order conditions.

4 Estimation Strategy

4.1 The Issue of Weak Instruments

A natural question to consider is the value of λ_j in practice, along with its variation across forecasters. Estimating (7) using OLS would not be feasible since $E_{jt+1-h}\pi(t+1)$ is not

³In related work, Dovern, *et. al* (2014) examine the presence of informational rigidities in a large panel of individual forecasters from 36 advanced and emerging markets. They find that, for real GDP growth forecasts, the estimates of the degree of informational rigidity in mean forecasts are substantially higher than in individual forecasts.

observed. Instrumental variable (IV) estimation would be an alternative if a valid instrument could be obtained. If forecasters have rational expectations, then actual inflation data could be used. In this case, $\pi(t+1) = E_{jt+1-h}\pi(t+1) + \varepsilon(j, t+1, h)$ and we could write (7) as:

$$\begin{aligned} & \pi_{jt+1-h}^a(t+1) \\ = & \omega(\lambda_j, h)[\pi(t+1) - \varepsilon(j, t+1, h)] + (1 - \omega(\lambda_j, h))\pi_{jt-h}^a(t+1) \\ = & \omega(\lambda_j, h)\pi(t+1) + (1 - \omega(\lambda_j, h))\pi_{jt-h}^a(t+1) - \omega(\lambda_j, h)\varepsilon(j, t+1, h), \end{aligned} \quad (8)$$

where $\varepsilon(j, t+1, h)$ is forecaster j 's h -period ahead forecast error. Under rational expectations, directly estimating (8) would be problematic as the forecast error $\varepsilon(j, t+1, h)$, would be correlated with actual inflation. Ang, Bekaert and Wei (2007) find that the *Survey of Professional Forecasters'* (SPF) inflation expectations data is one of the leading forecasts of inflation. By definition, however, we cannot use this as an instrument since the SPF data will be used to capture inflation announcements.

Consider two possible instruments for $\pi(t+1)$: 3-month U.S. Treasury Bill interest rates and the U.S. capacity utilization rate in manufacturing. The former instrument is closely linked to future inflation rates, while the choice of the latter instrument was motivated by Stock and Watson (1999), who find that when the rate of capacity utilization in manufacturing is used as a measure of real aggregate activity in a generalized Phillips curve, it is particularly effective in producing more accurate inflation forecasts than the conventional unemployment rate Phillips curve.⁴

I test the relevance of these instruments by examining several test statistics. First, consider the first-stage regression results obtained by a regression of the endogenous regressor, $\pi(t+1)$, on the full set of instruments. The Shea (1997) R^2 statistic is 0.10 which suggests that these instruments lack sufficient relevance to explain $\pi(t+1)$. Furthermore, Staiger and Stock (1997) formalized the definition of weak instruments suggesting that if the F -statistic on the excluded instruments in the first-stage is less than 10, it is a cause for concern. In our case this F -statistic is 5.83. The R^2 and F -statistics for the cases in which $h = 2, 3$ are even smaller. 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 $\pi(t+1)$ in a similar context has been documented in past research (for examples, see Stock and Watson (1999) and Hansen, Lunde and Nason (2011)).

⁴Note that past forecast errors may also be used as instruments. This would yield a specification similar to what is used in the following section.

4.2 The Anderson-Rubin Approach

I opt to take the Anderson-Rubin (AR) approach which allows a test of different values of λ_j (and hence $1 - \omega(\lambda_j, h)$). 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.

Start by taking the value of future inflation, $\pi(t + 1)$, to the left-hand side of (8) and then add a relevant auxiliary variable on the right-hand side:

$$\begin{aligned} \pi_{jt+1-h}^a(t + 1) - \omega(\lambda_j, h)\pi(t + 1) &= [1 - \omega(\lambda_j, h) - \alpha]\pi_{jt-h}^a(t + 1) \\ &\quad + \beta u(t + 1 - h) - \omega(\lambda_j, h)\varepsilon(j, t + 1, h), \end{aligned} \quad (9)$$

where $u(t + 1 - h)$ is a relevant auxiliary variable, such as the previous period's median inflation forecast.⁵ Then, we can choose a particular value of λ_j , labelled λ_j^0 , from a grid of possible values for λ_j , which would then yield a corresponding value $\omega(\lambda_j^0, h)$. While we cannot use this regression to estimate a value, it can be used to test any value of this weight on expected future inflation. To test whether $\lambda_j = \lambda_j^0$, the AR procedure involves an F -test of the null hypothesis that $\alpha = \beta = 0$, so the auxiliary variable is not statistically significant and the weight on the past announcement is given by $(1 - \omega(\lambda_j^0, h))$. The idea is that at the true value of λ_j , there should be no further role for the auxiliary variable, $u(t + 1 - h)$. The AR procedure yields an F -statistic known as the Anderson-Rubin (AR) statistic which follows a 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.

5 Forecast Data

The announced inflation forecast data comes from the *Survey of Professional Forecasters* (SPF), the oldest quarterly survey of macroeconomic forecasts in the US. The survey began in

⁵In the following sections, this will be used as an auxiliary variable along with short-term interest rates and the previous period's change in real oil prices.

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 which 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. Quarterly forecasts are annualized quarter-over-quarter percent changes. The particular quarterly forecasts series used are 1-quarter ahead through 5-quarter ahead forecasts of CPI inflation ("CPI2" to "CPI6"). The data runs from 1981:Q3 to 2010:Q1.

As an auxiliary variable, the corresponding median CPI inflation forecast from the SPF, denoted $u_1(t + 1 - h)$, is used as a measure of the consensus forecast. Since this was the consensus measure in the previous period for the same target date that forecasters will make announcements for in the current period, this consensus measure is known to forecasters at the time of their announcement. The 3-month U.S. Treasury Bill interest rate, denoted $u_2(t + 1 - h)$, will also be used as an auxiliary variable. This monthly series is averaged into a quarterly series to be consistent with the SPF data. The auxiliary variable denoted $u_3(t + 1 - h)$, is a lag of the log change in real WTI oil prices and is included to account for the correlation between the error term and current period variables.

The full sample contains observations from forecasters who have submitted enough forecasts to generate at least 6 revisions. This yields a panel of 124 forecasters. Estimation will be performed using the full sample of forecasters, but will be focused on the 10 forecasters in the sample that submitted at least 50 inflation announcements. Of these 10 professional forecasters, 4 are financial service providers, 5 are nonfinancial service providers and the industry of the 1 forecaster is unknown.

Data for actual inflation comes from FRED[®] (Federal Reserve Economic Data), a database of over 14,000 US economic time series maintained by the Federal Reserve Bank of St. Louis. The series used is "cpiaucsl", the monthly consumer price index for all items (seasonally adjusted). Monthly data running from 1981 to 2010 is averaged to quarterly data and is then converted to an annualized quarterly rate to be consistent with the expected inflation series.

The advantage of using CPI inflation is that the monthly US CPI releases by the Bureau of Labour Statistics are not subject to revisions. In contrast, GDP deflator data is usually subject to systematic revisions which could result in forecasts that differ systematically from announced values solely due to data revisions. To avoid this issue, all the results presented in this paper use CPI inflation.

6 Estimation Results

6.1 Forecaster-Specific Estimation

In this section, we will perform the AR test on (9) for various values of λ_j ($1 - \omega(\lambda_j, h)$) and report the values of λ_j at which the null $H'_0 : \alpha = \beta = 0$ is closest to holding. Recall, this involves minimizing the resulting F -statistic for a grid of possible values of λ_j . When $\lambda_j^0 = 0$ ($\omega(\lambda_j^0, h) = 1$), this is a test of the *no-sluggishness hypothesis* in which forecasters' sole objective is to announce accurate forecasts. Forecaster-specific estimates reveal that from our panel of 124 forecasters, 14 forecasters reject the no-sluggishness hypothesis at the 5% significance level when $h = 1$. When $h = 2$, and $h = 3$, 22 and 23 forecasters reject the no-sluggishness hypothesis at the 5% level, respectively. While these rejection rates for the no-sluggishness hypothesis do not seem high, it is important to highlight a limitation of the forecaster-specific estimation here: since most of the results follow from forecasters with small sample sizes, the power of our tests is quite low which implies a high probability of making Type II error (mistakenly failing to reject a false null hypothesis). Therefore, if we consider only the forecasters with relatively larger sample sizes who have submitted at least 50 inflation announcements, of these 10 forecasters, 2 reject the no-sluggishness hypothesis at the 5% significance level when $h = 1$, and 5 reject the no-sluggishness hypothesis at the 5% significance level both when $h = 2$ and when $h = 3$. Hence, it is possible that sluggishness in inflation announcements may be more prevalent than what is suggested when forecasters with small sample sizes are considered.

We continue to a more detailed analysis of these 10 forecasters who have at least 50 observations. These forecasters are labelled 1 through 10 in the order their ID number is listed in the survey, hence forecasters with the lowest ID numbers make up the relatively earlier participants in the survey. Recall that we wish to test the null hypothesis that $\alpha = \beta = 0$. This is a test of whether the conservatism model is a good guide or not.⁶ Here, the null hypothesis is tested for a range of values of λ_j between 0 and 10.

The estimation results test the null hypothesis that $\lambda_j = \lambda_j^0$ by estimating (9) and constructing the AR F -statistic for $H'_0 : \alpha = \beta_1 = \beta_2 = \beta_3 = 0$, for each forecaster j , where β_i is the parameter on indicator u_i , $i = 1, 2, 3$. The conservatism model is not rejected and results suggest the prevalence of substantial sluggishness in inflation announcements.

⁶The first-order conditions that follow from the conservative loss function for each inflation announcement are not altered by the sequence of forecasts reported by forecasters in the SPF. Specifically, in any given period, forecasters report a sequence of inflation forecasts: $\{\pi_{jt}^a(t+1), \pi_{jt}^a(t+2), \pi_{jt}^a(t+3), \pi_{jt}^a(t+4), \pi_{jt}^a(t+5)\}$, each of which forecast inflation for a *different* target date. Conversely, conservatism is measured by the reluctance to revise an inflation forecast for a *particular* target date, for example between $\pi_{jt}^a(t+1)$ and $\pi_{jt-1}^a(t+1)$ for $\pi(t+1)$.

This result holds both with and without the inclusion of auxiliary variables. The statistics that will be presented below consider the case in which the auxiliary variables (consensus inflation forecast, $u_1(t + 1 - h)$, interest rates, $u_2(t + 1 - h)$, and lag of log change in real oil prices $u_3(t + 1 - h)$) are included in estimation. The top panel in figure 2 represents the AR-statistics and their corresponding p -values are on the bottom panel for forecaster 1. The corresponding figures for the remaining 9 forecasters are presented in Appendix A.

The p -values reach their maximum when AR statistics reach their minimum. At the minimum AR statistic for each forecast horizon, we cannot reject the null that $\alpha = \beta_1 = \beta_2 = \beta_3 = 0$ at the corresponding significance level indicated by the p -value. For all forecasters, this joint test cannot be rejected at the 1% significance level at all forecast horizons.

The best fitting value of λ_j varies across forecasters, but there are certain patterns in the results. First, with the exception of forecaster 7, the minimum AR statistics highlight a higher degree of conservatism when $h = 3$, relative to the no-sluggishness hypothesis. There is also a high degree of conservatism when $h = 2$ with $\lambda_j \geq 0.7$ for all 10 forecasters. Excluding forecaster 7, there is more conservatism when $h = 2$ than when $h = 1$. The best fitting degree of conservatism when $h = 1$ is primarily values of $\lambda_j < 1$. These findings suggest that forecasters' priorities differ depending on the forecast horizon: in the shorter term, accuracy in their forecasts takes precedence, and in the medium term, their focus is on maintaining continuity in their forecasts.

Table 1 summarizes all of $(1 - \omega(\lambda_j, h))$ weights associated with the minimum- F λ values from figure 2 and Appendix A for each forecaster at each forecast horizon. Recall, this is the weight that forecaster j would place on their previous forecast announcement at horizon h . The table also indicates whether the forecaster is a financial (F) or nonfinancial (NF) industry forecaster or the forecaster's industry is unknown (U). Notice, there does not appear to be a distinct degree of conservatism depending on whether the forecaster is a financial or nonfinancial industry forecaster. Also note an interesting finding when the model is tested using the series of median forecasts: this series does not represent a sequence of forecasts made by a particular forecaster, and therefore we would not expect it to contain a high degree of conservatism. This is confirmed in Table 1 as $(1 - \omega(\lambda, h)) = 0.3$ when $h = 1$, and equals 0 when $h = 2, 3$. One conclusion that could be drawn from this finding is that to obtain a forecast that is focused on minimizing forecast inaccuracy, one should opt to use the median forecast over that of any individual forecaster.

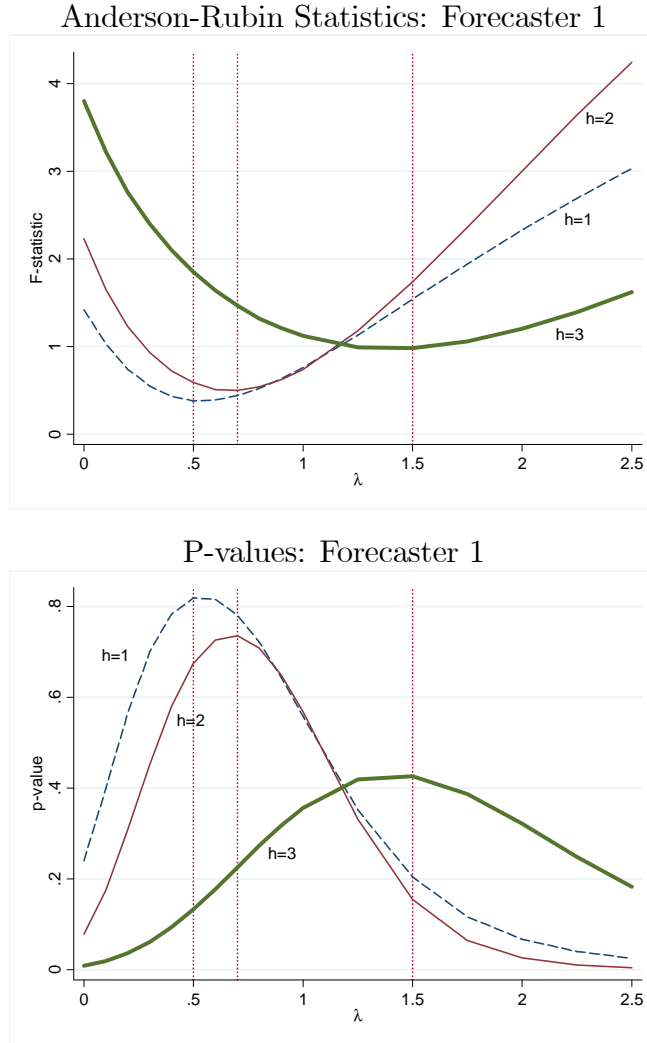
6.2 Pooled Estimation

Pooled estimation further supports the forecaster-specific estimates. The two panels in figure 3 show the AR statistics along with corresponding p -values when the data for the 10 forecasters is pooled, allowing each forecaster its own intercept. This yields a total of 628 observations when $h = 1, 2$ and 626 observations when $h = 3$. Both panels suggest that the degree of sluggishness that best fits the conservatism model is clearly distinct when $h = 1$ versus when $h = 2, 3$. Without the inclusion of auxiliary variables, the minimum AR statistic occurs when $\hat{\lambda} = 0.4$ when $h = 1$, $\hat{\lambda} = 1.50$ when $h = 2$, and $\hat{\lambda} = 2.25$ when $h = 3$. The corresponding p -values indicate that these minimum AR statistics do not reject the null hypothesis at the 5% significance level. When auxiliary variables are included, the minimum AR statistics under the null hypothesis $\alpha = \beta_1 = \beta_2 = \beta_3 = 0$ follow the same pattern. In the context of the conservatism model, achieving forecast accuracy is important in the short term (when $h = 1$) but conservatism in inflation announcements prevails at longer forecast horizons (when $h = 2, 3$).

The bottom of table 1 also indicates the minimum- F ($1 - \omega(\hat{\lambda}, h)$) values when data is pooled, along with a 95% confidence interval which gives the range of corresponding weights for which the AR-statistics in figure 3 fall below the α -percent critical value of the F distribution (or equivalently, the p -values lie above α).

Figure 2: The Anderson-Rubin Test

$$\begin{aligned} \pi_{jt+1-h}^a(t+1) - \omega(\lambda_j, h)\pi(t+1) &= \delta_0 + (1 - \omega(\lambda_j, h) - \alpha)\pi_{jt-h}^a(t+1) \\ &+ \beta_1 u_1(t+1-h) + \beta_2 u_2(t+1-h) + \beta_3 u_3(t+1-h) \\ &+ \xi(j, t+1, h) \end{aligned}$$




Notes: The top figure reports F-statistics for the joint test that $\alpha = \beta_1 = \beta_2 = \beta_3 = 0$. The minimum points on these curves indicate the value of λ at which the conservatism model is closest to holding for each forecast horizon. Corresponding p-values are depicted in the bottom figure. The values of $1 - \omega(\lambda, h)$ that are associated with the minimum AR statistics (or conversely, maximum p-values) are reported in Table 1 for all forecasters.


Figure 3: The Anderson-Rubin Test

$$\pi_{t+1-h}^a(t+1) - \omega(\lambda, h)\pi(t+1) = \delta_0 + \sum_{j=2}^{10} \delta_j + (1 - \omega(\lambda, h) - \alpha)\pi_{jt-h}^a(t+1) + \xi(t+1, h)$$

Anderson-Rubin Statistics: Pooled Estimation



P-values: Pooled Estimation



Notes: The top figure reports F-statistics for the joint test that $\alpha = 0$. The minimum points on these curves indicate the value of λ at which the conservatism model is closest to holding for each forecast horizon. Corresponding p-values are depicted in the bottom figure. The associated $1 - \omega(\lambda, h)$ values are reported in Table 1 along with a 95% confidence interval.

Table 1: Summary of Minimum- F values of $(1 - \omega(\lambda_j, h))$:

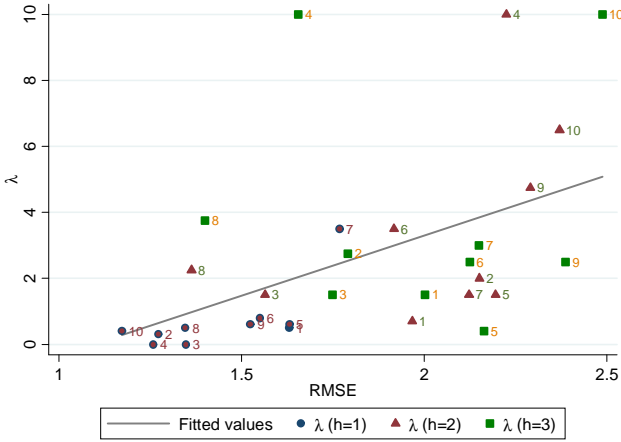
	$h = 1$	$h = 2$	$h = 3$
Forecaster 1 (NF)	0.33 (0.82)	0.36 (0.74)	0.52 (0.43)
Forecaster 2 (NF)	0.23 (0.02)	0.62 (0.04)	0.67 (0.08)
Forecaster 3 (F)	0 (0.58)	0.55 (0.38)	0.52 (0.13)
Forecaster 4 (F)	0 (0.02)	0.90 [†] (0.02)	0.90 [†] (0.02)
Forecaster 5 (F)	0.38 (0.41)	0.55 (0.39)	0.24 (0.15)
Forecaster 6 (F)	0.44 (0.67)	0.75 (0.02)	0.65 (0.09)
Forecaster 7 (F)	0.78 (0.56)	0.55 (0.64)	0.69 (0.77)
Forecaster 8 (NF)	0.33 (0.80)	0.65 (0.19)	0.74 (0.07)
Forecaster 9 (NF)	0.47 (0.49)	0.81 (0.60)	0.65 (0.11)
Forecaster 10 (U)	0.29 (0.76)	0.85 (0.23)	0.90 [†] (0.24)
Median Forecast	0.33 (0.01)	0 (0.01)	0 (0.02)
Pooled	0.29	0.55	0.62
95% c.i.	(0.12,0.39)	(0.41,0.63)	(0.50,0.70)

Notes: (1) p -values in parentheses. (2) [†] indicates that values of $(1 - \omega(\lambda_j, h))$ where λ_j may be greater than 10. (3) The median forecast used here is the same as what is reported by the SPF. (4) Acronyms used are: F: financial industry forecaster, NF: nonfinancial industry forecaster, U: industry unknown.

6.3 Forecast Accuracy

If conservatism is indeed present, then we might expect that the more conservative forecasters would also have less forecast accuracy on average, as these forecasters are more concerned with maintaining continuity in their forecasts. Figure 4 illustrates a scatter plot of conservatism and forecast accuracy as measured by the average RMSE, for each forecaster at different forecast horizons.

Figure 4: Conservatism and Forecast Accuracy



Note: The number to the left of each point in this figure indicates the forecaster.

There is indeed a positive relationship between conservatism in the announced forecasts and the RMSE. Notice that, for example, forecasters 2, 4 and 10 who have the lowest levels of conservatism for the current horizon ($h = 1$) from the group of forecasters also have the lowest levels of RMSEs. Forecaster 7, on the other hand, has the highest degree of conservatism when $h = 1$ and also the largest RMSE for that horizon. Forecasters 4, 9 and 10 have the largest degree of conservatism when $h = 2$ and they also have the largest RMSEs for that forecast horizon. Forecaster 10 has the highest degree of conservatism at $h = 3$, and also the highest RMSE in that forecast horizon.⁷

⁷An interesting case is forecaster 8 who has higher than the median levels of conservatism, yet still has lower than the median RMSEs. This forecaster is able to strike a balance between making accurate forecasts and maintaining their continuity.

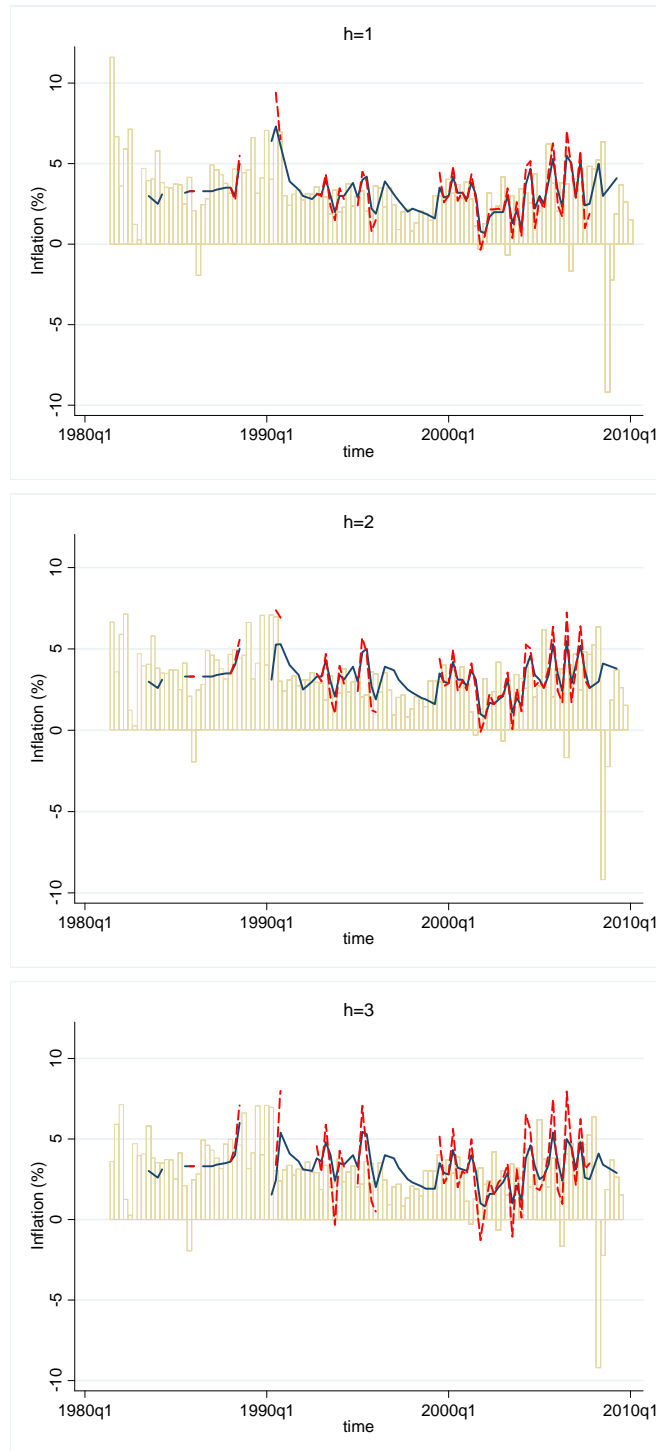
7 Measuring True Expectations

Notice that from the first-order conditions derived in section 2, once we have an estimate for the degree of conservatism, $\hat{\lambda}_j$, we can solve for forecaster j 's "true" or implicit inflation expectations since we have data on inflation announcements. That is, rearranging (7) and denoting true or implicit inflation expectations by $\hat{E}_{jt+1-h}\pi(t+1)$ yields:

$$\begin{aligned} \hat{E}_{jt+1-h}\pi(t+1) &= [\pi_{jt+1-h}^a(t+1) - (1 - \omega(\hat{\lambda}_j, h))\pi_{jt-h}^a(t+1)]/\omega(\hat{\lambda}_j, h) & (10) \\ h &= 1, 2, 3 \quad \text{and} \quad j = 1, \dots, 10 \end{aligned}$$

Figure 5 shows a comparison between inflation announcements reported in the SPF (captured by the blue solid line), and the corresponding implied inflation expectation series (captured by the red dashed line) for forecaster 1, and Appendix B shows similar graphs for each of the remaining 9 forecasters. For reference, the height of the bar shows the corresponding level of actual realized inflation throughout the sample period. Note that in most cases, true inflation expectations is a more volatile series than inflation announcements, as the model would suggest. In some cases, implicit inflation is very close to its corresponding announced value, and in some cases it is not. In some situations, notice the implicit inflation series is closer to the actual inflation series than announced values, though in some cases it is not. This implicit inflation expectations series does, however, come with some uncertainty as it is the calculated series without a confidence interval. For cases such as forecaster 4 or forecaster 10, the implicit inflation series shows much larger amounts of volatility which could also be an indication that these forecasters face a more complex loss function than the one presented here. In this paper, the degree of conservatism it is determined under the assumption that forecasters face symmetric, linear-quadratic loss functions. While a key starting point, for some forecasters, conservatism under an asymmetric loss function may be a more accurate representation of their forecasting problem and remains an important avenue for future research. Nonetheless, the implicit inflation series presented here does provide an important insight into better understanding the formation of inflation expectations: announced inflation forecasts may be slow to change relative to a forecaster's true inflation expectations.

Figure 5: Announced (blue) & Implicit (red) Forecasts: Forecaster 1



Notes: The implicit or "true" inflation expectations series (red, dashed) is calculated using (10). Gaps in this series occur if a forecaster does not submit inflation announcements for consecutive periods. The announced inflation forecast series (blue, solid) correspond to what forecasters report to the SPF. The bars indicate realized values of inflation corresponding to the forecasts.

8 Extension: The Multi-Period Case

Consider a multi-period version of the static conservative loss function presented in section 2:

$$\text{Min} \sum_{i=1}^h \gamma^{i-1} E_{jt+1-h} \{ [\pi(t+1) - \pi_{jt+i-h}^a(t+1)]^2 + \lambda_j [\pi_{jt+i-h}^a(t+1) - \pi_{jt-1+i-h}^a(t+1)]^2 \},$$

where γ is the discount factor, and as in our analysis above, we set $h = 3$. Now, forecasters will choose a sequence of forecasts based on their current information set (in this case given by time $t - 2$ information). Solving backwards, first-order conditions for this problem yield:

$$\begin{aligned} E_{jt-2} \pi_{jt}^a(t+1) &= \left(\frac{1}{1 + \lambda_j} \right) E_{jt-2} \pi(t+1) + \left(\frac{\lambda_j}{1 + \lambda_j} \right) E_{jt-2} \pi_{jt-1}^a(t+1) \\ E_{jt-2} \pi_{jt-1}^a(t+1) &= \left(\frac{1 + \lambda_j + \gamma \lambda_j}{1 + \lambda_j + \gamma \lambda_j + \lambda_j(1 + \lambda_j)} \right) E_{jt-2} \pi(t+1) \\ &\quad + \left(\frac{\lambda_j(1 + \lambda_j)}{1 + \lambda_j + \gamma \lambda_j + \lambda_j(1 + \lambda_j)} \right) \pi_{jt-2}^a(t+1) \\ \pi_{jt-2}^a(t+1) &= \left(\frac{(1 + \gamma \lambda_j)(1 + \lambda_j + \gamma \lambda_j) + \lambda_j(1 + \lambda_j)}{(1 + \lambda_j + \gamma \lambda_j)^2 + \lambda_j(1 + \lambda_j)^2} \right) E_{jt-2} \pi(t+1) \\ &\quad + \left(\frac{\lambda_j(1 + \lambda_j + \gamma \lambda_j) + \lambda_j^2(1 + \lambda_j)}{(1 + \lambda_j + \gamma \lambda_j)^2 + \lambda_j(1 + \lambda_j)^2} \right) \pi_{jt-3}^a(t+1). \end{aligned}$$

The optimal forecasts in the multi-period problem have a similar structure as in the static case except that the optimal forecasts are now based on forecasters' expected future announcements for which data is typically unavailable. The weighted averages for the current three-step ahead ($\pi_{jt-2}^a(t+1)$) and expected two-step ahead announcement ($E_{jt-2} \pi_{jt-1}^a(t+1)$) also differ. To further understand the difference between the multi-period and static versions of the problem, table 2 below presents the weights on the past forecast announcement (ie. $(1 - \omega(\lambda_j, h))$) corresponding to both problems for different values of λ_j and at all three forecast horizons.

Notice that the weight on past announcements remain the same between the multi-period and the static models when $h = 1$, since it is an (expected) one-step ahead forecast so the actual data would be realized in the following period. For $h = 2$ and $h = 3$, the weight on past announcements in the multi-period case are smaller than those in the static case. This occurs since, in the dynamic case, forecasters are concerned about both their future forecast errors and future forecast revisions. Conversely, if we allow the discount factor γ to be very small as shown in the bottom panel of table 2, the difference in the weights between

the multi-period and static cases also become very small.

Table 2: A Comparison Between Static and Multi-Period Conservatism

$\gamma = 0.95$

λ_j	Static Model			Multi-Period Model		
	$h = 1$	$h = 2$	$h = 3$	$h = 1$	$h = 2$	$h = 3$
0	0	0	0	0	0	0
2	0.67	0.62	0.59	0.67	0.55	0.52
4	0.80	0.78	0.75	0.80	0.69	0.65
6	0.86	0.84	0.83	0.86	0.77	0.72
8	0.89	0.88	0.87	0.89	0.81	0.77
10	0.91	0.90	0.90	0.91	0.84	0.80

$\gamma = 0.05$

λ_j	Static Model			Multi-Period Model		
	$h = 1$	$h = 2$	$h = 3$	$h = 1$	$h = 2$	$h = 3$
0	0	0	0	0	0	0
2	0.67	0.62	0.59	0.67	0.66	0.66
4	0.80	0.78	0.75	0.80	0.79	0.79
6	0.86	0.84	0.83	0.86	0.85	0.85
8	0.89	0.88	0.87	0.89	0.88	0.88
10	0.91	0.90	0.90	0.91	0.91	0.91

Note: The tables indicates values of $1 - \omega(\lambda_j, h)$ (i.e. the weight on past announcements) when $\gamma = 0.95$ (top panel) and when $\gamma = 0.05$ (bottom panel).

9 Conclusions


Given the influence that inflation expectations have on the level of actual inflation, it is vital to have a deeper understanding of the process of inflation forecasting. This paper consid-

ers the situation in which professional forecasters have rational expectations that may not be reflected in their inflation forecast announcements. In the conservatism model, this discrepancy between true inflation expectations and inflation forecast announcements is fuelled by forecasters' desire to maintain continuity in their forecasts, though it compromises their accuracy.


Using the Anderson-Rubin (1949) approach, the evidence supports the conservatism model, suggesting some conservatism in inflation announcements for the current forecast horizon and a larger degree of conservatism in inflation forecasts for future forecast horizons. This finding is consistent with past literature that has found departures from rationality and biases in inflation forecasts. From the point of view of a policy maker, this could be a sign of increased credibility as forecasters would react less to monetary policy shocks than the actual inflation data would suggest. From the perspective of a forecaster, they may 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. This paper also presents a simple method to undo this effect, making it possible to back out a true or implicit inflation expectations series. Future work will aim to examine conservatism under the assumption of an asymmetric loss function.

Appendix A: Anderson-Rubin Statistics and P-values⁸


Anderson-Rubin Statistics: Forecaster 2




Anderson-Rubin Statistics: Forecaster 3



P-values: Forecaster 2




P-values: Forecaster 3




⁸Please refer to explanatory notes in Figure 2.


Anderson-Rubin Statistics: Forecaster 4




Anderson-Rubin Statistics: Forecaster 5



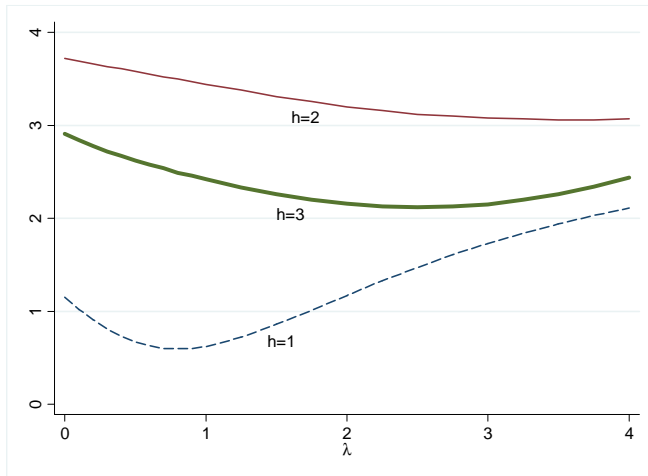
P-values: Forecaster 4




P-values: Forecaster 5




Anderson-Rubin Statistics: Forecaster 6



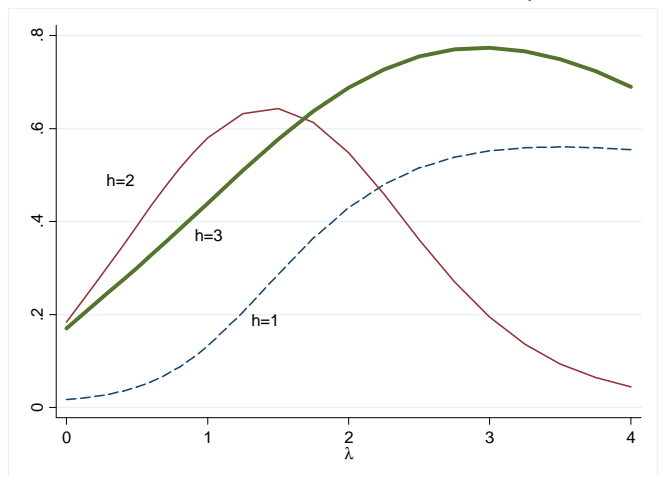
Anderson-Rubin Statistics: Forecaster 7




P-values: Forecaster 6




P-values: Forecaster 7




Anderson-Rubin Statistics: Forecaster 8




Anderson-Rubin Statistics: Forecaster 9




P-values: Forecaster 8




P-values: Forecaster 9



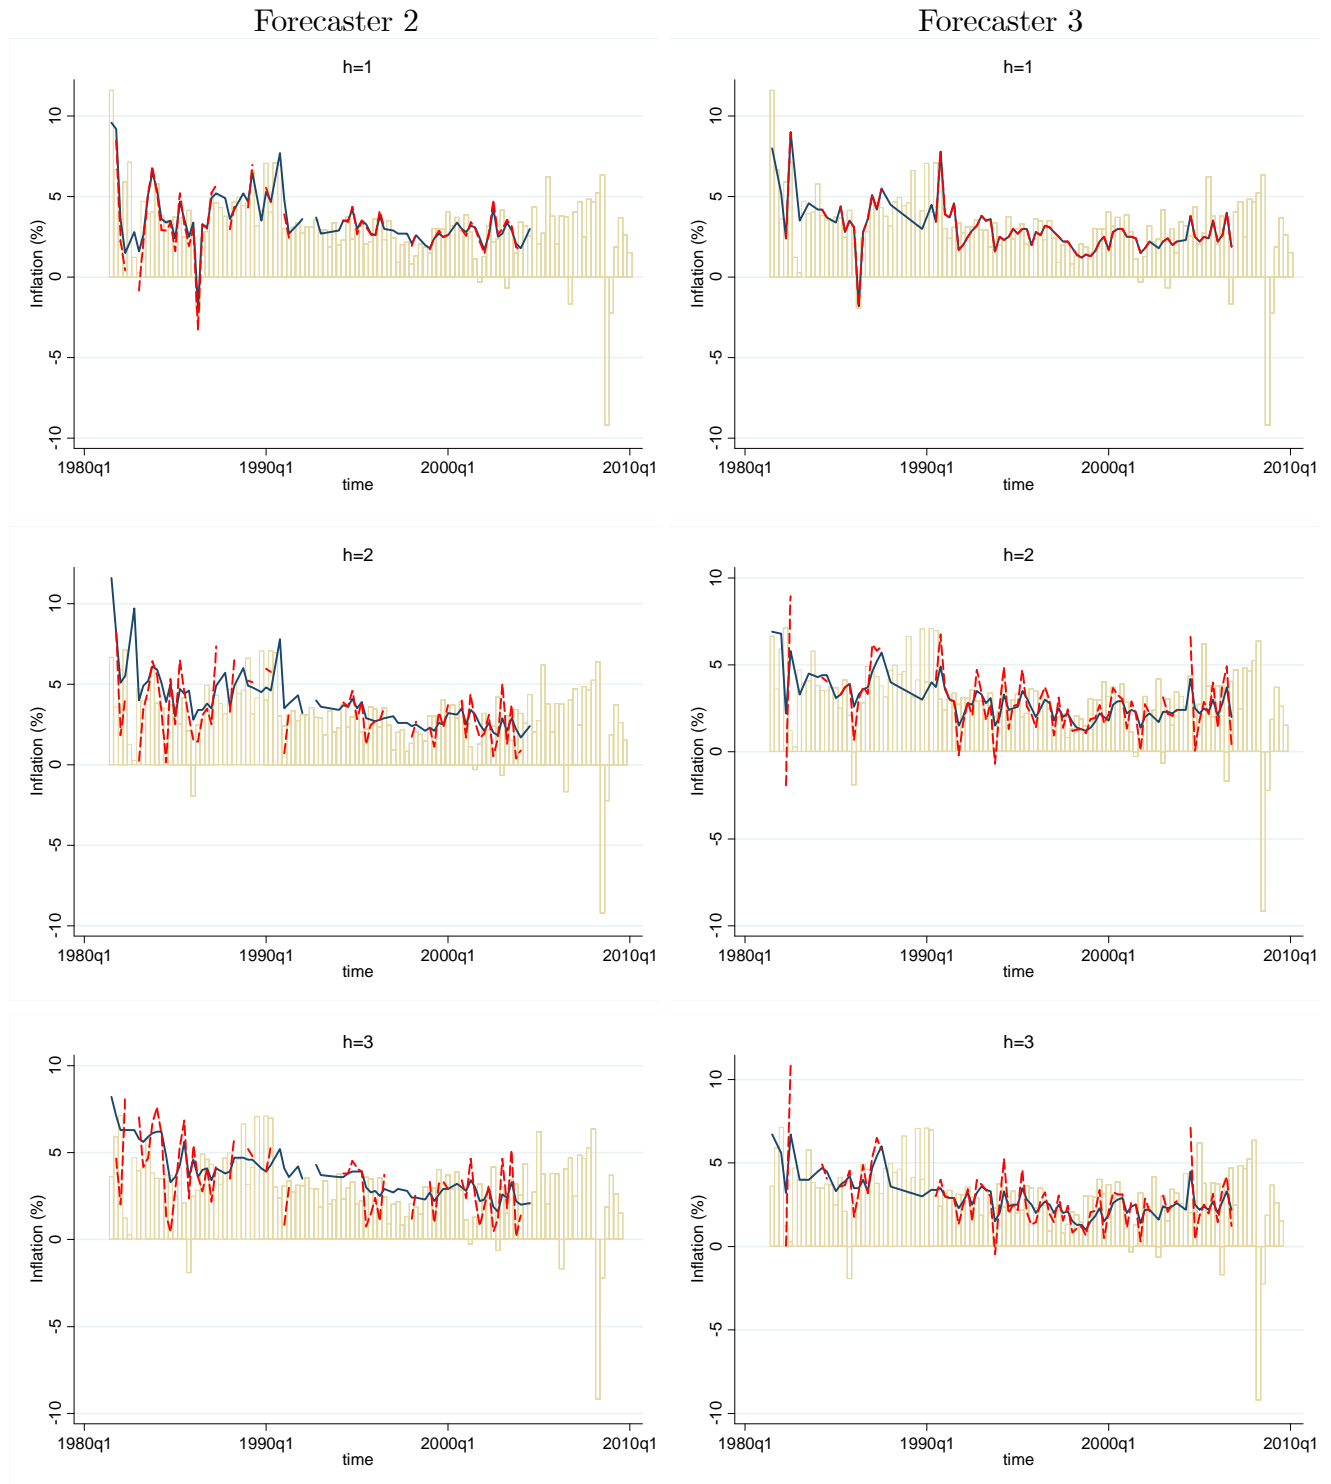
Anderson-Rubin Statistics: Forecaster 10



P-values: Forecaster 10




Appendix B: Announced (blue) & Implicit (red) Forecasts⁹




⁹Please refer to explanatory notes in Figure 5.


Forecaster 4




Forecaster 5




h=2



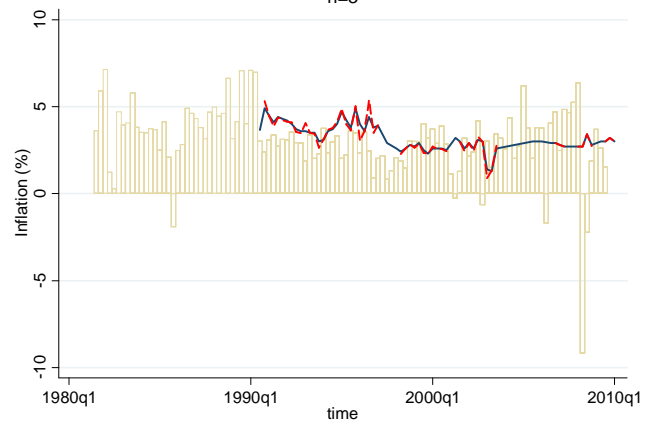
h=2



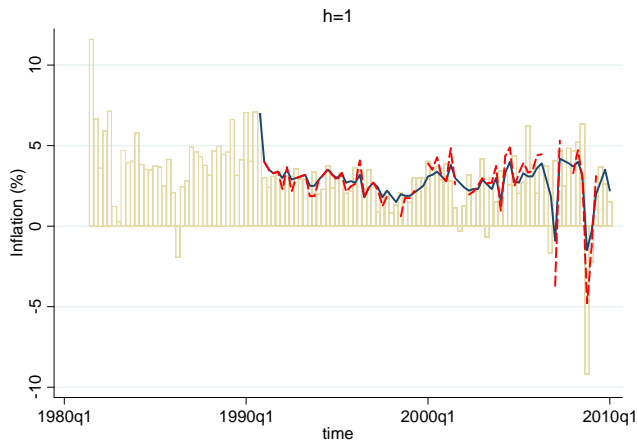
h=3



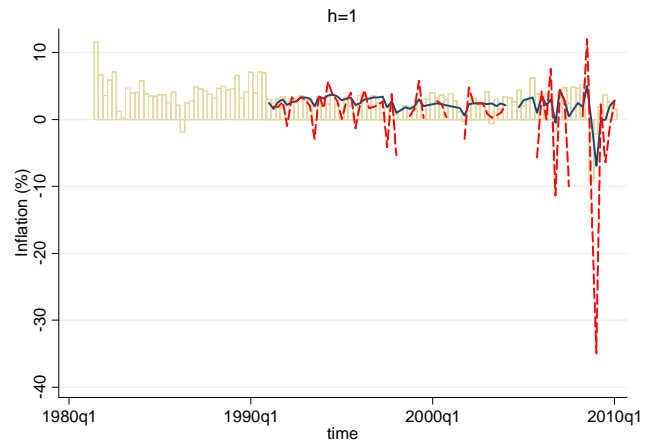
h=3




Forecaster 6




Forecaster 7




h=2




h=2




h=3




h=3




Forecaster 8




Forecaster 9




h=2




h=2






h=3



h=3



Forecaster 10



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] Batchelor, Roy. 2007. Bias in Macroeconomic Forecasts. *International Journal of Forecasting*, 23: 189–203
- [4] Batchelor, Roy and Pami Dua. 1992. Conservatism and Consensus-seeking Among Economic Forecasters. *Journal of Forecasting*, 11: 169-181.
- [5] Bennett, Paul, Geoum, In Sun and David Laster. 1999. Rational Bias in Macroeconomic Forecasts. *Quarterly Journal of Economics*, 114(1): 293-318.
- [6] Berkmen, Pelin, Reza Moghadam, and Ranjit Teja. 2014. Euro Area - "Deflation" Versus "Lowflation." iMFdirect - The IMF Blog. March 4.
- [7] Capistran, Carlos and Allan Timmermann. 2009. Disagreement and Biases in Inflation Expectations. *Journal of Money, Credit and Banking*, 41(2-3): 365-396.
- [8] Coibion, Olivier and Yuriy Gorodnichenko. 2012. Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts. *NBER Working Paper* 16537.
- [9] Doovern, Jonas, Ulrich Fritsche, Prakash Loungani, and Natalia Tamirisa. 2014. Informational Rigidities: Comparing Average and Individual Forecasts For a Large International Panel. RPF Working Paper No. 2014-001.
- [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, Maral Kichian. 2006. Inflation Dynamics and the New Keynesian Phillips Curve: An Identification-Robust Econometric Analysis. *Journal of Economic Dynamics and Control*, 30: 1707-27.
- [12] Gali, Jordi and Mark Gertler. 1999. Inflation Dynamics: A Structural Econometric Analysis. *Journal of Monetary Economics*, 44: 195-222.

- [13] Hansen, Lars Peter R, Asger Lunde, and James M. Nason. 2011. Model Confidence Sets for Forecasting Models. *Econometrica*, 79: 453-497
- [14] Keane, Michael F. and David E. Runkle. 1990. Testing the Rationality of Price Forecasts: New Evidence from Panel Data. *American Economic Review*, 80: 714-735.
- [15] Kirchgassner, Gebhard and Ulrich K. Muller. 2006. Are Forecasters Reluctant to Revise Their Predictions? Some German Evidence. *Journal of Forecasting*, 25: 401 – 413.
- [16] Maćkowiak, Bartosz and Mirko Wiederholt. 2009. Optimal Sticky Prices Under Rational Inattention. *American Economic Review*, 99: 769-803.
- [17] Mankiw, Gregory N., Ricardo Reis, and Justin Wolfers. 2003. Disagreement about Inflation Expectations. *NBER Macroeconomics Annual*, 209-248.
- [18] Manski, Charles F. 2010. When Consensus Choice Dominates Individualism: Jensen’s Inequality and Collective Decisions under Uncertainty. *Quantitative Economics*, 1: 187-202.
- [19] Messina, Jeff, Tara M. Sinclair, and Herman O. Stekler. 2014. What Can We Learn From Revisions to the Greenbook Forecasts? RPF Working Paper No. 2014-003.
- [20] Nason, James M., and Gregor W. Smith. 2008. Identifying the New Keynesian Phillips Curve. *Journal of Applied Econometrics*, 23: 525-51.
- [21] Nordhaus, William D. 1987. Forecasting Efficiency: Concepts and Applications. *Review of Economics and Statistics*, 69 (4): 667-74.
- [22] Shea, John. 1997. Instrument Relevance in Multivariate Linear Models: A Simple Measure. *Review of Economics and Statistics*, 79: 348-352.
- [23] Staiger, Douglas and James H. Stock. 1997. Instrumental Variables Regression with Weak Instruments. *Econometrica*, 65: 557-586.
- [24] Stock, James H. and Mark W. Watson. 1999. Forecasting Inflation. *Journal of Monetary Economics*, 44: 293-335.
- [25] Woodford, Michael. 2009. Information-Constrained State-Dependent Pricing. *Journal of Monetary Economics*, 56: S100-S124.
- [26] Zarnowitz, Victor, and Phillip Braun. 1993. Twenty-Two Years of the NBER-ASA Quarterly Economic Outlook Surveys: Aspects and Comparisons of Forecasting Performance,

in *Business Cycles, Indicators, and Forecasting*, James H. Stock and Mark W. Watson, eds., Chicago, IL: University of Chicago Press.