BIS Working Papers
No 700

The perils of approximating fixed-horizon inflation forecasts with fixed-event forecasts

by James Yetman

Monetary and Economic Department

February 2018

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

This publication is available on the BIS website (www.bis.org).

The perils of approximating fixed-horizon inflation forecasts with fixed-event forecasts

February 2018

James Yetman[1]

Abstract

A common practice in studies using inflation forecasts is to approximate fixed-horizon forecasts with fixed-event ones. Here we show that this may be problematic. In a panel of US inflation forecast data that allows us to compare the two, the approximation results in a mean absolute approximation error of around 0.2–0.3 percentage points (around 10% of the level of inflation), and statistically significant differences in both the variances and persistence of the approximate inflation forecasts relative to the actual forecasts. To reduce these problems, we propose an adjustment to the approximation, consistent with a model where longer-horizon forecasts are more heavily "anchored", while shorter-horizon forecasts more closely reflect current inflation levels.

Keywords: fixed-event forecasts, fixed-horizon forecasts, inflation expectations

JEL classification: C43, E31

1. Introduction

Macroeconomic forecasts can be characterised as either "fixed event" or "fixed horizon". With fixed-event forecasts, the outcome being forecast remains fixed while the forecast horizon changes from one forecast to the next – eg forecasts of 2017 inflation made in consecutive months. By contrast, with fixed-horizon forecasts, the timing of the outcome and the forecast move together, so that the period between them remains constant – eg forecasts of inflation over the coming 12 months made at different points in time.

Fixed-event and fixed-horizon forecasts are each better suited for different purposes. Fixed-event forecasts provide an effective means to assess how expectations are changing in response to developments: since the outcome being forecast is unchanged from one forecast to the next, forecast changes must be due to changed expectations of the outcome. However, because the forecast horizon varies from one forecast to the next, other properties of the forecasts (such as the degree of forecaster disagreement, or the persistence of forecasts) may not be readily comparable.

In contrast, fixed-horizon forecasts are generally preferable when using forecasts as a proxy for expectations in estimation exercises. For example, when estimating a forward-looking Phillip's curve,[2] there is presumably an optimal forecast horizon for expected inflation for which the model will fit best. Using fixed-event forecasts in this context without adequately controlling for the varying horizons would compromise econometrics performance. Perhaps not surprisingly, when central banks collect inflation forecasts, they generally focus on fixed-horizon forecasts.[3]

Given the relatively limited availability of high-quality surveys of fixed-horizon forecasts, especially outside some major advanced economies, researchers have employed several different approaches to use fixed-event forecasts in their place:

(i)   Group forecasts by horizons and treat those in each group as having the same properties – for example, grouping all forecasts for "next year" versus "this year" as in Johnson (2002). Along similar lines, Capistrán and Ramos-Francia (2010) group all forecasts made in a given year in their study of forecast dispersion and, apart from adding fixed effects for each month, assume that the data-generating process is the same across horizons. Especially at

---

[2]   See Mavroeidis et al (2014), especially Section 3.1.3, for a summary of this approach.

[3]   Moreno and Villar (2010) and Sousa and Yetman (2016) discuss the inflation forecast surveys collected by emerging market central banks. In their survey responses, where central banks indicated the nature of the forecasts they were collecting, these were generally fixed-horizon rather than fixed-event. In 2011 the Reserve Bank of New Zealand switched from collecting fixed-event to fixed-horizon inflation forecasts (Lees (2016)).

shorter horizons, this is likely to be problematic. Forecasts made in January for the outcome in that calendar year are likely to be very different from forecasts made in December for the same year, given that most of the information for the current year are likely to be known by the forecaster at the time of the shortest-horizon forecasts.

(ii)   Model the evolution of the forecasts across the different horizons explicitly. An early paper taking this approach was Davies and Lahiri (1995), who model forecast errors as being made up of bias, monthly shocks and idiosyncratic errors, where the monthly shocks accumulate across horizons and the distribution of the errors depends on the forecast horizon. Another approach is to allow forecasts to be more heavily influenced by a long-run anchor of inflation expectations, or endpoint, the longer is the forecast horizon (eg Kozicki and Tinsley (2012), Mehrotra and Yetman (2014) and Yetman (2017)). Modelling the evolution of forecasts across horizons may be computationally demanding compared with other options, however.

(iii)   Approximate fixed-horizon forecasts based on fixed-event forecasts. In this category, see Dovern et al (2012) and Siklos (2013, 2017), in studies of forecast disagreement; Johnson (2003), in a study of the effect of inflation targeting on the level of expected inflation; Gerlach (2007), who uses the approximation to construct a measure of inflation expectations when estimating the monetary policy reaction function of the European Central Bank; Kortelainen et al (2011), who generate approximations of expected inflation and GDP to use in estimating a three-equation New Keynesian model; and Davis (2014), who assesses the responses of inflation expectations to various shocks. Some authors cite Dovern and Fritsche (2008) in defence of this approach. They compare approximate fixed-horizon forecasts, constructed from fixed-event forecasts, with actual fixed-horizon forecasts constructed from quarterly forecasts made by the same panel of forecasts. They argue that the approximation results in levels of dispersion across forecasters that are highly correlated with the true level of dispersion.

This third approach involves an approximation error. One objective of this paper is to assess how large the approximation error is likely to be, and what it implies for other properties of the forecasts. This is generally difficult to determine because, while there are panels of fixed-horizon and fixed-event forecasts for the same variables, they generally differ along many dimensions (eg the timing of the forecasts and the composition of the forecaster panel) such that they cannot be readily compared.

However, there is one panel of forecaster-level inflation forecasts that we can use to construct both fixed-event and fixed-horizon forecasts from the same underlying data. This data set is the "Survey of Professional Forecasters" (SPF), collated by the Federal Reserve Bank of Philadelphia. It is the same data

set as Dovern and Fritsche (2008) explored.[4] The quarterly SPF includes forecasts of inflation for six quarters (the previous quarter, the current quarter and each of the following four quarters) as well as up to three calendar years (defined as the fourth-quarter-over-fourth-quarter growth rate; for the current year as well as the following two years).[5] We use this data to construct both actual "fixed-event" 12-month inflation (based on the quarterly inflation forecasts) and approximate 12-month inflation (based on the fixed-horizon annual inflation forecasts), and then compare their behaviour.

Using data beginning in 1990, we find that the approximation introduces distortions into the inflation measures that are often quantitatively large and statistically significant. The mean absolute approximation error is around 0.2–0.3 percentage points, or around 10% of the level of inflation. The approximation error also leads to statistically significant changes in the variance of the resulting inflation measures. Further, the persistence in the annual inflation forecast series is often statistically significantly different when the approximation is used, a problem that appears to be worse in the second half of our sample period than in the first.

We illustrate a possible partial fix for the approximation bias, drawing on the idea that inflation forecasts are more heavily "anchored" as the forecast horizon increases, as discussed in Kozicki and Tinsley (2012), Mehrotra and Yetman (2014) and Yetman (2017). The "fix" reduces the mean absolute error by 13% (and the mean squared error by 21%), and results in fewer significant differences between the actual and approximate inflation rates in terms of their variances and persistence as well. Further, higher order moments – skewness and kurtosis – of the resulting panel of approximate forecasts are more alike the actual data when our adjustment to the standard approximation is implemented.

In the next section, we briefly outline the most common approximation used to translate fixed-event forecasts into fixed-horizon ones. Section 3 summarises how we can use the SPF to produce both actual and approximate inflation forecasts. Section 4 then compares the resulting "approximate" and actual fixed-horizon forecasts. Section 5 discusses one way to improve on the standard approximation,

---

[4]   Additionally, it would be possible to examine the effect of using the approximation using data from Consensus Economics, but based on average forecasts across a panel. For the 27 economies included in their "Consensus Forecasts" or "Asia-Pacific Consensus Forecasts" publications, average quarterly inflation forecasts are published for each of the following seven to eight quarters on a quarterly basis. Under the assumption that the composition of the annual fixed-event and quarterly forecasts are the same, the quarterly forecasts could be used to compute actual average 12-month inflation forecasts which should be comparable with approximate forecasts based on average fixed-event annual forecasts. In their eastern European and Latin American publications, Consensus Economics includes average forecasts of monthly inflation for each of the following six months, so no such comparison with annual inflation is possible. However, in no cases could published Consensus Economics data be used to examine the effects of the approximation at the level of the individual forecaster, unlike the SPF forecast data examined here.

[5]   Forecasts from the SPF are all available at www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/.

and compares the performance of the "improved" approximation with the standard approximation. Section 6 then concludes.

2. Approximating 12-month inflation forecasts

When approximating fixed-horizon forecasts with fixed-event ones, the most common approach is as follows. Let $\hat{\pi}_{t+k|t}$ denote the forecast of annual inflation for the year ending k-months-ahead based on information available at time t. Some forecast surveys, such as from Consensus Economics,[6] include a pair of forecasts $\{\hat{\pi}_{t+k|t}, \hat{\pi}_{t+12+k|t}\}$ for each month, with horizons $k \in \{1, 2, \ldots, 12\}$. The value of k shortens as the year progresses, as all forecasts are for years ending in December, so in January $k = 12$, in February $k = 11$, and so on.

Fixed 12-month horizon forecasts are then approximated in the studies listed in the introduction as averages of the forecasts for the current and next calendar years, weighted by their shares in the forecast period:

$$\hat{\pi}_{t+12|t} = \frac{k}{12}\hat{\pi}_{t+k|t} + \frac{12-k}{12}\hat{\pi}_{t+12+k|t}.$$
(1)

This approach implies that the 12-month-ahead forecast for inflation made in October 2012 is approximated by the sum of $\hat{\pi}_{2012M12|2012M10}$ and $\hat{\pi}_{2013M12|2012M10}$, with weights $3/12$ and $9/12$ respectively, for example.

3. Calculating actual and approximate forecasts

The SPF provides us with quarterly and annual inflation forecasts at a quarterly frequency that we can use to construct both actual and approximate 12-month inflation forecasts for each quarter. For example, in the first quarter of each year, the available forecasts include the following:

---

[6]   http://www.consensuseconomics.com/.

(Forecasts one-through-six indicate forecasts of quarterly inflation, while A and B indicate forecasts of annual inflation. There is also a forecast C for two years ahead, but we do not make use of this.)

We use these forecasts in two ways. First, we confirm that the quarterly forecasts are approximately consistent with the annual forecasts: that is, annual forecast A is consistent with quarterly forecasts 2–5, linked by the following formula:

$$\hat{\pi}_A = 100\left[\left[\left(1+\frac{\hat{\pi}_2}{100}\right)\left(1+\frac{\hat{\pi}_3}{100}\right)\left(1+\frac{\hat{\pi}_4}{100}\right)\left(1+\frac{\hat{\pi}_5}{100}\right)\right]^{1/4}-1\right]. \tag{2}$$

Second, we compare the behaviour of actual four-quarter inflation implied by quarterly forecasts labelled 3–6 with approximate four-quarter inflation for the same period using the approximation outlined in Section 2. Based on the approximation were to hold, we would expect that:

$$100\left[\left[\left(1+\frac{\hat{\pi}_3}{100}\right)\left(1+\frac{\hat{\pi}_4}{100}\right)\left(1+\frac{\hat{\pi}_5}{100}\right)\left(1+\frac{\hat{\pi}_6}{100}\right)\right]^{1/4}-1\right] \approx \frac{3}{4}\hat{\pi}_A+\frac{1}{4}\hat{\pi}_B. \tag{3}$$

For the second quarter, the relevant available forecasts are:

For this quarter, our test that the annual and quarterly forecasts are consistent is that:

$$\hat{\pi}_A = 100\left[\left[\left(1+\frac{\hat{\pi}_1}{100}\right)\left(1+\frac{\hat{\pi}_2}{100}\right)\left(1+\frac{\hat{\pi}_3}{100}\right)\left(1+\frac{\hat{\pi}_4}{100}\right)\right]^{1/4}-1\right], \tag{4}$$

while we also have two possible tests of the approximation, based on the quarterly forecasts 2–5 and 3–6, as follows:

$$100\left[\left[\left(1+\frac{\hat{\pi}_2}{100}\right)\left(1+\frac{\hat{\pi}_3}{100}\right)\left(1+\frac{\hat{\pi}_4}{100}\right)\left(1+\frac{\hat{\pi}_5}{100}\right)\right]^{1/4}-1\right] \approx \frac{3}{4}\hat{\pi}_A+\frac{1}{4}\hat{\pi}_B; \tag{5}$$

$$100\left[\left[\left(1+\frac{\hat{\pi}_3}{100}\right)\left(1+\frac{\hat{\pi}_4}{100}\right)\left(1+\frac{\hat{\pi}_5}{100}\right)\left(1+\frac{\hat{\pi}_6}{100}\right)\right]^{1/4}-1\right] \approx \frac{1}{2}\hat{\pi}_A+\frac{1}{2}\hat{\pi}_B. \tag{6}$$

For quarter three, no tests of consistency are possible (since we do not have all four quarterly forecasts of any given calendar year), but we can consider three tests of the approximation, based on quarterly forecasts 1–4, 2–5 and 3–6 respectively, as follows:

$$100\left[\left[\left(1+\frac{\hat{\pi}_1}{100}\right)\left(1+\frac{\hat{\pi}_2}{100}\right)\left(1+\frac{\hat{\pi}_3}{100}\right)\left(1+\frac{\hat{\pi}_4}{100}\right)\right]^{1/4}-1\right] \approx \frac{3}{4}\hat{\pi}_A+\frac{1}{4}\hat{\pi}_B; \tag{7}$$

$$100\left[\left[\left(1+\frac{\hat{\pi}_2}{100}\right)\left(1+\frac{\hat{\pi}_3}{100}\right)\left(1+\frac{\hat{\pi}_4}{100}\right)\left(1+\frac{\hat{\pi}_5}{100}\right)\right]^{1/4}-1\right] \approx \frac{1}{2}\hat{\pi}_A+\frac{1}{2}\hat{\pi}_B; \tag{8}$$

$$100\left[\left[\left(1+\frac{\hat{\pi}_3}{100}\right)\left(1+\frac{\hat{\pi}_4}{100}\right)\left(1+\frac{\hat{\pi}_5}{100}\right)\left(1+\frac{\hat{\pi}_6}{100}\right)\right]^{1/4}-1\right] \approx \frac{1}{4}\hat{\pi}_A+\frac{3}{4}\hat{\pi}_B. \tag{9}$$

Finally, for the fourth quarter, consistency requires that forecasts 3-6 correspond to forecast B:

$$\hat{\pi}_B = 100\left[\left[\left(1+\frac{\hat{\pi}_3}{100}\right)\left(1+\frac{\hat{\pi}_4}{100}\right)\left(1+\frac{\hat{\pi}_5}{100}\right)\left(1+\frac{\hat{\pi}_6}{100}\right)\right]^{1/4}-1\right], \tag{10}$$

while the two available tests of the approximation are:

$$100\left[\left[\left(1+\frac{\hat{\pi}_1}{100}\right)\left(1+\frac{\hat{\pi}_2}{100}\right)\left(1+\frac{\hat{\pi}_3}{100}\right)\left(1+\frac{\hat{\pi}_4}{100}\right)\right]^{1/4}-1\right] \approx \frac{1}{2}\hat{\pi}_A+\frac{1}{2}\hat{\pi}_B; \tag{11}$$

$$100\left[\left[\left(1+\frac{\hat{\pi}_2}{100}\right)\left(1+\frac{\hat{\pi}_3}{100}\right)\left(1+\frac{\hat{\pi}_4}{100}\right)\left(1+\frac{\hat{\pi}_5}{100}\right)\right]^{1/4}-1\right] \approx \frac{1}{4}\hat{\pi}_A+\frac{3}{4}\hat{\pi}_B. \tag{12}$$

Combining all horizons, then, we have three observations in each year for each forecaster where the annual forecasts and the quarterly forecasts should be consistent with each other (based on equations 2, 4 and 10). In addition, we have eight observations in each year for each forecaster where we can compute both approximate and actual 12-month inflation to compare (based on equations 3, 5–9 and 11–12).[7]

---

[7] The panel is not balanced, with forecasters entering and leaving the panel and some quarterly forecasts missing, affecting the number of total number of observations.

In our analysis of the SPF forecasts, we focus on the Q1 1990–Q1 2017 period. We avoid earlier forecasts, given that the behaviour of inflation forecasts during relatively high inflation periods may be a poor guide to forecast performance today.[8] We construct the measures outlined above for each forecaster, in each period. We examine all observations for which both comparable measures are available.

Graph 1. Equivalent inflation forecasts
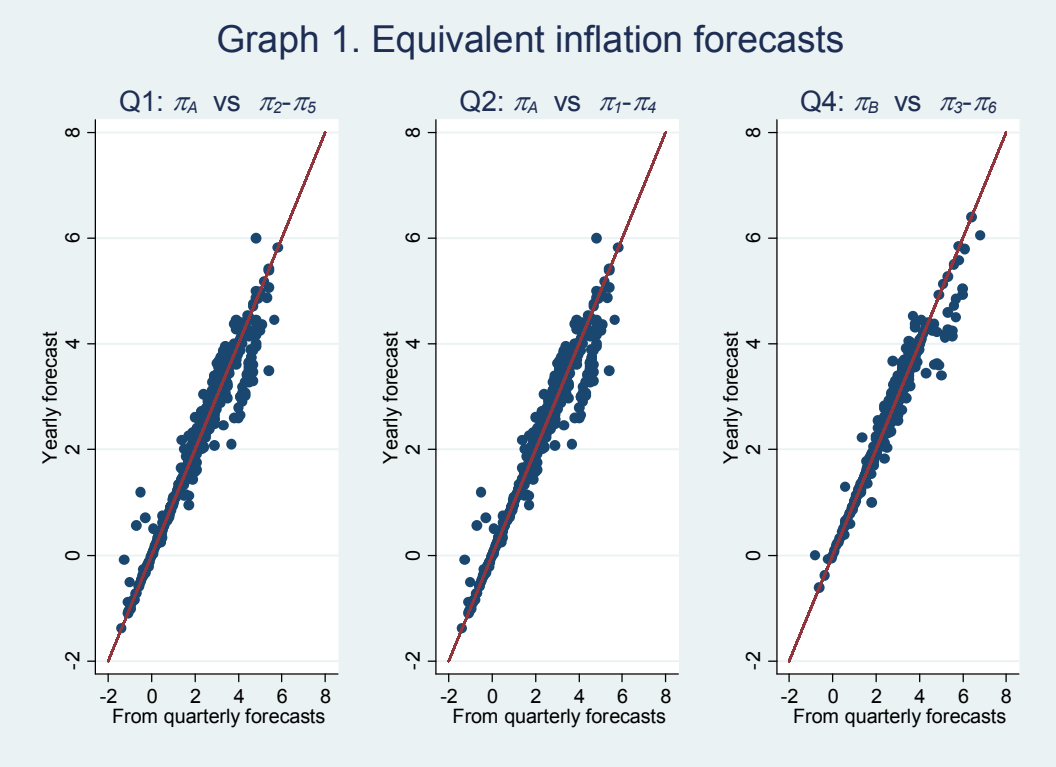
Notes: the panels display annual inflation forecasts and their equivalent made up from quarterly inflation forecasts. Deviations between the dots and the 45-degree line are a result of rounding errors and inconsistencies and inconsistencies between the quarterly and annual forecasts.

Graph 1 displays the three tests of consistency outlined above (in equations 2, 4 and 10). If the annual and quarterly forecasts were perfectly consistent, then we would expect them to lie on a 45-degree line drawn through the origin. Even allowing for rounding error (given that the forecasts are typically given as percentages to one decimal place), there are clearly some inconsistencies: in 4.5% of cases, the difference exceeds 0.4 percentage points, and in 1.4% of cases 1.0 percentage points. In the

---

[8]   In addition, the period we study roughly coincides with the period during which the survey was conducted by the Philadelphia Fed. Before Q2 1990, the survey was conducted by the American Statistical Association (ASA) and National Bureau of Economic Research (NBER). Background information on the survey available online notes that "in the surveys conducted by the NBER/ASA, the same identification number could represent different forecasters. In these surveys, we have noticed some occurrences in which an individual participates, suddenly drops out of the panel for a large number of periods, and suddenly re-enters, suggesting that the same identifier might have been assigned to different forecasters. Unfortunately, we cannot investigate the historical record of these individuals because we do not have hard-copy historical records from the early surveys."

following analysis, we will consider dropping all forecasts where the inconsistency is large as a robustness check.

4. Comparing actual and approximate forecasts

Our next step is to compare the actual and approximate 12-month inflation forecasts under each of the cases outlined above, and to highlight the differences between then. Some key results are summarised in Table 1 below, both for the full sample and dropping those forecasts where quarterly and annual inflation forecasts are inconsistent.[9] Shaded rows represent results where there is no approximation involved, and any difference is due to measurement or rounding error, or underlying inconsistency between the quarterly and annual forecasts. We use these as a benchmark to indicate the level of variation that attributable to random noise, as opposed to approximation error.

Column (3) contains the mean absolute difference between the approximation and the actual 12-month inflation forecasts. For (shaded) quarters in which the only source of difference is measurement and rounding error, these differences are less than 0.1 percentage points (and around 0.03 percentage points once we drop observations with larger inconsistencies from the sample). By contrast, these average differences are three to 12 times as large for other quarters, or around 0.2–0.3 percentage points, which is approximately 10% of the average inflation rate during the sample period. This increase in average error is highly statistically significant (p-value < 0.01 in all cases). The results in column (4) illustrate the same message based on mean squared differences.

Columns (5–7) display the variance of forecasts across the two measures of annual inflation, and the p-value from an F-test of equal variance. For (shaded) cases, in quarters where the only difference is rounding and measurement error, this difference is small and never statistically significant. However, for other cases, the difference is often statistically highly significant, with p-values smaller than 0.05 in nine out of the 16 cases examined. Note, however, that there is no simple pattern of differences: in some cases the approximation results in a significantly larger variance, and in others in a significantly smaller variance.

---

[9]   We define inconsistent as any time the difference displayed in Graph 1 exceeds 0.2 percentage points. For this exercise, we drop forecasts for a given forecaster for all quarters in any year where at least one of the three consistency checks is violated. (We use this condition because there are no consistency checks for quarter-3 forecasts.) We also examined alternative cut-off rules as robustness checks, such as excluding all forecasters if the largest inconsistency exceeds either 0.4 or 1.0 percentage points, and results are similar in each case.

(1) Forecast quarter	(2) Obs.	(3) MA diff.	(4) MS diff.	(5-7) Variance			(8-9) Skewness		(10-11) Kurtosis	
Table 1. Errors and moments										
				Approx.	Actual	p-val.	Approx.	Actual	Approx.	Actual
Approximation c.f. actual based on $\hat{\pi}_1$ - $\hat{\pi}_4$										
	Full sample									
1	0									
2	1012	0.09	0.05	1.03	0.98	0.47	-0.44	-0.62	3.83	3.74
3	904	0.31	0.17	0.85	0.74	0.04	0.01	-0.02	3.11	3.46
4	1022	0.36	0.22	0.75	0.95	0.00	0.91	1.03	5.36	7.58
	Drop forecasters in years with larger inconsistencies									
1	0									
2	805	0.03	0.00	0.92	0.93	0.88	-0.65	-0.65	3.70	3.66
3	724	0.29	0.15	0.77	0.74	0.51	0.00	0.09	2.78	3.48
4	808	0.36	0.23	0.52	0.63	0.00	0.21	0.06	3.46	6.50
Approximation c.f. actual based on $\hat{\pi}_2$ - $\hat{\pi}_5$										
	Full sample									
1	1021	0.08	0.05	0.90	0.83	0.17	-0.21	-0.29	4.68	4.38
2	923	0.29	0.16	0.86	0.69	0.00	-0.27	-0.24	3.65	3.91
3	904	0.29	0.16	0.67	0.68	0.81	0.10	0.08	3.45	3.87
4	1021	0.26	0.22	0.71	1.06	0.00	1.00	-0.03	6.10	7.01
	Drop forecasters in years with larger inconsistencies									
1	833	0.03	0.00	0.69	0.69	0.94	-0.51	-0.52	4.64	4.63
2	722	0.26	0.14	0.74	0.61	0.02	-0.48	-0.33	3.56	4.29
3	724	0.30	0.16	0.57	0.60	0.43	0.07	0.07	3.15	4.42
4	807	0.27	0.24	0.46	0.79	0.00	0.16	-0.89	4.06	7.56
Approximation c.f. actual based on $\hat{\pi}_3$ - $\hat{\pi}_6$										
	Full sample									
1	891	0.22	0.13	0.79	0.66	0.01	0.06	0.22	4.15	4.33
2	923	0.28	0.15	0.73	0.69	0.48	-0.01	0.22	3.89	4.17
3	895	0.20	0.09	0.59	0.66	0.09	0.15	0.02	3.83	3.61
4	1024	0.08	0.04	0.78	0.71	0.19	0.71	0.37	6.07	4.77
	Drop forecasters in years with larger inconsistencies									
1	724	0.19	0.10	0.58	0.52	0.11	-0.22	0.00	4.30	4.64
2	722	0.28	0.15	0.59	0.56	0.51	-0.22	0.15	4.05	5.24
3	718	0.20	0.08	0.47	0.56	0.03	0.13	-0.03	3.72	4.46
4	809	0.03	0.00	0.53	0.53	0.92	-0.09	-0.08	4.57	4.51

Notes: MA diff. (MS diff.) is the average mean absolute difference (mean squared difference) between the forecasts based on annual data (using the approximation in equation (1)) and quarterly data. Shaded rows are those where there is no approximation error involved.

Columns (8–9) illustrate differences in the degree of skewness across the two measures, and the final two columns display kurtosis. Here there is a pattern: kurtosis is nearly always greater in the 12-month forecasts constructed from quarterly forecasts than in the approximate measure based on annual forecasts.

We also illustrate the differences for one particular case, corresponding to the third row of results in the table, visually in Graph 2. For these forecasts, of inflation for the year ending in the fourth quarter that were made in the second quarter, we take the distribution of forecasts for each year and calculate the median (solid line) and quartiles (dotted lines) of the approximate (red) and actual (blue) inflation forecasts separately. The graph shows that the two distributions are very different at times. For example, for forecasts of 2001 inflation, the median approximate forecast is 0.6 percentage points higher than the actual forecast, and there is no overlap between the interquartile ranges of the two series. Similar sized gaps also exist in 2007–09, with the approximate forecasts much higher than the actual for the first two of those years, and then much lower for the third. Given that these years with the largest gaps all correspond to periods where the United States was in recession, it appears that the approximation is especially poor during economic downturns.

Graph 2: Comparison of distribution of actual and approximate forecasts: Quarter 2 forecasts of inflation for year ending quarter 4
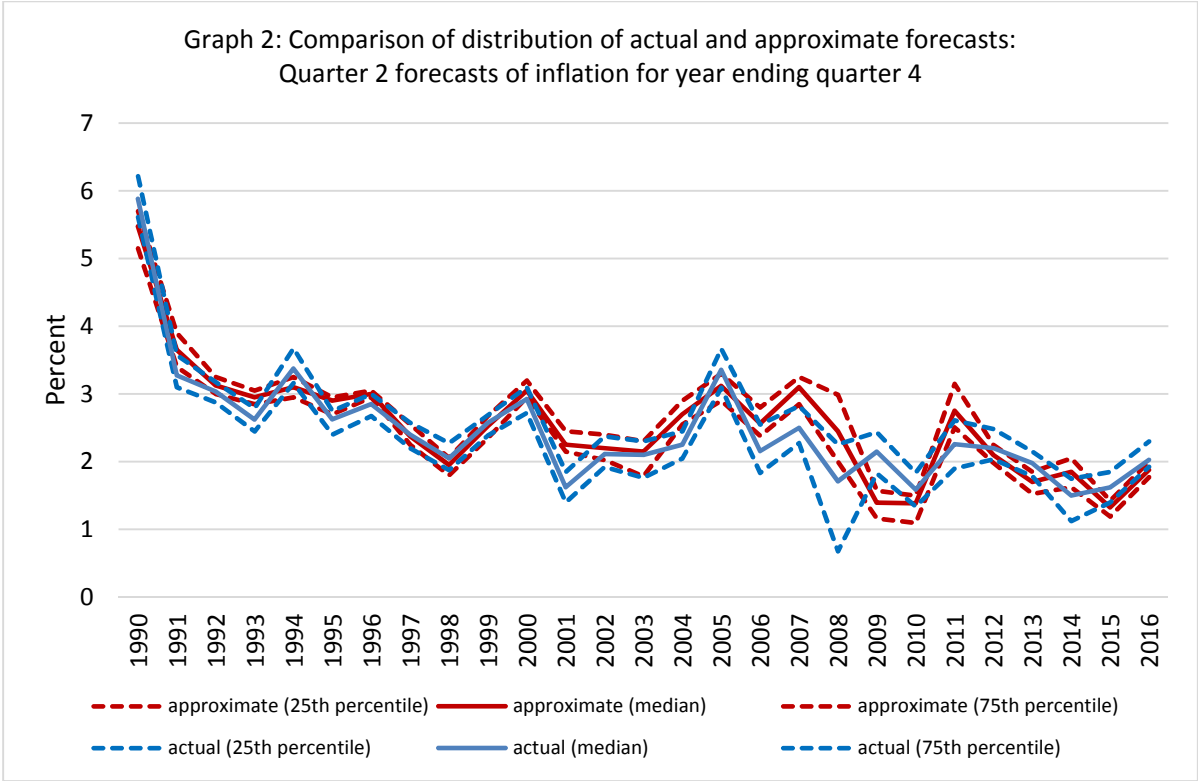
Taking these results together, the approximation results in a non-trivial amount of noise in the 12-month inflation forecasts, relative to forecasts constructed from quarterly data where no approximation is involved. In addition, it often results in a statistically significant change in the variance of the distribution and reduces the thickness of the distribution's tails. Given that the variance and kurtosis are potential measures of forecaster disagreement, which in turn is commonly interpreted as a measure of central bank credibility (eg Dovern et al (2012); Siklos (2013, 2017)), the approximation may lead to biased empirical analysis.

Another aspect of the forecasts that we assess is the degree of persistence. In particular, we estimate:

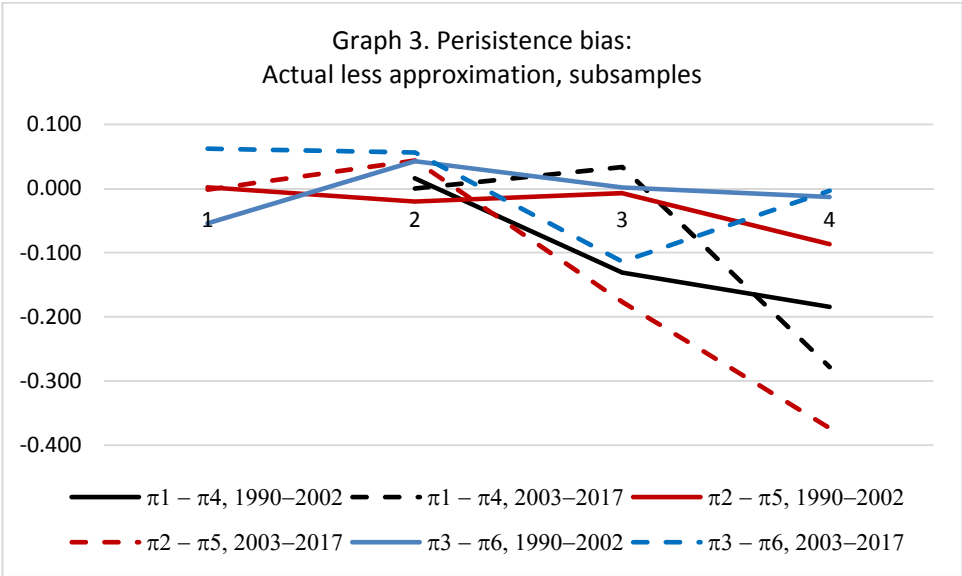$$\pi_{y,n}^i = \beta^i + \rho \pi_{y-1,n}^i + \varepsilon_{y,n}^i, \tag{13}$$

on a panel where forecasters are indexed by i, and compare the estimates of ρ between actual and approximate forecasts of 12-month inflation. This regression is estimated on data for each quarter (n) separately, so that there is one observation for each forecaster, per year in each regression.

(1) Forecast quarter	(2) Obs.	(3-5) ρ		
		Approx.	Actual	p-value
Table 2. Persistence				
Approximation c.f. actual based on $\hat{\pi}_1$ - $\hat{\pi}_4$				
	Full sample			
1	0			
2	741	0.14	0.15	0.84
3	603	0.23	0.26	0.52
4	749	0.32	0.15	0.00
	Drop forecasters in years with larger inconsistencies			
1	0			
2	534	0.12	0.12	0.95
3	437	0.21	0.22	0.84
4	533	0.30	0.06	0.00
Approximation c.f. actual based on $\hat{\pi}_2$ - $\hat{\pi}_5$				
	Full sample			
1	766	0.18	0.22	0.25
2	646	0.19	0.30	0.01
3	603	0.32	0.29	0.50
4	747	0.37	0.21	0.00
	Drop forecasters in years with larger inconsistencies			
1	558	0.18	0.18	0.99
2	455	0.16	0.21	0.31
3	437	0.28	0.16	0.03
4	531	0.36	0.09	0.00
Approximation c.f. actual based on $\hat{\pi}_3$ - $\hat{\pi}_6$				
	Full sample			
1	624	0.26	0.35	0.02
2	646	0.29	0.41	0.00
3	594	0.42	0.44	0.60
4	750	0.37	0.42	0.08
	Drop forecasters in years with larger inconsistencies			
1	453	0.24	0.29	0.35
2	455	0.23	0.30	0.20
3	432	0.36	0.31	0.33
4	532	0.34	0.34	0.97

Notes: Based on regressing inflation forecasts on a constant and lagged inflation forecasts. Shaded rows are those where there is no approximation error involved.

The results are contained in Table 2. For the quarters where the differences are only due to measurement and rounding error (the shaded rows), there is generally no statistically significant difference in the degree of persistence, especially once we remove observations in years with inconsistent forecasts. However, in other quarters, the differences are statistically significant in about half of all cases. There is no clear pattern to the differences: in some cases, the persistence is greater in the approximate forecasts, while in other cases it is smaller.

High levels of persistence in forecasts are commonly interpreted as indicators of poorly anchored expectations, since they imply that any departure from an inflation target is likely to self-perpetuate. Studies of central bank credibility sometimes focus on changes in the persistence of inflation forecasts over time (eg Corbo et al (2001)). Even more important than any bias in the level of persistence, therefore, would be changes in the amount of bias over time. To check for this, we divide the sample approximately in half and estimate the above equation on each subsample. Graph 3 displays the point estimates below.

Graph 3. Perisistence bias: Actual less approximation, subsamples

Legend: $\pi1 - \pi4$, 1990–2002 · $\pi1 - \pi4$, 2003–2017 · $\pi2 - \pi5$, 1990–2002 · $\pi2 - \pi5$, 2003–2017 · $\pi3 - \pi6$, 1990–2002 · $\pi3 - \pi6$, 2003–2017

Notes: The vertical axis is the difference in estimated persistence between actual and approximate data based on regressing inflation forecasts on a constant and lagged inflation forecasts on subsamples of the data. The horizontal axis is the quarter. Solid lines are based on the first half of the sample, and dashed lines on the second half.

We calculate the persistence in the actual forecasts less that in the approximate forecasts, for each of the first (solid lines) and second (dashed lines) halves of the sample. In most cases, this difference is negative, indicating that the approximate forecasts are more persistent than the actual forecasts. However, there is also a visible difference between the two subsamples. On average, the dashed lines lie below the solid lines of the same colour, indicating that the degree of bias in the estimates of

persistence based on approximate forecasts in the latter subsample is more severe than in the earlier subsample. The magnitude of this difference (as much as 0.3) is roughly the same magnitude as the underlying persistence in the forecasts themselves, suggesting that any studies of changing forecast persistence based on approximate data may be misleading.

5. Improving on the approximation

We have highlighted some problems with approximating fixed-horizon inflation forecasts using fixed-event forecasts. Now we offer a partial fix that improves on the standard approximation in terms of better matching the statistical properties of the actual forecast data.

Kozicki and Tinsley (2012), Mehrotra and Yetman (2014) and Yetman (2017) have suggested models of inflation forecasts where these are increasingly "anchored" as the forecast horizon increases. One formulation of this is the following:

$$\hat{\pi}_{t+k|t} = \alpha_k \pi^* + (1-\alpha_k)\pi_{t-1} + \varepsilon_t^k, \tag{14}$$

where π^* is the inflation anchor and π_{t-1} is the latest available inflation outcome at the time that the forecast is made (lagged by one month to allow for publication lags). We model α_k as a simple exponential function ($\alpha_k = 1 - \exp(-k/b)$ for some value of b, where k is the forecast horizon), and estimate π^* and b for each of the forecasters who appear frequently in our panel, across the available calendar-year inflation forecasts (labelled A and B in section 3), for all time periods. We then use these estimates to determine the appropriate weights to combine the calendar-year inflation forecasts to produce a 12-month fixed-horizon forecast. To be more precise, we have forecasts for two calendar years, $\hat{\pi}_{t+k|t}$ and $\hat{\pi}_{t+k+12|t}$, but we require forecasts for a fixed horizon of 12 months, $\hat{\pi}_{t+12|t}$. We wish to approximate the latter as:

$$\hat{\pi}_{t+12|t} = x_k \hat{\pi}_{t+k|t} + (1-x_k)\hat{\pi}_{t+k+12|t}, \tag{15}$$

for some value of x_k. Note that this can be viewed as a generalisation of the standard approximation given in equation (1). Now, by substituting from equation (14), we can rewrite equation (15) as:

$$\begin{aligned}
\hat{\pi}_{t+12|t} = &\left[x_k \alpha_k + (1-x_k)\alpha_{k+12}\right]\pi^* \\
&+ \left[x_k(1-\alpha_k) + (1-x_k)(1-\alpha_{k+12})\right]\pi_{t-1} \\
&+ \left[x_k \varepsilon_t^k + (1-x_k)\varepsilon_t^{k+12}\right].
\end{aligned} \tag{16}$$

From equation (14), we can also rewrite $\hat{\pi}_{t+12|t}$ as:

$$\hat{\pi}_{t+12|t} = \alpha_{12}\pi^* + (1-\alpha_{12})\pi_{t-1} + \varepsilon_t^{12}. \tag{17}$$

Equating the coefficients on π^* (or π_{t-1}) in (16) and (17), we obtain:

$$x_k = \frac{\alpha_{12} - \alpha_{k+12}}{\alpha_k - \alpha_{k+12}} = \frac{\exp(-[k+12]/b) - \exp(-12/b)}{\exp(-[k+12]/b) - \exp(-k/b)}. \tag{18}$$

Note that for $k=12$, $x=1$ so that the full weight goes on the shorter duration annual forecast, and for $k=0$, $x=0$ so that the full weight goes on the longer duration annual forecast. These extreme cases are analogous to the shaded rows in the previous tables of results, where there is no approximation error involved. More generally, $0 \le x_k \le 1$ for all values of $b \ge 0$. And for the limiting case of $b \to 0$, $x \to (12-k)/k$, which is the standard approximation given by equation (1) that we examined in the previous section.

To implement this adjustment to the approximation, we focus on the 65 forecasters who provided forecasts to the SPF at least 20 times within our sample. While these represent only 37% of the 174 forecasters in the SPF panel, their forecasts constitute 79% of the total forecast panel, as many of the other forecasters only enter a very limited number of times.[10] We then estimate equation (14) for each forecaster, obtaining estimates of b that range from 1.4 and 76.8.[11] Then, based on these estimates, we construct x_k for each forecaster, and use this to compute an "adjusted" approximate fixed-horizon forecast, using equations (15) and (18), that we can then compare with the standard measure. We also re-compute our standard measure for this comparison based on the same sample of forecast data (ie eliminating those forecasters who contribute fewer than 20 observations to the panel from our previous results).

The results for the first and second moments are given in Table 3, and for higher moments and persistence in Table 4. In comparing the adjustment with the standard approximation, we focus only on the unshaded portions of the table- ie the elements where (i) there is an approximation error inherent in the data and (ii) for $b > 0$, our adjusted approximation differs from the standard approximation. For each statistic, we are therefore comparing eight elements in our discussion.

Columns (3) and (4) of Table 3 display the mean absolute differences between the actual and approximate measures of inflation for the standard and the adjusted approximation, respectively. We can see that the adjustment results in a systematic improvement in the performance of the

---

[10]   18 forecaster codes are associated with only one forecast, for example.

[11]   For this exercise, π_{t-1} is the year-on-year change in the quarterly CPI (where quarterly CPI is the three-month average of monthly CPI) lagged by one quarter, where monthly CPI is taken from https://fred.stlouisfed.org/series/CPIAUCSL.

approximation: in seven of the eight cases, the difference declines, and the average decline in the mean absolute difference is by 13%. Columns (5) and (6) repeat the exercise for the mean squared difference, where there is some improvement across all eight cases, by an average of 21%.

Table 3. Comparison between standard and adjusted approximation: errors and variance											
(1) Forecast quarter	(2) Obs.	(3-4) MA difference		(5-6) MS difference		(7-12) Variance					
						Stand.			Adjust.		
		Stand.	Adjust.	Stand.	Adjust.	Approx.	Actual	p-val.	Approx.	Actual	p-val.
Approximation c.f. actual based on $\hat{\pi}_1 - \hat{\pi}_4$											
1	0										
2	790	0.08	0.08	0.03	0.03	0.96	0.94	0.76	0.96	0.94	0.76
3	712	0.30	0.31	0.16	0.16	0.84	0.75	0.16	0.72	0.75	0.51
4	813	0.36	0.32	0.24	0.19	0.65	0.81	0.00	0.63	0.81	0.00
Approximation c.f. actual based on $\hat{\pi}_2 - \hat{\pi}_5$											
1	816	0.07	0.07	0.04	0.04	0.78	0.72	0.27	0.78	0.72	0.27
2	723	0.28	0.25	0.16	0.12	0.77	0.63	0.01	0.67	0.63	0.37
3	712	0.31	0.25	0.17	0.13	0.64	0.66	0.77	0.60	0.66	0.21
4	812	0.27	0.25	0.24	0.20	0.60	0.96	0.00	0.64	0.96	0.00
Approximation c.f. actual based on $\hat{\pi}_3 - \hat{\pi}_6$											
1	725	0.21	0.19	0.12	0.09	0.66	0.55	0.01	0.59	0.55	0.32
2	723	0.29	0.23	0.16	0.11	0.62	0.59	0.47	0.57	0.59	0.69
3	706	0.21	0.16	0.10	0.08	0.56	0.64	0.09	0.58	0.64	0.19
4	815	0.06	0.06	0.02	0.02	0.68	0.65	0.55	0.68	0.65	0.55

Notes: MA difference (MS difference) is the average mean absolute difference (mean squared difference) between the forecasts based on annual data (generally involving an approximation) and quarterly data. Stand. indicates the standard approximation for converting fixed-event to fixed-horizon forecasts given in equation (1); Adjust. indicates the adjusted approximation given in equations (15) and (18). Shaded rows are those where there is no approximation error involved. Results are based on data for all forecasters who contribute to the SPF panel at least 20 times during the sample period.

The remainder of Table 3 compares the performance of the variance across the panel of forecasters. Once again, our adjusted approximation generally outperforms the standard approximation: in six of the eight cases, the variance is closer to the actual variance using our adjustment in place of the standard approximation, and the average decline in this gap across all eight cases is by 22%. Moreover, the difference between the variances is statistically significant in only two cases (at the 95% level) using our adjustment, versus four cases using the standard approximation.

Moving on to higher moments, here improvements are also evident (Table 4). The gap between the skewness of the distribution of approximate and actual forecasts (columns 3–6) declines using the adjusted approximation in five of the eight cases, and by an average of 9% overall. Moreover, the gap in the case of kurtosis (columns 7–10) decreases in all but one case, by an overall average of 23%.

Table 4. Comparison between standard and adjusted approximation: skewness, kurtosis and persistence

(1) Forecast quarter	(2) Obs.	(3–6) Skewness				(7–10) Kurtosis				(11–16) Persistence					
		Standard		Adjusted		Standard		Adjusted		Standard			Adjusted		
		Approx.	Actual	Approx.	Actual	Approx.	Actual	Approx.	Actual	Approx.	Actual	p-val.	Approx.	Actual	p-val.
Approximation c.f. actual based on $\hat{\pi}_1 - \hat{\pi}_4$															
1	0														
2	790	-0.48	-0.62	-0.48	-0.62	3.90	3.73	3.90	3.73	0.15	0.16	0.89	0.15	0.16	0.89
3	712	0.13	0.12	0.23	0.12	3.16	3.46	3.53	3.46	0.26	0.28	0.67	0.32	0.28	0.39
4	813	0.75	0.65	0.70	0.65	5.21	7.67	5.72	7.67	0.33	0.16	0.00	0.37	0.16	0.00
Approximation c.f. actual based on $\hat{\pi}_2 - \hat{\pi}_5$															
1	816	-0.45	-0.47	-0.45	-0.47	5.00	4.45	5.00	4.45	0.19	0.23	0.24	0.19	0.23	0.24
2	723	-0.30	-0.29	-0.11	-0.29	3.63	3.67	3.68	3.67	0.20	0.31	0.01	0.26	0.31	0.18
3	712	0.23	0.18	0.28	0.18	3.68	4.24	4.16	4.24	0.35	0.31	0.44	0.42	0.31	0.01
4	812	0.75	-0.45	0.53	-0.45	5.74	7.48	5.79	7.48	0.39	0.22	0.00	0.40	0.22	0.00
Approximation c.f. actual based on $\hat{\pi}_3 - \hat{\pi}_6$															
1	725	-0.09	0.07	0.16	0.07	4.38	4.00	4.21	4.00	0.27	0.36	0.02	0.32	0.36	0.37
2	723	0.01	0.13	0.19	0.13	3.82	4.25	4.12	4.25	0.29	0.44	0.00	0.37	0.44	0.09
3	706	0.28	0.10	0.20	0.10	4.26	4.01	4.60	4.01	0.45	0.46	0.75	0.48	0.46	0.67
4	815	0.37	0.19	0.37	0.19	5.55	4.71	5.55	4.71	0.39	0.44	0.16	0.39	0.44	0.16

Notes: Standard indicates the standard approximation for converting fixed-event to fixed-horizon forecasts given in equation (1); Adjusted indicates the adjusted approximation given in equations (15) and (18). Shaded rows are those where there is no approximation error involved. Based on data for all forecasters who contribute to the SPF panel at least 20 times during the sample period.

Finally, we compare the estimated persistence, in a manner analogous to the discussion of Table 2. Here the results are more modest. The gap between the persistence of the approximate inflation forecast and the actual inflation forecast declines by 5% with our adjustment. Further, the difference in persistence between the approximation and the actual data is statistically significant (at the 95% level) in three cases with the adjustment versus five cases without it. However, the level of the gap actually increases in more cases (five) than it decreases (three).

Taken together, our results indicate that our proposed adjustment improves the performance of a standard approximation that researchers have used to convert fixed-event inflation forecasts into fixed-horizon forecasts.

Graph 4: Weight on the anchor
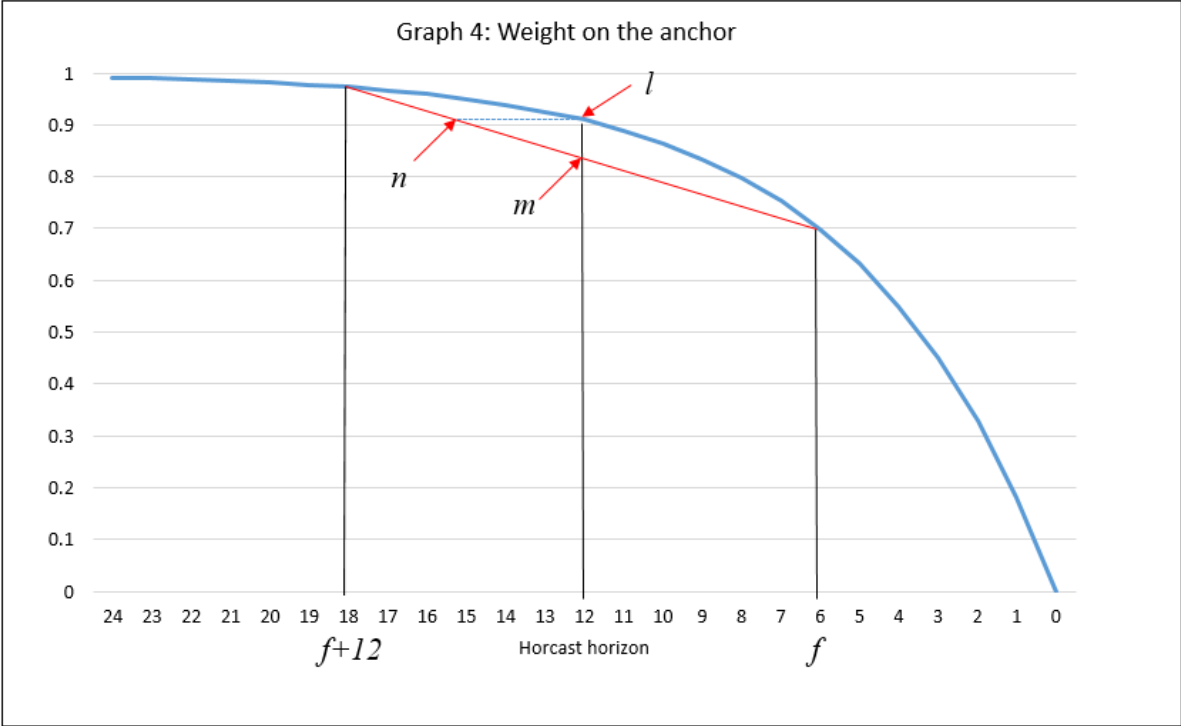
We can explain the effectiveness of our adjustment in the following way. Equation (17) implies that the weight on the inflation anchor varies with the forecast horizon in the manner displayed in Graph 4. Suppose that we wish to have a 12-month fixed-horizon forecast, but only have annual inflation forecasts, for two calendar years that each partially overlap with the 12-month horizon, indicated by f and $f+12$ on the graph, at our disposal. For the 12-month horizon inflation forecast, the weight on the anchor would be consistent with the displayed function where $h=12$ – ie at point l. Taking the standard approximation, the two annual forecasts are weighted by how far away f and $f+12$ are from 12. This results in the weight on the anchor indicated by m, on a straight line between the function evaluated at horizons f and $f+12$, evaluated at a horizon of 12 months. But because the

function is concave, this lies below l, which is the correct level for a 12-month horizon forecast. That is, the standard approximation results in too little weight on the anchor compared with what a 12-month fixed-horizon forecast would entail.

The effect of our adjustment is to use the estimate of the curvature of function (based on b) to move the weight on the anchor along the chord on the graph from point m to point n. That is, to increase the weight on the longer-horizon forecast (which is more heavily influenced by the anchor) and correspondingly decrease the weight on the shorter-horizon forecast. By doing so, we can remove the systematic bias present in the standard approximation. Our empirical example above illustrates how we can implement such an approach to improve the performance of the approximation, using real inflation forecast data.

Note, however, that this is not a complete fix for the problem of using the approximation, especially in terms of estimates of the persistence of inflation expectations. Clearly, our simple model of forecasts underlying the adjustment to the approximation is incomplete, and so our correction is only partial. One possibility for further work would be to model the error process in the expectations formation process, which might provide scope for further improvements. Additionally, our correction may be less effective if expectations are poorly anchored. In our model, that would be characterised by a low value of b and less curvature in the function in Graph 3, so that the gap between the standard approximation (point m) and our adjusted approximation (point n) would be more limited.

6. Conclusions

In this paper, we have shown that approximating fixed-horizon forecasts using fixed-event forecasts is problematic. In the SPF data that we examined, it results in a mean absolute approximation error of around 0.2–0.3 percentage points, or about 10% of the level of inflation. In addition, it results in statistically significant changes in both the variances and persistence of the resulting inflation forecast measures.

So what should researchers who desire a fixed-horizon forecast, but only have fixed-event forecasts available, do? We have shown one path to improve on the performance of the approximation, based on a model where inflation forecasts are a weighted average of an anchor and the latest available inflation forecast at the time that the forecasts are made, with the weight on the anchor declining as the forecast horizon shortens. Our model implies placing a relatively higher weight on the longer-horizon forecast than is the case using the standard approximation. Our proposed correction reduces the mean absolute error and means squared error of the approximation (by 13% and 21% respectively).

It also reduces the number of statistically significant differences between the approximate and actual inflation forecasts in terms of the variance and persistence. Further, the skewness and kurtosis of the distribution of adjusted approximate forecasts is more like that for the distribution of actual forecasts than when the standard approximation is used.

References

Capistrán, C and M Ramos-Francia (2010): "Does inflation targeting affect the dispersion of inflation expectations?" *Journal of Money, Credit and Banking* 42(1), 113-134.

Corbo, V, O Landerretche and K Schmitt-Hebbel (2001): "Assessing inflation targeting after a decade of world experience," *International Journal of Financial Economics* 6(4), 343-368.

Davies, A and K Lahiri (1995): "A new framework for analyzing survey forecasts using three-dimensional panel data," *Journal of Econometrics* 68(1), 205-227.

Davis, J (2014): "Inflation targeting and the anchoring of inflation expectations: cross-country evidence from consensus forecasts," *Federal Reserve Bank of Dallas Globalization and Monetary Policy Institute Working Paper* no 174.

Dovern, J and U Fritsche (2008): "Estimating fundamental cross-section dispersion form fixed event forecasts," *University of Hamburg Discussion Paper* no 1.

Dovern, J, U Fritsche and J Slacalek (2012): "Disagreement among forecasters in G7 countries," *Review of Economics and Statistics* 94(4), 1081-1096.

Gerlach, S (2007): "Interest rate setting by the ECB, 1999-2006: Words and deeds," *International Journal of Central Banking* 3(3), 1–45.

Heenan, G, M Peter and S Roger (2006): "Implementing inflation targeting: institutional arrangements, target design, and communications," *IMF Working Paper* no 06/278.

Johnson, D (2002): "The effect of inflation targeting on the behavior of expected inflation: evidence from an 11 country panel," *Journal of Monetary Economics* 49(8), 1521-1538.

Johnson, D (2003): "The effect of inflation targets on the level of expected inflation in five countries," *Review of Economics and Statistics* 85(4), 1076-1081.

Kortelainen, M, M Paloviita and M Viren (2011): "Observed inflation forecasts and the New Keynesian macro model," *Economics Letters* 112(1), 88–90.

Kozicki, S and P Tinsley (2012): "Effective use of survey inflation in estimating the evolution of expected inflation," *Journal of Money, Credit and Banking* 44(1), 145–169.

Lees, K (2016): "Assessing forecast performance," *Bulletin* 79(10), Reserve Bank of New Zealand, June.

Mavroeidis, S, M Plagborg-Møller and J Stock (2014): "Empirical evidence on inflation expectations in the New Keynesian Phillips Curve," *Journal of Economic Literature* 52(1), 124-188.

Mehrotra, A and J Yetman (2014): "Decaying expectations: what inflation forecasts tell us about the anchoring of inflation expectations," *BIS Working Papers* no 464.

Moreno, R and A Villar (2010): "Inflation expectations, persistence and monetary policy", *BIS Papers* no 49, 77–92.

Patton, A and A Timmermann (2010): "Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion," *Journal of Monetary Economics* 57(7), 803-820.

Siklos, P (2017): "What has publishing central bank forecasts accomplished? Central banks and Their competitors," Handbook of Central Banking, D Mayes, P Siklos and J Sturm (eds), Oxford University Press, forthcoming.

Siklos, P (2013): "Sources of disagreement in inflation forecasts: an international empirical investigation," *Journal of International Economics* 90(1), 218-231.

Sousa, R and J Yetman (2016): "Inflation expectations and monetary policy," *BIS Papers* no 89, 41-67.

Yetman, J (2017): "The evolution of inflation expectations in Canada and the US," *Canadian Journal of Economics* 50(3), 711-737.

Previous volumes in this series

All volumes are available on our website www.bis.org.