Inconsistent survey histograms and point forecasts revisited

Journal of Economic Behavior and Organization 236 (2025) 107097

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

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

Inconsistent survey histograms and point forecasts revisited[☆]

Michael P. Clements

ICMA Centre, Henley Business School, University of Reading, United Kingdom

ARTICLE INFO

Dataset link: https://data.mendeley.com/datasets/6zzwht77n4/1

JEL classification:
C23
C53
E66

Keywords:
Survey expectations
Histogram forecasts
Forecast revisions
Inflation
GDP growth

ABSTRACT

Past analyses of surveys of professional forecasters' histogram and point forecasts indicate that the two are not always consistent. The point forecasts are either systematically higher or lower than the corresponding histogram means, depending on whether we consider inflation or GDP growth. We consider whether inconsistencies are related to delayed updating of the histogram forecasts, or to the reaction of the two types of forecasts to new information, and whether inconsistent pairs typically imply less accurate point or histogram forecasts. We also re-consider explanations related to the complexity of the task on an extended dataset.

1. Introduction

Survey expectations of Professional Forecasters have been extensively used in the literature as a measure of otherwise unobserved agents' expectations. For example, they have contributed to the literature on the role of expectations shocks as sources of business-cycle variation in output and other macro-variables.[1] Much attention has also been directed to the impact of uncertainty shocks on economic activity, with surveys sometimes being used to measure uncertainty.[2] However, research on the properties of the survey expectations has uncovered some apparent anomalies between the survey respondents' subjective probability assessments, usually reported as histograms, and their point forecasts. Two apparent anomalies have been uncovered: "inconsistencies" between the histograms and point forecasts, and the so-called 'variance mismatch' between the histogram-based *ex ante* uncertainty and *ex post* uncertainty calculated using the realized values of the outturns. The anomalies do not necessarily impugn the value of the survey-based expectations and derived measures, such as the uncertainty series for business cycle analysis, but warrant further investigation into the "quality" of the Professionals' forecasts.[3] Zhao (2023) covers some of the ideas we consider in the

☆ An earlier version of this work was distributed as "Do professional macro-forecasters under or over-revise their histogram forecasts?". Helpful comments on that paper were gratefully received from participants in the Macro For Online Seminar December 2023, and in particular the comments of Alexander Glas are acknowledged. Helpful comments were also received from reviewers and the editors of this journal. All the computations were performed using code written in Gauss 21, other than the panel estimation which was done in Stata SE 18.5.

E-mail address: m.p.clements@reading.ac.uk.

[1] The literature has sought to distinguish expectations (and sentiment) shocks from news technology shocks and technology surprises using various identification schemes. See, e.g., Barsky and Sims (2012), Leduc and Sill (2013), Fève and Guay (2016), Levchenko and Pandalai-Nayar (2017) and Clements and Galvão (2021).

[2] See, for example, Bachmann et al. (2013), Haddow et al. (2013), Baker et al. (2016), Fajgelbaum et al. (2017) and Girardi and Reuter (2017), *inter alia.*

[3] Our focus is on professional forecasters, even though they may not be the most appropriate agents in any given instance. For example, for Phillips curve analysis Coibion and Gorodnichenko (2015) argue that firms' expectations are preferable to those of professional forecasters, and result in the Phillips curve

https://doi.org/10.1016/j.jebo.2025.107097
Received 5 September 2024; Received in revised form 26 May 2025; Accepted 3 June 2025

Available online 2 July 2025

context of the Survey of Consumer Expectations (SCE) inflation expectations, with a similar over-arching aim, of determining whether such agents' expectations are "reliable" and "accurate", and can "be used as a valuable source of information by monetary policy makers" (p.1734). These concerns are generally thought to have greater force for households/consumers, although recent work casts doubt on the value of the Professional Forecasters' histogram-based forecasts and risk assessments: whether they are meaningful or informative, sometimes in comparison to unconditional forecast distributions (see, e.g., Knüppel and Schultefrankenfeld (2012), Kenny et al. (2014), Clements (2018) and Knüppel and Schultefrankenfeld (2019)).

We provide an empirical analysis of the inconsistency anomaly.[4] This anomaly has been explored in the literature, and we update existing analyses to include more recent data. The more recent data from 2014:1–2019:4, and the change in histogram bin widths for these surveys, allows us to re-visit one possible explanation: computational complexity.

We consider potential explanations of the inconsistency motivated by the argument that histogram revisions are costly to produce, and have a lower profile than headline point projections, and so may not be updated as often as the point forecasts. The basic idea is consistent with Sim's notion of "rational inattention", that "people have limited information-processing capacity" (Sims, 2003), but the information rigidity literature on "sticky" and "noisy" information also results in delayed or partial updating (see, e.g., Coibion and Gorodnichenko (2012)).

We consider separately the properties of point forecasts and histograms corresponding to "small", "medium" and "large" histogram revisions, where revision size is determined by the Wasserstein distance between adjacent histograms of the same target, and suppose that small revisions are more likely to reflect delayed or incomplete updating. We use panel regressions with time fixed effects to control for confounding macro-effects.

We also make use of forecast error-revision regression used in the analysis of under- or over-reaction of (point) forecasts to new information to examine differences between the histogram mean forecasts, and the point forecasts, in terms of how they are revised in response to new information.

The final main contribution is an analysis of forecast inconsistencies and accuracy, again relying on panel regressions with time fixed effects to control for confounding effects.

A possible concern is that the inconsistencies arise because of measurement issues. The survey respondents provide their probability assessments in the form of histograms, that is, they assign probabilities to the variable of interest falling in different intervals. This is an incomplete representation of their underlying subjective distributions, and does not allow a unique determination of key statistics, such as moments and other quantities that might be of interest, without making additional assumptions.[5] As a consequence, the apparent inconsistencies we identify between the matched pairs of point and histogram forecasts might reflect the assumption made in calculating the histogram means.

We make use of the bounds approach to the determine inconsistencies between histogram means and point predictions of Engelberg et al. (2009). This does not require assumptions on the distribution of the probability mass within the histogram intervals, but at the cost of providing an interval estimate of the moment rather than a point estimate.[6]

The plan of the rest of the paper is as follows. Section 2 describes the forecast data set, and the approach taken to calculate moments from the reported histograms. Section 3 describes the measure of histogram revisions, given the relative novelty of the measure in the macro-forecasting literature. Section 4 describes the bounds approach to determining inconsistency. A matched pair of histogram and point forecasts are inconsistent when the point forecast is outside a bound on the mean calculated from the histogram, where the bound makes no assumptions about the underlying subjective distribution beyond those present in the histogram. We then provide estimates of inconsistency for the more recent historical period.[7] We also re-visit the role of computational complexity. Section 5 considers whether delayed updating of the histograms contributes to inconsistencies between the two forecasts of location. This is made operational by considering whether inconsistency is related to the size of the histogram revisions: whether smaller than average revisions, resulting from a failure to "'fully revise" the histograms, results in more inconsistencies. The evidence does not support this conjecture.

Rather than considering the whole histograms, Section 6 reports on the evidence for under- or over reaction of the two location estimates to news, and whether the two sets of forecasts are revised in line in response to news. Section 7 considers accuracy, and finds that the inconsistencies are generally due to inaccurate histogram forecasts rather than poor or inaccurate point forecasts. Finally, Section 8 offers some concluding remarks.

being better able to explain the 'missing disinflation' following the 2008–9 Financial Crisis. We consider professional forecasters, and in particular the U.S. SPF, because of the availability of histogram forecasts spanning a long time period.

[4] We do not consider the variance mis-match anomaly. Clements (2014a) showed that U.S. professional macro-forecasters are under-confident in their probability assessments concerning the outlook for inflation and real GDP growth at "within-year" horizons, that is, up to one-year ahead. The same has been reported by Knüppel and Schultefrankenfeld (2019) for the inflation forecasts of the Bank of England, the Banco Central do Brasil, the Magyar Nemzeti Bank and the Sveriges Riksbank.

[5] Bassetti et al. (2023) provide a critique of current approaches to inference on the probabilistic forecasts in surveys, including the approaches used in this paper. They highlight the failure to account for inference uncertainty: both the uncertainty about the form of the parametric distribution which is fitted to the histograms; and the uncertainty about the estimated parameters of the distribution. We leave an exploration of these issues to future research.

[6] Manski (2007) provides a detailed treatment of these issues.

[7] The early studies by Engelberg et al. (2009) and Clements (2009) necessarily only used survey data up to the mid 2000's.

2. Description of U.S. SPF data

The forecast data are taken from the U.S. Survey of Professional Forecasters (SPF)[8] (see, e.g., Croushore (1993), and the recent survey of professional forecasters by Clements et al. (2023)). The SPF is a quarterly survey of macroeconomic forecasters of the US economy that began in 1968, and is currently administered by the Philadelphia Fed. It is a staple for academic research on macro-expectations.

The survey constitutes a panel, albeit that each quarter there are joiners and leavers, as well as non-participation in certain periods (see, e.g., Figure 1 in Clements (2021).)

We use the 154 surveys from 1981:3 to 2019:4 inclusive. Although the survey starts in 1968:4, it is not clear what the target years are for many of the Q4 surveys in the 1970's.[9] In addition, there is a suspicion that some of the individual respondent identifiers used in the early years were subsequently used again for later participants. As a consequence, we exclude the 1970's. We also avoid the COVID-19 period, to ensure that the findings are not driven by those unprecedented times.

The survey elicits histogram forecasts of the annual rates of growth of real GDP and the GDP deflator, for the year of the survey relative to the previous year, as well as for the next year, relative to the current year. The 'fixed-event' nature of the histogram forecasts can be seen by noting that for annual inflation in 2016 (compared to 2015), for example, the 2016:1 survey yields a forecast of just under a year ahead, and the 2016:4 survey, a forecast with a horizon of just under a quarter. The next-year histograms in the surveys in 2015:1 to 2015:4 yield forecasts with approximate horizons of 8 down to 5 quarters of the same target — 2016 over 2015.

The location, size and number of the histogram bins to which the respondents attach probabilities has changed over the years, to reflect the changing ranges of values that respondents are expected to view as likely. (See the online documentation available at link to the SPF given above). A pertinent change for our purposes is the halving of the inflation bin widths to a half a percentage point for the 2014:1 to 2019:4 surveys, as discussed in Section 4.

Each survey also provides point predictions of quarterly growth rates for the current quarter and the next four quarters, providing a quarterly series of fixed-horizon forecasts, as well as point forecasts of the annual calendar year inflation and output growth rates on the same basis as the histogram forecasts. We make use of the point forecasts of the annual calendar year inflation and output growth rates. The matched point and density forecasts have been used by a number of authors, including (Engelberg et al., 2009; Clements, 2009, 2010, 2014b; Rich and Tracy, 2010, 2021), *inter alia*.

When we need a measure of the actual values of output growth and inflation, we make use of the "advance" estimates, available from the Philadelphia Fed Real-Time Data Set for Macroeconomists.[10]

We calculate individual mean estimates from the reported histograms either by calculating the means directly (see, e.g., Clements et al. (2023), p. 79–80) or by fitting normal distributions when the histogram has three or more bins with non-zero probability assigned, and symmetric (isosceles) triangle distributions when only one or two bins are used.[11] We use the normal rather than the generalized beta distribution suggested by Engelberg et al. (2009). The normal requires fewer choices than the generalized beta distribution, avoids attributing zero probability mass to some values (see, e.g., Bassetti et al. (2023), p. 447), and is a common choice. Another possibility would be the skew Student-t distribution applied to survey-based histograms by Ganics et al. (2024).

3. Measuring histogram revisions

A measure of distance between two densities $f(x)$ and $g(x)$ is given by Wasserstein distance (WD), defined as:

$$WD = \sqrt{\int_0^1 \left(F^{-1}(t) - G^{-1}(t)\right)^2 dt}$$

where $f(x)$ and $g(x)$ have inverse cumulative distribution functions $F^{-1}(t)$ and $G^{-1}(t)$. This can be calculated for the histograms directly, rather than for parametric distribution approximations to the histograms. This ensures that the estimate of the distance does not depend on extraneous assumptions (e.g., the choice of distribution). The calculation of WD does not require that the locations and sizes of the bins of the two histograms being compared match one another: this accommodates the changes that have occurred in the size and location of the U.S. SPF histogram bins over the period 1981:Q3 to 2019:Q4. The calculation of WD does assume that the probability mass is uniformly distributed within a bin. We closely follow (Verde and Irpino, 2013) in terms of notation and approach, and refer the reader to Arroyo and Maté (2009) and Verde and Irpino (2013), and the references therein, for further details.

8 https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/survey-of-professional-forecasters.

9 See the online documentation provided by the Philadelphia Fed, under "Documentation" from SPF line above. The problematic survey quarters are 1968.4, 1969.4, 1970.4, 1971.4, 1972.3, 1972.4, 1973.4, 1975.4, 1976.4, 1977.4, 1978.4, 1979.2, 1979.3, 1979.4. We have omitted the 1985.1, 1986.1 as there is some doubt about the targets in these surveys.

10 https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/real-time-data-set-for-macroeconomists.

11 When only one or two bins have mass, we follow the proposal of Engelberg et al. (2009) (Section 4.1, p.38), amended by Krüger and Pavlova (2024) to handle histograms with different bin widths. Engelberg et al. (2009) Fig. 2 implicitly assume the bin widths are equal at 1 for both bins. However, the US SPF widths were set at 2 prior to 1992:1 (before their sample period), and to a half for inflation from 2014:4 -, and to different widths for GDP growth from 2020:Q2. Krüger and Pavlova (2024) provide a simple generalization that allows for these possibilities.

Let us refer to histogram i by:

$$X_i = \left\{ \left(I_{1i}, p_{1i}\right), \left(I_{2i}, p_{2i}\right), \dots, \left(I_{h_i}, p_{h_i}\right) \right\}$$

where p_{si} is the probability assigned to bin s, $I_{si} = \left[a_{si}, b_{si}\right]$. Then define w_{si} as the cumulative probability up to and including bin s, so that $w_{0i} = 0$; and $w_{si} = \sum_{l=0}^{s} p_{li}$, for $s = 1, \dots, h_i$. The uniformity assumption implies the cdf:

$$F_i(x) = w_{(s-1)i} + p_{si}\frac{x - a_{si}}{b_{si} - a_{si}}$$

when $x \in I_{si}$ and $s > 0$, and $F_i(x) = 0$ if $x < a_{1i}$, and $F_i(x) = 1$ if $x > b_{hi}$. This implies the quantile function:

$$q_i(t) = F_i^{-1}(t) = a_{si} + \frac{t - w_{(s-1)i}}{w_{si} - w_{(s-1)i}}\left(b_{si} - a_{si}\right), \; 0 \le w_{(s-1)i} \le t \le w_{si} \le 1,$$

which is key to calculating the WD between histograms i and j.

Then, let $w = \left[w_0, \dots, w_m\right]$ contain the sorted and unique values of the union of $\left\{w_{s_i i}\right\}$, $s_i = 1, \dots, h_i$ and $\left\{w_{s_j j}\right\}$, $s_j = 1, \dots, h_j$. For each pair $\left(w_{l-1}, w_l\right)$, we use q_i and q_j to calculate 'new' bins, $I_{li}^* = \left[q_i\left(w_{l-1}\right); q_i\left(w_l\right)\right]$ and $I_{lj}^* = \left[q_j\left(w_{l-1}\right); q_j\left(w_l\right)\right]$. Computing the centres and radii of the new bins for each histogram yields:

$$\begin{aligned} c_{li}^* &= \left(q_i\left(w_l\right) + q_i\left(w_{l-1}\right)\right)/2, \quad r_{li}^* = \left(q_i\left(w_l\right) - q_i\left(w_{l-1}\right)\right)/2 \\ c_{lj}^* &= \left(q_j\left(w_j\right) + q_j\left(w_{l-1}\right)\right)/2, \quad r_{lj}^* = \left(q_j\left(w_l\right) - q_j\left(w_{l-1}\right)\right)/2. \end{aligned}$$

The squared WD between histograms i and j is:

$$WD^2\left(X_i, X_j\right) = \sum_{l=1}^{m}\left(w_l - w_{l-1}\right)\left[\left(c_{li}^* - c_{lj}^*\right)^2 + \frac{1}{3}\left(r_{li}^* - r_{lj}^*\right)^2\right].$$

For the U.S. SPF histograms described in Section 2, we adopt the notation that $H_{i,t,h}$ denotes the histogram of respondent i for the year t (i.e., the inflation rate, or output growth in year t, relative to year $t-1$) made in survey quarter $5-h$ of the target year t, for $h = \{1, 2, 3, 4\}$, and in survey quarter $9-h$ of the previous year for $h = \{5, 6, 7, 8\}$. We calculate Wasserstein distance to measure histogram revisions whenever a respondent i files a return to two adjacent surveys for the same target. We let $WD_{i,t,h}$ denote the revision between histograms $H_{i,t,h}$ and $H_{i,t,h+1}$.[12] Note that $WD_{i,t,h}$ draws on 'current-year' histogram forecasts for $h = 1, 2$ and $h = 3$; $WD_{i,t,4}$ compares the Q1 current-year forecast ($H_{i,t,4}$) to the Q4 survey in year $t-1$ next-year forecast ($H_{i,t,5}$); and for $h = 5$ to 7, the distance estimates are between next-year forecasts made in the year $t-1$.

There are few applications of WD in the macro-forecasting literature. (Rich and Tracy, 2021; Cumings-Menon et al., 2020) are notable exceptions. They use the Wasserstein metric to measure disagreement between respondents in terms of histogram forecasts. A novelty of our study is the use of WD to measure the revisions to individual respondents' histograms.

4. Inconsistencies between histograms and point forecasts: the Engelberg et al. (2009) bounds approach

In this section we set out the bounds approach of Engelberg et al. (2009). This approach is used to establish the extent of the inconsistencies, and in Section 4.1 we update previous estimates in the literature, and in Section 4.2 re-visit some of the explanations which have been proposed.

Engelberg et al. (2009) show that the matching point forecasts and histograms generally correspond, in that the point forecast is *usually* consistent with a measure of central tendency from the histogram. Determining whether the probability distributions (in the form of histograms) and point forecasts are consistent with one another is not straightforward, since it is unclear which, if any, of the three usual measures of central tendency the point forecast might be, and because the histogram provides an incomplete description of the individual's subjective probability density function. (Engelberg et al., 2009) calculate lower and upper bounds on the value of the histogram mean, avoiding the need to make any assumptions about how the histogram relates to the underlying distribution. If the point forecast falls within the bound, it is deemed consistent with the histogram mean. As an illustration, suppose the histogram attaches a probability of 0.5 that the variable will be between 2 and 3, 0.3 between 3 and 4, and 0.2 between 4 and 5. For the mean, the lower (upper) bound is calculated by assuming that all the probability lies at the lower (upper) limit of the histogram intervals. Then the lower bound on the mean is given by $0.5 \times 2 + 0.3 \times 3 + 0.2 \times 4 = 2.7$, and the upper bound by $0.5 \times 3 + 0.3 \times 4 + 0.2 \times 5 = 3.7$. (All the mean bounds are necessarily of length one when the bin widths are one). Similar bounds can be calculated for the mode and median. We begin by updating the earlier work on the U.S. SPF.[13]

12 Hence in terms of the notation at the beginning of this section, generic histograms X_i and X_j are replaced by $H_{i,t,h}$ and $H_{i,t,h+1}$ in order to calculate the size of the revision, $WD_{i,t,h}$.

13 Clements (2009, 2010, 2014b) for the U.S. SPF, García and Manzanares (2007) for the European Central Bank's Survey of Professional Forecasters, and Boero et al. (2008) for the Bank of England Survey of External Forecasters. See Clements (2019) (Chapter 6) for a review of the literature. More recently, Zhao (2023) has analysed the internal consistency of US households' inflation expectations reported as point and density forecasts in the New York Fed's Survey of Consumer Expectations. The Handbook of Economic Expectations by Bachmann et al. (2023) reviews the literature on the expectations of different agents.

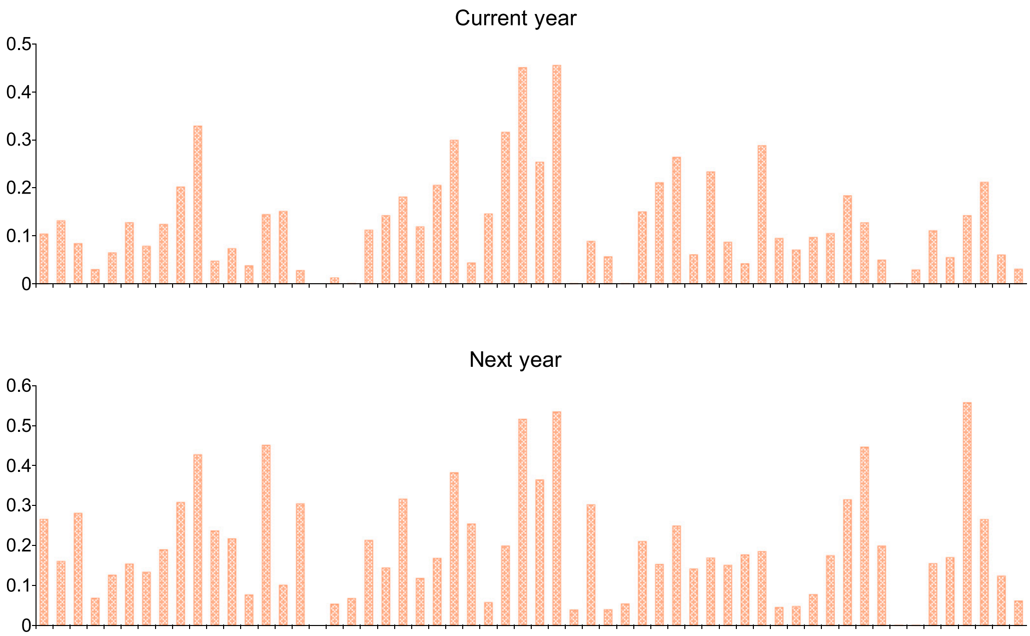

Fig. 1. Proportion of each respondent's pairs of within-year output growth forecasts which are inconsistent.

4.1. Evidence of inconsistencies

Table 1 updates the evidence for (in)consistency based on the bounds approach by including all the surveys from 1981:Q3 to 2019:Q4. The table reports the percentages of point forecasts which are within, below and above the bounds calculated on the mean, median and mode. The results aggregate over all respondents and all time periods, but are presented separately by horizon.

At the longest horizon $h = 8$, corresponding to a forecast from a first quarter survey of the following year, in excess of one fifth of forecast pairs for both variables are inconsistent with the interpretation that the point forecast corresponds to the mean. A preponderance of the inconsistencies are in the direction of the point forecasts being more 'optimistic' - the point forecast being above the upper bound on the histogram mean for output growth, and below the lower bound for inflation. The degree of inconsistency diminishes as the horizon shortens from $h = 8$ to $h = 1$, but the optimistic 'bias' remains evident, and even at $h = 1$ the inflation forecast is lower than is consistent with the histogram mean for over 1 in 10 pairs of forecasts.

The findings change a little for the mode and median, but we will focus on the mean, because the macro-literature typically adopts this interpretation of the point forecasts. The results in Table 1 are broadly in line with the findings in the literature. The findings of inconsistency are not much affected by including the post Financial Crisis period.

4.2. Existing explanations of inconsistency

Engelberg et al. (2009) and Clements (2009, 2010, 2011, 2014b) consider various explanations, including the role of the reporting practice — the rounding of probabilities either to convey ambiguity or to simplify communication, and that the point forecasts may be generated by respondents with asymmetric loss functions. It transpires that these do not explain the inconsistency, and are not considered further here.

One possibility is that the aggregate inconsistencies reported in Table 1 are primarily from a small number of respondents, or reflect particular atypical time periods. That is, the inconsistencies are not typical of the behaviour of the set of forecasters or of the period under study. Figs. 1 and 2 report the proportion of inconsistent forecasts for each individual respondent, and Figs. 3 and 4 the proportion of inconsistent pairs across respondents for each survey quarter. In each case, the top panel is for current-year forecasts, and the bottom panel next-year forecasts.

It is apparent that inconsistent pairs are not confined to a few respondents or survey quarters, albeit that there is variation in both dimensions.[14]

[14] There are a few outliers, in the sense of time periods which have a very large proportion of inconsistent pairs. One such is 90% for the current-year inflation forecasts made in 1990:Q1, while for the remaining 153 survey quarters the fraction is less than a half. In 1990:Q1 the number of active participants had fallen to a low level, before the Philadelphia Fed took over the survey from 1990:Q3, and there were only 13 matched pairs of forecasts.

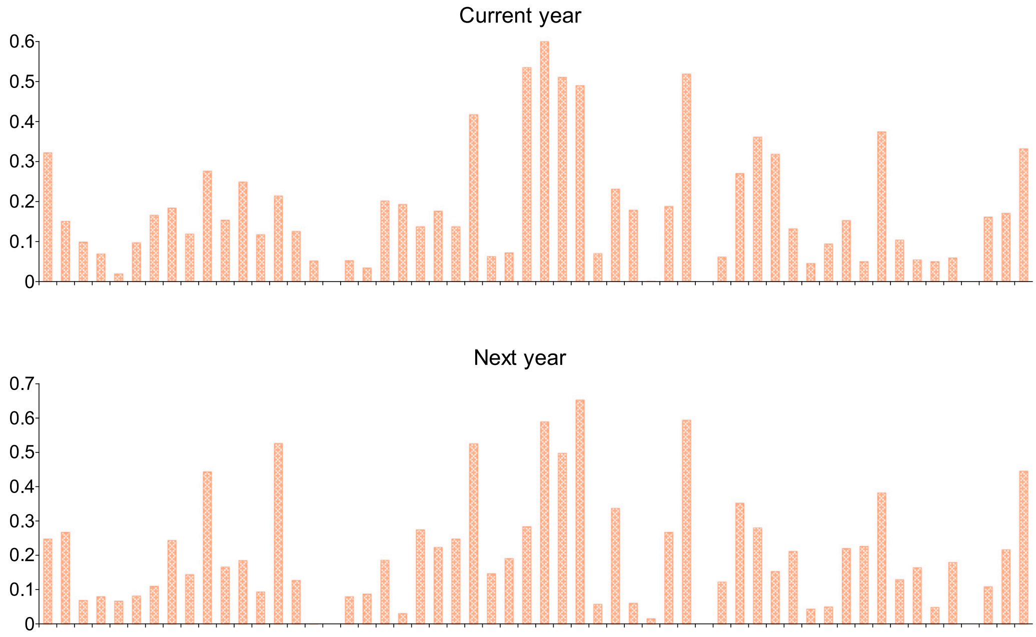

Fig. 2. Proportion of each respondent's pairs of within-year inflation forecasts which are inconsistent.

Clements (2010) suggests that respondents might be more likely to record inconsistent pairs of forecasts when the complexity of the task is greater. It seems reasonable to suppose that consistency is easier to achieve when the outcome is less uncertain, so that large probabilities are assigned to a narrowly defined range consisting of a small number of bins, rather than when outcomes are perceived as more uncertain and small probabilities are assigned to a wider range of values.[15] The SPF bin widths have changed over the historical period, becoming narrower with time, suggesting that the computational complexity of the task would have increased as typically more bins will be needed *other things being equal*. More precisely, from 1981:3 to 1991:Q4 the output growth bin widths were 2 percentage points, and this was halved to one percentage point from 1992:Q1 onwards.[16] For inflation, the 2 percentage point interval widths from 1981:Q3 to 1991:Q4 were also reduced to one percentage point in 1992:Q1, but were further reduced to half a percentage point over 2014:Q1 to 2019:Q4.

However, the historical changes to the histogram interval widths were not made exogenously, but reflect adaptation of the survey to the changing economic environment. While narrower bins might increase computational complexity by bringing more bins in to play, likely concurrent reductions in uncertainty would reduce the perceived range of values under consideration. Glas and Hartmann (2022) suggest some survey respondents may tend to 'report probability numbers for a fixed number of bins regardless of their precise definition', (p.990) suggesting that a finer gradation of bins might not tend to increase computational complexity much, but would reduce the scale of the histograms.

Figs. 3 and 4 provide some evidence that the degree of inconsistency has increased over the period, at least for inflation, although this is partly camouflaged by the seasonality — the typical decline in inconsistency from Q1 to Q4 surveys as the horizons shorten (for both current and next year targets). Table 2 perhaps provides a clearer picture, by reporting the proportions of mean inconsistencies (separately by survey quarter; equivalently, forecast horizon) for the sub-periods corresponding to different histogram widths, as well as reporting the average number of bins to which non-zero probabilities were assigned.

For neither variable do we observe an increase in the number of bins being used which is commensurate with the halving of bin widths. For output growth, the average number of bins used increased from around 3 to around 4.5, for all but the shortest horizon, when the widths halved. At the same time, the degree of consistency is lower in the later period for horizons 2 to 7 inclusive: see Table 2 Panel A. For inflation, the average number of bins used increased by just under 1 in 1992–2013, relative to 1981–1991, when the widths halved, and then again by a similar amount in 2014–2019 when they halved again: Table 2 Panel B. The table suggests the degree of inconsistency is higher during 2014:1 to 2019:4 compared to the two earlier periods — it is roughly twice as high for the half a percentage point intervals of 2014:1 to 2019:4 compared to the two-percentage point histogram intervals of 1981:3 to 1991:4. This suggests that increased complexity may have played a role as the bin widths have been reduced, not withstanding some support for the suggestion of Glas and Hartmann (2022).

[15] To underscore this point, Clements (2010) supposes perceived uncertainty about future outcomes is low, and in the extreme, all the probability mass is located in a single histogram bin. It then seems less likely that the histogram and point forecast would be inconsistent than if probabilities are attached to many bins. Inconsistency would require reporting a point forecast value that lies outside the range of the single bin.

[16] Further changes were made in response to the COVID-19 period, but our sample ends in 2019:Q4.

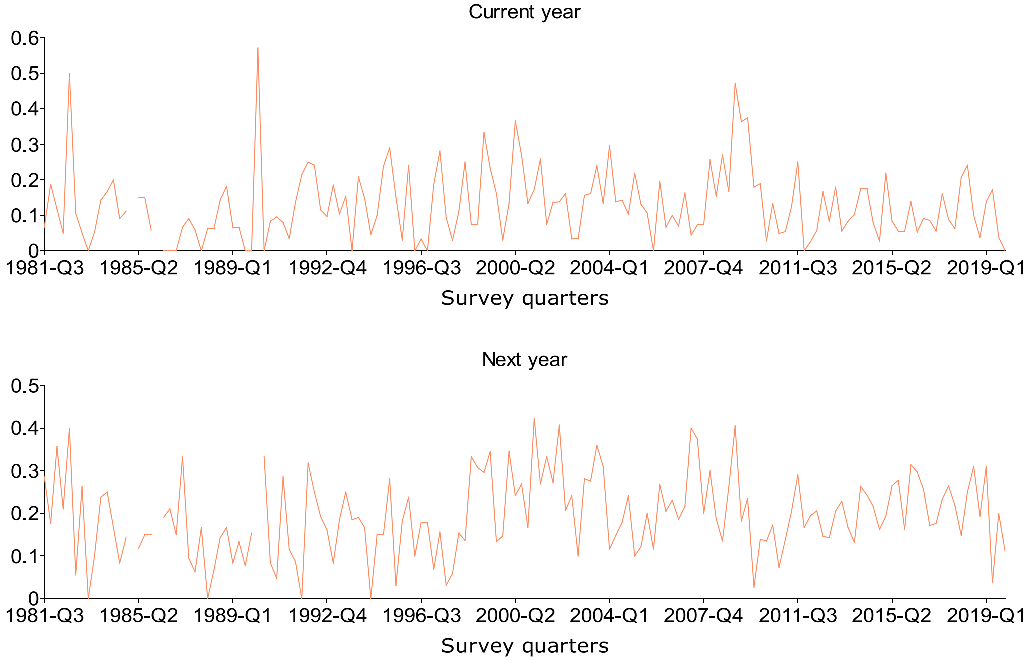

Fig. 3. Proportion of pairs of within-year output growth forecasts which are inconsistent at each survey quarter.

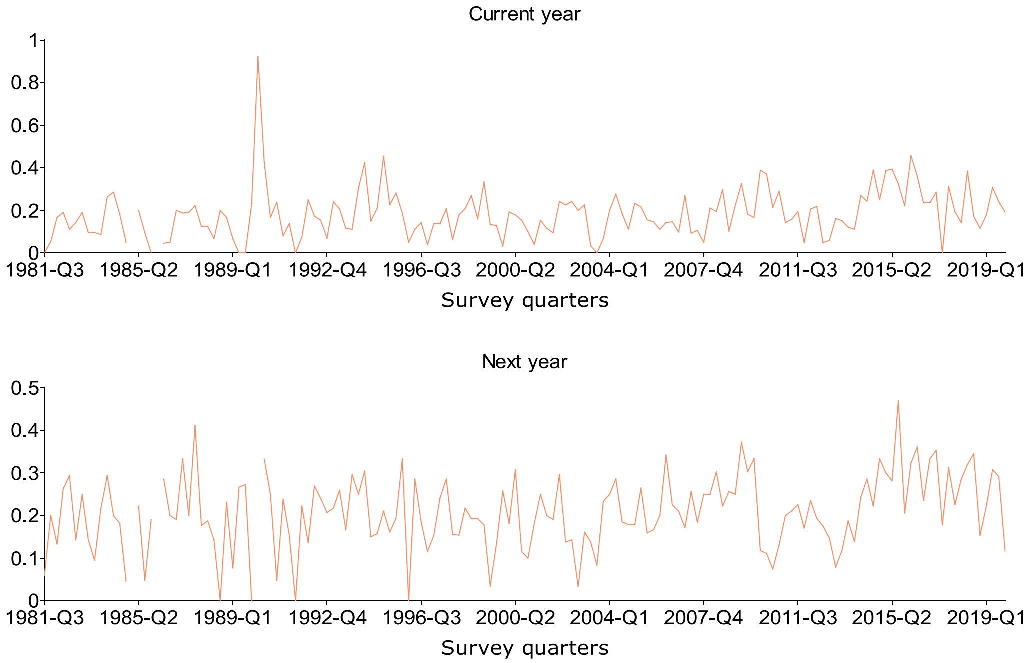

Fig. 4. Proportion of pairs of within-year inflation forecasts which are inconsistent at each survey quarter.

We conclude that computational complexity may play a role. The role of computational complexity can be rationalized by rational inattention ((Sims, 2003) and Woodford (2002)). That is, that consistent pairs of forecasts could be made by survey respondents, but doing so is costly — it is rational to report inconsistent pairs of forecasts. Consideration of the costs of producing the different types of forecasts suggests they may be revised differently over time as new information accrues, and in the following section we explicitly focus on revisions.

Table 1
Consistency of histograms and point forecasts: Bounds approach.

h	No.	Mean			Median			Mode		
		Within	Below	Above	Within	Below	Above	Within	Below	Above
Panel A. Output growth										
1	1130	92.6	3.4	4.1	91.6	4.3	4.2	91.8	4.3	3.9
2	1075	87.3	4.8	7.8	87.5	5.6	6.9	88.0	6.4	5.6
3	1132	85.8	5.2	9.0	85.8	6.9	7.2	88.0	7.3	4.7
4	1051	81.9	6.5	11.6	82.2	8.5	9.3	84.8	9.1	6.1
5	1137	84.1	4.2	11.7	86.3	6.0	7.7	88.2	6.8	5.0
6	1063	79.6	3.9	16.6	82.4	5.5	12.1	86.4	6.9	6.7
7	1068	79.9	4.1	16.0	81.4	6.9	11.6	86.4	8.2	5.4
8	955	76.9	4.7	18.4	80.6	5.2	14.2	87.1	6.2	6.7
Panel B. Inflation										
1	1105	87.0	10.7	2.4	86.3	11.0	2.7	86.3	10.7	3.0
2	1056	82.7	14.2	3.1	84.2	12.7	3.0	85.1	11.9	3.0
3	1086	79.6	18.1	2.3	81.7	16.2	2.1	83.6	14.2	2.2
4	1010	76.7	20.3	3.0	77.2	19.1	3.7	79.4	16.9	3.7
5	1108	80.6	16.4	3.0	83.0	14.1	2.9	84.7	12.2	3.1
6	1035	79.7	16.0	4.3	80.6	14.9	4.5	82.0	13.6	4.4
7	1040	76.2	18.8	5.0	78.5	16.1	5.4	80.6	14.6	4.8
8	927	78.9	16.0	5.2	78.7	14.5	6.7	80.3	13.2	6.4

The entries show the percentage of forecast pairs in each category. "Within" refers to the percentage of pairs for which the point forecast is within the bounds on the given histogram mean moment. "Below" refers to the percentage of pairs for which the point forecast is below the lower bound on the given histogram mean moment. "Above" refers to the percentage of pairs for which the point forecast is above the upper bound on the given histogram mean moment.

Table 2
Consistency of histograms and point forecasts: Mean-Bounds approach. Sub-sample analysis.

Panel A. Output growth

h	1981:3 to 1991:4					1992:1 to 2019:4				
	No.	Bins	Within	Below	Above	No.	Bins	Within	Below	Above
1	198	2.6	89.4	4.5	6.1	932	3.0	93.2	3.1	3.6
2	178	3.0	88.2	7.3	4.5	897	3.9	87.2	4.3	8.5
3	190	3.5	93.2	2.1	4.7	942	4.5	84.3	5.8	9.9
4	140	2.1	86.4	8.6	5.0	911	4.8	81.2	6.1	12.6
5	209	3.3	87.6	4.3	8.1	928	4.5	83.3	4.2	12.5
6	178	3.3	84.8	2.8	12.4	885	4.6	78.5	4.1	17.4
7	178	3.3	87.1	3.4	9.6	890	4.7	78.4	4.3	17.3
8	112	3.4	76.8	8.0	15.2	843	4.8	76.9	4.3	18.9

Panel B. Inflation

h	1981:3 to 1991:4					1992:1 to 2013:4					2014:1 to 2019:4				
	No.	Bins	W	B	A	No.	Bins	W	B	A	No.	Bins	W	B	A
1	211	2.2	88.2	9.0	2.8	705	2.8	88.7	9.5	1.8	189	3.7	79.4	16.9	3.7
2	189	2.7	90.5	8.5	1.1	677	3.5	83.6	13.9	2.5	190	4.4	71.6	21.1	7.4
3	192	2.9	84.9	13.0	2.1	708	3.8	80.2	17.2	2.5	186	4.6	71.5	26.9	1.6
4	143	3.2	77.6	20.3	2.1	673	4.0	78.3	19.5	2.2	194	4.9	70.6	23.2	6.2
5	219	3.0	84.5	13.2	2.3	703	3.6	80.5	17.6	1.8	186	4.2	76.3	15.6	8.1
6	184	3.0	84.8	13.6	1.6	662	3.8	81.4	16.0	2.6	189	4.7	68.8	18.5	12.7
7	180	3.0	78.9	18.9	2.2	678	3.9	76.8	20.5	2.7	182	4.5	70.9	12.6	16.5
8	117	3.0	77.8	18.8	3.4	620	3.8	81.5	16.3	2.3	190	4.7	71.1	13.2	15.8

The entries show the percentage of forecast pairs in each category. "Within" refers to the percentage of pairs for which the point forecast is within the bounds on the given histogram mean moment. "Below" refers to the percentage of pairs for which the point forecast is below the lower bound on the given histogram mean moment. "Above" refers to the percentage of pairs for which the point forecast is above the upper bound on the given histogram mean moment.
The column headed Bins denotes the average number of bins to which non-zero probabilities are attached, averaging across time periods and respondents.

5. Inconsistency and forecast revisions

In this section we focus on the role of forecast revisions or updating of the forecasts. Our prior belief is that histogram forecasts may not be revised in as timely a fashion as the point forecasts, because their production requires greater effort. Point forecasts are reported and discussed in the media, whereas histogram forecasts are not, and updating them from one survey quarter to the next may be costly if the production of such forecasts is not a routine part of the respondents' day jobs (and is undertaken solely at the behest of the SPF). That the production of histogram forecasts requires greater effort is consistent with findings in the literature pointing to real differences in accuracy across forecasters in terms of histogram forecasts, but not point forecasts. D'Agostino et al. (2012) report that there are no real differences between forecasters in terms of point forecast accuracy: 'better' forecasters are

Table 3
Proportion of times histograms are not changed.

h	Output growth	Inflation	Both
1	0.125	0.151	0.066
2	0.081	0.168	0.040
3	0.085	0.143	0.056
4	0.057	0.084	0.026
5	0.106	0.147	0.063
6	0.128	0.182	0.069
7	0.141	0.184	0.087

The table reports the proportion of histogram revisions with a WD value of 0 (no change on the previous survey submission of the same target year). Note: Failure to submit a histogram would result in a "missing" forecast: the previous period's forecast is not carried over.

better by 'chance'. However, adapting the approach of D'Agostino et al. (2012) to the U.S. SPF histogram forecasts, Clements (2020) concludes that some forecasters are superior to others.

If the histogram forecasts are updated less often, we would expect the histogram means and point forecasts to be out of synch when the point forecast is revised to reflect new information, but the histogram is not. Infrequent or delayed updating of the histogram forecasts relative to the point forecasts might be expected to result in inconsistencies between the two types of forecasts, but it is less clear that this would of itself explain the systematic tendency to relative optimism of the point forecasts compared to the histograms. For inflation, for example, this would require lower actual rates than the expected long-run inflation in the sample period, such that the inflation histogram forecasts, when they are revised, are revised down. The suggestion in Malmendier and Nagel (2016) that agents' inflation expectations may overweight their lifetime experiences of inflation is a potential explanation. According to which, for some agents the high inflation in the 1970s caused long-run expectations to be too high, with histogram forecasts indicating higher inflation because they are revised less often. A similar argument would need to be made for output growth, but in the reverse direction.

Other things being equal, we would also expect the histogram forecasts which are not updated to be less accurate than histogram forecasts which are updated (and are part of consistent pairs).

We explore these conjectures with a variety of approaches, either averaging over individuals and time periods, or using panel regressions to exploit the variation over individuals, while controlling for changing economic conditions with time fixed effects. Point forecast revisions are simply the difference between forecasts of the same target made in different (adjacent) survey quarters, whereas histogram revisions are measured by the Wasserstein distance between histograms of the same target, as explained in Section 3. As explained in Section 2, the U.S. SPF quarterly surveys provide forecasts of the current and next year annual growth rates for both the point forecasts and histograms, so that by a judicious matching of the current and next year forecasts, we obtain 8 forecasts, and 7 revisions, for each of the target years, 1982 to 2019.

5.1. Consistency and histogram revisions — Aggregate-level evidence

Before considering the relationship between revisions and inconsistency, Table 3 reports the proportion of histogram revisions with a WD value of 0, indicating the histogram is the same as it was in the previous period. The forecaster may have considered revising the histogram but decided no change is warranted, or simply filled in the same probabilities as in their previous return without reflection. Either way, the fewest no-change histograms occur at $h = 4$, corresponding to revisions between the Q1 survey current-year forecasts and the previous year's Q4 survey ("next year" forecasts). All but 6% and 8% of the histograms of output growth and inflation are revised at $h = 4$. The proportion not being revised increases to 12% and 15% at $h = 1$, and also increases as the horizon lengthens. The final column reports the proportion of times (over individuals and times periods, for a given horizon), that neither is revised. This is a relatively small proportion of the time (at a minimum of less than 3% for $h = 4$ and a maximum of 9% for $h = 7$), but several times higher than expected if not updating the output growth and inflation histograms were independent events. There is a tendency for "inattentiveness" to extend to both variables.

The histogram revisions reflect changes in both scale and location. This is evident in the Spearman rank correlations between the squared WD values and the squared changes in histogram standard deviations and squared changes in point forecasts (where the variation is across time and respondent for a given h).[17] Table 4 shows that the histogram revisions are highly correlated with the changes in the histogram means for all horizons, for both variables. The WD values are less correlated with the changes to the histogram standard deviations, but these correlations are nevertheless highly statistically significant. As expected, the WD values reflect changes in histogram location and scale.

Table 5 separately tabulates the proportion of point forecasts within the mean-bounds, below the lower bound, and above the upper bound, for small, medium and large histogram revisions (measured by WD). Small revisions are WD values less than the lower

[17] Verde and Irpino (2013) (eqn. (20)) formally show that the squared WD can be decomposed into the squared difference between the means of the two histograms, plus the squared difference between the standard deviations of two distributions, and a second-moment-related term. For this reason we look at the squared changes. We use the rank correlation coefficient to allow for non-linear dependence between the variables of interest.

Table 4
Spearman rank correlations between WD^2 and the squared revisions in the histogram standard deviations, and in the squared histogram mean revisions.

h	Output growth		Inflation	
	Rank correlation between WD^2 and:			
	Std dev change	Hist mean change	Std dev change	Hist mean change
1	0.686	0.925	0.720	0.922
	0.000	0.000	0.000	0.000
2	0.653	0.921	0.654	0.937
	0.000	0.000	0.000	0.000
3	0.614	0.935	0.628	0.928
	0.000	0.000	0.000	0.000
4	0.407	0.930	0.554	0.911
	0.000	0.000	0.000	0.000
5	0.508	0.933	0.654	0.931
	0.000	0.000	0.000	0.000
6	0.533	0.909	0.665	0.935
	0.000	0.000	0.000	0.000
7	0.679	0.871	0.712	0.868
	0.000	0.000	0.000	0.000

The correlation statistics (first row of each pair) are between WD^2 and the squared revisions in the histogram standard deviations, and in the squared histogram mean revisions. The second row elements for each horizon are the right tail probabilities of the Spearman test statistics, assuming the statistic is standard normal. As these are all zero to three decimal places, the statistics are all significant at conventional levels.

Table 5
Consistency of histogram means and point forecasts, by WD.

h	WD < Lower quartile			WD in IQ			WD > Upper quartile		
	Within	Below	Above	Within	Below	Above	Within	Below	Above
Panel A. Output growth									
1	93.7	2.7	3.6	97.3	1.8	0.9	82.4	7.2	10.4
2	88.0	6.5	5.6	90.3	3.2	6.5	80.2	8.3	11.5
3	83.8	7.5	8.8	88.5	3.7	7.8	84.1	7.0	8.8
4	82.7	5.3	11.9	85.5	3.6	11.0	74.9	15.1	10.1
5	86.5	3.6	9.9	85.7	2.7	11.6	77.6	7.9	14.5
6	81.1	2.8	16.1	81.6	4.4	14.1	77.7	4.3	18.0
7	79.9	5.1	15.0	81.4	3.2	15.4	75.5	5.7	18.9
Panel B. Inflation									
1	87.4	10.7	1.9	90.1	8.4	1.6	82.9	12.5	4.6
2	77.6	17.6	4.8	87.6	10.1	2.3	78.9	17.9	3.2
3	78.3	19.4	2.3	84.0	14.0	2.0	76.5	21.2	2.2
4	77.3	20.5	2.3	79.4	18.7	1.8	71.8	21.5	6.7
5	78.6	17.6	3.8	82.8	14.2	3.0	77.7	18.6	3.7
6	74.3	19.3	6.4	84.7	12.5	2.9	75.4	19.0	5.7
7	80.0	13.8	6.2	82.0	13.7	4.3	66.7	27.5	5.8

quartile (across all surveys and individuals) for a given h, medium are those between the lower and upper quartiles, and large are those in excess of the upper quartile. From Table 3 we know that more than a half of the lower quartile WD values for inflation at all horizons other than $h = 4$ are zero, indicating unchanged histograms.

There is no evidence for either output growth or inflation to support the hypothesis that a failure to update the histograms generates inconsistency, if we assume that "small" histogram revisions signify a failure to update to incorporate new information. Small WD values are not associated with less consistency compared to large values for output growth, and generally show consistency patterns similar to medium WD values. For output growth, we find that the proportion of consistent pairs is lowest for the large WD revisions (for all horizons bar one). For inflation, small revisions are generally no more likely than large revisions to be associated with inconsistency. For inflation, the degrees of inconsistency for the small and large revisions are similar (for all horizons other than $h = 7$), and generally higher than for medium-sized revisions.

Broadly, the ratio of optimistic to pessimistic point forecasts, relative to the histograms means, is not much affected by the size of the histogram revisions. That is, the majority of output growth inconsistent pairs result from the point forecast exceeding the upper bound, and for inflation, from the point forecast being below the lower bound.

The aggregate-level evidence suggests "small" histogram revisions are not correlated with more inconsistency. In the next sub-section we consider individual-level evidence relating to the relationship between inconsistency and histogram revisions, to control for the macro-outlook.

Table 6
Consistency of histogram means and point forecasts, by squared point forecast revisions. Output growth.

h	SQPR < Lower quartile			SQPR in IQ			SQPR > Upper quartile		
	Within	Below	Above	Within	Below	Above	Within	Below	Above
Panel A. Output growth									
1	95.2	3.1	1.8	93.8	2.2	4.0	88.9	5.3	5.8
2	90.1	4.3	5.6	88.8	4.7	6.5	80.7	7.7	11.6
3	88.3	4.0	7.6	88.8	3.7	7.5	78.8	10.2	11.1
4	87.8	3.2	9.0	83.0	5.0	12.0	78.0	10.6	11.5
5	86.6	1.8	11.6	84.6	2.9	12.6	80.7	9.4	9.9
6	83.0	0.5	16.5	81.9	2.6	15.5	74.3	9.5	16.2
7	78.3	3.4	18.2	81.8	2.5	15.8	78.6	9.0	12.4
Panel B. Inflation									
1	89.5	9.1	1.4	89.9	8.9	1.1	83.7	11.3	5.0
2	84.1	13.6	2.3	84.2	13.3	2.5	79.8	15.2	4.9
3	82.6	15.5	1.9	79.4	19.5	1.2	80.6	14.7	4.6
4	80.9	17.2	1.9	79.6	18.3	2.1	71.1	22.9	6.0
5	85.3	12.8	1.8	83.1	13.2	3.7	70.5	25.5	4.1
6	81.1	14.1	4.9	81.8	14.1	4.1	72.0	22.3	5.7
7	82.3	12.1	5.6	78.8	16.3	4.8	71.5	23.0	5.5

SQPR denotes "squared point revision".

5.2. Consistency and histogram revisions - Individual-level evidence

Section 5.1 suggests that small histogram revisions are not more likely to result in inconsistent pairs of histograms and point forecasts. However, we have not controlled for factors which might influence both revision size and inconsistency. It is plausible that when the macroeconomic outlook becomes more uncertain, respondents might make larger histogram revisions (to produce a histogram which reflects the greater uncertainty, as well as to revise the location of the distribution) and the tendency to generate inconsistent forecasts might also increase (possibly for the reasons mentioned in Section 4.2, such as the increased computational complexity of the task). In fact, if we tabulate the proportion of point forecasts within the mean-bounds, below the lower bound, and above the upper bound, by small, medium and large (squared) *point forecast revisions*, rather than histogram revisions, we find small point revisions are also less likely to result in inconsistency: see Table 6. The aggregate-level analysis may reflect the effects of a more uncertain macro-outlook rather than delayed updating.

To counter this, we report conditional time fixed effects logistic regressions in Table 7. The inclusion of the time effects means that it is only cross-sectional variation at each point in time which is captured in the estimates, not common variation in revision size and inconsistency induced by changes in the macro environment.

The dependent variable *hit* takes the value 1 for consistent pairs, and zero otherwise. The table shows that larger histogram revisions (measured by WD) reduce the probability of a consistent pair — the coefficients are generally negative. The absolute magnitudes vary over forecast horizon when we interact revision size with quarterly dummies but are negative for all horizons for inflation, and for all horizons but one for output growth. The coefficients are statistically significant when constrained to be equal across (within, or next-year) horizons in columns (1) and (3).

6. Point forecasts and histogram mean forecasts: over or under-reaction to news.

In this section we analyse the point forecasts and histogram mean forecasts with the commonly-used forecast error-forecast revision regressions, to determine whether the forecasts under- or over-react to news. We then adapt this framework to consider the discrepancy between the point forecast and the histogram mean forecast, in place of the forecast error. The approach is based on regressions such as (1) below. Such regressions have been run on consensus or aggregate expectations and also on individual-level expectations, to estimate the reaction to "news", the latter being captured by the forecast revision. For example, at the individual level, for a "rational" forecaster i who receives a (possibly) noisy signal, the population value of the slope parameter β_i in the forecast error-revision regression:

$$y_{t+h} - E_{i,t}\left[y_{t+h}\right] = \beta_i\left(E_{i,t}\left[y_{t+h}\right] - E_{i,t-1}\left[y_{t+h}\right]\right) + u_{i,t+h} \quad (1)$$

is $\beta_i = 0$. Here, y_{t+h} is the actual value at time $t+h$, where $E_{i,t}[.]$ is the forecast made at time t by individual i, and the use of the expectations operator indicates the conditional mean. The forecaster makes optimal use of her possibly noisy signal in determining $E_{i,t}\left[y_{t+h}\right]$ such that the resulting forecast error is not systematically related to the forecast revision between time $t-1$ and t, $E_{i,t}\left[y_{t+h}\right] - E_{i,t-1}\left[y_{t+h}\right]$.

Table 7
Conditional time fixed effect logit regressions of "hit" on WD, for output growth, and for inflation.

	Output growth				Inflation			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
	hit	*hit*	hit_{nx}	hit_{nx}	*hit*	*hit*	hit_{nx}	hit_{nx}
WD	−0.72*8				−0.626*			
	(−5.59)				(−5.08)			
WD × Q1		−0.292				−0.394		
		(−1.18)				(−1.71)		
WD × Q2		0.0142				−0.614*		
		(0.05)				(−2.72)		
WD × Q3		−0.817*				−0.408		
		(−2.93)				(−1.73)		
WD × Q4		−2.039*				−1.308*		
		(−5.88)				(−4.35)		
WD_{nx}			−0.573*				−0.773*	
			(−4.40)				(−5.69)	
WD × $Q2_{nx}$				−0.791*				−1.051*
				(−3.68)				(−4.47)
WD × $Q3_{nx}$				−0.349				−0.511*
				(−1.49)				(−2.10)
WD × $Q4_{nx}$				−0.536*				−0.719*
				(−2.25)				(−3.04)
Observations	2284	2284	1858	1858	2323	2323	1778	1778

t statistics in parentheses.
The dependent variable *hit* takes the value 1 for consistent pairs, and zero otherwise. The columns headed *hit* report results for current-year forecasts, namely, forecasts of the current-year annual average growth rate made in one of the four quarters of the current year. The columns headed hit_{nx} are for next-year growth rates relative to the current year survey quarters.
* Denotes statistically significantly different from zero at the 5% level.

In models such as this, estimated by pooling or fixed effects panel regression, such that the slope is the same across all respondents, both (Bordalo et al., 2020; Broer and Kohlhas, 2021) report finding that forecasters *over*-react to new information: $\beta < 0$.[18] In this literature, the "forecast" is typically the point forecast.

We begin by estimating (1) for both point forecasts and histogram mean forecasts separately. In the former case, the forecast error is based on the point forecast, and the explanatory variable is the revision to the point forecast, and in the latter case the forecast error is based on the histogram mean forecast, and the revision is to the histogram mean point forecast. We use panels with time-effects, and relative to (1) allow the slope coefficient to depend on the survey forecast, generating results for a range of forecast horizons h.

Table 8 reports the results for output growth and inflation, for point forecasts, and Table 9 is the corresponding tables for histogram mean forecasts. Matching the literature cited above, we find over-reaction, $\beta < 0$ for output growth and inflation point forecasts, with slope estimates of around −0.3 for "within-year forecasts" (corresponding to horizons of 1 to 4 quarters ahead), and of around −0.4 for "next-year" forecasts (corresponding to horizons of 5 to 8 quarters ahead). All are highly statistically significant.

The results for the histogram mean forecasts similarly indicate over-reaction for inflation, but less so for output growth, and especially for the within-year forecasts. The slope is statistically insignificant for the two shortest horizons (the current-quarter Q3 and Q4 surveys), and statistically significantly positive and negative the Q2 and Q1 horizons. The factors that cause over-reaction exert less of an influence for the output growth histogram forecasts.

A more direct test of whether point forecasts and histogram mean forecasts are revised in-step in response to new information is to replace the forecast error in (1) by the difference between the point forecast and histogram mean — the discrepancy. That is, we consider whether the degree of inconsistency is systematically related to the arrival of new information. This can be proxied either by the point forecast revision, or by the histogram mean revision.[19] In either case, the revision measures news regarding the 'location', and of interest is whether this is simultaneously reflected in both the respondent's histogram mean and point forecasts to the same extent.

Table 10 indicates that the revision to the point forecast (Rev) is positively related to the discrepancy ($Disc$) measured as the point forecast minus the histogram mean. The expected change in the location, as measured by the point forecast, is only partially

[18] Broer and Kohlhas (2021) attribute the over-reaction to private information as being due to 'absolute' over-confidence, and the over-reaction to public information as being due to 'relative' over-confidence. Bordalo et al. (2020) explain the over-reaction with a model of 'diagnostic' expectations (following Bordalo et al. (2019)). See also Kohlhas and Walther (2021) and Angeletos et al. (2020).

[19] Clements (2009, 2010) runs similar regressions, with the same dependent variable, that is, the difference between the point forecasts and histogram means. One of the explanatory variables he includes is the histogram variance (or standard deviation), rather than the revision, to determine whether the inconsistencies result from the point forecasts being deliberately "biased" because the respondents have asymmetric loss functions. His findings do not support asymmetric loss as an explanation of the inconsistencies.

Table 8
Point forecast error regressed on point forecast revisions. Output growth & Inflation.

	Output Growth				Inflation			
	(1) FE	(2) FE	(3) FE_{nx}	(4) FE_{nx}	(1) FE	(2) FE	(3) FE_{nx}	(4) FE_{nx}
Rev	−0.293* (−18.54)*				−0.257* (−17.12)			
Rev × Q1		−0.404* (−14.41)*				−0.295* (−12.11)		
Rev × Q2		−0.148* (−5.75)				−0.118* (−4.46)		
Rev × Q3		−0.405* (−12.61)				−0.343* (−10.27)		
Rev × Q4		−0.202* (−3.79)				−0.376* (−8.31)		
Rev_{nx}			−0.402* (−16.51)				−0.431* (−19.04)	
Rev × $Q2_{nx}$				−0.341* (−7.67)				−0.490* (−12.29)
Rev × $Q3_{nx}$				−0.480* (−12.19)				−0.446* (−13.15)
Rev × $Q4_{nx}$				−0.367* (−8.51)				−0.322* (−6.86)
Constant	0.00166 (0.30)	−0.00257 (−0.47)	−0.388* (−32.69)	−0.387* (−32.43)	−0.0373* (−7.87)	−0.0421* (−8.65)	−0.262* (−24.29)	−0.258* (−23.75)
Observations	2924	2924	2103	2103	2758	2758	1979	1979

t statistics in parentheses.
Time fixed effects are included. These account for the absence of level quarterly dummies. *FE* is point forecast error. *Rev* is the point forecast minus the point forecast of the same target made in the previous quarter. The $_{nx}$ subscript indicates next-year forecasts.
* Denotes statistically significantly different from zero at the 5% level.

reflected in the histogram forecast. There is an *under*-reaction of the histogram mean relative to the response of the point forecast. The explanatory variable *Rev* is multiplied by a quarter of the year dummy variable, apart from in columns (2) and (4), to allow the effects to depend on the forecast horizon (equivalently, the survey quarter). The estimates vary across survey quarter — but are always positive and statistically significant, except for output growth in Q4 (which is negative, but not significantly different from zero). A coefficient of 0.42 for the current-year Q4 inflation forecasts suggests a one percentage point increase in the annual inflation point forecast, relative to the forecast made in the third quarter, is associated with an increased discrepancy of 0.42% percentage points. Only for output growth in Q4, current year, do we not reject the hypothesis that the two forecasts are in lockstep. Apart from for Q4, the estimates for output growth are generally larger than for inflation, suggesting a higher degree of inconsistency for output growth than for inflation.

In an Appendix (Table 13) we report the results of measuring news by the histogram mean revision: these are similar to those for the point forecast revision in Table 10.

The time fixed effects mean the relationship between the discrepancy and point forecast revision is determined by variation across respondents in these variables at each point time. That is, a positive association results if those with a larger than (cross-section) average discrepancy at time t also have larger than average upward revisions.

The discrepancy — forecast revision regression results indicate that the larger the revisions, the larger the discrepancy, and the more inconsistencies. In the next section we consider whether the larger revisions and inconsistencies are more detrimental to the point forecasts or the histograms.

7. Accuracy of the histogram mean forecasts and point forecasts and inconsistencies

In this section we compare the histogram mean forecasts and point forecasts in terms of forecast accuracy. We have considered the consistency of the two, but a natural question is whether the inconsistencies we have identified tend to reflect poor or inaccurate histogram mean forecasts, or poor point forecasts.

Aggregate evidence on the accuracy of the two sets of forecasts is provided by Table 11. Unconditionally, that is, for all forecasts irrespective of whether they belong to consistent or inconsistent pairs, the point forecasts are markedly more accurate than the histogram means at the shorter horizons. For example, at $h = 1$ the ratio of the two is 0.3 for output growth, and 0.1 for inflation. The two align more closely at larger h, but for inflation the point forecasts remain over 10% more accurate. If we condition on consistent pairs, the relative advantage of the point forecasts is diminished.

Table 9
Histogram mean forecast error regressed on histogram mean forecast revisions. Output growth & Inflation.

	Output Growth				Inflation			
	(1) FE	(2) FE	(3) FE_{nx}	(4) FE_{nx}	(1) FE	(2) FE	(3) FE_{nx}	(4) FE_{nx}
Rev	−0.00157 (−0.06)				−0.433* (−28.53)			
Rev × Q1		−0.203* (−4.41)				−0.502* (−17.78)		
Rev × Q2		0.164* (3.84)				−0.375* (−12.78)		
Rev × Q3		0.0519 (0.93)				−0.428* (−13.68)		
Rev × Q4		−0.121 (−1.26)				−0.417* (−12.64)		
Rev_{nx}			−0.127* (−4.72)				−0.469* (−22.50)	
Rev × $Q2_{nx}$				−0.206* (−4.20)				−0.481* (−13.30)
Rev × $Q3_{nx}$				−0.101* (−2.33)				−0.414* (−11.21)
Rev × $Q4_{nx}$				−0.0839 (−1.73)				−0.508* (−14.41)
Constant	0.0117 (1.29)	0.0119 (1.30)	−0.284* (−21.20)	−0.282* (−20.86)	−0.210* (−29.28)	−0.213* (−29.47)	−0.452* (−42.62)	−0.452* (−42.41)
Observations	2688	2688	1927	1927	2613	2613	1931	1931

t statistics in parentheses.
Time fixed effects are included. These account for the absence of level quarterly dummies. *FE* is the histogram mean forecast error. *Rev* is the histogram mean forecast minus the histogram mean forecast of the same target made in the previous quarter. The $_{nx}$ subscript indicates next-year forecasts.
* Denotes statistically significantly different from zero at the 5% level.

This suggests the inconsistencies are caused by aberrant or poor histogram forecasts. Further evidence is provided by panel regressions with time fixed effects, as in Section 6. A measure of accuracy, S_{ith} is regressed on an indicator of consistency, hit_{ith}, which as before is multiplied by a quarter of the year dummy variable. hit_{ith} takes the value one if respondent i's histogram and point forecast are consistent, and zero otherwise.

For the point forecasts, S is the squared-forecast error. For the histograms, accuracy is measured by the ranked probability score (RPS: Epstein (1969)).[20]

The inclusion of time fixed effects controls for the possibility that less accurate forecasts may eventuate at more challenging times when inconsistencies are also more likely to arise.

Table 12 shows statistically significant negative coefficients on the hit indicator in all instances bar one for the histogram accuracy (ln RPS) regressions for output growth.[21] The evidence for the point forecasts is much weaker. For example, 'hit' is not significant at conventional levels for Q1 or Q4 current-year forecasts, and is only significant at the 1% level for Q3 next year forecasts. The findings are broadly similar for inflation.

Our interpretation of these findings is that histograms of inconsistent pairs are less accurate than those of consistent histograms (controlling for the macroeconomic conditions, via time effects). But the point forecasts of inconsistent pairs are little worse than those of consistent pairs. That is, inconsistent pairs reflect poor histograms rather than poor point forecasts.

8. Conclusions

Extending the historical coverage of earlier studies, we show that, over the nearly 40 year period of quarterly surveys from 1981 to 2019, there is a persistent tendency for survey respondents to report inconsistent histogram and point forecasts.

[20] The RPS is given by:

$$RPS = \sum_{k=1}^{K} \left(P^k - Y^k\right)^2 \tag{2}$$

where P^k is the cumulative probability assigned to bin k (i.e., $P^k = \sum_{s=1}^{k} p^s$, where p^s is the bin s probability). Similarly Y^k cumulates y^s, where y^s takes the value 1 when the actual value is in bin s, and zero otherwise. Relative to the quadratic probability score (QPS: (Brier, 1950)), RPS penalizes less severely density forecasts with probability close to the bin containing the actual value.

[21] *hit* is one for consistent forecast pairs, so a negative coefficient indicates consistent pairs are more accurate.

Table 10
Point forecast and histogram mean discrepancy regressed on point forecast revisions. Output growth & Inflation.

	Output Growth				Inflation			
	(1) Disc.	(2) Disc.	(3) Disc.$_{nx}$	(4) Disc.$_{nx}$	(1) Disc.	(2) Disc.	(3) Disc.$_{nx}$	(4) Disc.$_{nx}$
Rev	0.253*				0.186*			
	(10.92)				(8.04)			
Rev × Q1		0.197*				0.173*		
		(4.87)				(4.62)		
Rev × Q2		0.307*				0.133*		
		(8.21)				(3.22)		
Rev × Q3		0.365*				0.171*		
		(7.43)				(3.29)		
Rev × Q4		−0.113				0.411*		
		(−1.33)				(5.88)		
Rev$_{nx}$			0.270*				0.212*	
			(13.65)				(10.11)	
Rev × Q2$_{nx}$				0.117*				0.237*
				(3.31)				(6.35)
Rev × Q3$_{nx}$				0.381*				0.224*
				(12.12)				(7.18)
Rev × Q4$_{nx}$				0.280*				0.155*
				(7.92)				(3.56)
Constant	0.0116	0.0176*	0.105*	0.106*	−0.147*	−0.146*	−0.152*	−0.154*
	(1.46)	(2.19)	(10.76)	(10.92)	(−19.83)	(−19.16)	(−15.14)	(−15.19)
Observations	2747	2747	1963	1963	2591	2591	1840	1840

t statistics in parentheses.
Time fixed effects are included. These account for the absence of level quarterly dummies. *Disc* is point forecast minus the histogram mean. *Rev* is the point forecast minus the point forecast of the same target made in the previous quarter. The $_{nx}$ subscript indicates next-year forecasts.
* Denotes statistically significantly different from zero at the 5% level.

Table 11
Histogram mean and point forecast accuracy.

h				Conditional on consistency		
	Point	H.mean	Ratio	Point	H.mean	Ratio
	Output Growth					
1	0.010	0.032	0.325	0.010	0.023	0.444
2	0.046	0.056	0.819	0.043	0.046	0.948
3	0.064	0.122	0.522	0.063	0.092	0.688
4	0.282	0.328	0.858	0.290	0.305	0.951
5	0.468	0.501	0.935	0.450	0.465	0.968
6	0.563	0.594	0.947	0.525	0.546	0.961
7	0.810	0.794	1.020	0.796	0.795	1.001
8	0.906	0.984	0.921	0.829	0.835	0.992
	Inflation					
1	0.004	0.042	0.097	0.004	0.025	0.160
2	0.018	0.066	0.269	0.018	0.040	0.440
3	0.044	0.104	0.419	0.046	0.077	0.595
4	0.102	0.155	0.659	0.106	0.125	0.849
5	0.240	0.257	0.932	0.240	0.244	0.982
6	0.330	0.377	0.875	0.341	0.350	0.975
7	0.380	0.435	0.873	0.380	0.410	0.925
8	0.371	0.435	0.853	0.381	0.410	0.930

The table shows the forecast accuracy of the histogram means and point forecasts (on squared error), where we report the medians of the squared errors over time periods and individuals for a given horizon *h*.

Table 12
Panel regressions of forecast accuracy and inconsistency — Output growth & inflation.

	Output growth				Inflation			
	(1) lnRPS	(2) lnRPS_{nx}	(3) Ln e^2	(4) Ln $e^2{}_{nx}$	(1) lnRPS	(2) lnRPS_{nx}	(3) Ln e^2	(4) Ln $e^2{}_{nx}$
hit × Q1	−0.435* (−3.29)		−0.205 (−1.11)		−0.838* (−6.71)		−0.238 (−1.27)	
hit × Q2	−0.950* (−6.72)		−0.396* (−2.02)		−1.065* (−8.53)		−0.179 (−0.95)	
hit × Q3	−1.388* (−9.12)		−0.488* (−2.32)		−1.687* (−12.23)		−0.626* (−2.98)	
hit × Q4	−2.618* (−13.33)		−0.242 (−0.90)		−2.301* (−14.87)		−0.384 (−1.65)	
hit × Q1_{nx}		−0.120 (−1.72)		−0.331* (−2.12)		−0.163 (−1.87)		0.149 (0.81)
hit × Q2_{nx}		−0.323* (−4.63)		−0.324* (−2.07)		−0.422* (−5.18)		0.428* (2.47)
hit × Q3_{nx}		−0.218* (−3.07)		−0.716* (−4.50)		−0.399* (−4.55)		−0.112 (−0.61)
hit × Q4_{nx}		−0.457* (−6.04)		−0.235 (−1.39)		−0.495* (−5.78)		0.0435 (0.24)
Constant	−1.008* (−13.99)	−0.645* (−20.14)	−3.120* (−30.55)	−0.512* (−7.15)	−1.084* (−17.90)	−0.921* (−24.44)	−3.776* (−40.71)	−1.847* (−23.04)
Observations	3012	3044	3165	3052	2819	2875	2937	2892

t statistics in parentheses.
All columns include time fixed effects. Columns (1) and (2) are for current-year and next-year histogram forecast (RPS) accuracy, respectively, and columns (3) and (4) are for current and next-year point forecast (squared error) accuracy.
* Denotes statistically significantly different from zero at the 5% level.

We re-visit earlier explanations, such as the computational complexity involved in producing histogram forecasts. We find that computational complexity plays a role. If we consider the period 2014:1 to 2019:4 for inflation, for example, the bin widths were half a percentage point, down from one percentage point. The number of bins used increased, but did not double, as might have been expected. This was associated with more inconsistency, on average, at a time when the change to the bin widths likely reflected lower expected macro uncertainty.

We also investigate inconsistency and forecast revisions, and whether the inconsistencies reflect "poor" histogram forecasts or point forecasts (or neither).

Our prior belief was that failure to update the histogram forecasts may explain the inconsistencies between histogram means and point forecasts. Updating of histograms may not occur because producing a new histogram forecast is costly compared to issuing a point forecast.

Using a measure of histogram revision, we find that looking across all individuals and surveys for a given forecast horizon, it is not the case that the smaller histogram revisions are associated with more inconsistency. At least at the shorter forecast horizons, larger histogram revisions tend to be associated with a higher probability of histogram-point forecast inconsistency.

It may be the case that larger revisions are associated with more inconsistency because both revisions and inconsistency increase when the outlook becomes more uncertain. We report logit panel regressions which indicate that the probability of a consistent pair is reduced as the size of the histogram revision increases. By including time fixed effects we control for changes in the macro outlook.

Moving beyond the estimation of bounds on the histogram mean, we consider the under- / over-reaction of histogram mean forecasts and point forecasts, and find clear evidence of over-reaction of both for inflation. We show that the point forecast and histogram mean estimates do not move in lockstep, and the discrepancy between the two is correlated with the forecast revisions. We show that inconsistencies tend to be caused by poor histograms, rather than poor point forecasts, in the sense that the histograms of inconsistent pairs are less accurate than those of consistent pairs, but that the same is not true for point forecasts.

Histogram forecasts are more informative about an individual's beliefs, relative to the point forecasts, indicating how likely various ranges of outcomes are perceived to be. However, this appears to come at a cost in terms of the accuracy of the estimate of the conditional expectation.

The empirical analysis uses Wasserstein distance to measure the size of histogram revisions and these are then conditioned on in subsequent analyses. An interesting direction for future research would be to model the histogram revisions (possibly measured by Wasserstein distance) possibly with some of the variables suggested by Manzan (2021) in his analysis of histogram variance revisions.

Table 13
Point forecast and histogram mean discrepancy regressed on histogram mean forecast revisions. Output growth & Inflation.

	Output growth				Inflation			
	(1) Disc.	(2) Disc.	(3) Disc.$_{nx}$	(4) Disc.$_{nx}$	(1) Disc.	(2) Disc.	(3) Disc.$_{nx}$	(4) Disc.$_{nx}$
Rev	−0.392* (−27.54)				−0.355* (−24.86)			
Rev × Q1		−0.354* (−13.32)				−0.382* (−14.17)		
Rev × Q2		−0.271* (−9.24)				−0.313* (−11.32)		
Rev × Q3		−0.328* (−10.81)				−0.320* (−10.94)		
Rev × Q4		−0.583* (−21.77)				−0.410* (−13.37)		
Rev$_{nx}$			−0.274* (−15.54)				−0.319* (−16.61)	
Rev × Q2$_{nx}$				−0.281* (−9.16)				−0.306* (−9.24)
Rev × Q3$_{nx}$				−0.215* (−7.25)				−0.237* (−6.96)
Rev × Q4$_{nx}$				−0.332* (−10.63)				−0.405* (−12.52)
Constant	0.00117 (0.17)	0.00475 (0.68)	0.0588* (6.22)	0.0579* (6.11)	−0.174* (−25.57)	−0.176* (−25.66)	−0.176* (−18.01)	−0.177* (−18.07)
Observations	2729	2729	1988	1988	2585	2585	1883	1883

t statistics in parentheses.
Time fixed effects are included. These account for the absence of level quarterly dummies. *Disc* is point forecast minus the histogram mean. *Rev* is the histogram mean forecast minus the histogram mean forecast of the same target made in the previous quarter. The $_{nx}$ subscript indicates next-year forecasts.
* Denotes statistically significantly different from zero at the 5% level.

Declaration of competing interest

Author declares no conflict of interest.

Appendix

Table 13 is the companion table to Table 10 in the main text. Relative to Table 10, we use the histogram mean revision instead of the point forecast revision. The results are similar but there is a sign switch, because we are still measuring the "discrepancy" in the same way, that is, as the point forecast minus the histogram mean forecast.

Data availability

The data and code are available in the Mendeley Data Repository: https://data.mendeley.com/datasets/6zzwht77n4/1.

References

Angeletos, G.-M., Huo, Z., Sastry, K.A., 2020. Imperfect Macroeconomic Expectations: Evidence and Theory. In: NBER Macroeconomics Annual 2020. Vol. 35, National Bureau of Economic Research, Inc, pp. 1–86.
Arroyo, J., Maté, C., 2009. Forecasting histogram time series with k-nearest neighbours methods. Int. J. Forecast. 25 (1), 192–207.
Bachmann, R., Elstner, S., Sims, E.R., 2013. Uncertainty and economic activity: Evidence from business survey data. Am. Econ. J.: Macroecon. 5 (2), 217–249.
Bachmann, R., Topa, G., Klaauw, W., 2023. Handbook of Economic Expectations. Academic Press, Elsevier.
Baker, S.R., Bloom, N., Davis, S.J., 2016. Measuring economic policy uncertainty. Q. J. Econ. 131 (4), 1593–1636.
Barsky, R.B., Sims, E.R., 2012. Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence. Am. Econ. Rev. 102 (4), 1343–1377.
Bassetti, F., Casarin, R., Del Negro, M., 2023. Inference on probabilistic surveys in macroeconomics with an application to the evolution of uncertainty in the survey of professional forecasters during the COVID pandemic. In: Bachmann, R., Topa, G., Klaauw, W. (Eds.), Handbook of Economic Expectations. Academic Press, Elsevier, pp. 443–476, (Chapter 15).
Boero, G., Smith, J., Wallis, K.F., 2008. Evaluating a three-dimensional panel of point forecasts: The Bank of England survey of economic forecasters. Int. J. Forecast. 24, 354–367.
Bordalo, P., Gennaioli, N., Ma, Y., Shleifer, A., 2020. Overreaction in Macroeconomic Expectations. Am. Econ. Rev. 110 (9), 2748–2782.
Bordalo, P., Gennaioli, N., Porta, R.L., Shleifer, A., 2019. Diagnostic Expectations and Stock Returns. J. Financ. 74 (6), 2839–2874.
Brier, G.W., 1950. Verification of forecasts expressed in terms of probability. Mon. Weather Rev. 75, 1–3.

Broer, T., Kohlhas, A., 2021. Forecaster (Mis-)Behavior. CEPR Discussion Papers 12898/2018, C.E.P.R. Discussion Papers.

Clements, M.P., 2009. Internal consistency of survey respondents' forecasts: Evidence based on the survey of professional forecasters. In: Castle, J.L., Shephard, N. (Eds.), The Methodology and Practice of Econometrics. A Festschrift in Honour of David F. Hendry. Oxford University Press, Oxford, pp. 206–226, (Chapter 8).

Clements, M.P., 2010. Explanations of the inconsistencies in survey respondents forecasts. Eur. Econ. Rev. 54 (4), 536–549.

Clements, M.P., 2011. An empirical investigation of the effects of rounding on the SPF probabilities of decline and output growth histograms. J. Money Credit. Bank. 43 (1), 207–220.

Clements, M.P., 2014a. Forecast uncertainty - Ex Ante and Ex Post: US inflation and output growth. J. Bus. Econom. Statist. 32 (2), 206–216. http://dx.doi.org/10.1080/07350015.2013.859618.

Clements, M.P., 2014b. Probability distributions or point predictions? Survey forecasts of US output growth and inflation. Int. J. Forecast. 30 (1), 99–117. http://dx.doi.org/10.1016/j.ijforecast.2013.07.010.

Clements, M.P., 2018. Are macroeconomic density forecasts informative? Int. J. Forecast. 34, 181–198.

Clements, M.P., 2019. Macroeconomic Survey Expectations. In: Palgrave Texts in Econometrics, Palgrave Macmillan, http://dx.doi.org/10.1007/978-3-319-97223-7.

Clements, M.P., 2020. Are some Forecasters' Probability Assessments of Macro Variables Better Than those of Others? Econometrics 8, 16, doi:10.3390/econometrics8020016.

Clements, M.P., 2021. Do survey joiners and leavers differ from regular participants? The US SPF GDP growth and inflation forecasts. Int. J. Forecast. 37 (2), 634–646.

Clements, M.P., Galvão, A.B., 2021. Measuring the effects of expectations shocks. J. Econom. Dynam. Control 124 (C).

Clements, M.P., Rich, R., Tracy, J., 2023. Surveys of professionals. In: Handbook of Economic Expectations. Academic Press, Elsevier, pp. 71–106, (Chapter 3).

Coibion, O., Gorodnichenko, Y., 2012. What can survey forecasts tell us about information rigidities? J. Polit. Econ. 120 (1), 116–159.

Coibion, O., Gorodnichenko, Y., 2015. Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation. Am. Econ. J.: Macroecon. 7 (1), 197–232.

Croushore, D., 1993. Introducing: The survey of professional forecasters. Fed. Reserv. Bank Phila. Bus. Rev. Novemb. 3–15.

Cumings-Menon, R., Shin, M., Sill, K., 2020. Measuring Disagreement in Probabilistic and Density Forecasts. Working Papers 21–03, Federal Reserve Bank of Philadelphia.

D'Agostino, A., McQuinn, K., Whelan, K., 2012. Are some forecasters really better than others? J. Money Credit. Bank. 44 (4), 715–732.

Engelberg, J., Manski, C.F., Williams, J., 2009. Comparing the point predictions and subjective probability distributions of professional forecasters. J. Bus. Econom. Statist. 27 (1), 30–41.

Epstein, E.S., 1969. A scoring system for probability forecasts of ranked categories. J. Appl. Meteorol. 8, 985–987.

Fajgelbaum, P.D., Schaal, E., Taschereau-Dumouchel, M., 2017. Uncertainty traps. Q. J. Econ. 132 (4), 1641–1692.

Fève, P., Guay, A., 2016. Sentiments in SVARs. Working Papers TSE-656, Toulouse School of Economics (TSE).

Ganics, G., Rossi, B., Sekhposyan, T., 2024. From fixed-event to fixed-horizon density forecasts: Obtaining measures of multihorizon uncertainty from survey density forecasts. J. Money Credit. Bank. 56 (7), 1675–1704.

García, J.A., Manzanares, A., 2007. Reporting Biases and Survey Results: Evidence from European professional forecasters. ECB Working Paper No. 836, European Central Bank, Frankfurt.

Girardi, A., Reuter, A., 2017. New uncertainty measures for the EURO area uisng survey data. Oxf. Econ. Pap. 69, 278–300.

Glas, A., Hartmann, M., 2022. Uncertainty measures from partially rounded probabilistic forecast surveys. Quant. Econ. 13 (3), 979–1022. http://dx.doi.org/10.3982/QE1703.

Haddow, A., Hare, C., Hooley, J., Shakir, T., 2013. Macroeconomic uncertainty: what is it, how can we measure it and why does it matter? Bank Engl. Q. Bull. 100–109, 2013 Q2.

Kenny, G., Kostka, T., Masera, F., 2014. How Informative are the Subjective Density Forecasts of Macroeconomists? J. Forecast. 33 (3), 163–185.

Knüppel, M., Schultefrankenfeld, G., 2012. How Informative Are Central Bank Assessments of Macroeconomic Risks? Int. J. Cent. Bank. 8 (3), 87–139.

Knüppel, M., Schultefrankenfeld, G., 2019. Assessing the uncertainty in central banks' inflation outlooks. Int. J. Forecast. 35, 1748–1769.

Kohlhas, A.N., Walther, A., 2021. Asymmetric Attention. Am. Econ. Rev. 111 (9), 2879–2925.

Krüger, F., Pavlova, L., 2024. Quantifying subjective uncertainty in survey expectations. Int. J. Forecast. 40 (2), 796–810.

Leduc, S., Sill, K., 2013. Expectations and Economic Fluctuations: An Analysis Using Survey Data. Rev. Econ. Stat. 95 (4), 1352–1367.

Levchenko, A.A., Pandalai-Nayar, N., 2017. TFP, News, and 'Sentiments': The International Transmission of Business Cycles. Working Papers, University of Michigan.

Malmendier, U., Nagel, S., 2016. Learning from inflation experiences. Q. J. Econ. 131 (1), 53–87.

Manski, C.F., 2007. Identification for Prediction and Decision. Harvard Univerity Press, Cambridge, Massachusetts.

Manzan, S., 2021. Are professional forecasters Bayesian? J. Econom. Dynam. Control 123 (C).

Rich, R., Tracy, J., 2010. The relationships among expected inflation, disagreement, and uncertainty: Evidence from matched point and density forecasts. Rev. Econ. Stat. 92 (1), 200–207.

Rich, R., Tracy, J., 2021. A Closer Look at the Behavior of Uncertainty and Disagreement: Micro Evidence from the Euro Area. J. Money Credit. Bank. 53 (1), 233–253.

Sims, C.A., 2003. Implications of rational inattention. J. Monet. Econ. 50, 665–690.

Verde, R., Irpino, A., 2013. Multiple linear regression for histogram data using least squares of quantile functions: a two-components model. In: Lechevallier, Y., Saporta, G., Guan, R., Wang, H. (Eds.), Advances in Theory and Applications of High Dimensional and Symbolic Data Analysis. In: Revue des Nouvelles Technologies de l'Information, Vol. RNTI-E-25.

Woodford, M., 2002. Imperfect common knowledge and the effects of monetary policy. In: Aghion, P., Frydman, R., Stiglitz, J., Woodford, M. (Eds.), Knowledge, Information, and Expectations in Modern Macroeconomics: in Honor of Edmund Phelps. Princeton University Press, pp. 25–58.

Zhao, Y., 2023. Internal consistency of household inflation expectations: Point forecasts vs. density forecasts. Int. J. Forecast. 39 (4), 1713–1735.