

A Bayesian Approach for Inference on Probabilistic Surveys

Marco Del Negro*

*Federal Reserve Bank
of New York*

Roberto Casarin

Ca' Foscari University of Venice

Federico Bassetti

Polytechnic University of Milan

July 1, 2023

Abstract

We propose a nonparametric Bayesian approach for conducting inference on probabilistic surveys. We use this approach to study whether US Survey of Professional Forecasters density projections for output growth and inflation are consistent with the noisy rational expectations hypothesis. We find that in contrast to theory for horizons close to two years there is no relationship whatsoever between subjective uncertainty and forecast accuracy for output growth density projections, both across forecasters and over time, and only a mild relationship for inflation projections. As the horizons shortens, the relationship becomes one-to-one, as the theory would predict.

JEL CLASSIFICATION: C1, C11, C13, C15, C32, C58, G12, G13, G15

KEY WORDS: Bayesian inference, Bayesian Nonparametric, Survey of Professional Forecasters, Noisy Rational Expectations.

*Correspondence: Marco Del Negro (marco.delnegro@ny.frb.org): Research Department, Federal Reserve Bank of New York, 33 Liberty Street, New York NY 10045. Other Correspondence: Roberto Casarin (r.casarin@unive.it) and Federico Bassetti (federico.bassetti@polimi.it). This project was begun with Francesco Ravazzolo, whom we thank for many useful conversations. We thank Alissa Johnson for terrific research assistance. We are also grateful for insightful feedback to the participants of several seminars, workshops, and conferences, and Michael Clements, Mikkel Plagborg-Møller, and Christopher Sims, for their useful feedback. The views expressed in this paper do not necessarily reflect those of the Federal Reserve Bank of New York or the Federal Reserve System.

I Introduction

The pioneering work of Manski (2004) highlighted the benefits of probabilistic surveys compared to surveys that only ask respondents for their point projections: probabilistic surveys simply provide a wealth of information that is not included in point projections.¹ As Potter (2016) writes, “in a world characterized by pervasive uncertainty, density forecasts provide a comprehensive representation of respondents’ views about possible future outcomes for the variables of interest.” Given the respondents’ density forecasts, the econometrician can compute numerous objects of interest, such as the mean, the median, the variance, the skewness, the interquartile range, *et cetera*.

Except that survey respondents do not provide us with density forecasts. For most surveys concerning continuous variables, they only provide the percent chance that the variable of interest (e.g., inflation over the next year) would fall within different pre-specified contiguous ranges or bins. That is, the information we have consists in the integral of the forecast density over these bins, or equivalently, in a few points of the cumulative density function (CDF). In order to extract most quantities of interest, standard practice consists in postulating a parametric form for the forecast distribution and computing its parameters by minimizing the distance between the observed CDF points and those implied by the assumed distribution, which is often either a step-wise uniform (Zarnowitz and Lambros, 1987), a Gaussian (Giordani and Soderlind, 2003), or a generalized beta distribution (Engelberg et al., 2009).²

In this paper we propose a Bayesian nonparametric approach for the estimation of the survey respondents’ forecast densities.³ The approach starts by making parametric assump-

¹Indeed, a number of recent surveys, including the Federal Reserve Bank of New York’s Survey of Consumer Expectations and Survey of Primary Dealers and Market Participants, the Atlanta Fed’s Survey of Business Uncertainty and Business Inflation Expectations Survey, the ECB’s Survey of Professional Forecasters, and the the Bank of England’s Survey of External Forecasters, rely heavily on probabilistic questions.

²For a few quantities of interest, such as the median, one can compute nonparametric bounds as in Engelberg et al. (2009). We are well aware that this perspective on how forecasters respond to probabilistic surveys—namely, that they assign probabilities to bins using an underlying predictive distribution, possibly adding some noise because of rounding or other reasons—may be challenged in favor of more a behavioral alternative. However, this is explicitly or implicitly the assumption made by the existing literature when extracting an underlying distribution, whether a normal or a beta, and constructing measures of uncertainty. Moreover, it is the assumption that underlies the (noisy) rational expectations hypothesis, which we test in our application.

³In economics, the Bayesian nonparametric approach so far has applied to the analysis of treatment effects

tions on the mapping between the predictive distribution of forecasters and the bin probabilities they report, where this mapping explicitly allows for the introduction of noise in the reporting (due to, e.g., rounding toward zero). We then relax this parametric model by embedding it into the more general Bayesian nonparametric framework, thereby amending the potential misspecification associated with the parametric assumptions. This is because, loosely speaking, Bayesian nonparametrics replaces any parametric model with a potentially infinite mixture of such models, attaining more flexibility while at the same time using the information from the cross-section of forecasters to estimate the parameters of the mixture components. Intuitively, each mixture component corresponds to a forecaster type (e.g., low/high variance; optimists/pessimists; low/high noise; *et cetera*, and combinations thereof). As long as the number of types grows more slowly than the number of forecasters, there is enough information to estimate the parameters corresponding to each type.

Our approach differs from existing methods in a few important dimensions. First, it allows for full-fledged inference regarding the mapping between data and objects of interest, in the sense that it generates a posterior probability for these objects. While current approaches provide point estimates for, say, measures of the scale of the predictive densities like the variance, they do not supply any assessment of the uncertainty surrounding these estimates, which is often large given the limited information provided by the survey responses. Second, inference conducted using a specific parametric distribution can be naturally sensitive to the choice of the distribution, or the choice of the mapping between the distribution and the reported bin probabilities. The nonparametric nature of our approach provides some robustness to misspecification regarding these parametric assumptions. Last, our approach conducts inference jointly across survey respondents, that is, using the entire cross-section instead of being applied to each respondent separately. As hinted above, this joint inference allows for partial information pooling across forecasters thereby improving the precision of the inference, making it possible to obtain some consistency results when the number of forecasters grows to infinity.

We use this approach to address the question of whether US Survey of Professional Forecasters (SPF) density forecasts are consistent with the noisy rational expectations hypothesis (see, for instance, Coibion and Gorodnichenko, 2012, 2015). According to this hypothesis,

 (Chib and Hamilton, 2002), autoregressive panel data (Hirano, 2002; Gu and Koenker, 2017; Liu, 2021), time series (Bassetti et al., 2014), stochastic production frontiers models (Griffin and Steel, 2004), unemployment duration (Burda et al., 2015), and finance (Griffin, 2011, and Jensen and Maheu, 2010). Griffin et al. (2011) provide an intuitive description of the approach and a survey of this literature.

forecasters receive both public and private signals about the state of the economy. The precision of forecasters' signals, both public and private, ought to be reflected in equal measure in their density forecasts and, under rational expectations, in their *ex-post* forecast accuracy, both in the cross-section and over time. For example, if the economy becomes more uncertain and the precision deteriorates, this should be reflected in both higher subjective uncertainty and worse *ex-post* forecast errors. In fact, we find that for horizons close to two years there is no relationship whatsoever between subjective uncertainty and *ex-post* forecast accuracy for output growth density projections, and only a very mild relationship for inflation projections. As the horizons shortens, the relationship becomes one-to-one, in accordance with the theory. These findings suggest that forecasters do not correctly anticipate periods of macroeconomic uncertainty, except for very short horizons. Notably, this finding is robust to the exclusion of the Covid period.

The outline of the paper is as follows. Section II presents the inference problem, briefly describes current approaches, and formally discusses our Bayesian nonparametric approach. Section III first provides a few examples of how our approach differs from current practice and then discusses the relationship between subjective uncertainty and forecast accuracy. Section IV concludes pointing out some of the limitations of the analysis and discussing avenues for further research.


II Inference for Probabilistic Surveys

In this section we start by providing a short introduction to probabilistic survey data focusing on those features that are relevant for this analysis, and in the process describe the SPF data used in our application. Then we briefly discuss the approaches used so far for translating the information provided by the respondents into subjective predictive distributions and objects of interest, such as measures of uncertainty. The rest of the section is devoted to the description of our approach.

II.A The inference problem and current approaches

Probabilistic forecasts such as those elicited by the Philadelphia Fed as part of the SPF take the form of probabilities assigned to bins: the percent chance that the variable of interest, e.g., inflation, falls within different contiguous ranges, where these ranges are pre-specified

Figure 1: Probability forecasts for selected examples



Note: Each panel displays the forecast probabilities z_{ij} , $j = 1, \dots, J$ (step-wise solid lines) for a given forecaster i (forecaster number shown in parentheses) and the bin bounds (black ticks, horizontal axis).

by the survey designer.⁴ Formally, for each forecaster $i = 1, \dots, n$ the data consist of a vector of probabilities $\mathbf{z}_i = (z_{i1}, \dots, z_{iJ})$, with $z_{ij} \geq 0$ and $\sum_{j=1}^J z_{ij} = 1$, measuring the predictive likelihood that the continuous variable y falls within the respective bin. Reflecting the fact that the bins are mutually exclusive and contiguous, and generally cover the entire real line, we denote the bins by $(y_{j-1}, y_j]$, $j = 1, \dots, J$, with $y_0 < y_1 < \dots < y_J$, where y_0 and y_J are equal to $-\infty$ (left open bin) and $+\infty$ (right open bin), respectively. Figure A-3 in the Appendix displays the evolution of the bin ranges from 1982, the beginning of our sample,

⁴Some recent surveys, such as the Atlanta Fed's Survey of Business Uncertainty, only specify the number of bins and let the respondents determine their boundaries.

until the end in 2021, for both output growth and inflation, and shows that bins were changed in 1992, 2009, and 2020 for output growth surveys, and in 1985, 1992, and 2014 for inflation surveys. The SPF is conducted at a quarterly frequency (answers are collected in the middle of each quarter, right after GDP figures for the previous quarter have been released) and asks about probabilistic predictions for current and the following year year-over-year growth rates in real output (GDP) and the price level, as measured by the GDP deflator. Stark (2013) describes in detail the features of the SPF survey, and the Philadelphia Fed’s site provides a manual for interpreting the data.⁵

Figure 1 provides a few examples of survey responses that illustrate a number of common features of the SPF data. The top two panels show the probabilistic forecasts for output growth in 2020 made in 2019Q2 by respondents 532 and 584, while the bottom two panels show the forecasts for inflation in 2009 made in 2008Q4 by respondents 516 and 560 (respondents are anonymous). The probabilities \mathbf{z}_i ’s are displayed as histograms, while the black ticks on the horizontal axis mark the boundaries of the bins.

A first feature that emerges from Figure 1 is that probabilistic forecasts are very heterogeneous. For each row the respondents are forecasting the same object, and yet their probabilistic predictions are very different. Second, forecasters often assign zero probability to some if not most bins. Forecaster 532 for instance places zero probability on output growth being between -1 and 1 percent, but positive probability on output being between -2 and -1 percent, and between 1 and 3 percent. The econometrician could interpret this information literally, or as an indication that this respondent has a bimodal forecast distribution with some probability on a recession, a larger probability on an expansion, and very small likelihood of in-between outcomes. Other forecasters, such as respondent 584, place positive mass on almost all bins, however. A third feature of the data is that almost all probabilities in Figure 1 are round numbers, with responses for forecaster 584 being again the only exception. Fourth, forecasters do place mass on open bins and sometimes, as is the case for respondent who in 2008 was fearing deflation in 2009, most of the mass. Figures A-6 and A-7 in the Appendix show for each output growth and inflation survey the percentage of respondents placing positive probability on either one open bin or both. These percentages

⁵Figure A-4 in the Appendix displays the number of respondents n for output growth surveys conducted in Q1, Q2, Q3, and Q4 of each year (the numbers for inflation are essentially the same). n is about 35 in the early 1980s, and then drops steadily over time until 1992 when the Philadelphia Fed begins to manage the survey; n hovers around 35 until the mid-2000s and then starts to increase reaching a peak of about 50 during the Great Recession; it declines steadily thereafter and is about 30 in 2021. Figure A-5 shows survey participation by respondent, and provides a visual description of the panel’s composition.

are as high as 70 for output and 90 percent for inflation before 1992, when the bins were changed, but are on average about 20 percent, with peaks of 40 percent or higher, even after 1992. Finally, many of these predictive densities appear asymmetric. These examples display a left skew for output and, at least for forecaster 560, a right skew for inflation.

The econometrician’s problem is to use the information given by the elements of the survey probability vector \mathbf{z}_i of the i -th forecaster to address a number of questions of interest: What is the mean prediction for forecaster i ? How uncertain are they? Is there skewness in their predictive densities? The approach predominantly used so far in the literature concerning macroeconomic surveys has been to implicitly or explicitly assume that each forecaster i assigns the bin probabilities \mathbf{z}_i using a given predictive probability distribution $F_i(y)$. The task of the econometrician is then to estimate $F_i(y)$ from the data \mathbf{z}_i , and then use it to answer the questions of interest. Most existing literature has accomplished this task by fitting a given *parametric distribution* to the cumulative distribution function (CDF) implied by the bin probabilities, respondent by respondent, that is fitting $Z_{ij} = z_{i1} + \dots + z_{ij}$ $j = 1, \dots, J$, $i = 1, \dots, n$ using a parametric family of distributions $\{F(y|\boldsymbol{\theta}) : \boldsymbol{\theta} \in \Theta\}$. The type of the parametric distribution varies across studies, from a mixture of uniforms/piecewise linear CDF (that is, assuming that the probability is uniformly distributed within each bin; Zarnowitz and Lambros, 1987), to a Gaussian (Giordani and Soderlind, 2003), a generalized beta (Engelberg et al., 2009),⁶ and a skew-t distribution (e.g., Ganics et al., 2020), with the generalized beta assumption arguably being the most popular approach. The parameters of each distribution are usually estimated using nonlinear least squares, respondent by respondent; that is, $F_i(y) = F(y|\hat{\boldsymbol{\theta}}_i)$, where

$$\hat{\boldsymbol{\theta}}_i = \underset{\boldsymbol{\theta}_i}{\operatorname{argmin}} \sum_{j=1}^J \left| Z_{ij} - F(y_j|\boldsymbol{\theta}_i) \right|^2. \quad (1)$$

These approaches arguably have some limitations that are well understood in the literature (see Clements et al., forthcoming). First, the assumed parametric distribution may be misspecified—it may not fit well the individual responses. Moreover, since the width of the bins can be large (as is obviously the case when the respondent places probability on open bins), even if the distributions fit the Z_{ij} ’s, the inference results on moments and quantiles can be sensitive to the distributional assumption. Second, and related, existing approaches ignore inference uncertainty, even that concerning $\boldsymbol{\theta}_i$ for a given parametric assumption,

⁶Engelberg et al. (2009) use a triangular distribution, whenever the number of (adjacent) bins with positive probability is two or fewer.

let alone the uncertainty about the shape of $F_i(\cdot)$. This omission implies that confidence bands and hypothesis testing procedures cannot be derived.⁷ Third, bounded distributions such as the beta or the mixture of uniforms take literally the z_{ij} that are zero, in that they place no probability mass on bins where the respondents place no mass. More broadly, the approach outlined in expression (1) does not directly address the issue of rounding: it solves the minimization problem taking all the Z_{ij} 's literally even though the respondent may be reporting approximate probabilities (Dominitz and Manski, 1996; D'Amico and Orphanides, 2008; Boero et al., 2008, 2014; Engelberg et al., 2009; Manski and Molinari, 2010; Manski, 2011; Giustinelli et al., 2020; Glas and Hartmann, 2022, among others, discuss the issue of rounding; Binder, 2017 uses rounding to measure uncertainty).⁸ There have been attempts to address some of these issues, e.g., the potential misspecification, by choosing more flexible families of distributions such as the skew-Student-t distribution (e.g., Ganics et al., 2020). The possibility of misspecification remains, however. Moreover, if the econometrician does not account for inference uncertainty this flexibility comes at the price of overparameterization.

In the following two sections, we propose a Bayesian model that attempts to overcome some of these limitations. We first introduce a *parametric* model for the data. This model follows the literature in assuming that each forecaster uses a specific predictive distribution $F(\cdot)$ to assign probabilities ν to the bins, but differs from the literature in that it states that the data \mathbf{z} are noisy versions of the ν 's, where again the noise follows a parametric distribution. We then relax this parametric model by embedding it into the more general Bayesian *nonparametric* approach, thereby amending the potential misspecification associated with the parametric assumptions.

⁷Researchers have of course understood the presence of an inference issue especially when the information provided by the respondent is very limited. The proposed solutions generally amount to either choosing less heavily parameterized distributions or discarding the respondent: Clements (2010) for instance simply discards respondents with fewer than three bins, while Engelberg et al. (2009) use a triangle distribution in these cases. Liu and Sheng (2019) propose a maximum likelihood estimation approach in order to account for parameter uncertainty for given parametric assumptions.

⁸Manski and Molinari (2010) and Giustinelli et al. (2020) address the issue of rounding by considering interval data and using a person's response pattern across different questions to infer her or his rounding practice. The inferential approach pursued by these researchers is very different from the one followed by much of the literature and addresses different questions.

II.B A parametric model

We assume that the probability vector \mathbf{z}_i reported by forecaster i is a noise-ridden measurement of an unobserved vector of subjective probabilities over the J bins $\boldsymbol{\nu}_i = (\nu_{i1}, \dots, \nu_{iJ})$, with $\nu_{ij} \geq 0$ and $\nu_{i1} + \dots + \nu_{iJ} = 1$. We assume that the bin probabilities ν_{ij} are computed using forecaster i 's subjective probability distribution $F_i(\cdot)$:

$$\nu_{ij} = \nu_{ij}(\boldsymbol{\theta}_i) = F(y_j|\boldsymbol{\theta}_i) - F(y_{j-1}|\boldsymbol{\theta}_i), \quad j = 1, \dots, J, \quad (2)$$

where $\boldsymbol{\theta}_i \in \Theta$ is the vector of all parameters which includes those describing the CDF $F_i(\cdot) = F(\cdot|\boldsymbol{\theta}_i)$. For concreteness, in our application the subjective distribution $F(\cdot|\boldsymbol{\theta})$ is a mixture of two Gaussian distributions, that is

$$F(y|\boldsymbol{\theta}) = (1 - \omega)\Phi(y|\mu, \sigma_1^2) + \omega\Phi(y|\mu + \mu_\delta, \sigma_2^2), \quad (3)$$

but the approach accommodates any other choice for $F(\cdot|\boldsymbol{\theta})$.

The probability distribution $h(\cdot)$ captures the noise in the mapping between $\boldsymbol{\nu}_i$ and \mathbf{z}_i due to approximations or mistakes in reporting:

$$\mathbf{z}_i = (z_{i1}, \dots, z_{iJ}) \stackrel{ind}{\sim} h(\mathbf{z}_i|\boldsymbol{\nu}_i(\boldsymbol{\theta}_i), \boldsymbol{\theta}_i). \quad (4)$$

In choosing $h(\cdot)$, one needs to account for the fact that the elements of \mathbf{z}_i are positive and sum up to one (\mathbf{z}_i belongs to the simplex). A convenient choice for a distribution on the simplex is the Dirichlet distribution. A drawback of this distribution is that it assigns zero probability to \mathbf{z}_i 's that have some elements equal to zero, when in fact for the vast majority of forecasters some z_{ij} 's are zero. To specify $h(\cdot)$, we then follow Zadora et al. (2010) and use a distribution which allows for values of the random vector on the boundary of the simplex.

In order to describe this distribution it is useful to introduce the equivalent representation of \mathbf{z}_i given by the couple $(\mathbf{z}_{\boldsymbol{\xi}_i}, \boldsymbol{\xi}_i)$, where $\boldsymbol{\xi}_i = (\xi_{i1}, \dots, \xi_{iJ})$ with $\xi_{ij} = 1$ if and only if $z_{ij} = 0$ and $\xi_{ij} = 0$ otherwise. $\mathbf{z}_{\boldsymbol{\xi}_i}$ is the set of strictly positive \mathbf{z}_i —that is, $\mathbf{z}_{\boldsymbol{\xi}_i} = [z_{ij} : j \in \mathcal{J}_i^*]$ for $\mathcal{J}_i^* = \{j = 1, \dots, J : \xi_{ij} = 0\}$ —and takes values in the open $|\mathcal{J}_i^*|$ -dimensional simplex $\Delta^{|\mathcal{J}_i^*|}$, where $|\mathcal{J}_i^*|$ is the number of bins with positive mass. Using these definitions we can write the $h(\cdot)$ distribution as

$$h(\mathbf{z}_i|\boldsymbol{\theta}_i) = \frac{1}{c(\boldsymbol{\theta}_i)} \prod_{j=1}^J \varrho_j(\boldsymbol{\theta}_i)^{\xi_{ij}} (1 - \varrho_j(\boldsymbol{\theta}_i))^{1-\xi_{ij}} h(\mathbf{z}_{\boldsymbol{\xi}_i}|\boldsymbol{\theta}_i), \quad (5)$$

where $\boldsymbol{\varrho}(\boldsymbol{\theta}_i) = (\varrho_1(\boldsymbol{\theta}_i), \dots, \varrho_J(\boldsymbol{\theta}_i))$ are the probabilities that a forecaster will report zero probability on bin 1 through J , and $c(\boldsymbol{\theta}_i) = 1 - (\varrho_1(\boldsymbol{\theta}_i) \cdot \dots \cdot \varrho_J(\boldsymbol{\theta}_i))$ is a normalizing

constant. $h(\mathbf{z}_{\xi_i}|\boldsymbol{\theta}_i)$ is the standard Dirichlet distribution defined on the elements of \mathbf{z}_i that are non zero:

$$h(\mathbf{z}_{\xi_i}|\boldsymbol{\theta}_i) = \frac{\Gamma\left(\sum_{j \in \mathcal{J}_i^*} \phi(\boldsymbol{\theta}_i) \nu_j(\boldsymbol{\theta}_i)\right)}{\prod_{j \in \mathcal{J}_i^*} \Gamma(\phi(\boldsymbol{\theta}_i) \nu_j(\boldsymbol{\theta}_i))} \prod_{j \in \mathcal{J}_i^*} z_j^{\phi(\boldsymbol{\theta}_i) \nu_j(\boldsymbol{\theta}_i) - 1}, \quad (6)$$

where $\phi(\boldsymbol{\theta}_i) \kappa_i$ is the rescaled precision with $\kappa_i = \sum_{j \in \mathcal{J}_i^*} \nu_j(\boldsymbol{\theta}_i)$, and $\nu_j(\boldsymbol{\theta}_i)/\kappa_i$, $j \in \mathcal{J}_i^*$, are the renormalized $\nu(\boldsymbol{\theta}_i)$'s, which take into account the fact that if a forecaster decides to report zero probability for one or more bins, they need to adjust the probabilities associated with the other bins so that they still sum up to one.⁹

Omitting for ease of notation the dependence on the forecaster index i , the probability of reporting zero mass in the j -th bin is modeled as $\varrho_j(\boldsymbol{\theta}) = \varrho(\nu_j(\boldsymbol{\theta}), \epsilon(\boldsymbol{\theta}))$ where the function $\varrho(\nu, \epsilon)$ is decreasing in ν such that $\varrho \rightarrow 1$ for $\nu \rightarrow 0$ and $\varrho \rightarrow 0$ for $\nu \rightarrow 1$. $\epsilon(\boldsymbol{\theta})$ measures the sensitivity of $\varrho_j(\boldsymbol{\theta})$ to ν (that is, $\varrho \rightarrow 1$ for $\epsilon \rightarrow 0$ and $\varrho \rightarrow 0$ for $\epsilon \rightarrow 1$). In practice, the probabilities $\boldsymbol{\varrho}$ are parameterized as:

$$\varrho_j(\boldsymbol{\theta}) = \int_0^{\epsilon(\boldsymbol{\theta})} b(x|\nu_j(\boldsymbol{\theta}), r) dx \quad (7)$$

$j = 1, \dots, J$, where $b(x|m, r)$ is the PDF of a beta distribution $\mathcal{B}e(m, r)$ with mean m and precision r parameters.¹⁰ We assume r is fixed at 100 and $\epsilon(\boldsymbol{\theta}) = \epsilon$. The parameter vector of $h(\mathbf{z}|\boldsymbol{\theta})$ in the new parametrization is $\boldsymbol{\theta} = (\mu, \mu_\delta, \sigma_1, \sigma_2, \omega, \phi, \epsilon)$, where we set $\phi(\boldsymbol{\theta}) = \phi$.

Some of the parametric assumptions outlined above are less palatable than others. For instance, the assumption that the noise around the non-zero z_{ij} 's takes the form of a Dirichlet distribution is at odds with the observation on the prevalence of rounding. And even when the parametric assumption may be more palatable, it can still be wrong. Embedding these parametric assumptions into a more general nonparametric model arguably protects us, at least to some extent, from misspecification. We describe this approach in the next section.

⁹Note that the conditional Dirichlet satisfies some relevant properties of the unconditional Dirichlet, the marginal conditional means $\mathbb{E}(z_{ij}|\boldsymbol{\xi}_i) = \frac{\nu_j(\boldsymbol{\theta}_i)}{\sum_{j \in \mathcal{J}_i^*} \nu_j(\boldsymbol{\theta}_i)}$, $j \in \mathcal{J}_i^*$ sum up to one, and their marginal conditional variances $\mathbb{V}(z_{ij}|\boldsymbol{\xi}_i) = \frac{\nu_j(\boldsymbol{\theta}_i)(\kappa_i - \nu_j(\boldsymbol{\theta}_i))}{\kappa_i^2(\phi \sum_{j \in \mathcal{J}_i^*} \nu_j(\boldsymbol{\theta}_i) + 1)}$, $j \in \mathcal{J}_i^*$ go to zero with $\phi(\boldsymbol{\theta}_i) \rightarrow \infty$.

¹⁰We chose the beta distribution because it is the marginal of a Dirichlet, but we could have chosen any other distribution satisfying the above requirements. Our parametrization of the beta distribution is $b(x|m, r) = \frac{1}{B(mr, (1-m)r)} x^{mr-1} (1-x)^{(1-m)r-1}$ with $x \in (0, 1)$, $m \in (0, 1)$ and precision $r > 0$.

II.C A Bayesian nonparametric model

The Bayesian nonparametric hierarchical model works as follows. We assume that the parameter vector $\boldsymbol{\theta}_i$ is sampled from a mixture of forecaster “types” (for concreteness, let us think of low versus high uncertainty; low versus high mean; left versus right-skewed; low versus high reporting noise; a combination of all the above, *et cetera*). For now, imagine that the number of types K is finite. At the first stage of the hierarchy $\boldsymbol{\theta}_i$ is drawn from

$$\boldsymbol{\theta}_i \stackrel{iid}{\sim} \begin{cases} \boldsymbol{\theta}_1^* & \text{with probability } w_1 \\ \vdots \\ \boldsymbol{\theta}_K^* & \text{with probability } w_K \end{cases} \quad (8)$$

where the weights w_k are positive and sum to one. At the second stage of the hierarchy, the unknown parameters $\boldsymbol{\theta}_k^*$ characterizing the types, which are referred to as atoms, are sampled from a common distribution $\boldsymbol{\theta}_k^* \stackrel{iid}{\sim} G_0$, called base measure, which can be viewed as the probability distribution generating the types. The type probabilities w_k are drawn from the prior distribution

$$(w_1, \dots, w_K) \sim \text{Dir}\left(\frac{\psi_0}{K}, \dots, \frac{\psi_0}{K}\right), \quad (9)$$

where $\text{Dir}(\cdot)$ is a Dirichlet distribution whose K parameters are all identical and equal to $\frac{\psi_0}{K}$.

Now let the number of types K go to infinity. When this happens, expression (8) is replaced by the discrete random measure

$$G(\boldsymbol{\theta}) = \sum_{k=1}^{\infty} w_k \delta(\boldsymbol{\theta} - \boldsymbol{\theta}_k^*) \quad (10)$$

where $\delta(x)$ denotes a point mass distribution located at 0, the atoms $\boldsymbol{\theta}_k^*$ are drawn from G_0 as before, and the random weights w_k are generated by the stick-breaking representation $SB(\psi_0)$ given by

$$w_k = v_k \prod_{l=1}^{k-1} (1 - v_l) \quad (11)$$

where the stick-breaking components v_l are i.i.d. random variables from a beta distribution $\mathcal{Be}(1, \psi_0)$ (e.g., see Pitman, 2006). The random measure G is a Dirichlet process $\mathcal{DP}(\psi_0, G_0)$ (Ferguson, 1973) and our hierarchical model is a Dirichlet process prior $\boldsymbol{\theta}_i \stackrel{iid}{\sim} G$, $G \sim \mathcal{DP}(\psi_0, G_0)$. The precision parameter ψ_0 determines how uneven the weights are in the stick-breaking representation: when $\psi_0 \rightarrow 0$ all forecasters are assumed to be of the same

type ($w_1 \rightarrow 1$) while when $\psi_0 \rightarrow +\infty$ the inference is done forecaster by forecaster (using the same prior). Outside of this latter limiting case, the Dirichlet process prior generates a priori dependence among the forecaster-specific parameters $\boldsymbol{\theta}_i$'s via the formation of clusters of forecasters of the same type.¹¹

Sethuraman (1994)'s constructive representation (10) implies that our model has the infinite mixture representation

$$\mathbf{z}_i \stackrel{iid}{\sim} h_G(\mathbf{z}) = \int h(\mathbf{z}|\boldsymbol{\theta})G(d\boldsymbol{\theta}) = \sum_{k=1}^{\infty} w_k h(\mathbf{z}|\boldsymbol{\theta}_k^*), \quad (12)$$

where the weights w_k come from the same prior distribution (11) for all forecasters. Each forecaster is modeled a priori as a potentially infinite mixture of types each encoded by the parametric distribution $h(\mathbf{z}|\boldsymbol{\theta}_k^*)$. The Bayesian nonparametric model is therefore quite flexible. As such, it can overcome the inherent misspecification implied by the use of specific parametric assumptions, as shown below in section II.E. The weights of the mixture are given by (11), implying that some degree of pooling is imposed: most forecasters come from the same relatively few (depending on ψ_0) types. Such pooling mitigates overfitting. Moreover, the number of types grows naturally as more data becomes available and depends on the degree of heterogeneity in the sample.¹² A posteriori, both the unknown atoms/types $\boldsymbol{\theta}_k^*$ and the weights w_k are estimated, as described next.

II.D Posterior inference

Since expression (12) is a mixture distribution, it can be reparameterized using auxiliary allocation variables d 's, which are equal to k if $\boldsymbol{\theta}_i$ is sampled from the k^{th} mixture component:

$$\mathbf{z}_i \stackrel{iid}{\sim} \sum_{k=1}^{\infty} \mathbb{I}\{d = k\} h(\mathbf{z}|\boldsymbol{\theta}_d^*), \quad Pr\{d = k\} = w_k, \quad (13)$$

¹¹As shown in Pitman (2006), the predictive distribution of $\boldsymbol{\theta}_{i+1}$ conditional on $(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_i)$ can be represented as a Polya's urn process $\boldsymbol{\theta}_{i+1}|\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_i \sim \frac{\psi_0}{\psi_0 + i} G_0(\boldsymbol{\theta}_{i+1}) + \frac{1}{\psi_0 + i} \sum_{k=1}^i \delta(\boldsymbol{\theta}_k - \boldsymbol{\theta}_{i+1})$. With probability $\frac{\psi_0}{\psi_0 + i}$ the new draw $\boldsymbol{\theta}_{i+1}$ is generated from G_0 , but it is otherwise equal to one of the previous i draws. When $\psi_0 \rightarrow \infty$ we have the same parametric model for each forecaster: $\mathbf{z}_i \sim h(\cdot|\boldsymbol{\theta}_i)$ where the $\boldsymbol{\theta}_i$'s are drawn independently from G_0 .

¹²Given a population of n forecasters', their distribution can be characterized using N_n different clusters, where N_n is a random variable with prior mean $\mathbb{E}[N_n] \approx \psi_0 \log\left(\frac{\psi_0 + n}{\psi_0}\right)$.

Thus the posterior distribution of $(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_n, G)$ given $(\mathbf{z}_1, \dots, \mathbf{z}_n)$ can be expressed in terms of the posterior distribution $\Pi(d_1, \dots, d_n, \boldsymbol{\theta}_1^*, \boldsymbol{\theta}_2^*, \dots, w_1, w_2, \dots | \mathbf{z}_1, \dots, \mathbf{z}_n)$, where a *posteriori* the allocation variables are naturally forecaster specific. If the mixture (13) were finite, Bayesian inference would be straightforward. The slice Gibbs sampler algorithm of Walker (2007) and Kalli et al. (2011) surmounts the issue of infinity using data augmentation, as explained in detail in Appendix B. The MCMC samples

$$(d_1^{(m)}, \dots, d_n^{(m)}, \boldsymbol{\theta}_1^{*(m)}, \boldsymbol{\theta}_2^{*(m)}, \dots, w_1^{(m)}, w_2^{(m)}, \dots)$$

over $m = 1, \dots, M$ iterations are used to approximate the posterior distribution for any quantity of interest. For example, the set of posterior draws $\{F(y|\boldsymbol{\theta}_i^{(m)}) : y \in \mathcal{Y}, m = 1, \dots, M\}$, with $\boldsymbol{\theta}_i^{(m)} := \boldsymbol{\theta}_{d_i^{(m)}}^{*(m)}$ approximate the posterior distribution of the subjective CDF $F_i(\cdot)$ (see Figure 2 below). Analogously, the posterior mean of the standard deviation of the predictive distribution of the i -th forecaster is approximated by

$$\hat{\sigma}_i = \frac{1}{M} \sum_{m=1}^M \sigma(\boldsymbol{\theta}_i^{(m)})$$

where $\sigma(\boldsymbol{\theta})$ is the standard deviation of $F(\cdot|\boldsymbol{\theta})$. Finally, quantities involving the whole population of forecasters can be approximated in a similar way. For example, the posterior mean of the cross-sectional standard deviation of the individual standard deviations is given by

$$\hat{\sigma} = \frac{1}{M} \sum_{m=1}^M \frac{1}{n} \sum_{i=1}^n \left(\sigma(\boldsymbol{\theta}_i^{(m)}) - \bar{\sigma}^{(m)} \right)^2 \quad \text{with} \quad \bar{\sigma}^{(m)} = \frac{1}{n} \sum_{i=1}^n \sigma(\boldsymbol{\theta}_i^{(m)}).$$

II.E Posterior consistency

In this section we discuss asymptotic properties of the posterior distribution as the number of forecasters goes to infinity. We only state the main result on consistency, leaving all the details, proofs, and some additional results to Section D of the Appendix.

We formalize asymptotic convergence using the notion of weak consistency of the posterior distribution (Ghosh and Ramamoorthi, 2003), which provides a widely accepted minimal requirement for large sample behavior of Bayesian nonparametric models (e.g., Norets and Pelenis, 2012; Pelenis, 2014; Norets and Pelenis, 2014; Bassetti et al., 2018). Roughly speaking, posterior consistency means that in a frequentist experiment with a given data generating density, the posterior distribution concentrates around this density as the sample

size (number of forecasters) increases. More formally, let \mathcal{H} be the set of all possible data generating densities (with respect to a dominating measure) on the simplex Δ^J where the data \mathbf{z} lives. Given a prior Π on \mathcal{H} , the posterior is said to be *weakly consistent* at h_0 if for every i.i.d. sequence $\mathbf{z}_1, \dots, \mathbf{z}_n$ of random vectors with common density h_0 the posterior probability $\Pi(U|\mathbf{z}_1, \dots, \mathbf{z}_n)$ converges a.s. to 1 as $n \rightarrow +\infty$ for every weak neighborhood U of h_0 . In our model, the prior is $\Pi(U) = P\{h_G \in U\}$, where h_G is defined in (12). In order to prove weak consistency we use the Schwartz theorem (see e.g. Chapter 4 in Ghosh and Ramamoorthi, 2003), which states that weak consistency at a true density h_0 holds if the prior assigns positive probabilities to Kullback-Leibler neighborhoods of h_0 .

Before stating the main theorem, it is helpful to clarify the definition of Kullback-Leibler divergence for densities over Δ^J with possible zero elements. We define a σ -finite measure on Δ^J by $\lambda(d\mathbf{z}) = c(d\xi) \otimes \mathcal{L}_\xi(d\mathbf{z}_\xi)$ where c is the counting measure on the space $\{\xi \in \{0, 1\}^J, \text{ s.t. } \xi_1 + \dots + \xi_J < J\}$, and \mathcal{L}_ξ is the Lebesgue measure on $\Delta^{|\mathcal{J}^*|}$ (recall that \mathbf{z}_ξ, ξ , and \mathcal{J}^* were defined in section II.B). The set of all possible data generating densities \mathcal{H} is the set of all the densities $g(\mathbf{z}) = g(\mathbf{z}_\xi, \xi)$ absolutely continuous with respect to λ . Given two densities h_0 and g in \mathcal{H} the Kullback-Leibler divergence between h_0 and g is then defined as

$$KL(h_0, g) = \int_{\Delta^J} h_0(\mathbf{z}) \log \left(\frac{h_0(\mathbf{z})}{g(\mathbf{z})} \right) \lambda(d\mathbf{z}). \quad (14)$$

Call \mathcal{M}^* the set of finite mixtures of densities (5) that define the parametric component of our model, and \mathcal{H}_0^* the set of densities that can be approximated in the Kullback-Leibler sense by densities in \mathcal{M}^* , i.e.

$$\mathcal{H}_0^* = \{h_0 \text{ density w.r.t. } \lambda: \forall \epsilon > 0 \exists g \in \mathcal{M}^* \text{ s.t. } KL(h_0, g) \leq \epsilon \}.$$

Theorem 1. Assume that $\boldsymbol{\theta} \mapsto (\varrho_1(\boldsymbol{\theta}), \dots, \varrho_J(\boldsymbol{\theta}), \phi(\boldsymbol{\theta})\nu_1(\boldsymbol{\theta}), \dots, \phi(\boldsymbol{\theta})\nu_J(\boldsymbol{\theta}))$ is a continuous function such that $\nu_j(\boldsymbol{\theta}) > 0$ and $0 < \varrho_j(\boldsymbol{\theta}) < 1$ for every $j = 1, \dots, J$. If G_0 has full support, then the posterior is weakly consistent at any density h_0 in \mathcal{H}_0^* such that

$$\int_{\Delta^J} \left| \log \left(\prod_{j: z_j > 0} z_j \right) \right| h_0(\mathbf{z}) \lambda(d\mathbf{z}) < +\infty. \quad (15)$$

The result guarantees that the posterior distribution concentrates around the true process generating the histogram data \mathbf{z} in the SPF cross-section as the number of forecasters grows to infinity. In particular, it shows that the Bayesian nonparametric approach is robust to deviations from the specific parametric assumptions, such as the notion that forecasters's

noise is distributed according to a Dirichlet distribution (as opposed to rounding toward integer numbers) or the particular choice of the $F(\cdot)$ predictive CDF. Hence, even if the specific form of $h(\mathbf{z}|\boldsymbol{\theta})$ is not correct, the true distribution h_0 is recovered in the limit as long as h_0 belongs to the very broad class of models \mathcal{H}_0^* , which includes all the models that are not “too far” from any finite mixture of $h(\mathbf{z}|\boldsymbol{\theta})$. This property is not shared by any of the current outstanding approaches for inference on probabilistic surveys.

A couple of observations are in order. First, as is always the case for Bayesian non-parametrics, the consistency results do not apply to *individual* forecasters, but only to the data generating process for the entire distribution of forecasters. Concretely, this means that they apply to any object that involves a suitably large number of forecasters, such as the consensus distribution. Second, it must be clear that the consistency holds for the true distribution h_0 on the available data \mathbf{z} and *not* for the underlying predictive CDF $F(\cdot)$ over the entire domain of y . This is due to the fact that the available data do not provide enough information to fully recover the CDF of y since the number of bins J is taken as fixed (and finite), even when n goes to infinity: loosely speaking, we can claim consistency for the value of the predictive CDF $F(\cdot)$ at the bin edges y_1, \dots, y_J , but do not have enough information about the value of $F(\cdot)$ for $y \in (y_j, y_{j+1}]$. This identification issue is overcome in the case where the number of bins J goes to infinity and the bin size goes to zero, as we show in Appendix C. Specifically, we show that under these conditions when the number of forecasters n also goes to infinity the consistency result discussed above applies also to estimates of the predictive distribution $F(\cdot)$. Since these results are of limited interest for our application where the number of bins is limited and non negligible mass is often placed on the open bins, they are relegated to the Appendix.¹³

In practice in our application both the number of bins, as just mentioned, and of observations n is not large (e.g., n around 30 in 2020 and the width of some of the bins is as large as 6 percent for output growth). Still, an advantage of the Bayesian approach is that lack of information is reflected in the posterior credible intervals. Section III.A below for instance shows that when forecasters place a large amount of mass on the open bins, the estimates of their predictive CDF $F(\cdot|\boldsymbol{\theta}_i)$ becomes more uncertain. Still, one needs to be aware that in these situations the choice of the distribution family $F(\cdot)$, $h(\cdot)$, and the prior distribution can impact the results. Therefore, a robustness check with respect to these choices should

¹³Arguably, output growth SPF surveys between the Great Recession and the Covid episode, which displayed fairly narrow bins and little mass on the open bins, are the only ones for which these conditions come close to applying

be included in all applications of our method.

II.F Priors

In this section we discuss the prior settings used in our application. The parameter vector $\boldsymbol{\theta}$ is composed of $(\mu, \mu_\delta, \sigma_1, \sigma_2, \omega, \phi, \epsilon)$. The first four parameters pertain to the $F(\cdot)$ function—the mixture of two normals (3), which we repeat here for convenience: $F(y|\boldsymbol{\theta}) = (1 - \omega)\Phi(y|\mu, \sigma_1^2) + \omega\Phi(y|\mu + \mu_\delta, \sigma_2^2)$. The parameters ϕ and ϵ are used to specify $h(\cdot)$ and $\varrho(\cdot)$ in (6) and (7), respectively. We should stress that we use the same priors for both output growth and inflation and for all years in our sample.

The location of the first mixture component is $\mu \sim \mathcal{N}(2, 5^2)$, where the standard deviation of 5 implies that this is a very loose prior. The scales of the mixture components follow $\sigma_j \sim \mathcal{IGa}(a_\sigma, b_\sigma)\mathcal{I}(\sigma_1)_{(0,10)}$, $j = 1, 2$ where a_σ, b_σ are chosen so that the standard deviation has mean $E[\sigma_j] = 2$ and a variance $V[\sigma_j] = 4$, and where we truncate the distribution at 10 for numerical reasons. The parameters μ_δ captures the deviation of the mean of the second mixture component relative to the first one. Its prior is centered at zero (implying that the second mixture a priori mainly captures fat tails) and has a standard deviation of 1: $\mu_\delta \sim \mathcal{N}(0, 1^2)$. Finally, the prior for ω , the weight on the second component of the mixture, is $\omega \sim \mathcal{Be}(0.5, 3)$. Its mode is zero, implying that the prior favors models with one mixture only. The prior places roughly 20 percent probability on $\{\omega \geq 0.25\}$.

For the precision parameter ϕ of the Dirichlet distribution (6) that determines the amount of noise around the underlying probabilities ν we assume a gamma distribution $\mathcal{Ga}(a_\phi, b_\phi)$, where a_ϕ and b_ϕ are chosen so that both the mean and the variance of ϕ are equal to 100. The left panel of Figure A-2 in the Appendix shows the 50 and 90 percent a-priori coverage intervals for the noise associated with three different values of ν : 0.1, 0.6 and 0.3. The 50 and 90 percent intervals are about 5 and 10 percent wide, respectively. As regards the prior for the rounding-off-to-zero parameter ϵ we assume a gamma distribution $\mathcal{Ga}(a_\epsilon, b_\epsilon)$ and set a_ϵ, b_ϵ such that in expectations the probability of reporting a zero, ϱ_j , is close to one when the true mass on the bin ν_j is less than 1 percent, declining to very small values for any ν_j larger than 5 percent. The right panel of Figure A-2 shows the mean, the 50, and the 90 percent coverage intervals of $\varrho_j(\boldsymbol{\theta})$ as a function of $\nu_j(\boldsymbol{\theta})$. The a-priori uncertainty is such that when ν_j is 2 percent the 90 percent interval for ϱ_j ranges from 0 to 25 percent. Finally, the concentration hyperparameter ψ_0 of the Bayesian nonparametric prior, which determines

the prior number of clusters, we follow the standard choice and set $\psi_0 = 1$. This implies that the expected number of clusters for a cross-section with $n = 30$ is roughly 4.

III Results

In this section we first discuss the application of the nonparametric Bayesian approach to the few selected examples mentioned at the beginning of Section II, so that the reader becomes familiar with how the approach works in practice. Next, we document the evolution from 1982 to 2021 of individual measures of uncertainty obtained using our approach. This analysis sets the stage for the analysis in the following section where we study the relationship between subjective uncertainty and ex-post forecast errors, and assess whether SPF predictive densities conform with the noisy rational expectations hypothesis.

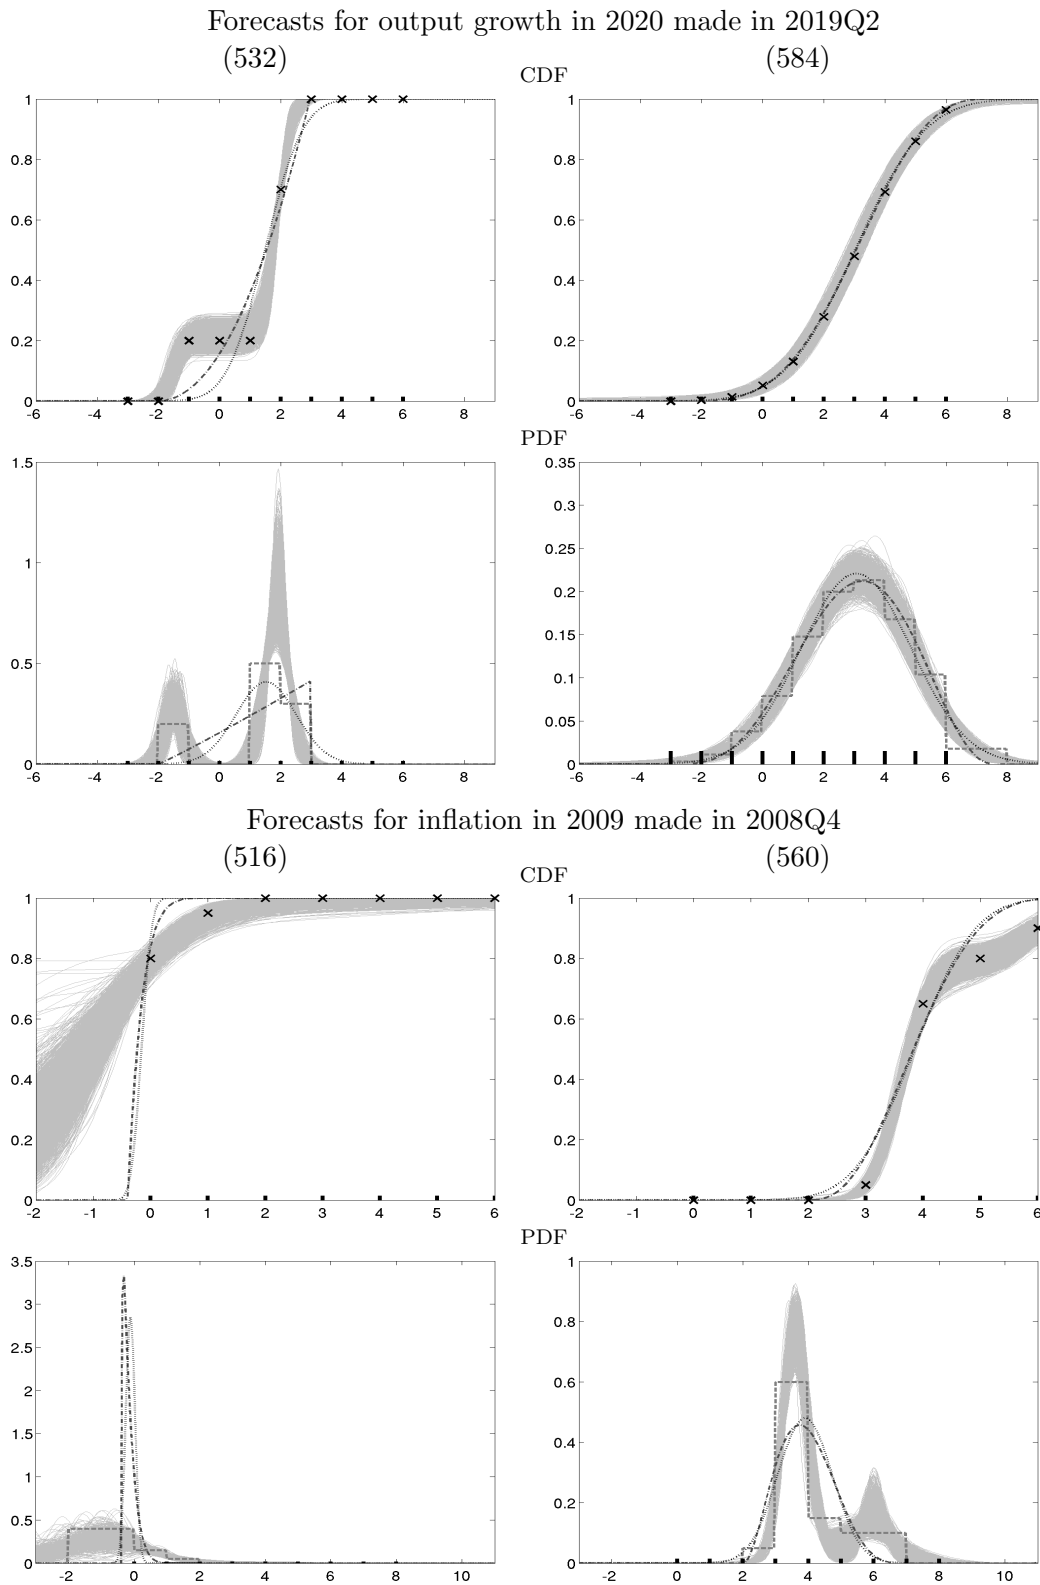
III.A Examples

Figure 2 shows the inference results for the four SPF respondents shown in Figure 1. For each forecaster we show posterior draws (thin gray lines) from the BNP model for the subjective CDF $F(y|\theta_i)$ (top) and PDF (bottom), and compare it with the results under the generalized beta (black, dash-and-dotted) and Gaussian (black, dotted) approaches. The CDF plots also show the observed cumulative probabilities Z_{ij} (crosses), while the PDF plots show the step-wise uniform PDF (gray dashed lines) implied by the histogram probabilities z_{ij} .

Figure 2 illustrates a few aspects of the Bayesian nonparametric approach. For starters, the observed cumulative probabilities (the Z_{ij} 's; crosses) belongs to the high posterior density region for all these respondents, suggesting that the approach is flexible enough to capture a variety of arguably challenging cases. In contrast, the beta and the normal approaches do not fit the Z_{ij} 's well in these examples and their CDFs and PDFs do not belong to the high posterior density region obtained from the BNP approach, with the exception of respondent (584). As a consequence there can be important differences in the objects of interests, such as the measure of uncertainty, or quantiles, implied by the different approaches.¹⁴ Figure 2 also shows that the BNP approach delivers wider posterior coverage intervals that reflect the higher degree of uncertainty whenever there is less information from the respondent. The

¹⁴Bassetti et al. (forthcoming) show that inference using the BNP approach differs from that obtained using standard approaches for several other examples obtained during the recent Covid episode.

Figure 2: Inference using the Bayesian nonparametric approach: CDFs and PDFs for selected examples



Note: For each forecaster the top and bottom panels show the subjective CDF and PDF, respectively, estimated using the BNP approach (posterior draws; light gray), as well as the least-squares estimates obtained under the normal (gray, dashed line) and the beta (black, dash-and-dotted line) parametric assumptions. The CDF panels also show the observed cumulated histogram probabilities Z_{ij} $j = 1, \dots, J$ (crosses), while the PDF panel show the step-wise uniform PDF (gray dashed lines) implied by such probabilities.

case of respondent 516, who placed 80 percent probability on the left open bin (see Figure 1), is exemplary. The posterior coverage intervals for both the BNP CDF and PDF reflect the fact that we know very little about the left-tail behavior of this forecaster, as evidenced by the fact that the gray lines for both the CDF and the PDF are far less concentrated for forecaster 576 in the left tail than in other regions or for other forecasters.

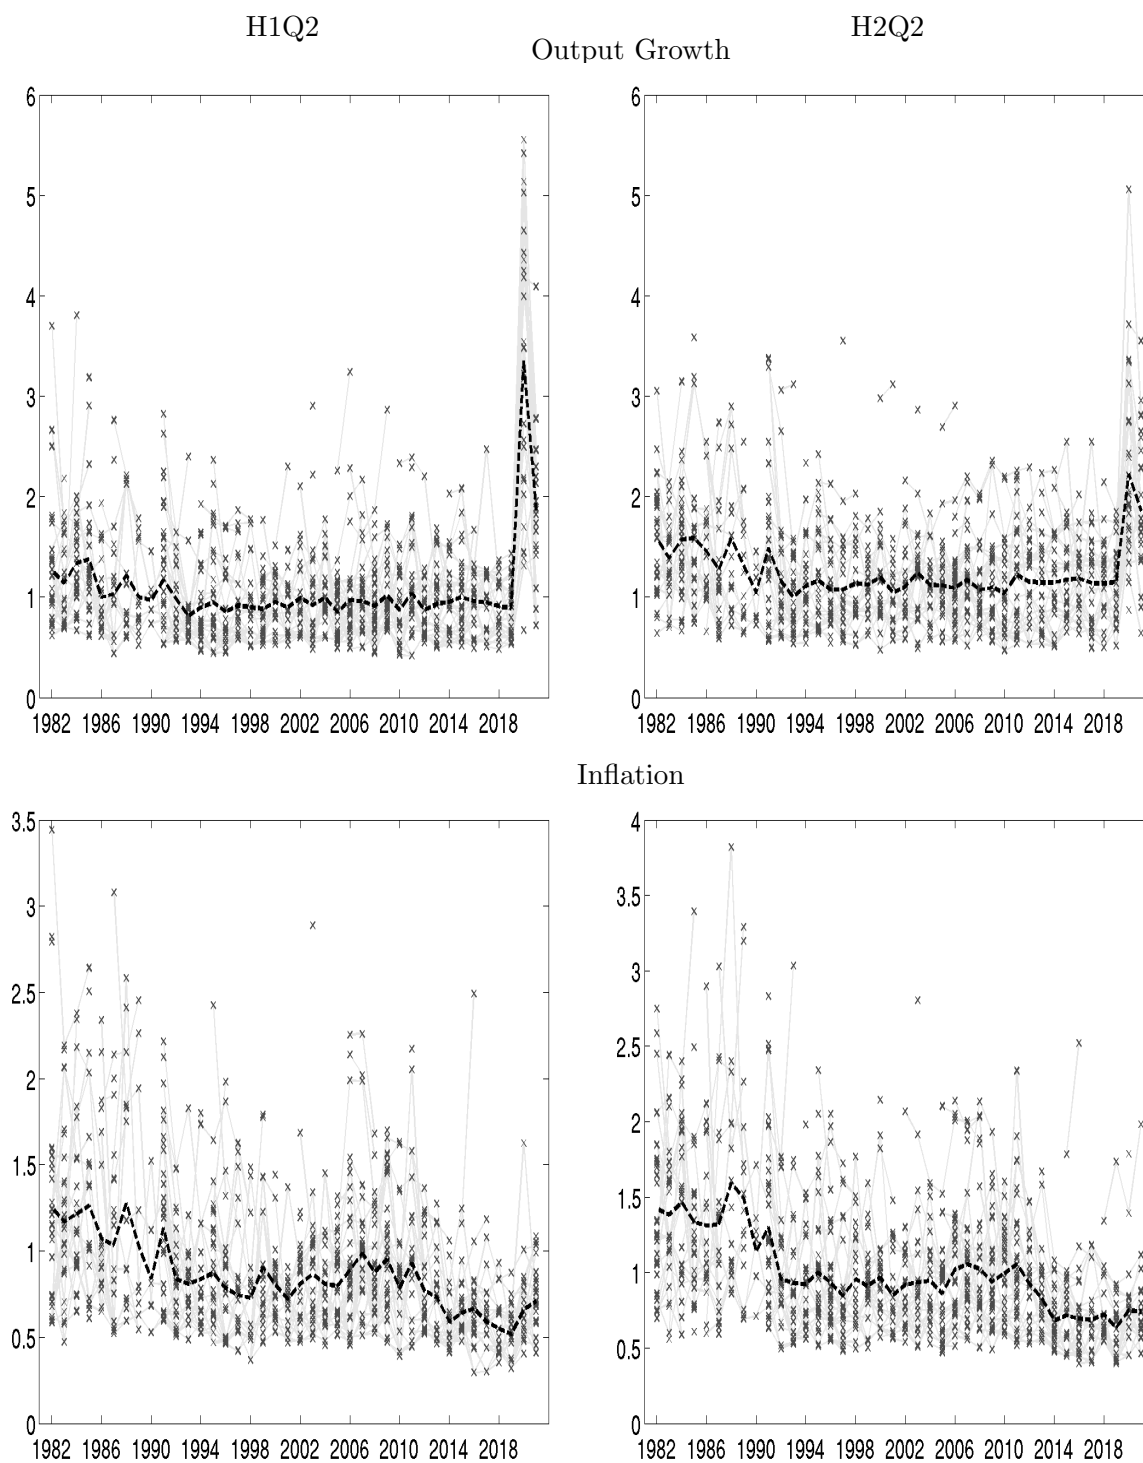
III.B Heterogeneity in subjective uncertainty

In this section we document the evolution of individual measures of uncertainty obtained using our approach in the 1982-2021 sample. We do this for two reasons. First, we set the stage for the analysis in the next section, where we study the relationship between subjective uncertainty and ex-post forecast errors. In particular, we follow the literature and show that professional forecasters differ significantly in terms assessment of uncertainty, and that these differences vary over time.¹⁵ We also show and that while these differences are persistent, forecasters do change their mind from period to period about their subjective uncertainty—a variation that we will exploit later. Second, we take advantage of our inference-based approach and test the extent to which these differences are significant.

Figure 3 shows the evolution of subjective uncertainty by individual respondent for output growth (top) and inflation (bottom). The left and right panels display uncertainty for the current and the next year projections, respectively, made in the second quarter of each year (the Appendix shows that results for other quarters are qualitatively similar). In each panel the crosses indicate the posterior mean of the standard deviation of the individual predictive distribution. We use the standard deviation (as opposed to the variance) because its units are easily grasped quantitatively and are comparable with alternative measures of uncertainty such as the interquartile range. Thin gray lines connect the crosses across periods when the respondent is the same, providing information on both whether respondents

¹⁵Heterogeneity in macroeconomic probabilistic forecasts was noted long ago. While much of the early literature focused on disagreement in point projections or central tendencies (see Mankiw et al., 2003; Capistrán and Timmermann, 2009; Patton and Timmermann, 2010, 2011; Andrade and Le Bihan, 2013, and other work mentioned in the recent survey by Clements et al., forthcoming) more recent work documents the fact that forecasters disagree about uncertainty and that these differences are long-lasting (Lahiri and Liu, 2006; D’Amico and Orphanides, 2008; Bruine De Bruin et al., 2011; Boero et al., 2014; Rich and Tracy, 2021, among others). Kozeniauskas et al., 2018, discuss the conceptual differences between macroeconomic uncertainty and disagreement using a model where forecasters have private information and update their beliefs using Bayes’ law).

Figure 3: Subjective uncertainty by individual respondent



Note: Each panel displays the posterior mean of the standard deviation of the subjective predictive distribution by individual respondent (light gray crosses, connected by thin gray line whenever the respondent appears in consecutive surveys), and the cross-sectional average of the individual standard deviations (dashed black line). Top and bottom panels correspond to output growth and inflation projections. The left and right column correspond to current and next year projections.

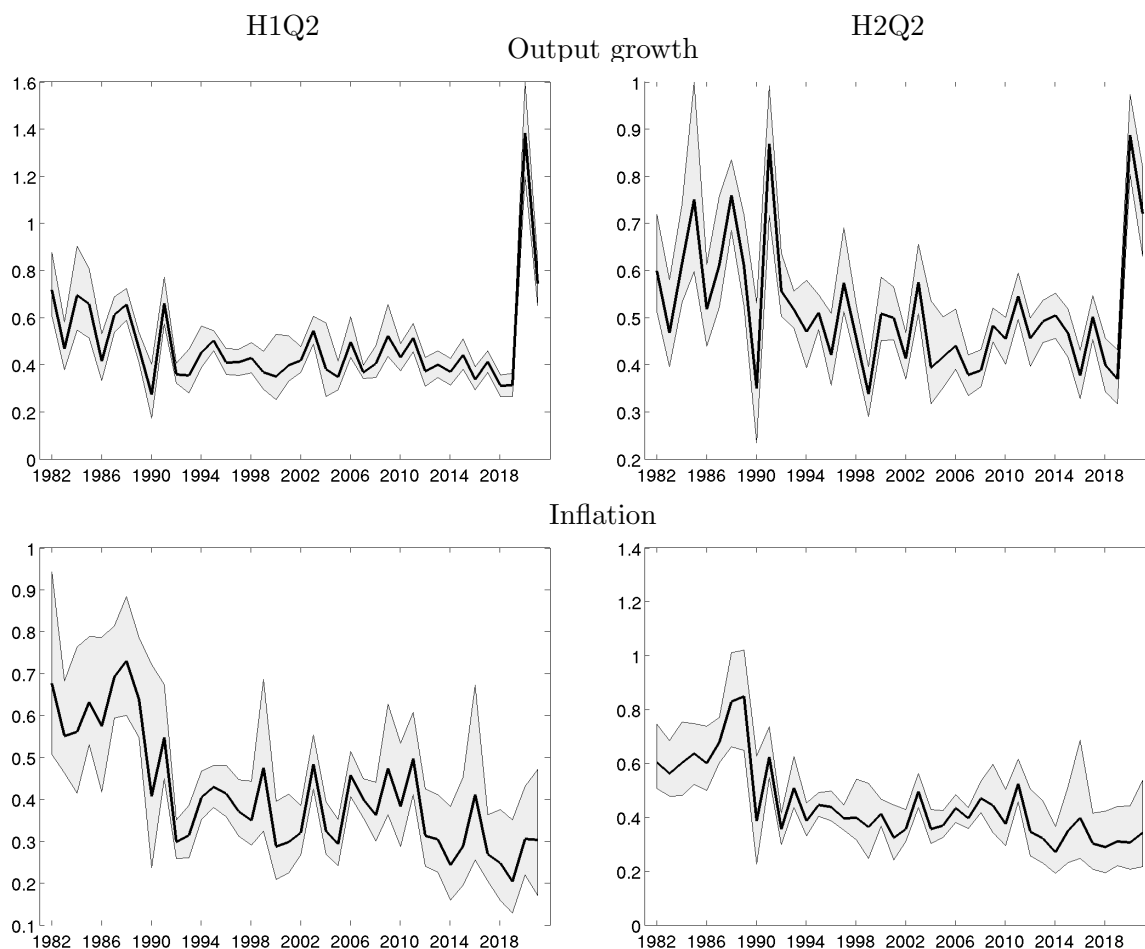
change their view on uncertainty and whether the composition of the panel affects the cross-sectional average measure of uncertainty, which is shown by a black dashed line (Manski, 2018, stresses the extent to which the literature has often ignored compositional changes when discussing the evolution of consensus or average measures). Figure 4 provides a time series of the differences in individual uncertainty, as measured by the cross-sectional standard deviation of the individual standard deviations. The solid black line displays the posterior mean of this measure, while the shaded areas represent the 90 percent posterior coverage.

Figure 3 shows that on average uncertainty for current year output growth projections declined from the 1980s to the early 1990s, likely reflecting a gradual learning about the Great Moderation, and then remained fairly constant up to 2020 when the Covid pandemic struck, and average uncertainty grew threefold. Average uncertainty for next year projections tends to be in general higher than for current year projections. It follows a similar pattern, except that it displays a small but very steady upward shift in the aftermath of the Great Recession. It appears unlikely that changes in survey design, and particularly in the bins, affect these patterns: for output growth these changes take place in 1992, 2009, and 2020. Except for 2020, where much of the change in uncertainty is arguably attributed to Covid, there are no evident breaks associated with the bin changes. Interestingly, we do not see any upticks in average subjective uncertainty in the run up to recessions, even for current year forecasts, with the exception of the Covid crisis. Using the interquartile range to measure uncertainty, as done in Figure A-11 in the Appendix, leads to very similar results. Using the generalized beta approach to fit histograms (see Figure A-12) also produces similar overall patterns, although perhaps not surprisingly this approach leads to lower estimates of subjective uncertainty relative to our approach.

Cross-sectional differences in individual uncertainty are very large, and quantitatively trump any time variation in average uncertainty. The standard deviation of low uncertainty individuals remains below 1 throughout the sample, with the sole exception of the Covid period, while that of high uncertainty individuals is often higher than two. More formally, the cross-sectional standard deviation of individual standard deviations, shown in Figure 4, hovers between 0.4 and 0.8 throughout the sample, and then jumps during the Covid period. The cross-sectional standard deviation is quite tightly estimated indicating that differences across individuals are significant. The level and the dispersion of uncertainty appear to be tightly linked, in that the cross-sectional standard deviation is high when the average is high. From Figure 3 this seems due to the fact that it is mostly high uncertainty respondents who change their mind about the confidence in their projections, thereby driving both the average

and the cross-sectional standard deviations. While differences in subjective uncertainty are persistent, forecasters do revise their assessment of unpredictability and their relative ranking varies as indicated by the fact that the thin gray lines very often cross one another.

Figure 4: Cross-sectional standard deviation of individual uncertainty—Q2 survey



Note: Posterior mean of the cross-sectional standard deviation of the individual standard deviations (solid black line) The shaded areas display the 90 percent posterior credible intervals. Top panels: output growth. Bottom panels: inflation. Left column: current year; right column: next year.

The bottom panels of Figure 3 shows that on average subjective uncertainty for inflation in both for current (left) and following (right) year declined from the 1980s to the mid-1990s and then was roughly flat up until the mid-2000s. Average uncertainty rose in the years surrounding the Great Recession, but then declined again quite steadily starting in 2011 reaching a lower plateau around 2015. Interestingly, average uncertainty did not really rise much in 2020 and 2021 in spite of the Covid related disturbances, and in spite of the fact that for most respondents expected inflation rose sharply, as documented in Figure A-14 in the Appendix. In the case of inflation changes in the bins, which took place in 1985, 1992, and 2014 (see Figure A-3 in the Appendix), may have played some role as we see that the

average standard deviation drops markedly in both 1992 and 2014. At the same time it is arguably not the only explanation since such drops appear to be the continuation of a trend that had started before the change in survey design.

As was the case for output growth, also for inflation cross-sectional differences in individual uncertainty are very large. The cross-sectional standard deviation of individual standard deviations (Figure 4) follows the same pattern of the average standard deviation: it starts around 0.6 percent in the 1980s, drops to around 0.4 percent in the 1990s, and then drops a bit further in the 2010s. This measure of cross-sectional heterogeneity in uncertainty is tightly estimated and its fluctuations appear to be statistically significant. As was the case for output, high uncertainty respondents becoming less uncertain are mostly driving both the average and the cross-sectional standard deviations.

III.C Subjective uncertainty and forecast accuracy: Testing the noisy information hypothesis for density forecasts

In this section we use our approach to assess whether SPF predictive densities conform with the noisy rational expectations hypothesis (see Coibion and Gorodnichenko, 2012, 2015, for instance). We do that by subjecting predictive densities to three types of tests. The first two tests concern the scale of the forecasters' predictive distribution, while the last test concerns its location. According to the noisy rational expectations hypothesis, forecasters receive both public and private signals about the state of the economy, which they do not observe. Under this hypothesis heterogeneity in the signals, and in their precision, explains the heterogeneity in both mean predictions $E_{t-q,i}[y_t]$ and in their subjective uncertainty $\sigma_{t|t-q,i}^2 = E_{t-q,i}[(y_t - E_{t-q,i}[y_t])^2]$, where i denotes the forecaster, q the horizon of the forecast, and $E_{t-q,i}[\cdot]$ the expectation operator under forecaster i 's information set at time $t - q$. In the time series, changes in the precision of either the private or public signals—the latter due, say, to changes in policy or the structure of the economy, a recession approaching, or some other major event like Covid-19—will be reflected in $\sigma_{t|t-q,i}^2$. In the cross-section, if forecaster i has a more precise signal than forecaster j , then $\sigma_{t|t-q,i}^2$ ought to be lower than $\sigma_{t|t-q,j}^2$. We have seen in the previous section that $\sigma_{t|t-q,i}^2$ varies substantially both over time and in the cross-section.

Regardless, if expectations are rational there needs to be a correspondence between the

subjective uncertainty $\sigma_{t|t-q,i}$ and the *ex-post* forecast error $|y_t - E_{t-q,i}[y_t]|$. Define

$$\eta_{i,t|t-q} = (y_t - E_{t-q,i}[y_t]) / \sigma_{t|t-q,i}, \quad (16)$$

the standardized forecast error. Under rational expectations (that is, if the predictive distribution is consistent with the data generating process for y_t) it has to be the case that

$$E[\eta_{i,t|t-q}^2] = 1. \quad (17)$$

We will test whether $\eta_{i,t|t-q}^2 = (y_t - E_{t-q,i}[y_t])^2 / \sigma_{t|t-q,i}^2$, is equal to 1 on average, and refer to this test as a *scale* test. Next, taking logs of the absolute value of both sides of equation (16) we obtain:

$$\ln |y_t - E_{t-q,i}[y_t]| = \ln \sigma_{t|t-q,i} + \ln \eta_{i,t|t-q}. \quad (18)$$


By regressing $\ln |y_t - E_{t-q,i}[y_t]|$ on $\ln \sigma_{t|t-q,i}$ we will test whether absolute forecast error change proportionally to the subjective uncertainty both in the time series and in the cross-section, and refer to this test as a *variation* test. While both tests hinge on rational expectations, they are different. The *scale* test tells us whether subjective distributions are properly scaled on average, while the *variation* test inquires whether variations in subjective uncertainty map into variations in forecast errors. Aside from testing the noisy rational expectation hypothesis, this latter test is interesting in itself, as it sheds light on the relationship between the *ex-ante* uncertainty expressed by survey respondents and their *ex-post* ability to predict macroeconomic outcomes, both in the time series and in the cross-section.

Finally, under rational expectations it has to be the case that the mean projection $E_{t-q,i}[y_t]$ better lead to smaller forecast errors on average than any other forecast, because the mean projection minimizes the expected squared forecast error under $E_{t-q,i}[\cdot]$ and therefore, under rational expectations, under $E[\cdot]$ as well:

$$E[(y_t - E_{t-q,i}[y_t])^2] \leq E[(y_t - y_{t,t-q,i}^{pp})^2] \text{ for any } y_{t,t-q,i}^{pp}. \quad (19)$$

We refer to this test as a *location* test, as it assesses whether the predictive densities' mean fulfills its properties under rational expectations. A strand of literature has investigated whether point forecasts coincide with means and, to the extent that they do not, whether this reflects a forecaster's loss function that is not quadratic (e.g., Engelberg et al., 2009; Clements, 2010; Elliott et al., 2008; Patton and Timmermann, 2007). Our test is based on the notion that regardless of the forecaster's loss function, the mean better minimize the square loss if forecasters are rational.

Figure 5: Do forecasters over or under-estimate uncertainty? A scale test



Note: Black dots correspond to OLS estimates of α_q from regression (20) for $q = 8, \dots, 1$. Solid black whiskers indicate 90 percent posterior coverage intervals based on robust standard errors.

III.C.1 A scale test: Do forecasters over or under-estimate uncertainty?

We can assess the hypothesis in (17) by testing whether $\alpha_q = 1$ in the panel regression


$$(y_t - E_{t-q,i}[y_t])^2 / \sigma_{t|t-q,i}^2 = \alpha_q + \epsilon_{t,i,q}, \quad t = 1, \dots, T, \quad i = 1, \dots, N. \quad (20)$$

where we use the posterior means of $E_{t-q,i}[y_t]$ and $\sigma_{t|t-q,i}^2$ from our approach.¹⁶ Estimates of α_q that are significantly greater/lower than 1 indicate that forecasters under/over estimate uncertainty. Figure 5 shows estimates of α_q for different horizons, ranging from $q = 8$ (H2Q1) to $q = 1$ (H1Q4) (recall that the variables being forecasted are the year-over-year growth rates of output or the price level). The dots indicate the OLS point estimates and the whiskers the two-standard deviations posterior intervals, computed using robust standard errors (see Müller, 2013, for a discussion of the use of the sandwich matrix in a Bayesian context). The figure shows that for horizons between two and one-and-a-half years (e.g., $q = 6, 7$ or 8) α_q is significantly larger than 1 for both output growth and inflation. In fact, for output growth α_q is about 3, indicating that forecasters grossly underestimate uncertainty, in line with the literature on overconfidence (Daniel and Hirshleifer, 2015; Malmendier and Taylor, 2015). For horizons closer to one year ($q = 5, 4$) α_q remains well above 1 for output, but

¹⁶In line with most of the literature evaluating forecasting accuracy from surveys in this section we use the so-called “first final” estimate as a measure for y_t , where the first final is the third estimate of GDP or the GDP deflator. Results obtained using the latest revision are shown in the Appendix. These results are very similar for longer horizons, but tend to be different for shorter horizons arguably because the revised series contain methodological changes in measuring GDP that forecasters simply cannot anticipate.

is not significantly different from 1 for inflation. For shorter horizons (q lower than 3) α_q is significantly lower than 1, indicating that forecasters overestimate uncertainty. The overestimation is sizable for inflation, with estimates hovering around .5, but less so for output. For $q = 1$ the estimate of α_q is barely significantly below 1. Figure A-19 in the Appendix shows that these results do not change much across different sub-samples (eg, excluding the Covid period and/or the period 1982-1991 when the Philadelphia Fed was not in charge of the survey).

The idea behind the regression in (20) borrows heavily from existing literature. Clements (2014) in particular computes values for $(y_t - E_{t-q,i}[y_t])^2 / \sigma_{t|t-q,i}^2$ using the point predictions in place of the mean $E_{t-q,i}[y_t]$, and estimates of $\sigma_{t|t-q,i}^2$ obtained from fitting a generalized beta distribution. Clements then computes α_{iq} using a time series regression for each forecaster i , tests the hypothesis $\alpha_{iq} > 1$ and $\alpha_{iq} < 1$, and reports the fraction of forecasters for which each hypothesis is rejected.¹⁷




Note: Black dots correspond to OLS estimates of α_q from regression (20) for $q = 8, \dots, 1$ using the posterior means for $E_{t-q,i}[y_t]$ and $\sigma_{t|t-q,i}^2$ from our approach. Gray crosses correspond to OLS estimates when these objects are obtained using the generalized beta approach. Whiskers indicate 90 percent posterior coverage intervals based on robust standard errors.

Broadly speaking, Clements (2014)’s findings are in line with those reported above: at longer horizons forecasters generally tend to be overconfident, and this overconfidence diminishes as the horizon gets shorter. The benefit of running a panel regression as in (20)

¹⁷This exercise is conducted for US SPF surveys for output growth and inflation from 1981Q3 to 2010Q4. Giordani and Soderlind (2003) also find that forecasters underestimate inflation uncertainty. Similarly, Diebold et al. (1999) and Rossi and Sekhposyan (2015, 2018) investigated whether uncertainty is over or underestimated but focus on the consensus (that is, average) predictive density.

is twofold. First, we explicitly test whether predictive distributions are correctly scaled using the entire panel, rather than forecaster by forecaster, thereby getting a clear answer on whether the rational expectation hypothesis is rejected or not for the SPF. Second, we obtain quantitative estimates of the average degree of over/under-confidence that are not marred by the small sample problem affecting individual forecasters' regressions. The finding that at longer horizons forecasters are as much as 3 times and one and half times as confident as they should be for output growth and inflation, respectively, for instance, was not known to our knowledge. Also, previous literature mostly used point forecasts, while of course under rational expectations equation (17) holds for the mean, but not necessarily for the point forecast if this differs from the mean (Figure A-17 in the Appendix shows that the results for the point forecasts are not very different at long horizons, but can be quite different at short horizons). Finally, Figure 6 shows that it makes a big difference whether one uses the posterior mean of $\sigma_{t|t-q,i}^2$ from our approach or that obtained from fitting a generalized beta distribution, especially at long horizons where forecasters place more probability on the open bins.¹⁸


Figure 7: Do differences in subjective uncertainty map into differences in forecast accuracy? A variation test



Note: Black dots correspond to OLS estimates of $\beta_{1,q}$ from regression (21) for $q = 8, \dots, 1$. Solid black whiskers indicate 90 percent posterior coverage intervals based on robust standard errors.

¹⁸Glas and Hartmann (2022) conduct Clement's exercise distinguishing between rounders (respondents who round up to 0 and/or to round numbers) and non-rounders, and find that rounders tend to underestimate uncertainty, especially for long horizons, while non-rounders do not. Our approach takes rounding explicitly into account and in fact we find that the underestimation of uncertainty at long horizons is about half that implied by the generalized beta approach.

Figure 9: A variation test—accounting for inference uncertainty (baseline vs weighted OLS)



Note: Black dots correspond to OLS estimates of $\beta_{1,q}$ from regression (21) for $q = 8, \dots, 1$. Solid black whiskers indicate 90 percent posterior coverage intervals based on robust standard errors. Gray crosses correspond to weighted OLS estimates, where the weights are inversely proportional to inference uncertainty as measured by the interquartile range of the posterior distribution of $\sigma_{t|t-q,i}$. Whiskers indicate 90 percent posterior coverage intervals based on robust standard errors.

III.C.2 A variation test: Do differences in subjective uncertainty map into differences in forecast accuracy?

Next, we explore a different implication of the noisy rational expectations hypothesis: subjective uncertainty and forecast accuracy should co-move, both across forecasters and over time, as highlighted by equation (18). We test this hypothesis by testing whether in the panel regression

$$\ln |y_t - E_{t-q,i}[y_t]| = \beta_{0,q} + \beta_{1,q} \ln \sigma_{t|t-q,i} + \epsilon_{t,i,q}, \quad t = 1, \dots, T, \quad i = 1, \dots, N, \quad (21)$$

the coefficient $\beta_{1,q}$ is equal to one. As before, equation (21) is estimated via OLS where $E_{t-q,i}[y_t]$ and $\sigma_{t|t-q,i}$ are measured using the posterior mean of the standard deviation estimated using our approach, and robust standard errors are computed. Figure 7 plots the point estimates of $\beta_{1,q}$ for different horizons (crosses) and the whiskers denote the two-standard deviations posterior intervals.

We observe that, quite strikingly, for output growth there is essentially no relationship between subjective uncertainty and the size of the *ex-post* forecast error for horizons above one year. As the forecast horizon shortens the relationship becomes tighter, and for $q = 1$ one cannot reject the hypothesis that $\beta_{1,1} = 1$. For inflation the estimates of $\beta_{1,q}$ hover around 0.5 for longer horizons, but increase toward 1 as the horizon shortens, with $\beta_{1,1}$ that is also not significantly different from 1.

Figure 8 shows the estimates of $\beta_{1,q}$ controlling for time, forecaster, and both time and forecaster fixed effects in order to ascertain whether the results in Figure 7 are mostly due to differences across forecasters or over time. The results with time fixed effects (top panels) indicate that for output growth it is generally not the case that for longer horizons forecasters with lower/higher subjective uncertainty have lower/higher absolute forecast errors on average, although for short horizons the correspondence between the two becomes tighter. At longer horizons there is little relationship also for inflation, although for $q \leq 2$ the coefficient $\beta_{1,q}$ is one or very close to. Similarly, the results with forecaster fixed effects (middle panels) suggest that when forecasters change their subjective uncertainty, possibly because the quality of their private signal has changed, on average this maps one-to-one into corresponding changes in the absolute forecast errors for horizons close to one quarter, but not for longer horizons. The bottom panels of Figure 8 show that these findings still hold when we include both forecaster and time fixed effects. The Appendix shows that all these results are broadly robust to different samples.

To our knowledge, both the idea of testing the noisy rational expectations hypothesis using the variation test and this set of results are new to the literature. Clements (2014) computes time series averages of $\sigma_{t|t-q,i}$ for each forecaster and plots them against the corresponding predictive root mean square error (RMSE) computed during the same period.¹⁹ Clements concludes that “there is little evidence that more (less) confident forecasters are more (less) able forecasters.” This exercise compares to our model with time fixed effects, where we study whether forecasters that are more uncertain also have higher absolute forecast errors. Our results agree with Clements for output growth and inflation at long horizons, but differ at short horizons. One reason for the difference is that Clements uses point forecasts while expression (18) only holds for the mean: if absolute forecast errors are computed using predictions other than the mean, there is no *a priori* reason why there should be a correspondence with the subjective standard deviation, even under rational expectations. In fact Figure A-20 in the Appendix shows that when we use the point predictions the correspondence between subjective uncertainty and forecast error vanishes at short horizons.²⁰

¹⁹Clements adjusts for the unbalanced nature of the sample—that is, the fact that each forecaster’s average is computed for a different time period—by constructing weighted averages where the weights reflect the average forecast error or subjective uncertainty during that period.

²⁰A more proper comparison with Clements cross sectional results is in Figure A-21 in the Appendix where we show the results with time fixed effects and point forecasts. Indeed we find that most coefficients are not significantly different from 0 for point forecasts. Figure A-22 shows the results obtained using the generalized beta approach, which are similar to those shown in Figure 7.

Most important, the purely cross-sectional comparison undertaken so far by the literature misses the time dimension of our regression, where we investigate whether changes in subjective uncertainty over time actually map into changes in forecasting performance. This aspect is particularly important as it sheds light on whether forecasters correctly anticipate periods of macroeconomic uncertainty. The finding that in the time dimension (that is, using forecaster fixed effects and both time and forecaster fixed effects) at longer horizons the mapping between subjective uncertainty and forecast accuracy is just not there for output, and is only partial for inflation, but is in line with the noisy rational expectations model for both output and inflation at short horizons, is entirely novel.

Last, one benefit of our approach is that we can measure inference uncertainty about $\sigma_{t|t-q,i}$. We can therefore assess to what extent such uncertainty may be driving the results in Figure 7. We do so by running a weighted OLS panel regression, where the weights are inversely proportional to inference uncertainty as measured by the interquantile range of the posterior distribution of $\sigma_{t|t-q,i}$. Figure 9 shows that the weighted OLS results are nearly identical to the results in 7, assuaging concerns of a bias driven by inference uncertainty. The Appendix shows that the weighted results are very similar to the unweighted ones also whenever we use fixed effects and/or different samples.

III.C.3 A location test: The relative accuracy of mean and point predictions

Last, we turn to the *location* test, where we use the point forecast as an alternative to the mean projections in testing (19). The top panels of Figure 10 shows OLS estimates of the coefficient γ_q in the panel regression

$$\ln \frac{(y_t - E_{t-q,i}[y_t])^2}{(y_t - y_{t,t-q,i}^{pp})^2} = \gamma_q + \epsilon_{t,i,q}, \quad t = 1, \dots, T, \quad i = 1, \dots, N, \quad (22)$$

where $y_{t,t-q,i}^{pp}$ is the point forecast for y_t made by forecaster i in period $t - q$. Estimates of γ_q significantly greater than zero indicate that on average point fare worse than mean projections in terms of mean squared error. In fact, these estimates can be interpreted as the percentage improvement/worsening in forecast accuracy for point relative to mean projections. For horizons longer than one year estimates of γ_q are not significantly different from zero for output growth, and only slightly positive for inflation. This result may partly reflect the fact that for these horizons point and mean predictions are not very different (see

Engelberg et al., 2009). As the horizon gets shorter the estimates tend to become much larger and significantly positive for both output growth and inflation.

The result that point forecasts perform better than mean forecasts in terms of mean squared error for short horizons is not new to the literature: Clements (2010) reports mean squared forecast errors for horizons shorter than one year and find that these are lower for point than for mean projections. As in Clements (2010), we interpret these results explicitly as an indirect test of the rationality of density projections: under rational expectations, it better be that the mean of the predictive distribution produces a lower mean squared error than any other point prediction regardless of the forecasters' loss function. The fact that for short horizons this is clearly not the case casts some doubt on explanations for the divergence between mean and point forecasts that rely on the forecasters' loss function (e.g., Patton and Timmermann, 2007; Elliott et al., 2008; Lahiri and Liu, 2009).

As a further test of the rationality of mean projections, we also run the Fair and Shiller (1990) regression

$$y_t = \delta_{0,q} + \delta_{1,q}E_{t-q,i}[y_t] + \delta_{2,q}y_{t-q,i}^{pp} + \epsilon_{t,i,q}, \quad t = 1, \dots, T, \quad i = 1, \dots, N. \quad (23)$$

The rationality of density projections would imply $\delta_{0,q} = 0$, $\delta_{1,q} = 1$, and $\delta_{2,q} = 0$. If point projections $y_{t,t-q,i}^{pp}$ coincide with mean forecasts then the two regressors are multicollinear. The middle panels of Figure 10 report estimates of $\delta_{1,q}$ (black crosses) and $\delta_{2,q}$ (gray diamonds) for different horizons q , while the bottom panels report estimates for the constant $\delta_{0,q} = 1$. For horizons longer than one year, estimates of $\delta_{1,q}$ are generally larger than those for $\delta_{2,q}$. Estimates for $\delta_{1,q}$ are significantly below 1, and estimates for the constant are significantly different from zero. As the horizon shortens, estimates for the constant become closer to 0, in line with rational expectations, but estimates of $\delta_{2,q}$ rise toward 1 while estimates of $\delta_{1,q}$ fall to zero, indicating that point predictions are much closer to actual outcomes than mean forecasts.

III.C.4 Summing up: Are SPF density forecasts consistent with the noisy rational expectation hypothesis?

The body of evidence collected in this section suggests that the answer is no. For horizons close to two years there is strong evidence that 1) forecasters are overconfident, and 2) there is virtually no relationship between differences in subjective uncertainty both across forecasters and over time and differences in forecasting performance. For horizons close to

one year we cannot reject that inflation density forecasts are correctly scaled, while output growth density forecasts flip from being overconfident to being underconfident. The mapping between *ex-ante* uncertainty and ex-post forecast errors is far from one, however, for both output and inflation projections. For very short horizons, density forecasts are correctly scaled for output growth, and slightly underconfident for inflation. For both variables there is (almost) a one-to-one mapping between subjective and ex-post uncertainty, both across forecasters and over time, in accordance with the noisy rational expectation hypothesis. But while the second moments of the density projections seems to line up with theory at short horizons, the first moments do not: mean projections deliver higher mean squared errors than point projections. In sum, we reach a similar conclusion for density projections as Patton and Timmermann (2010) reach for point forecasts, namely that differences across forecasters (and in our case also over time) cannot be explained by differences in information sets.

IV Conclusions

In this paper we presented a novel approach for conducting inference on data from probabilistic surveys, and used it to investigate whether U.S. Survey of Professional Forecasters density projections for output growth and inflation are consistent with the noisy rational expectations hypothesis. We find that for horizons close to two years there is no correspondence between subjective uncertainty and forecast accuracy for output growth density projections, both across forecasters and over time, and only a very mild correspondence for inflation projections, in contrast to what rational expectations would predict. As the horizons shortens, the relationship becomes one-to-one, in accordance with the theory.

While the inference approach we propose arguably several advantages relative to current practice—for starters, the fact that we explicitly conduct inference—it is important to point out some limitations of our analysis. We provided some consistency results that take advantage of the nonparametric nature of the approach, but these only apply to the model as a data generating process for the data that we observe—the bin probabilities. Regarding the objects we are truly interested in—the underlying continuous predictive densities—consistency results are only available in the unrealistic case that the number of bins goes to infinity and the bin width goes to zero. When these conditions are not met, the limited information provided by forecasters implies that posterior uncertainty regarding the objects of interest remains even when the number of forecasters goes to infinity, simply because there is not

enough information to identify the underlying predictive densities. This implies that the results obtained with our approach may be sensitive to the choice of the base function and of priors. More work needs to be done in this dimension.

In addition, the approach proposed in this paper deals with one survey (one cross-section) and one forecast variable at the time. It would be interesting to extend the approach to a panel context, which would permit joint inference across surveys for any object of interest. We leave this extension to future research.

References

- Andrade, Philippe and Hervé Le Bihan**, “Inattentive professional forecasters,” *Journal of Monetary Economics*, 2013, 60 (8), 967–982.
- Andrieu, Christophe and Johannes Thoms**, “A tutorial on adaptive MCMC,” *Statistics and Computing*, Dec 2008, 18 (4), 343–373.
- Barrientos, Andrés F., Alejandro Jara, and Fernando A. Quintana**, “Bayesian density estimation for compositional data using random Bernstein polynomials,” *J. Statist. Plann. Inference*, 2015, 166, 116–125.
- Bassetti, Federico, Roberto Casarin, and Fabrizio Leisen**, “Beta-product dependent Pitman–Yor processes for Bayesian inference,” *Journal of Econometrics*, 2014, 180 (1), 49–72.
- , – , **and Francesco Ravazzolo**, “Bayesian Nonparametric Calibration and Combination of Predictive Distributions,” *Journal of the American Statistical Association*, 2018, 113 (522), 675–685.
- , – , **and Marco Del Negro**, “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 R, ed., *Handbook of Economic Expectations*, Elsevier, forthcoming.
- Binder, Carola C**, “Measuring uncertainty based on rounding: New method and application to inflation expectations,” *Journal of Monetary Economics*, 2017, 90, 1–12.

Boero, Gianna, Jeremy Smith, and Kenneth F. Wallis, “Uncertainty and disagreement in economic prediction: the Bank of England Survey of External Forecasters,” *The Economic Journal*, 2008, *118* (530), 1107–1127.

– , – , **and** – , “The Measurement and Characteristics of Professional Forecasts’ Uncertainty,” *Journal of Applied Econometrics*, 2014, *7* (30), 1029–1046.

Bruine De Bruin, Wändi, Charles F. Manski, Giorgio Topa, and Wilbert Van Der Klaauw, “Measuring consumer uncertainty about future inflation,” *Journal of Applied Econometrics*, 2011, *26* (3), 454–478.

Burda, Martin, Matthew Harding, and Jerry Hausman, “A Bayesian semiparametric competing risk model with unobserved heterogeneity,” *Journal of Applied Econometrics*, 2015, *30* (3), 353–376.

Capistrán, Carlos and Allan Timmermann, “Disagreement and biases in inflation expectations,” *Journal of Money, Credit and Banking*, 2009, *41* (2-3), 365–396.

Chib, Siddhartha and Barton H. Hamilton, “Semiparametric Bayes analysis of longitudinal data treatment models,” *Journal of Econometrics*, 2002, *110* (1), 67–89.

Clements, Michael P., “Explanations of the inconsistencies in survey respondents’ forecasts,” *European Economic Review*, 2010, *54* (4), 536–549.

– , “Forecast Uncertainty—Ex Ante and Ex Post: US Inflation and Output Growth,” *Journal of Business & Economic Statistics*, 2014, *32* (2), 206–216.

– , **Robert Rich, and Joseph Tracy**, “Surveys of Professionals,” in R, ed., *Handbook of Economic Expectations*, Elsevier, forthcoming.

Coibion, Olivier and Yuriy Gorodnichenko, “What can survey forecasts tell us about information rigidities?,” *Journal of Political Economy*, 2012, *120* (1), 116–159.

– **and** – , “Information rigidity and the expectations formation process: A simple framework and new facts,” *American Economic Review*, 2015, *105* (8), 2644–78.

D’Amico, Stefania and Athanasios Orphanides, “Uncertainty and disagreement in economic forecasting,” Technical Report, Board of Governors of the Federal Reserve System (US) 2008.

- Daniel, Kent and David Hirshleifer**, “Overconfident investors, predictable returns, and excessive trading,” *Journal of Economic Perspectives*, 2015, 29 (4), 61–88.
- Diebold, Francis X, Anthony S Tay, and Kenneth F Wallis**, “Evaluating density forecasts of inflation: the Survey of Professional Forecasters,” in R.F. Engle and H. White, eds., *Cointegration, Causality, and Forecasting: A Festschrift in Honour of Clive WJ Granger*, Oxford University Press, 1999.
- Dominitz, Jeff and Charles F. Manski**, “Eliciting student expectations of the returns to schooling,” *Journal of Human resources*, 1996, pp. 1–26.
- Elliott, Graham, Ivana Komunjer, and Allan Timmermann**, “Biases in macroeconomic forecasts: irrationality or asymmetric loss?,” *Journal of the European Economic Association*, 2008, 6 (1), 122–157.
- Engelberg, Joseph, Charles F. Manski, and Jared Williams**, “Comparing the point predictions and subjective probability distributions of professional forecasters,” *Journal of Business & Economic Statistics*, 2009, 27 (1), 30–41.
- Fair, Ray C and Robert J Shiller**, “Comparing information in forecasts from econometric models,” *The American Economic Review*, 1990, pp. 375–389.
- Ferguson, T. S.**, “A Bayesian analysis of some nonparametric problems,” *Annals of Statistics*, 1973, 1, 209–230.
- Ganics, Gergely, Barbara Rossi, and Tatevik Sekhposyan**, “From fixed-event to fixed-horizon density forecasts: obtaining measures of multi-horizon uncertainty from survey density forecasts,” 2020.
- Ghosh, J. K. and R. V. Ramamoorthi**, *Bayesian nonparametrics* Springer Series in Statistics, Springer-Verlag, New York, 2003.
- Ghoshal, S., J. K. Gosh, and R. V. Ramamoorthi**, “Consistent semiparametric Bayesian inference about a location parameter,” *Journal of Statistical Planning and Inference*, 1999, 77 (2), 181–193.
- Giordani, Paolo and Paul Soderlind**, “Inflation forecast uncertainty,” *European Economic Review*, 2003, 47, 1037–1059.

- Giustinelli, Pamela, Charles F. Manski, and Francesca Molinari**, “Tail and center rounding of probabilistic expectations in the health and retirement study,” *Journal of Econometrics*, 2020.
- Glas, Alexander and Matthias Hartmann**, “Uncertainty Measures from Partially Rounded Probabilistic Forecast Surveys,” *Quantitative Economics*, 2022, *13*, 979 – 1022.
- Griffin, Jim E.**, “Inference in infinite superpositions of non-Gaussian Ornstein–Uhlenbeck processes using Bayesian nonparametric methods,” *Journal of Financial Econometrics*, 2011, *9* (3), 519–549.
- **and Mark F.J. Steel**, “Semiparametric Bayesian inference for stochastic frontier models,” *Journal of econometrics*, 2004, *123* (1), 121–152.
- **, Fernando Quintana, and Mark F.J. Steel**, “Flexible and Nonparametric Modeling,” in John Geweke, Gary Koop, and Herman K. van Dijk, eds., *Handbook of Bayesian Econometrics*, Oxford University Press, 2011.
- Gu, Jiaying and Roger Koenker**, “Unobserved heterogeneity in income dynamics: An empirical Bayes perspective,” *Journal of Business & Economic Statistics*, 2017, *35* (1), 1–16.
- Hirano, Keisuke**, “Semiparametric Bayesian Inference in Autoregressive Panel Data Models,” *Econometrica*, 2002, *70*, 781–799.
- Jensen, Mark J. and John M. Maheu**, “Bayesian semiparametric stochastic volatility modeling,” *Journal of Econometrics*, 2010, *157* (2), 306–316.
- Kalli, Maria, Jim E. Griffin, and Stephen G. Walker**, “Slice sampling mixture models,” *Statistics and Computing*, 2011, *21*, 93–105.
- Kozeniauskas, Nicholas, Anna Orlik, and Laura Veldkamp**, “What are uncertainty shocks?,” *Journal of Monetary Economics*, 2018, *100*, 1–15.
- Lahiri, Kajal and Fushang Liu**, “Modelling multi-period inflation uncertainty using a panel of density forecasts,” *Journal of Applied Econometrics*, 2006, *21* (8), 1199–1219.
- **and –**, “On the use of density forecasts to identify asymmetry in forecasters’ loss function,” *Business and Economic Statistics Section-JSM*, 2009, pp. 2396–2408.

- Liu, Laura**, “Density forecasts in panel data models: A semiparametric Bayesian perspective,” *Journal of Business & Economic Statistics*, 2021, (just-accepted), 1–42.
- Liu, Yang and Xuguang Simon Sheng**, “The measurement and transmission of macroeconomic uncertainty: Evidence from the U.S. and BRIC countries,” *International Journal of Forecasting*, 2019, *35* (3), 967–979.
- Malmendier, Ulrike and Timothy Taylor**, “On the verges of overconfidence,” *Journal of Economic Perspectives*, 2015, *29* (4), 3–8.
- Mankiw, N Gregory, Ricardo Reis, and Justin Wolfers**, “Disagreement about inflation expectations,” *NBER macroeconomics annual*, 2003, *18*, 209–248.
- Manski, Charles F.**, “Measuring expectations,” *Econometrica*, 2004, *72* (5), 1329–1376.
- , “Interpreting and combining heterogeneous survey forecasts,” in Michael P. Clements and D. F. Hendry, eds., *Oxford handbook of economic forecasting*, Vol. 85, Oxford University Press, 2011, pp. 457–472.
- , “Survey measurement of probabilistic macroeconomic expectations: progress and promise,” *NBER Macroeconomics Annual*, 2018, *32* (1), 411–471.
- **and Francesca Molinari**, “Rounding probabilistic expectations in surveys,” *Journal of Business & Economic Statistics*, 2010, *28* (2), 219–231.
- Müller, Ulrich K**, “Risk of Bayesian inference in misspecified models, and the sandwich covariance matrix,” *Econometrica*, 2013, *81* (5), 1805–1849.
- Norets, Andriy and Justinas Pelenis**, “Bayesian modeling of joint and conditional distributions,” *Journal of Econometrics*, 2012, *168* (332-346).
- **and** – , “Posterior consistency in conditional density estimation by covariate dependent mixtures,” *Econometric Theory*, 2014, *30* (3), 606–646.
- Patton, Andrew J and Allan Timmermann**, “Properties of optimal forecasts under asymmetric loss and nonlinearity,” *Journal of Econometrics*, 2007, *140* (2), 884–918.
- **and** – , “Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion,” *Journal of Monetary Economics*, 2010, *57* (7), 803–820.

- **and** –, “Predictability of output growth and inflation: A multi-horizon survey approach,” *Journal of Business & Economic Statistics*, 2011, 29 (3), 397–410.
- Pelenis, Justinas**, “Bayesian regression with heteroscedastic error density and parametric mean function,” *Journal of Econometrics*, 2014, 178, 624–638.
- Pitman, Jim**, *Combinatorial Stochastic Processes*, Vol. 1875, Springer-Verlag, 2006.
- Potter, Simon**, “The advantages of probabilistic survey questions: remarks at the IT Forum and RCEA Bayesian Workshop, keynote address, Rimini, Italy, May 2016,” Technical Report, Federal Reserve Bank of New York 2016.
- Rich, Robert and Joseph Tracy**, “A closer look at the behavior of uncertainty and disagreement: Micro evidence from the euro area,” *Journal of Money, Credit and Banking*, 2021, 53 (1), 233–253.
- Rossi, Barbara and Tatevik Sekhposyan**, “Macroeconomic Uncertainty Indices Based on Nowcast and Forecast Error Distributions,” *American Economic Review*, May 2015, 105 (5), 650–55.
- **and** –, “Alternative tests for correct specification of conditional predictive densities,” *Journal of Econometrics*, 2018.
- Sethuraman, Jayaram**, “A constructive definition of Dirichlet priors,” *Statistica Sinica*, 1994, 4, 639–650.
- Stark, Tom**, “SPF panelists’ forecasting methods: A note on the aggregate results of a November 2009 special survey,” *Federal Reserve Bank of Philadelphia*, 2013.
- Walker, Stephen G.**, “Sampling the Dirichlet mixture model with slices,” *Communications in Statistics – Simulation and Computation*, 2007, 36, 45–54.
- Wu, Yuefeng and Subhashis Ghosal**, “Correction to: “Kullback Leibler property of kernel mixture priors in Bayesian density estimation” [MR2399197],” *Electron. J. Stat.*, 2009, 3, 316–317.
- **and** –, “Kullback Leibler property of kernel mixture priors in Bayesian density estimation,” *Electron. J. Stat.*, 2009, 2, 298–331.
- Zadora, G., T. Neocleous, and Aitken C.**, “A Two-Level Model for Evidence Evaluation in the Presence of Zeros,” *Journal of Forensic Sciences*, 2010, 55 (2), 371–384.

Zarnowitz, Victor and Louis A. Lambros, "Consensus and Uncertainty in Economic Prediction," *Journal of Political Economy*, 1987, (95), 591 – 621.

Appendix

A Data description

We focus on the Survey of Professional Forecasters, managed since 1992 by the Federal Reserve Bank of Philadelphia, and previously by the American Statistical Association and the National Bureau of Economic Research. The panel of forecasters include university professors and private-sector macroeconomic researchers, and the composition of the panel changes gradually over time. The survey, which is performed quarterly, is mailed to panel members the day after the government release of quarterly data on the national income and product accounts. We restrict our attention to the two variables for which the SPF has probabilistic questions, namely year-over-year GDP growth and GDP deflator inflation over the sample 1982Q1-2021Q4.

B The Gibbs Sampler

The infinite mixture model is

$$h_G(\mathbf{z}) = \int h(\mathbf{z}|\boldsymbol{\theta})G(d\boldsymbol{\theta}) = \sum_{k=1}^{\infty} w_k h(\mathbf{z}|\boldsymbol{\theta}_k^*). \quad (\text{A-1})$$

Our Gibbs sampler applied to the cross section of \mathbf{z}_i , $i = 1, \dots, n$ uses the convenient approach proposed by Walker (2007) and Kalli et al. (2011). For each forecaster i , conditional on the sequence of weights w_k 's ($w_{1:\infty}$) and the sequence of atoms $\boldsymbol{\theta}_k^*$'s ($\boldsymbol{\theta}_{1:\infty}$), expression (A-1) can be written as the marginal distribution of

$$h(\mathbf{z}_i, u_i | w_{1:\infty}, \boldsymbol{\theta}_{1:\infty}^*) = \sum_{k=1}^{\infty} \mathbb{I}(u_i < w_k) h(\mathbf{z}_i | \boldsymbol{\theta}_k^*) \quad (\text{A-2})$$

with respect to u_i , where u_i is uniformly distributed over the interval $[0, 1]$, and independent across i , and $\mathbb{I}(\cdot)$ is an indicator function. This implies that the conditional distribution of \mathbf{z}_i given u_i , the weights and the atoms, is

$$h(\mathbf{z}_i | u_i, w_{1:\infty}, \boldsymbol{\theta}_{1:\infty}^*) = \frac{1}{h(u_i | w_{1:\infty})} \sum_{k \in A(u_i | w_{1:\infty})} h(\mathbf{z}_i | \boldsymbol{\theta}_k^*), \quad (\text{A-3})$$

where the set $A(u_i|w_{1:\infty})$ includes all the atoms with a weight w_k larger than u_i ($A(u_i|w_{1:\infty}) = \{k : u_i < w_k\}$), and the marginal $h(u_i|w_{1:\infty}) = \sum_{k=1}^{+\infty} \mathbb{I}(u_i < w_k)$ since each $h(\cdot|\boldsymbol{\theta}_k^*)$ integrates to one. Unlike expression (A-1), expression (A-3) is a *finite* mixture where each component has probability $\frac{1}{h(u_i|w_{1:\infty})}$, which is straightforward to draw from using standard methods. Specifically, we will use the auxiliary indicators d_i 's, which are equal to k if we draw from the k^{th} mixture component (note that, given u_i , the k^{th} component will only be drawn if it belongs to the set $A(u_i|w_{1:\infty})$). The resulting complete-data likelihood function is

$$L(\mathbf{z}_{1:n}|u_{1:n}, d_{1:n}, v_{1:\infty}, \boldsymbol{\theta}_{1:\infty}) = \prod_{i=1}^n \mathbb{I}_{\{u_i < w_{d_i}\}} h(\mathbf{z}_i|\boldsymbol{\theta}_{d_i}^*) \quad (\text{A-4})$$

with $d_i \in \{k : u_i < w_k\}$, where $v_{1:\infty}$ is the infinite dimensional sequence containing the stick-breaking components which map into the weights via expression (11).

Let $\mathcal{D}_k = \{i : d_i = k\}$ denote the set of indexes of the observations allocated to the k -th component of the mixture. Let $\mathcal{D} = \{k : \mathcal{D}_k \neq \emptyset\}$ denote the set of indexes of the non-empty mixture components (in the sense that at least one i is using the k^{th} component) and $\bar{d} = \max \mathcal{D}$ the overall number of stick-breaking components used. The Gibbs sampler works as follows:

1. $v_{1:\infty}, u_{1:n}|d_{1:n}, \boldsymbol{\theta}_{1:\infty}^*, \psi, \mathbf{z}_{1:n}$

Call $v_{1:\bar{d}}$ the stick-breaking elements associated with the mixture components that are being used (conditional on $d_{1:n}$). Following Kalli et al. (2011), drawing from the joint posterior of $v_{1:\bar{d}}, v_{\bar{d}+1:\infty}$, and $u_{1:n}$, conditional on all other parameters, is accomplished by drawing sequentially from: (a) the marginal distribution of $v_{1:\bar{d}}$, (b) the conditional distribution of $u_{1:n}$ given $v_{1:\bar{d}}$, and (c) from the conditional distribution of $v_{\bar{d}+1:\infty}$ given $u_{1:n}$ and $v_{1:\bar{d}}$.

- (a) $v_{1:\bar{d}}|d_{1:n}, \boldsymbol{\theta}_{1:\infty}^*, \psi, \mathbf{z}_{1:n}$.

After integrating out the u_i 's, the posterior of $v_{1:\infty}$ is proportional to

$$\begin{aligned} p(v_{1:\infty}|d_{1:n}, \boldsymbol{\theta}_{1:\infty}^*, \psi, \mathbf{z}_{1:n}) &\propto \left(\prod_{i=1}^n w_{d_i} h(\mathbf{z}_i|\boldsymbol{\theta}_{d_i}^*) \right) \left(\prod_{l=1}^{\infty} (1 - v_l)^{\psi-1} \right) \\ &\propto \left(\prod_{i=1}^n \left(v_{d_i} \prod_{l=1}^{d_i-1} (1 - v_l) \right) h(\mathbf{z}_i|\boldsymbol{\theta}_{d_i}^*) \right) \left(\prod_{l=1}^{\infty} (1 - v_l)^{\psi-1} \right). \end{aligned}$$

Now note that since $v_{\bar{d}+1:\infty}$ do not enter the likelihood (A-4) – that is, the term within the first parenthesis – they can be easily integrated out resulting in

$$p(v_{1:\bar{d}}|d_{1:n}, \boldsymbol{\theta}_{1:\infty}^*, \psi, \mathbf{z}_{1:n}) \propto \left(\prod_{i=1}^n \left(v_{d_i} \prod_{l=1}^{d_i-1} (1 - v_l) \right) h(\mathbf{z}_i | \boldsymbol{\theta}_{d_i}^*) \right) \left(\prod_{l=1}^{\bar{d}} (1 - v_l)^{\psi-1} \right).$$

Therefore samples for $v_{1:\bar{d}}$ are obtained by drawing each v_k independently from

$$\pi(v_k | u_{1:n}, d_{1:n}, \dots) \propto (1 - v_k)^{\psi + b_k - 1} v_k^{a_k} \quad (\text{A-5})$$

where $a_k = \sum_{i=1}^n \mathbb{I}(d_i = k)$ and $b_k = \sum_{i=1}^n \mathbb{I}(d_i > k)$, that is, v_k is drawn from a $Beta(a_k + 1, b_k + \psi)$.

(b) $u_{1:n} | v_{1:\bar{d}}, d_{1:n}, \boldsymbol{\theta}_{1:\infty}^*, \psi, \mathbf{z}_{1:n}$.

The likelihood (A-4), seen as a function of each u_i , $i = 1, \dots, n$, is simply a uniform distribution over $[0, w_{d_i}]$. Hence

$$\pi(u_i | \dots) \propto \frac{1}{w_{d_i}} \mathbb{I}(u_i < w_{d_i}). \quad (\text{A-6})$$

(c) $v_{\bar{d}+1:\infty} | u_{1:n}, v_{1:\bar{d}}, d_{1:n}, \boldsymbol{\theta}_{1:\infty}^*, \psi, \mathbf{z}_{1:n}$.

Again, $v_{\bar{d}+1:\infty}$ do not enter the likelihood (A-4), so samples from those v_k with $k > \bar{d}$ are simply obtained by drawing from the prior $Beta(1, \psi)$:

$$\pi(v_k | u_{1:n}, d_{1:n}, \dots) \propto (1 - v_k)^{\psi-1}. \quad (\text{A-7})$$

Of course, even if it is straightforward to execute, we do not want to generate an infinite number of draws. Fortunately we do not need to, as explained in Walker (2007). Inspection of (A-4) reveals that those mixtures for which $w_k < u_i$ will never be used, at least given the the draw for u_i . Let \bar{n}_i the smallest integer such that $\sum_{k=1}^{\bar{n}_i} w_k > 1 - u_i$. Since by construction $\sum_{k=1}^{\infty} w_k = 1$, it must be that $\sum_{k=1}^{\bar{n}_i} w_k < u_i$ and therefore, a fortiori, $w_k < u_i$ for $k > \bar{n}_i$. Now define $\bar{n} = \max\{\bar{n}_i, i = 1, \dots, n\}$. Conditional on $u_{1:n}$, at most we will use \bar{n} mixture components in the estimation. Hence we only need to draw $v_{\bar{d}+1:\bar{n}}$.

2. $\boldsymbol{\theta}_{1:\infty}^* | v_{1:\infty}, u_{1:n}, d_{1:n}, \psi, \mathbf{z}_{1:n}$

For the same argument given above, we actually do not have to draw an infinite number of atoms, but only as many as they may possibly be used (at least given the current draw of $u_{1:n}$) – that is, at most \bar{n} . Note also that given the way the u_i 's are drawn (from a uniform distribution over $[0, w_{d_i}]$), if $k \in \mathcal{D}$ then $k \leq \bar{n}$.

(a) For $k \in \mathcal{D}$ draws of $\boldsymbol{\theta}_k^*$ are obtained from

$$\pi(\boldsymbol{\theta}_k^* | \dots) \propto \left(\prod_{i \in \mathcal{D}_k} h(\mathbf{z}_i | \boldsymbol{\theta}_k^*) \right) G_0(\boldsymbol{\theta}_k^*) \quad (\text{A-8})$$

Since the joint distribution is not tractable, samples have been generated by Adaptive Metropolis Hastings (AMH) proposed in Andrieu and Thoms (2008). More specifically, at the j -th iteration of the AMH for a parameter $\boldsymbol{\theta}^*$ of dimension p the proposal distribution is

$$\boldsymbol{\theta}^{new} \sim \mathcal{N}(\boldsymbol{\theta}^{(j-1)*}, \Upsilon^{(j)}) \quad (\text{A-9})$$

with covariance matrix $\Upsilon^{(j)} = \exp\{\xi^{(j)}\} I_p$ where $\xi^{(j)}$ is adapted over the iterations as follows

$$\xi^{(j)} = \xi^{(j-1)} + \gamma^{(j)}(\hat{\alpha}^{(j-1)} - \bar{\alpha}) \quad (\text{A-10})$$

where $\bar{\alpha} = 0.3$ represents the desired level of acceptance probability, and $\hat{\alpha}^{(j-1)}$ is the previous iteration estimate of the acceptance probability (i.e. the acceptance rate). The diminishing adaptation condition is satisfied by choosing $\gamma^{(j)} = j^{(-a)}$. In the application we set $a = 0.7$.

(b) For $k \notin \mathcal{D}$, $k \leq \bar{n}$ draws of $\boldsymbol{\theta}_k^*$ are obtained via independent draws from the base measure G_0 .

We therefore obtained a sequence of draws $\boldsymbol{\theta}_{1:\bar{n}}^*$, which we will use in the next Gibbs sampler step.

3. $d_{1:n} | v_{1:\infty}, u_{1:n}, \boldsymbol{\theta}_{1:\infty}^*, \psi, \mathbf{z}_{1:n}$

Draws for each d_i , $i = 1, \dots, n$, are obtained by drawing from a multinomial with weights proportional to

$$\pi(d_i | \dots) \propto \mathbb{I}(u_i < w_{d_i}) h(\mathbf{z}_i | \boldsymbol{\theta}_{d_i}^*) \quad (\text{A-11})$$

with $d_i \in \{1, \dots, \bar{n}_i\}$. Note that in this draw we consider all possible mixture components from 1 to \bar{n}_i , not only those used so far (that is, those in \mathcal{D}). They will be drawn proportionally to their ability to fit of the data, as measured by $h(\mathbf{z}_i | \boldsymbol{\theta}_k^*)$.

C Further theoretical results

C.A Model properties

In this section, we present some properties which illustrate the flexibility of our nonparametric random histogram model. The behaviour of the model as the number of bins goes to infinity shows that our framework is theoretically sound since it can be used to approximate any subjective distribution when (2) holds.

Let $\mathbf{z}_i, i = 1, \dots, n$ be i.i.d. samples from $h(\mathbf{z}|\boldsymbol{\theta})$ and assume the forecasters never report zero probabilities (that is, conditional on $\xi_{ij} = 0 \forall j$), then in expectation z_{ij} coincides with ν_j : $\mathbb{E}[z_{ij}|\boldsymbol{\theta}] = \nu_j(\boldsymbol{\theta})$. Expression (A-1) then implies that the distribution of each z_{ij} will be centered at the infinite mixture of the bin probabilities ν_j 's implied by each mixture component $F(\cdot|\boldsymbol{\theta}_k)$:

$$\mathbb{E}[z_{ij}|G] = \sum_{k=1}^{\infty} w_k \nu_j(\boldsymbol{\theta}_k) = \sum_{k=1}^{\infty} w_k (F(y_j|\boldsymbol{\theta}_k) - F(y_{j-1}|\boldsymbol{\theta}_k)). \quad (\text{A-12})$$

We show that our random histogram (prior) model converges to an infinite dimensional (prior) model approximating any subjective distribution in the topology of weak convergence. This flexibility implies that the nonparametric prior alleviates possible misspecification issues.

Introduce a latent Dirichlet process $Z_{i,\infty}(\cdot)|\boldsymbol{\theta}_i \sim \mathcal{DP}(\phi(\boldsymbol{\theta}_i), F(\cdot|\boldsymbol{\theta}_i))$ with parameters $\phi(\boldsymbol{\theta}_i)$ and $F(\cdot|\boldsymbol{\theta}_i)$, given $\boldsymbol{\theta}_i$ from G . This process defines a random measure on the observation space \mathcal{Y} of the variable of interest (inflation), that is the support set of the subjective distribution $F(\cdot|\boldsymbol{\theta})$, and admits the equivalent stick breaking representation

$$Z_{i,\infty}(y) = \sum_{j=1}^{\infty} w_{ij} \mathbb{I}\{y_{ij} \leq y\} \quad (\text{A-13})$$

where $y_{ij} \ j = 1, 2, \dots$ are i.i.d. random variables with common distribution $F(\cdot|\boldsymbol{\theta}_i)$ and $w_{ij} \ j = 1, 2, \dots$ are obtained by using a sequence of i.i.d. $\mathcal{Be}(1, \phi(\boldsymbol{\theta}_i))$ random variables.

Proposition 1. If $\varrho_j = 0$ for $j = 1, \dots, J$, the Bayesian model

$$\begin{aligned} \mathbf{z}_i|G &\stackrel{ind}{\sim} h_G(\mathbf{z}), \quad i = 1, \dots, n \\ G &\sim \mathcal{DP}(\psi, G_0) \end{aligned}$$

where $\mathbf{z}_i = (z_{i1}, \dots, z_{iJ})$ admits the following stochastic representation:

$$\begin{aligned} (z_{i1}, \dots, z_{iJ}) &:= (Z_{i,\infty}(y_1), Z_{i,\infty}(y_2) - Z_{i,\infty}(y_1), \dots, 1 - Z_{i,\infty}(y_{J-1})) \quad i = 1, \dots, n \\ Z_{i,\infty} &\stackrel{\text{ind}}{\sim} \mathcal{DP}(\phi(\boldsymbol{\theta}_i), F(\cdot|\boldsymbol{\theta}_i)) \quad i = 1, \dots, n \\ \boldsymbol{\theta}_i &\stackrel{\text{i.i.d.}}{\sim} G \quad i = 1, \dots, n \\ G &\sim \mathcal{DP}(\psi, G_0). \end{aligned}$$

The previous Proposition suggests the following interpretation: given the true subjective probability distribution $F(\cdot|\boldsymbol{\theta}_i)$ of the i -th forecaster and its level of noise $\phi(\boldsymbol{\theta}_i)$, the forecaster reports the weights (z_{i1}, \dots, z_{iJ}) corresponding to the increments of a “noisy” version $Z_{i,\infty}$ of $F(\cdot|\boldsymbol{\theta}_i)$. This “noisy” version is the CDF obtained by a Dirichlet process with base measure $F(\cdot|\boldsymbol{\theta}_i)$ and concentration parameter $\phi(\boldsymbol{\theta}_i)$.

By (A-13), the latent Dirichlet process $Z_{i,\infty}$ is a random discrete CDF with infinite number of discontinuity points. To exemplify we depict $Z_{i,\infty}$ by the red stepwise line in Figure A-1. Despite of its discreteness, the process $Z_{i,\infty}$ ensures that our *prior model* gives positive probability to any weak neighbourhood of any distribution defined on the support set of $F(\cdot|\boldsymbol{\theta}_i)$. A combination of Proposition 1 and Theorem 3.2.4 of Ghosh and Ramamoorthi (2003) gives the following result.

Corollary 1. Assume that $\mathcal{Y} \subset \mathbb{R}$ is the support set of $F(\cdot|\boldsymbol{\theta})$ for any $\boldsymbol{\theta}$. Let $F(\cdot)$ be a distribution function with support subset of \mathcal{Y} , then $P(\{Z_{i,\infty} \in U_F\}) > 0$ for any weak neighbourhood U_F of $F(\cdot)$.

The random process $Z_{i,\infty}$ can be seen as the limit of the histograms \mathbf{z}_i when the number of bins goes to infinity. To show this formally, we associate the random histogram \mathbf{z}_i to a random CDF Z_{iJ} . For any J we consider the partition $\mathcal{P}_J = \{y_0^J = -\infty < y_1^J < \dots < y_J^J = +\infty\}$ and define the following one-to-one mapping between \mathbf{z}_i and the CDF Z_{iJ} . Without loss of generality, we assign to the middle point of each interval the bin probability mass, and account for the two open bins (first and last) by introducing two auxiliary points y_-^J, y_+^J , such that $-\infty < y_-^J < y_1 < y_{J-1} < y_+^J < +\infty$. With this position we define the process $Z_{i,J}(y)$ (black line in Figure A-1):

$$Z_{iJ}(y) = \begin{cases} 0 & \text{if } y < y_-^J \\ z_{i1} & \text{if } y_-^J \leq y < (y_1^J + y_2^J)/2 \\ z_{i1} + \dots + z_{ij} & \text{if } y \in [(y_{j-1}^J + y_j^J)/2, (y_j^J + y_{j+1}^J)/2) \text{ for } 2 < j \leq J-2 \\ z_{i1} + \dots + z_{iJ-1} & \text{if } y \in [(y_{J-2}^J + y_{J-1}^J)/2, y_+^J) \\ 1 & \text{if } y \geq y_+^J \end{cases}$$

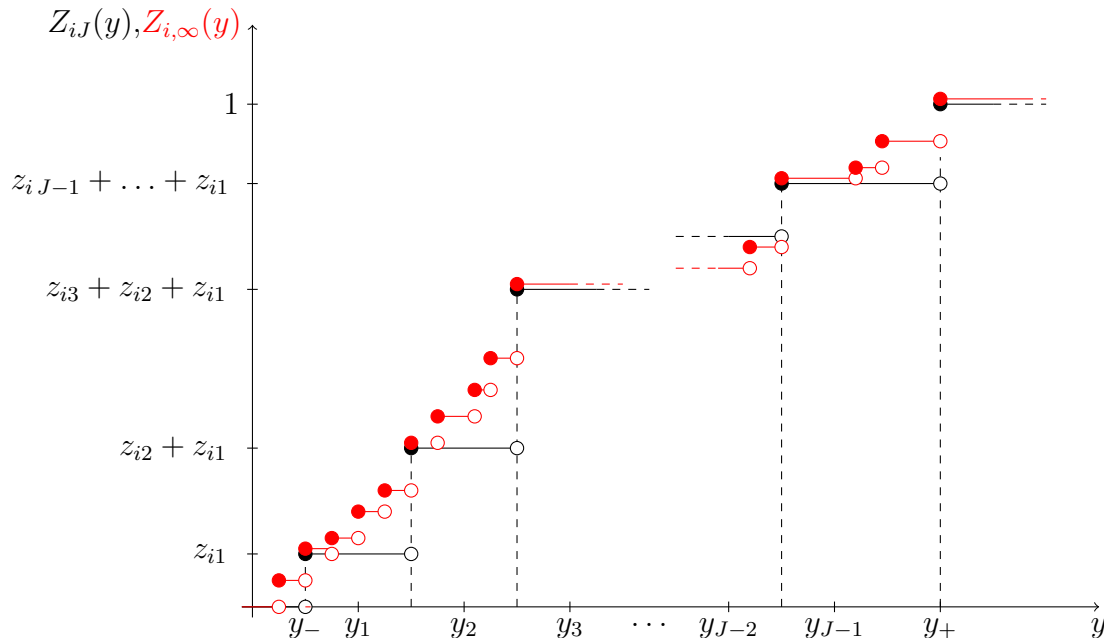


Figure A-1: Mapping between z_{ij} , $j = 1, \dots, J$, $Z_{iJ}(y)$ and $Z_{i,\infty}(y)$.

The next theorem shows that Z_{iJ} converges to $Z_{i,\infty}$ with probability one in the topology of the weak convergence. Moreover, under continuity assumptions, the asymptotic mean of Z_{iJ} , conditionally on θ_i , coincides with the true subjective distribution. Note that, conditionally on θ_i , the mean of $Z_{i,\infty}$ is the true subjective distribution, i.e. $\mathbb{E}[Z_{i,\infty}(\cdot)|\theta_i] = F(\cdot|\theta_i)$.

Theorem 2. Assume that $\varrho_j = 0$ for all j and the sequence of partitions $(\mathcal{P}_J)_J$ is such that $y_1 \rightarrow -\infty$, $y_{J-1} \rightarrow +\infty$ and $\max\{|y_{j+1} - y_j| : 1 \leq j \leq J - 2\} \rightarrow 0$ for $J \rightarrow +\infty$. Then,

$$P\left\{\lim_{J \rightarrow +\infty} Z_{iJ}(y) = Z_{i,\infty}(y) \text{ for any } y \text{ point of continuity of } Z_{i,\infty}\right\} = 1.$$

If $F(\cdot|\theta_i)$ is a continuous CDF, then

$$\lim_{J \rightarrow +\infty} \mathbb{E}[Z_{iJ}(y)|\theta_i] = \mathbb{E}[Z_{i,\infty}(y)|\theta_i] = F(y|\theta_i) \quad a.s.$$

C.B Further asymptotics

C.B.1 Posterior consistency

If $\varrho_j = 0$ for $j = 1, \dots, J$, i.e. forecasters give non-zero probability to each bin, \mathcal{H} is the set densities (absolutely continuous respect to the Lebesgue measure) on Δ^J , that is with

probability one $z_j > 0$ for all j . In this case, Kullback-Leibler divergence between two distribution h_0, g on Δ^J is easily defined as

$$KL(h_0, g) := \int_{\mathbf{z}} h_0(\mathbf{z}) \log \left(\frac{h_0(\mathbf{z})}{g(\mathbf{z})} \right) d\mathbf{z}.$$

As a corollary of the main theorem, we get a simpler result for the case in which $\varrho_j(\boldsymbol{\theta}) = 0$ for all $j = 1, \dots, J$. In this case \mathcal{M}^* is replaced by the set \mathcal{M} of finite mixtures of

$$h(\mathbf{z}|\boldsymbol{\theta}) = \frac{\prod_{j=1}^J \Gamma(\phi(\boldsymbol{\theta})\nu_j(\boldsymbol{\theta}))}{\Gamma\left(\sum_{j=1}^J \phi(\boldsymbol{\theta})\nu_j(\boldsymbol{\theta})\right)} \prod_{j=1}^{J-1} z_j^{\phi(\boldsymbol{\theta})\nu_j(\boldsymbol{\theta})-1} \left(1 - \sum_{j=1}^{J-1} z_j\right)^{\phi(\boldsymbol{\theta})\nu_J(\boldsymbol{\theta})-1}.$$

and \mathcal{H}_0^* by the set \mathcal{H}_0 of densities on Δ^J that can be approximated in the Kullback-Leibler sense by densities in \mathcal{M} , i.e.

$$\mathcal{H}_0 = \{h_0 \text{ density on } \Delta^J: \forall \epsilon > 0 \exists g \in \mathcal{M} \text{ s.t. } KL(h_0, g) \leq \epsilon \}.$$

Theorem 3. Let Θ be an open subset of \mathbb{R}^m for some m and $\varrho_j(\boldsymbol{\theta}) = 0$ for all $j = 1, \dots, J$. Assume that $\boldsymbol{\theta} \mapsto (\phi(\boldsymbol{\theta})\nu_1(\boldsymbol{\theta}), \dots, \phi(\boldsymbol{\theta})\nu_J(\boldsymbol{\theta}))$ is a continuous function on \mathbb{R}_+^J such that $\phi(\boldsymbol{\theta})\nu_j(\boldsymbol{\theta}) > 0$ for every $j = 1, \dots, J$. If G_0 has full support, then the posterior is weakly consistent at any density h_0 in \mathcal{H}_0 such that

$$\int_{\Delta^J} \left| \log \left(\prod_{j=1}^{J-1} z_j \left(1 - \sum_{j=1}^{J-1} z_j\right) \right) \right| h_0(\mathbf{z}) d\mathbf{z} < +\infty. \quad (\text{A-14})$$

Remark 1. If $\varrho_j(\boldsymbol{\theta}) = 0$ for all $j = 1, \dots, J$, $\phi(\boldsymbol{\theta}) = \phi$, and a mixture of normal distributions is assumed for the subjective distribution, that is

$$F(y|\boldsymbol{\theta}) = \sum_{i=1}^M \omega_i \Phi(y|\mu_i, \sigma_i^2) \quad (\text{A-15})$$

then the parameter vector is $\boldsymbol{\theta} = (\mu_1, \dots, \mu_M, \sigma_1^2, \dots, \sigma_M^2, \omega_1, \dots, \omega_M, \phi)$. If G_0 has full support, then the posterior is weakly consistent at any h_0 in \mathcal{H}_0 satisfying (A-14). Indeed, in this case $(\phi\nu_1(\boldsymbol{\theta}), \dots, \phi\nu_J(\boldsymbol{\theta}))$ is a continuous function on \mathbb{R}_+^J and $\phi\nu_j(\boldsymbol{\theta}) > 0$ for every $j = 1, \dots, J$.

The next Proposition gives some conditions ensuring that any continuous density function belongs to \mathcal{H}_0 .

Proposition 2. Assume $\varrho_j(\boldsymbol{\theta}) = 0$ for all $j = 1, \dots, J$ and that $\boldsymbol{\theta} \mapsto (\phi(\boldsymbol{\theta})\nu_1(\boldsymbol{\theta}), \dots, \phi(\boldsymbol{\theta})\nu_J(\boldsymbol{\theta}))$ is a continuous function on \mathbb{R}_+^J such that $\phi(\boldsymbol{\theta})\nu_j(\boldsymbol{\theta}) > 0$ for every $j = 1, \dots, J$. If for every $\mathbf{a} = (a_1, \dots, a_J) \in [1, +\infty)^J$ and $\delta > 0$, there is $\boldsymbol{\theta}_\delta$ in Θ such that $\|\mathbf{a} - \mathbf{a}_\delta\|_\infty \leq \delta$ with $\mathbf{a}_\delta = \phi(\boldsymbol{\theta}_\delta)(\nu_1(\boldsymbol{\theta}_\delta), \dots, \nu_J(\boldsymbol{\theta}_\delta))$, then any continuous density function on Δ^J belongs to \mathcal{H}_0 .

Remark 2. Note that combining Theorem 3 and Proposition 2 one gets that, under the assumptions of Proposition 2, if G_0 has full support, then the posterior is weakly consistent at any h_0 which is continuous on Δ^J and satisfies (A-14). An example in which all the conditions of Proposition 2 are met is the fully nonparametric case

$$F(y|\boldsymbol{\theta}) = \sum_{j=1}^J \varphi_j \mathbb{I}_{A_j}(y) \quad (\text{A-16})$$

where $A_j = [y_j, +\infty)$, $j = 1, \dots, J-1$, $A_J = [y^+, +\infty]$ and $\nu_j(\boldsymbol{\theta}) = \varphi_j$, $j = 1, \dots, J$. Conditions in Proposition 2 are satisfied also in the Gaussian mixture case of (A-15) with $M = J-1$.

C.B.2 Posterior consistency of the consensus distribution

The aggregate subjective distribution, also known as consensus distribution, is defined as

$$\bar{F}(y) = \frac{1}{n} \sum_{i=1}^n F_i(y)$$

where $F_i(y)$ is the forecast-specific subjective probability. In what follows, F_{n+1} denotes the posterior predictive distribution of y , defined as

$$F_{n+1}(y) := P\{y_{n+1} \leq y | \mathbf{z}_i, i = 1, \dots, n\}.$$

The next proposition shows the connection between the two quantities in our model.

Proposition 3. The distributions \bar{F} and F_n are related by

$$F_{n+1}(y) = \frac{n}{n + \psi_0} \bar{F}(y) + \frac{\psi_0}{n + \psi_0} \int F(y|\theta) G_0(d\theta).$$

Using the previous relation one obtain a useful asymptotic properties of the consensus distribution.

Proposition 4. Under the same assumptions of Theorem 3,

$$\lim_{n \rightarrow +\infty} (F_{n+1}(y_i) - F_{n+1}(y_{i-1})) = \lim_{n \rightarrow +\infty} (\bar{F}(y_i) - \bar{F}(y_{i-1})) = \int z_i h_0(\mathbf{z}) d\mathbf{z} \quad a.s.$$

for $i = 1, \dots, J$. Hence, if there exists F^* such that $\int z_i h_0(\mathbf{z}) = F^*(y_i) - F^*(y_{i-1})$, then

$$\lim_{n \rightarrow +\infty} F_{n+1}(y_i) = \lim_{n \rightarrow +\infty} \bar{F}(y_i) = F^*(y_i) \quad a.s..$$

As in Subsection C.A, we consider set of nested partitions $\mathcal{P}_J = \{y_0^J = -\infty < y_1^J < \dots < y_J^J = +\infty\}$ in such a way \mathcal{P}_{J+1} is a refinement of \mathcal{P}_J . We assume that observations $\mathbf{z}_1^J, \dots, \mathbf{z}_n^J$ are available with a “true” distribution $h_0 = h_0^J$ in \mathcal{M} , i.e. $h_0(\mathbf{z}) = \sum_{i=1}^M w_{i,0} h(\mathbf{z}|\boldsymbol{\theta}_{i,0})$ for suitable integer M , positive weights $(w_{1,0}, \dots, w_{M,0})$ and parameters $\boldsymbol{\theta}_{1,0}, \dots, \boldsymbol{\theta}_{M,0}$ in Θ . Note that with these hypotheses $\mathbf{z}_1^J, \dots, \mathbf{z}_n^J$ are consistent in J , that is if $J' > J$ then $z_i^J = \sum_{j \in I(i)} z_j^{J'}$ if the i -th bin in \mathcal{P}_J correspond the the union of the bins $j \in I(i)$ in $\mathcal{P}_{J'}$. This allows to consider limit jointly in the number of observations ($n \rightarrow +\infty$) and in the number of bins ($J \rightarrow +\infty$). Note also that for every J and every bin $(y_{i-1}, y_i]$ in \mathcal{P}_J

$$\int z_i h_0^J(\mathbf{z}) = F^*(y_i) - F^*(y_{i-1}).$$

for

$$F^*(y) := \sum_{i=1}^M w_{i,0} F(y|\boldsymbol{\theta}_{i,0}).$$

Proposition 5. In the setting described above, under the same assumptions of Theorem 2 on \mathcal{P}_J , then

$$\lim_{J \rightarrow +\infty, n \rightarrow +\infty} F_{n+1}(y) = \lim_{J \rightarrow +\infty, n \rightarrow +\infty} \bar{F}(y) = F^*(y) \quad a.s.$$

for every y point of continuity of F^* .

D Proofs

D.A Proofs of Theorem 1 and 3

The proof of Theorem 1 is based on an application of Theorem 1 and Lemma 3 Wu and Ghosal (2009b,a). In order to prove Theroem we need a slight generalization of these results. For the shake of clarity we state and prove this generalization.

In what follows, we denote with $supp(\mu)$ the weak support of a probability measure μ . We assume that \mathcal{X}_0 is a subset the finite set $\mathcal{X} = \{\boldsymbol{\xi} \in \{0, 1\}^J : |\boldsymbol{\xi}| < J\}$. Let \mathcal{Z} be the sample space, i.e. the set of all the pairs $(\boldsymbol{\xi}, \mathbf{z}_\boldsymbol{\xi})$ where $\boldsymbol{\xi} = (\xi_1, \dots, \xi_J)$, $\xi_i = \mathbb{I}\{z_i = 0\}$ and $\mathbf{z}_\boldsymbol{\xi}$ are the non-null elements of \mathbf{z} . In what follows we assume that $\mathbf{z}_\boldsymbol{\xi}$ takes values in an open subset $\mathcal{Z}_\boldsymbol{\xi}$ of $\mathbb{R}^{J-|\boldsymbol{\xi}|-1}$. In our application $\mathcal{Z}_\boldsymbol{\xi} = \Delta^{J-|\boldsymbol{\xi}|}$. On the sample space \mathcal{Z} , one defines the σ -finite measure $\lambda(d\mathbf{z}) = c(d\boldsymbol{\xi}) \otimes \mathcal{L}_\boldsymbol{\xi}(d\mathbf{z}_\boldsymbol{\xi})$ where c is the counting measure on \mathcal{X} and,

given $\boldsymbol{\xi}$, \mathcal{L}_ξ is the Lebesgue measure on $\mathcal{Z}_\xi \subset \mathbb{R}^{J-|\xi|-1}$. Let \mathcal{H} be the set of all the densities with respect to λ and note that the densities g factorize as $g(\mathbf{z}) = g(\boldsymbol{\xi})g(\mathbf{z}_\xi|\boldsymbol{\xi})$. We also assume that the kernel $h(\mathbf{z}|\boldsymbol{\theta})$ factorizes in the same way, i.e.

$$h(\mathbf{z}|\boldsymbol{\theta}) = h(\boldsymbol{\xi}|\boldsymbol{\theta})h(\mathbf{z}_\xi|\boldsymbol{\xi}, \boldsymbol{\theta}).$$

Finally, given a probability measure G on Θ , we write

$$h_G(\mathbf{z}) = \int_{\Theta} h(\mathbf{z}|\boldsymbol{\theta})G(d\boldsymbol{\theta}). \quad (\text{A-17})$$

and we assume that Π is the prior on \mathcal{H} induced by the map (A-17) when G has prior $\hat{\Pi}$.

In our application, $h_G(\mathbf{z})$ is given by (5) and $\hat{\Pi}$ is the Dirichlet process prior $\mathcal{DP}(\psi, G_0)$.

Given two densities h_0 and g in \mathcal{H} the Kullback-Leibler divergence between h_0 and g is defined as

$$KL(h_0, g) = \int_{\mathcal{Z}} h_0(\boldsymbol{\xi}, \mathbf{z}) \log \left(\frac{h_0(\boldsymbol{\xi}, \mathbf{z})}{g(\boldsymbol{\xi}, \mathbf{z})} \right) d\lambda.$$

Hence, writing $h_0(\mathbf{z}) = h_0(\boldsymbol{\xi})h_0(\mathbf{z}_\xi|\boldsymbol{\xi})$ and $g(\boldsymbol{\xi}, \mathbf{z}) = g(\boldsymbol{\xi})g(\mathbf{z}_\xi|\boldsymbol{\xi})$, by Fubini Theorem one can re-arrange the previous expression as

$$\begin{aligned} & \sum_{\boldsymbol{\xi} \in \mathcal{X}} h_0(\boldsymbol{\xi}) \int_{\mathcal{Z}_\xi} h_0(\mathbf{z}_\xi|\boldsymbol{\xi}) \log \left(\frac{h_0(\mathbf{z}_\xi|\boldsymbol{\xi})h_0(\boldsymbol{\xi})}{g(\mathbf{z}_\xi|\boldsymbol{\xi})g(\boldsymbol{\xi})} \right) d\mathbf{z}_\xi \\ &= \sum_{\boldsymbol{\xi} \in \mathcal{X}} h_0(\boldsymbol{\xi}) \left(\log \left(\frac{h_0(\boldsymbol{\xi})}{g(\boldsymbol{\xi})} \right) + \int_{\mathcal{Z}_\xi} h_0(\mathbf{z}_\xi|\boldsymbol{\xi}) \log \left(\frac{h_0(\mathbf{z}_\xi|\boldsymbol{\xi})}{g(\mathbf{z}_\xi|\boldsymbol{\xi})} \right) d\mathbf{z}_\xi \right). \end{aligned}$$

where for simplicity we write $\mathcal{L}_\xi(d\mathbf{z}_\xi)$ simply as $d\mathbf{z}_\xi$

Theorem 4. Let Θ be a Polish space and h_0 a density in \mathcal{H} . If for any $\varepsilon > 0$ there is a probability measure $G_\varepsilon \in \text{supp}(\hat{\Pi})$ and a closed set D_ε in Θ such that

$$(H1) \quad KL(h_0, h_{G_\varepsilon}) = \sum_{\boldsymbol{\xi} \in \mathcal{X}} h_0(\boldsymbol{\xi}) \int_{\mathcal{Z}_\xi} \log \left(\frac{h_0(\mathbf{z}_\xi|\boldsymbol{\xi})h_0(\boldsymbol{\xi})}{h_{G_\varepsilon}(\mathbf{z}_\xi|\boldsymbol{\xi})h_{G_\varepsilon}(\boldsymbol{\xi})} \right) h_0(\mathbf{z}_\xi|\boldsymbol{\xi},) d\mathbf{z}_\xi < \varepsilon;$$

(H2) D_ε contains $\text{supp}(G_\varepsilon)$ in its interior and for every $\boldsymbol{\xi}$

$$\int_{\mathcal{Z}_\xi} \log \left(\frac{h_{G_\varepsilon}(\mathbf{z}_\xi|\boldsymbol{\xi})h_{G_\varepsilon}(\boldsymbol{\xi})}{\inf_{\boldsymbol{\theta} \in D_\varepsilon} h(\mathbf{z}_\xi|\boldsymbol{\xi}, \boldsymbol{\theta})h(\boldsymbol{\xi}|\boldsymbol{\theta})} \right) h_0(\mathbf{z}_\xi|\boldsymbol{\xi}) d\mathbf{z}_\xi < +\infty;$$

(H3) $\inf_{\mathbf{z}_\xi \in C_\xi} \inf_{\boldsymbol{\theta} \in D_\varepsilon} h(\boldsymbol{\xi}|\boldsymbol{\theta})h(\mathbf{z}_\xi|\boldsymbol{\xi}, \boldsymbol{\theta}) > 0$ for every $\boldsymbol{\xi}$ and every compact set C_ξ in \mathcal{Z}_ξ ;

(H4) $\{\boldsymbol{\theta} \mapsto h(\boldsymbol{\xi}|\boldsymbol{\theta})h(\mathbf{z}_\xi|\boldsymbol{\xi}, \boldsymbol{\theta}) : \mathbf{z}_\xi \in C_\xi\}$ is uniformly equicontinuous on D_ε , for every $\boldsymbol{\xi}$ and every compact set C_ξ in \mathcal{Z}_ξ ;

then $\Pi\{KL(h_0, h_G) \geq \varepsilon\} > 0$ for every $\varepsilon > 0$ and hence Π is weakly consistent at h_0 .

Assumption (H1) corresponds to (A1) in Theorem 1 of Wu and Ghosal (2009b). Assumptions (H2)-(H3) correspond to assumptions (A7)-(A8) of Lemma 3 of Wu and Ghosal (2009b), while (H4) is slightly different from the original assumption (A9), see Wu and Ghosal (2009a). The theorem reduces to Theorem 1 and Lemma 3 of Wu and Ghosal (2009b) when \mathcal{X}_0 is the single point $\boldsymbol{\xi} = (0, \dots, 0)$.

Proof of Theorem 4. One has

$$\begin{aligned} KL(h_0, h_G) &= KL(h_0, h_{G_\varepsilon}) + \sum_{\boldsymbol{\xi} \in \mathcal{X}} h_0(\boldsymbol{\xi}) \int_{\Delta^{J-|\boldsymbol{\xi}|}} \log \left(\frac{h_{G_\varepsilon}(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi})h_{G_\varepsilon}(\boldsymbol{\xi})}{h_G(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi})h_G(\boldsymbol{\xi})} \right) h_0(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi}) d\mathbf{z}_\boldsymbol{\xi} \\ &\leq \varepsilon + \sum_{\boldsymbol{\xi} \in \mathcal{X}} h_0(\boldsymbol{\xi}) \int_{\Delta^{J-|\boldsymbol{\xi}|}} \log \left(\frac{h_{G_\varepsilon}(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi})h_{G_\varepsilon}(\boldsymbol{\xi})}{h_G(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi})h_G(\boldsymbol{\xi})} \right) h_0(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi}) d\mathbf{z}_\boldsymbol{\xi} =: \varepsilon + A_\varepsilon(G). \end{aligned}$$

If we show that there is an open neighbourhood V of G_ε such that for every G in V one has $A_\varepsilon(G) \leq \varepsilon$, then $\Pi\{KL(h_0, h_G) \geq 2\varepsilon\} > 0$ for every $\varepsilon > 0$. To prove the claim, for every $\boldsymbol{\xi}$ by (H2) we find a compact set $C_\boldsymbol{\xi}$ such that

$$\int_{C_\boldsymbol{\xi}^c} \log \left(\frac{h_{G_\varepsilon}(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi})h_{G_\varepsilon}(\boldsymbol{\xi})}{\inf_{\boldsymbol{\theta} \in D_\varepsilon} h(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi}, \boldsymbol{\theta})h(\boldsymbol{\xi}|\boldsymbol{\theta})} \right) h_0(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi}) d\mathbf{z}_\boldsymbol{\xi} \leq \frac{\varepsilon}{4}$$

and

$$\int_{C_\boldsymbol{\xi}^c} h_0(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi}) d\mathbf{z}_\boldsymbol{\xi} \leq \frac{\varepsilon}{4 \log(2)}.$$

Let $V_0 := \{G : G(D_\varepsilon) > 1/2\}$. Since $G_\varepsilon(D_\varepsilon) = 1$, by Portmanteau Theorem V is an open neighbourhood of G_ε . Now

$$h_G(\boldsymbol{\xi}, \mathbf{z}_\boldsymbol{\xi}) = \int_{D_\varepsilon} h(\boldsymbol{\xi}, \mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\theta})G(d\boldsymbol{\theta}) \geq \inf_{\boldsymbol{\theta} \in D_\varepsilon} h(\boldsymbol{\xi}|\boldsymbol{\theta})h(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi}, \boldsymbol{\theta})G(D_\varepsilon),$$

hence, for every G in V_1 ,

$$\begin{aligned} &\int_{C_\boldsymbol{\xi}^c} \log \left(\frac{h_{G_\varepsilon}(\boldsymbol{\xi}, \mathbf{z}_\boldsymbol{\xi})}{h_G(\boldsymbol{\xi}, \mathbf{z}_\boldsymbol{\xi})} \right) h_0(\boldsymbol{\xi}, \mathbf{z}_\boldsymbol{\xi}) d\mathbf{z}_\boldsymbol{\xi} \\ &\leq \int_{C_\boldsymbol{\xi}^c} \log \left(\frac{h_{G_\varepsilon}(\boldsymbol{\xi}, \mathbf{z}_\boldsymbol{\xi})}{\inf_{\boldsymbol{\theta} \in D_\varepsilon} h(\boldsymbol{\xi}|\boldsymbol{\theta})h(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi}, \boldsymbol{\theta})} \right) h_0(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi}) d\mathbf{z}_\boldsymbol{\xi} + \log(2) \int_{C_\boldsymbol{\xi}^c} h_0(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi}) d\mathbf{z}_\boldsymbol{\xi} \leq \frac{\varepsilon}{2}. \end{aligned} \tag{A-18}$$

By condition (H4), for every $\boldsymbol{\xi}$ there are $\mathbf{z}_\boldsymbol{\xi}^{(i)} \in C_\boldsymbol{\xi}$ $i = 1, \dots, m$, such that for every $\mathbf{z}_\boldsymbol{\xi} \in C_\boldsymbol{\xi}$ there is i for which

$$\sup_{\boldsymbol{\theta} \in D_\varepsilon} |h(\boldsymbol{\xi}|\boldsymbol{\theta})h(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi}, \boldsymbol{\theta}) - h(\boldsymbol{\xi}, \mathbf{z}_\boldsymbol{\xi}^{(i)}|\boldsymbol{\theta})| \leq \frac{c\varepsilon}{12}$$

where $c := \inf_{\mathbf{z}_\xi \in C_\xi} \inf_{\boldsymbol{\theta} \in D_\varepsilon} h(\boldsymbol{\xi}|\boldsymbol{\theta})h(\mathbf{z}_\xi|\boldsymbol{\xi}, \boldsymbol{\theta}) > 0$ by (H3). Since $G_\varepsilon(\partial D_\varepsilon) = 0$, the set

$$V_\xi := \left\{ G : \left| \int_{D_\varepsilon} h(\boldsymbol{\xi}, \mathbf{z}_\xi^{(i)}|\boldsymbol{\theta})G_\varepsilon(d\boldsymbol{\theta}) - \int_{D_\varepsilon} h(\boldsymbol{\xi}, \mathbf{z}_\xi^{(i)}|\boldsymbol{\theta})G(d\boldsymbol{\theta}) \right| < \frac{c\varepsilon}{12}; \quad i = 1, \dots, m \right\}$$

is a weak neighbourhood of G_ε . Hence, for G in V_ξ

$$\left| \int_{D_\varepsilon} h(\boldsymbol{\xi}, \mathbf{z}_\xi|\boldsymbol{\theta})G_\varepsilon(d\boldsymbol{\theta}) - \int_{D_\varepsilon} h(\boldsymbol{\xi}, \mathbf{z}_\xi|\boldsymbol{\theta})G(d\boldsymbol{\theta}) \right| \leq \frac{c\varepsilon}{4} \quad (\text{A-19})$$

Since $\text{supp}(G_\varepsilon) \subset D_\varepsilon$,

$$\int_{C_\xi} \log \left(\frac{h_{G_\varepsilon}(\boldsymbol{\xi}, \mathbf{z}_\xi)}{h_G(\boldsymbol{\xi}, \mathbf{z}_\xi)} \right) h_0(\mathbf{z}_\xi|\boldsymbol{\xi}) d\mathbf{z}_\xi \leq \int_{C_\xi} \log \left(\frac{\int_{D_\varepsilon} h(\boldsymbol{\xi}, \mathbf{z}_\xi|\boldsymbol{\theta})G_\varepsilon(d\boldsymbol{\theta})}{\int_{D_\varepsilon} h(\boldsymbol{\xi}, \mathbf{z}_\xi|\boldsymbol{\theta})G(d\boldsymbol{\theta})} \right) h_0(\mathbf{z}_\xi|\boldsymbol{\xi}) d\mathbf{z}_\xi.$$

Hence, using $\log(x+1) \leq x$ and (A-19), for G in $V_0 \cap V_\xi$ one obtains

$$\int_{C_\xi} \log \left(\frac{h_{G_\varepsilon}(\boldsymbol{\xi}, \mathbf{z}_\xi)}{h_G(\boldsymbol{\xi}, \mathbf{z}_\xi)} \right) h_0(\mathbf{z}_\xi|\boldsymbol{\xi}) d\mathbf{z}_\xi \leq \frac{\varepsilon}{2}. \quad (\text{A-20})$$

At this stage, combining (A-18) and (A-20), one obtains that $A_\varepsilon(G) \leq \varepsilon$ for every G in $V = V_0 \cap (\cap_\xi V_\xi)$. \square

We can now prove both Theorem 1 and Theorem 3. We start with the second theorem because the proof is easier.

Proof of Theorem 3. The proof follows from an application of Theorem 4 for $\mathcal{X}_0 = \{(0, \dots, 0)\}$.

Let

$$\tilde{\boldsymbol{\nu}}(\boldsymbol{\theta}) := (\tilde{\nu}_1(\boldsymbol{\theta}), \dots, \tilde{\nu}_J(\boldsymbol{\theta})) = (\phi(\boldsymbol{\theta})\nu_1(\boldsymbol{\theta}), \dots, \phi(\boldsymbol{\theta})\nu_J(\boldsymbol{\theta})) \quad (\text{A-21})$$

and

$$Z_\boldsymbol{\theta} = \frac{\prod_{j=1}^J \Gamma(\tilde{\nu}_j(\boldsymbol{\theta}))}{\Gamma\left(\sum_{j=1}^J \tilde{\nu}_j(\boldsymbol{\theta})\right)}.$$

Verification of (H1) of Theorem 4. By hypothesis, for every $\varepsilon > 0$ there is $g_\varepsilon(\mathbf{z}) = \sum_{i=1}^{M_\varepsilon} w_{i,\varepsilon} h(\mathbf{z}|\boldsymbol{\theta}_{i,\varepsilon})$ in \mathcal{M} such that $KL(h_0, g_\varepsilon) \leq \varepsilon$. To see that (H1) is satisfied, write $g_\varepsilon(\mathbf{z}) = \int h(\mathbf{z}|\boldsymbol{\theta})G_\varepsilon(d\boldsymbol{\theta}) = h_{G_\varepsilon}(\mathbf{z})$ for $G_\varepsilon(d\boldsymbol{\theta}) = \sum_{i=1}^{M_\varepsilon} w_{i,\varepsilon} \delta_{\boldsymbol{\theta}_{i,\varepsilon}}(d\boldsymbol{\theta})$. Now $\text{supp}(G_\varepsilon) = \cup_{i=1}^{M_\varepsilon} \{\boldsymbol{\theta}_{i,\varepsilon}\}$. To conclude recall that if Π is $\mathcal{DP}(\psi, G_0)$ and $\text{supp}(G_\varepsilon) \subset \text{supp}(G_0)$, then $G_\varepsilon \in \text{supp}(\Pi)$; see, for instance, Theorem 3.2.4 of Ghosh and Ramamoorthi (2003).

Verification of (H2) of Theorem 4. Given G_ε as above, one can find a compact set D_ε in Θ such that D_ε contains $\cup_{i=1}^{M_\varepsilon} \{\boldsymbol{\theta}_{i,\varepsilon}\} = \text{supp}(G_\varepsilon)$ in its interior.

Now

$$\begin{aligned} I_\varepsilon(\mathbf{z}) &:= \inf_{\boldsymbol{\theta} \in D_\varepsilon} h(\mathbf{z}|\boldsymbol{\theta}) \\ &= \inf_{\boldsymbol{\theta} \in D_\varepsilon} \frac{1}{Z_\boldsymbol{\theta}} \prod_{j=1}^{J-1} z_j^{\tilde{\nu}_j(\boldsymbol{\theta})-1} \left(1 - \sum_{j=1}^{J-1} z_j\right)^{\tilde{\nu}_j(\boldsymbol{\theta})-1} \\ &\geq C_{1,\varepsilon} \prod_{j=1}^{J-1} z_j^{\mu_{j,\varepsilon}-1} \left(1 - \sum_{j=1}^{J-1} z_j\right)^{\mu_{j,\varepsilon}-1} =: I_\varepsilon^*(\mathbf{z}) \end{aligned}$$

where $C_{1,\varepsilon} = \inf_{\boldsymbol{\theta} \in D_\varepsilon} Z_\boldsymbol{\theta}^{-1}$, $\mu_{j,\varepsilon} := \sup\{\tilde{\nu}_j(\boldsymbol{\theta}) : \boldsymbol{\theta} \in D_\varepsilon\}$. Now one has that $C_{1,\varepsilon} > 0$ and $\mu_{j,\varepsilon} > 0$, since D_ε is compact and the $\nu_j(\boldsymbol{\theta})$ s are continuous and strictly positive.

On the one hand $h_{G_\varepsilon}(\mathbf{z}) \geq I_\varepsilon(\mathbf{z})$ and hence $\log(h_{G_\varepsilon}(\mathbf{z})/I_\varepsilon(\mathbf{z})) \geq 0$, on the other hand

$$\begin{aligned} \int \log\left(\frac{h_{G_\varepsilon}(\mathbf{z})}{I_\varepsilon(\mathbf{z})}\right) h_0(\mathbf{z}) d\mathbf{z} &\leq \int \log\left(\frac{h_{G_\varepsilon}(\mathbf{z})}{I_\varepsilon^*(\mathbf{z})}\right) h_0(\mathbf{z}) d\mathbf{z} \\ &\leq \int \left| \log\left(\frac{g_\varepsilon(\mathbf{z})}{\prod_{j=1}^{J-1} z_j^{\mu_{j,\varepsilon}-1} \left(1 - \sum_{j=1}^{J-1} z_j\right)^{\mu_{j,\varepsilon}-1}}\right) \right| h_0(\mathbf{z}) d\mathbf{z} + |\log(C_{1,\varepsilon})|. \end{aligned}$$

Since

$$C_{2,\varepsilon} \prod_{j=1}^{J-1} z_j^{A_{j,\varepsilon}-1} \left(1 - \sum_{j=1}^{J-1} z_j\right)^{A_{j,\varepsilon}-1} \leq g_\varepsilon(\mathbf{z}) \leq C_{3,\varepsilon} \prod_{j=1}^{J-1} z_j^{B_{j,\varepsilon}-1} \left(1 - \sum_{j=1}^{J-1} z_j\right)^{B_{j,\varepsilon}-1}$$

for suitable constants $C_{2,\varepsilon}, C_{3,\varepsilon}, A_{1,\varepsilon}, \dots, B_{1,\varepsilon}, \dots, B_{J,\varepsilon}$, it follows that

$$\begin{aligned} \left| \log\left(\frac{g_\varepsilon(\mathbf{z})}{\prod_{j=1}^{J-1} z_j^{\mu_{j,\varepsilon}-1} \left(1 - \sum_{j=1}^{J-1} z_j\right)^{\mu_{j,\varepsilon}-1}}\right) \right| &\leq C_{4,\varepsilon} \left[1 + \sum_{j=1}^{J-1} |\log(z_j)| + \left| \log\left(1 - \sum_{j=1}^{J-1} z_j\right) \right| \right] \\ &\leq C_{4,\varepsilon} \left[1 + \left| \log\left(\prod_{j=1}^{J-1} z_j \left(1 - \sum_{j=1}^{J-1} z_j\right)\right) \right| \right] \end{aligned}$$

Combining all the estimates, one gets

$$\int \log\left(\frac{h_{G_\varepsilon}(\mathbf{z})}{I_\varepsilon(\mathbf{z})}\right) h_0(\mathbf{z}) d\mathbf{z} \leq C_{5,\varepsilon} \left[1 + \int \left| \log\left(\prod_{j=1}^{J-1} z_j \left(1 - \sum_{j=1}^{J-1} z_j\right)\right) \right| h_0(\mathbf{z}) d\mathbf{z} \right] < +\infty$$

by assumption (A-14). Hence

$$0 < \int \log\left(\frac{h_{G_\varepsilon}(\mathbf{z})}{\inf_{\boldsymbol{\theta} \in D_\varepsilon} h(\mathbf{z}|\boldsymbol{\theta})}\right) h_0(\mathbf{z}) d\mathbf{z} < +\infty.$$

Verification of (H3) of Theorem 4. It follows immediately that, for every compact set C in the open simplex Δ^J ,

$$\inf_{\mathbf{z} \in C} \inf_{\boldsymbol{\theta} \in D_\varepsilon} h(\mathbf{z}|\boldsymbol{\theta}) \geq \inf_{\mathbf{z} \in C} I_\varepsilon^*(\mathbf{z})$$

and the right hand side is strictly positive.

Verification of (H4) of Theorem 4. Under the hypotheses, the function $(\boldsymbol{\theta}, \mathbf{z}) \mapsto h(\mathbf{z}|\boldsymbol{\theta})$ is continuous and hence uniformly continuous on the compact set $C \times D_\varepsilon$. It follows that the family $\{(\boldsymbol{\theta}, \mathbf{z}) \mapsto h(\mathbf{z}|\boldsymbol{\theta}) : \mathbf{z} \in C\}$ is uniformly equicontinuous on D_ε . □

Proof of Theorem 1. The proof consists in an application of Theorem 4 for $\mathcal{X}_0 = \mathcal{X}$ and follows the same line of the proof of Theorem 3. In the present case, everything has an extra dependence on the fixed $\boldsymbol{\xi}$ in \mathcal{X} . In place of $I_\varepsilon(\mathbf{z})$ one has

$$I_\varepsilon(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi}) := \inf_{\boldsymbol{\theta} \in D_\varepsilon} \frac{1}{c(\boldsymbol{\theta})} \prod_{j=1}^J \varrho_j(\boldsymbol{\theta})^{\xi_j} (1 - \varrho_j(\boldsymbol{\theta}))^{1-\xi_j} \frac{1}{Z_\boldsymbol{\theta}(\boldsymbol{\xi})} \prod_{j \in \mathcal{J}^*} z_j^{\tilde{\nu}_j(\boldsymbol{\theta})-1}$$

where

$$Z_\boldsymbol{\theta}(\boldsymbol{\xi}) = \frac{\prod_{j \in \mathcal{J}^*(\boldsymbol{\xi})} \Gamma(\tilde{\nu}_j(\boldsymbol{\theta}))}{\Gamma\left(\sum_{j \in \mathcal{J}^*} \tilde{\nu}_j(\boldsymbol{\theta})\right)}.$$

Moreover,

$$I_\varepsilon(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi}) \geq C_{1,\varepsilon}(\boldsymbol{\xi}) \prod_{j \in \mathcal{J}^*} z_j^{\mu_{j,\varepsilon}-1} =: I_\varepsilon^*(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\xi})$$

where

$$C_{1,\varepsilon}(\boldsymbol{\xi}) = \inf_{\boldsymbol{\theta} \in D_\varepsilon} \frac{1}{c(\boldsymbol{\theta})} \prod_{j=1}^J \varrho_j(\boldsymbol{\theta})^{\xi_j} (1 - \varrho_j(\boldsymbol{\theta}))^{1-\xi_j} Z_\boldsymbol{\theta}^{-1}(\boldsymbol{\xi}),$$

and $\mu_{j,\varepsilon} := \sup\{\tilde{\nu}_j(\boldsymbol{\theta}) : \boldsymbol{\theta} \in D_\varepsilon\}$. Also in this case, $C_{1,\varepsilon} > 0$ and $\mu_{j,\varepsilon} > 0$, since D_ε is compact, $\nu_j(\boldsymbol{\theta})$ and $\varrho_j(\boldsymbol{\theta})$ are continuous, $0 < \varrho_j(\boldsymbol{\theta}) < 1$ and $\nu_j(\boldsymbol{\theta}) > 0$, $j = 1, \dots, J$. Finally,

$$C_{2,\varepsilon}(\boldsymbol{\xi}) \prod_{j \in \mathcal{J}^*} z_j^{A_{j,\varepsilon}-1} \leq h_{G_\varepsilon}(\boldsymbol{\xi}, \mathbf{z}) \leq C_{3,\varepsilon}(\boldsymbol{\xi}) \prod_{j \in \mathcal{J}^*} z_j^{B_{j,\varepsilon}-1}$$

for suitable constants $C_{2,\varepsilon}(\boldsymbol{\xi}), C_{3,\varepsilon}(\boldsymbol{\xi}), A_{1,\varepsilon}, \dots, B_{1,\varepsilon}, \dots, B_{J,\varepsilon}$. With this minor modifications, the verification of (H1) and (H2) is exactly as in the proof of Theorem 3. Assumption (H3) is true since

$$\inf_{\mathbf{z}_\boldsymbol{\xi} \in C_\boldsymbol{\xi}} \inf_{\boldsymbol{\theta} \in D_\varepsilon} h(\boldsymbol{\xi}|\boldsymbol{\theta}) h(\mathbf{z}_\boldsymbol{\xi}|\boldsymbol{\theta}) \geq \inf_{\mathbf{z} \in C_\boldsymbol{\xi}} I_\varepsilon^*(\mathbf{z}|\boldsymbol{\xi})$$

and the right hand side is strictly positive by the assumptions on the $\nu_j(\boldsymbol{\theta})$ s and $\varrho_j(\boldsymbol{\theta})$ s. Analogously,

$$(\boldsymbol{\theta}, \mathbf{z}_\xi) \mapsto h(\boldsymbol{\xi}|\boldsymbol{\theta})h(\mathbf{z}_\xi|\boldsymbol{\theta})$$

is uniformly continuous on the compact set $C_\xi \times D_\varepsilon$ and hence (H4) follows. \square

D.B Proof of Proposition 2

The proof of Proposition 2 is divided in various Lemmata. For the sake of notational simplicity set

$$D(\mathbf{z}; a_1, \dots, a_J) = \frac{\Gamma\left(\sum_{j=1}^J a_j\right)}{\prod_{j=1}^J \Gamma(a_j)} \prod_{j=1}^{J-1} z_j^{a_j-1} \left(1 - \sum_{j=1}^{J-1} z_j\right)^{a_J-1}.$$

Note that

$$h(\mathbf{z}|\boldsymbol{\theta}) = D(\mathbf{z}; \tilde{\boldsymbol{\nu}}(\boldsymbol{\theta}))$$

where $\tilde{\boldsymbol{\nu}}(\boldsymbol{\theta})$ is defined in (A-21).

Lemma 1. [Barrientos et al. (2015)] Let g_0 be a continuous density on Δ^J . Then, for every $\varepsilon > 0$ there is a density $g_\varepsilon(\mathbf{z}) = \sum_{i=1}^{M_\varepsilon} q_{i,\varepsilon} D(\mathbf{z}; a_{i,1,\varepsilon}, \dots, a_{i,J,\varepsilon})$ where $a_{i,j,\varepsilon} \geq 1$ for every i and j , such that

$$\|g_0 - g_\varepsilon\|_\infty \leq \varepsilon.$$

Lemma 2. Let $a = (a_1, \dots, a_J) \in [1, +\infty)^J$. If for any $\delta > 0$ there is $\boldsymbol{\theta}_\delta \in \Theta$ such that $\|a - \tilde{\boldsymbol{\nu}}(\boldsymbol{\theta}_\delta)\|_\infty \leq \delta$ then for any $\varepsilon > 0$ there is $\boldsymbol{\theta}_\varepsilon \in \Theta$ such that

$$\|D(\cdot; a_1, \dots, a_J) - D(\cdot; \tilde{\nu}_1(\boldsymbol{\theta}_\varepsilon), \dots, \tilde{\nu}_J(\boldsymbol{\theta}_\varepsilon))\|_\infty \leq \varepsilon.$$

Proof. The Proof is left to the reader. \square

Lemma 3. Assume that, for every $a = (a_1, \dots, a_J) \in [1, +\infty)^J$ and every $\delta > 0$ there is $\boldsymbol{\theta}_\delta \in \Theta$ such that $\|a - \tilde{\boldsymbol{\nu}}(\boldsymbol{\theta}_\delta)\|_\infty \leq \delta$. Then, for every continuous density g_0 on Δ^J and for every $\varepsilon > 0$, there is a density $\tilde{g}_\varepsilon(\mathbf{z}) = \sum_{i=1}^{M_\varepsilon} q_{i,\varepsilon} D(\mathbf{z}; \tilde{\boldsymbol{\nu}}(\boldsymbol{\theta}_{i,\varepsilon}))$ in \mathcal{M} such that

$$\|g_0 - \tilde{g}_\varepsilon\|_\infty \leq \varepsilon.$$

Proof. By Lemma 1, there is a density $g_\varepsilon(\mathbf{z}) = \sum_{i=1}^{M_\varepsilon} q_{i,\varepsilon} D(\mathbf{z}; a_{i,1,\varepsilon}, \dots, a_{i,J,\varepsilon})$ where $a_{i,j,\varepsilon} \geq 1$ for every i and j , such that $\|g_0 - g_\varepsilon\|_\infty \leq \varepsilon/2$. Now, by Lemma 2, there are $\boldsymbol{\theta}_{i,\varepsilon}$ such that $\|D(\cdot; a_{i,1,\varepsilon}, \dots, a_{i,J,\varepsilon}) - D(\cdot; \tilde{\nu}_1(\boldsymbol{\theta}_{i,\varepsilon}), \dots, \tilde{\nu}_J(\boldsymbol{\theta}_{i,\varepsilon}))\|_\infty \leq \varepsilon/2$. Hence, setting $\tilde{g}_\varepsilon(\mathbf{z}) := \sum_{i=1}^{M_\varepsilon} q_{i,\varepsilon} D(\mathbf{z}; \tilde{\nu}_1(\boldsymbol{\theta}_{i,\varepsilon}), \dots, \tilde{\nu}_J(\boldsymbol{\theta}_{i,\varepsilon}))$, one gets

$$\begin{aligned} \|g_0 - \tilde{g}_\varepsilon\|_\infty &\leq \|g_0 - g_\varepsilon\|_\infty \\ &\quad + \sum_{i=1}^{M_\varepsilon} q_{i,\varepsilon} \|D(\cdot; a_{i,1,\varepsilon}, \dots, a_{i,J,\varepsilon}) - D(\cdot; \tilde{\nu}_1(\boldsymbol{\theta}_{i,\varepsilon}), \dots, \tilde{\nu}_J(\boldsymbol{\theta}_{i,\varepsilon}))\|_\infty \leq \varepsilon. \end{aligned}$$

□

Lemma 4. For every densities g_1 and g_2 in Δ^J

$$KL(g_1, g_2) \leq \frac{\sup_{\mathbf{z}} |g_1(\mathbf{z}) - g_2(\mathbf{z})|^2}{\inf_{\mathbf{z}} g_2(\mathbf{z})}$$

Proof. By Jensen inequality

$$KL(g_1, g_2) \leq \log \left(\int \frac{g_1^2}{g_2} \right).$$

Now, since $\log(1+x) \leq x$ for every $x > 0$

$$\log \left(\int \frac{g_1^2}{g_2} \right) = \log \left(\int \left(\frac{(g_1 - g_2)^2}{g_2} + 1 \right) \right) \leq \int \frac{(g_1 - g_2)^2}{g_2} \leq \frac{\sup_{\mathbf{z}} |g_1(\mathbf{z}) - g_2(\mathbf{z})|^2}{\inf_{\mathbf{z}} g_2(\mathbf{z})}$$

□

Proof of Proposition 2. We need to prove that, if h_0 is a continuous density on Δ^J , then, for every $\eta > 0$, there is a density g_η in \mathcal{M} such that

$$KL(h_0, g_\eta) \leq \eta.$$

Let $h_\varepsilon(\mathbf{z}) = \max(\varepsilon, h_0(\mathbf{z})) C_\varepsilon^{-1}$ where $C_\varepsilon := \int \max(\varepsilon, h_0(\mathbf{z})) d\mathbf{z} \leq 1 + \varepsilon$. Clearly $h_\varepsilon > \varepsilon$ and $h_0 \leq C_\varepsilon h_\varepsilon$. Hence, by Lemma 5.1. in Ghoshal et al. (1999), for any density g

$$KL(h_0, g) \leq (2 + \varepsilon) \log(1 + \varepsilon) + (1 + \varepsilon) [KL(h_\varepsilon, g) + \sqrt{KL(h_\varepsilon, g)}]. \quad (\text{A-22})$$

By Lemma 3 there is a density \tilde{g}_ε in \mathcal{M} such that $\|h_\varepsilon - \tilde{g}_\varepsilon\|_\infty \leq \varepsilon/2$. From the previous inequality it follows that $\tilde{g}_\varepsilon \geq h_\varepsilon - \varepsilon/2 \geq \varepsilon/2$. Hence, by Lemma 4

$$KL(h_\varepsilon, \tilde{g}_\varepsilon) \leq \varepsilon.$$

The thesis follows by taking $\eta = (2 + \varepsilon) \log(1 + \varepsilon) + (1 + \varepsilon)(\varepsilon + \sqrt{\varepsilon})$ and $g_\eta = \tilde{g}_\varepsilon$. □

D.C Proofs of Proposition 1 and Theorem 2

Proof of Proposition 1. Recall that since $Z_{i,\infty}(dy)$ is a Dirichlet process with concentration parameter ϕ_i and base measure $F(dy|\boldsymbol{\theta}_i)$, then for any finite partition B_1, \dots, B_J of \mathbb{R} it follows that $(Z_{i,\infty}(B_1), \dots, Z_{i,\infty}(B_J))$ has a Dirichlet distribution on Δ^J of parameters $(\phi(\boldsymbol{\theta}_i)F(B_1|\boldsymbol{\theta}_i), \dots, \phi(\boldsymbol{\theta}_i)F(B_J|\boldsymbol{\theta}_i))$. Hence, the random vector $\mathbf{z}_i = (z_{i1}, \dots, z_{iJ}) := (Z_{i,\infty}(y_1) - Z_{i,\infty}(y_0), \dots, Z_{i,\infty}(y_J) - Z_{i,\infty}(y_{J-1}))$ has the Dirichlet distribution on the simplex Δ^J of parameters $(\phi(\boldsymbol{\theta}_i)\nu_1(\boldsymbol{\theta}_i), \dots, \phi(\boldsymbol{\theta}_i)\nu_J(\boldsymbol{\theta}_i))$. When $\varrho_j(\cdot|\epsilon) = 0$ for $j = 1, \dots, J$, our Bayesian mode is

$$\begin{aligned} (z_{i1}, \dots, z_{iJ}) &\sim \text{Dir}_J(\phi(\boldsymbol{\theta}_i)\nu_1(\boldsymbol{\theta}_i), \dots, \phi(\boldsymbol{\theta}_i)\nu_J(\boldsymbol{\theta}_i)) \\ \boldsymbol{\theta}_i &\stackrel{i.i.d.}{\sim} G \\ G &\sim \mathcal{DP}(\psi, G_0), \end{aligned}$$

and the thesis follows. □

Proof of Theorem 2. The thesis is easily deduced from Proposition 1. □

D.D Proofs of Propositions 3 and 4

Proof of Proposition 3. Note that

$$F_{n+1}(y) = E[F(y|\boldsymbol{\theta}_{d_{n+1}}^*)|\mathbf{z}_i, i = 1 \dots, n]$$

which yields

$$\begin{aligned} E[F(y|\boldsymbol{\theta}_{d_{n+1}}^*)|\mathbf{z}_i, i = 1 \dots, n] &= E[E[F(y|\boldsymbol{\theta}_{d_{n+1}}^*)|\boldsymbol{\theta}_{d_i}^*, \mathbf{z}_i, i = 1 \dots, n]|\mathbf{z}_i, i = 1, \dots, n] \\ &= E[E[F(y|\boldsymbol{\theta}_{d_{n+1}}^*)|\boldsymbol{\theta}_{d_i}^*, i = 1, \dots, n]|\mathbf{z}_i, i = 1, \dots, n] \end{aligned}$$

By Proposition 1, $\boldsymbol{\theta}_i := \boldsymbol{\theta}_{d_i}^*$ are drawn from a $\mathcal{DP}(\psi, G_0)$, hence the predictive distribution of $\boldsymbol{\theta}_{d_{n+1}}^*$ given $\boldsymbol{\theta}_{d_i}^*, i = 1, \dots, n$ is

$$G_{n+1}(\cdot) = \frac{n}{n + \psi} \sum_{i=1}^n \delta_{\boldsymbol{\theta}_{d_i}^*}(d\boldsymbol{\theta}) + \frac{\psi}{n + \psi} G_0(\cdot),$$

Hence by the law of iterated expectations

$$\begin{aligned} E[F(y|\boldsymbol{\theta})|\boldsymbol{\theta}_{d_i}^*, i = 1 \dots, n] &= \int F(y|\boldsymbol{\theta})G_{n+1}(d\boldsymbol{\theta}) \\ &= \frac{n}{n + \psi} \frac{1}{n} \sum_{i=1}^n F(y|\boldsymbol{\theta}_{d_i}^*) + \frac{\psi}{n + \psi} \int F(y|\boldsymbol{\theta})G_0(d\boldsymbol{\theta}) \end{aligned}$$

Since

$$E\left[\frac{1}{n}\sum_{i=1}^n F(y|\boldsymbol{\theta}_{d_i})|\mathbf{z}_i, i = 1, \dots, n\right] = \bar{F}(y)$$

we obtain the result

$$F_{n+1}(y) := P\{Y_{n+1} \leq y|\mathbf{z}_i, i = 1 \dots, n\} = \frac{n}{n + \psi} \bar{F}(y) + \frac{\psi}{n + \psi} \int F(y|\theta)G_0(d\theta)$$


□

Proof of Proposition 4. Recall that posterior consistency yields predictive consistency, see e.g. Theorem 4.2.1 in Ghosh and Ramamoorthi (2003) since $\phi(\mathbf{z}) = z_i$ is a bounded and continuous function on the simplex the thesis follows. □

E Additional Results

E.A Priors

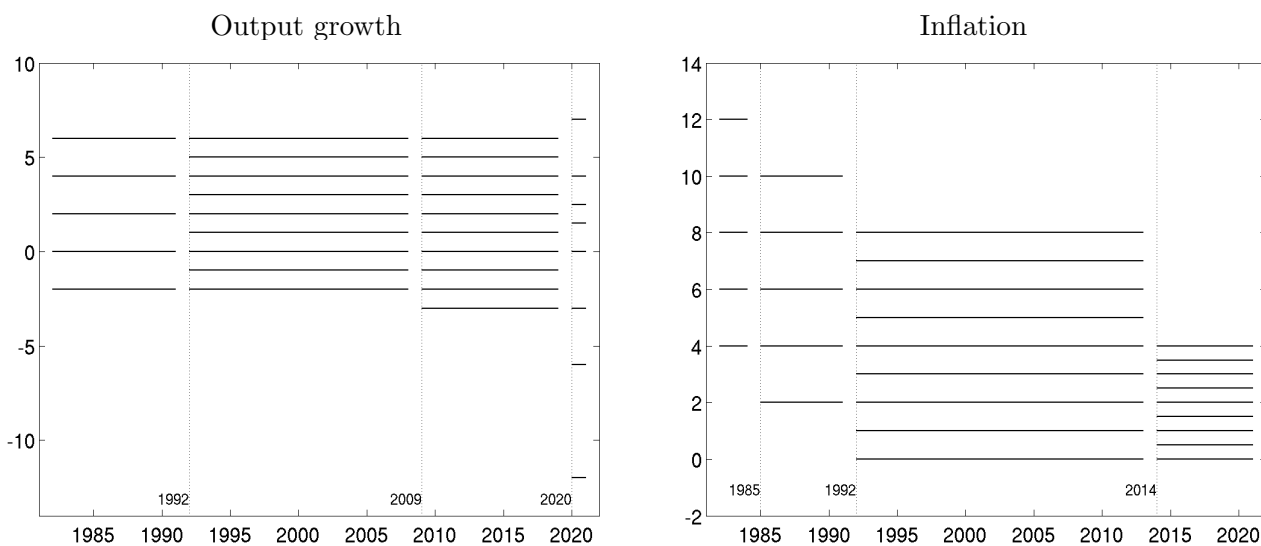
Figure A-2: Noise and zero probability



Note: The left panel of Figure A-2 shows the 50 and 90 percent a-priori coverage intervals for the noise associated with three different values of ν : 0.1, 0.6 and 0.3.

E.B Survey design

Figure A-3: Survey bins



Note: Solid lines show the survey bins. Vertical dotted lines highlight years when bins changed.

Figure A-4: Number of respondents for H1 Output growth surveys

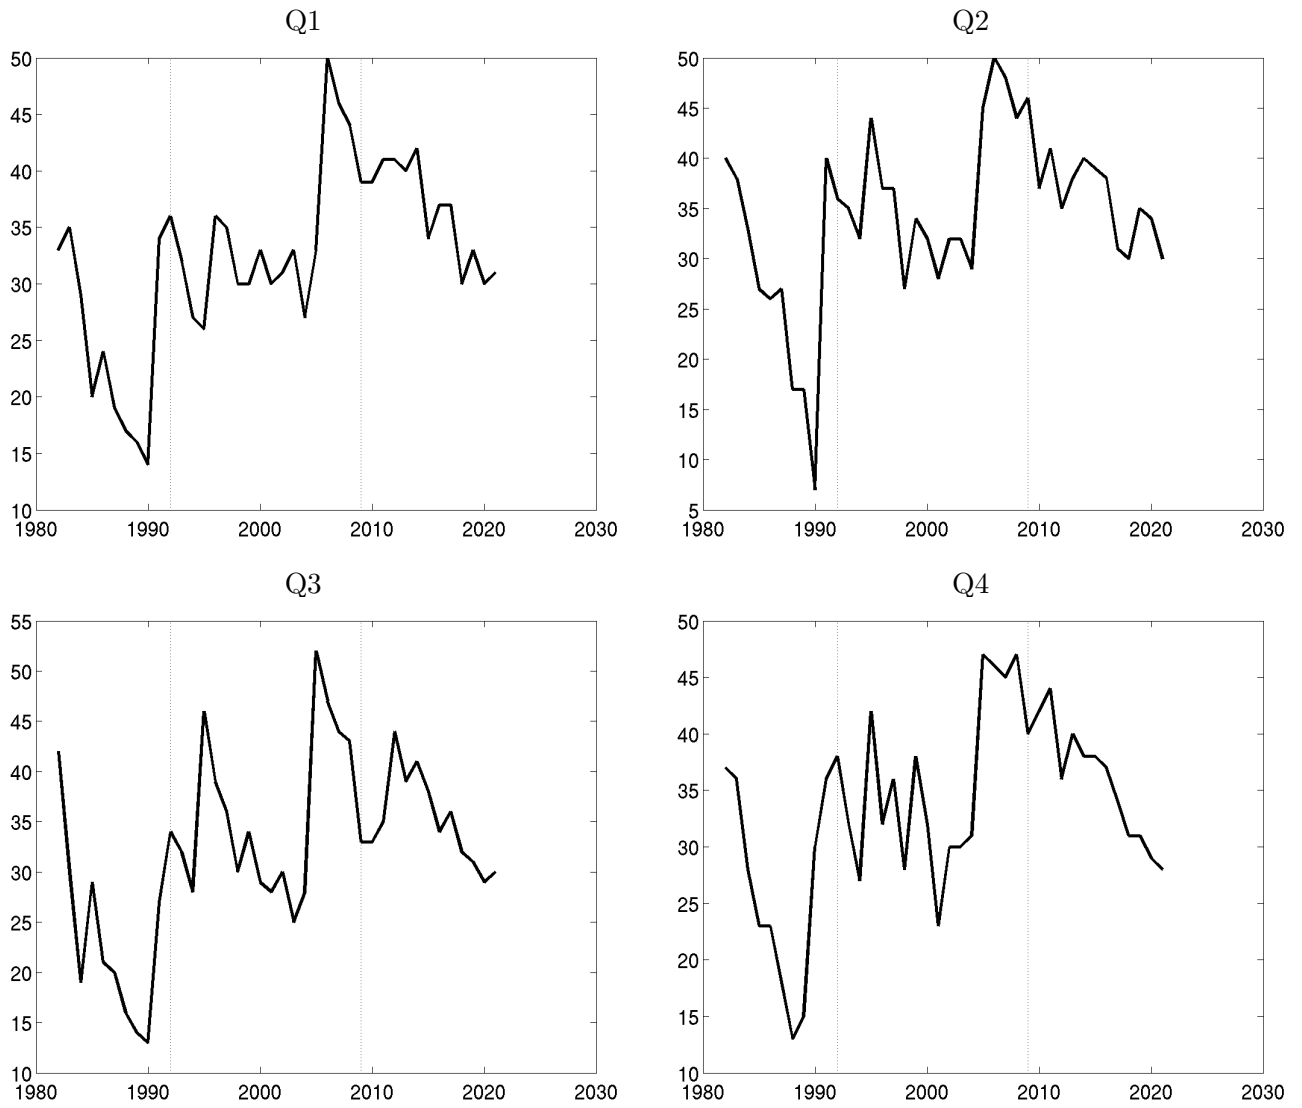



Figure A-5: SPF survey participation by respondent



Note: The light gray crosses indicate when respondents participate in a survey, and are connected by a thin dotted gray line whenever the respondent appears in consecutive surveys. Respondents are indexed by a number increasing in the year they joined the survey (y axis).

Figure A-6: Percentage of respondents for H2 Output growth surveys placing positive probability on either one open bin (solid line) or both (dash-and-dotted line)

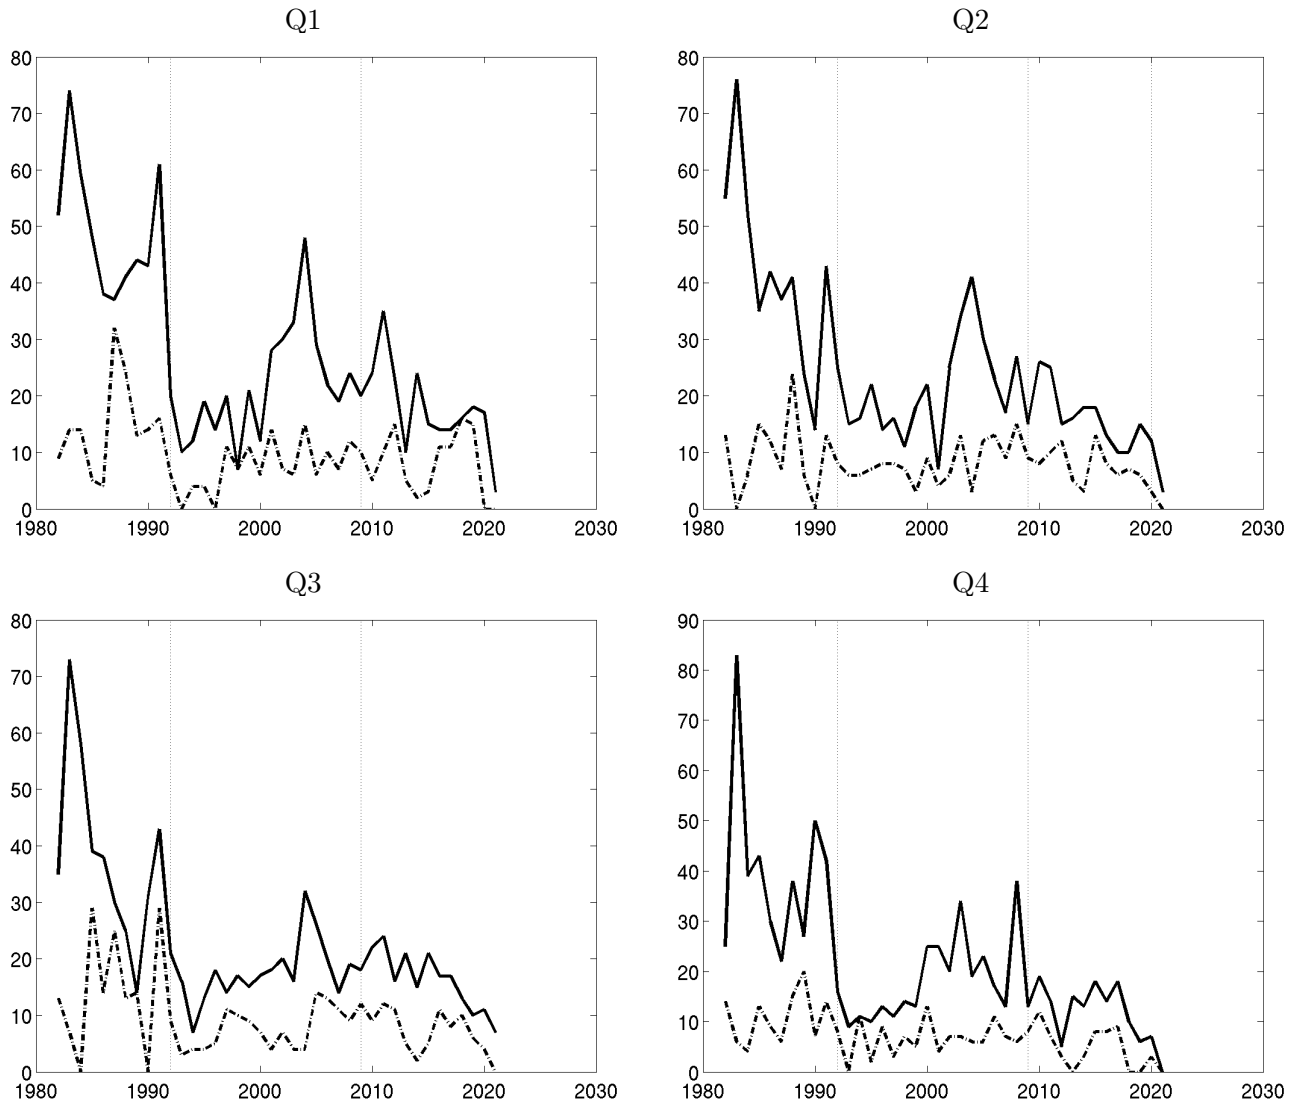
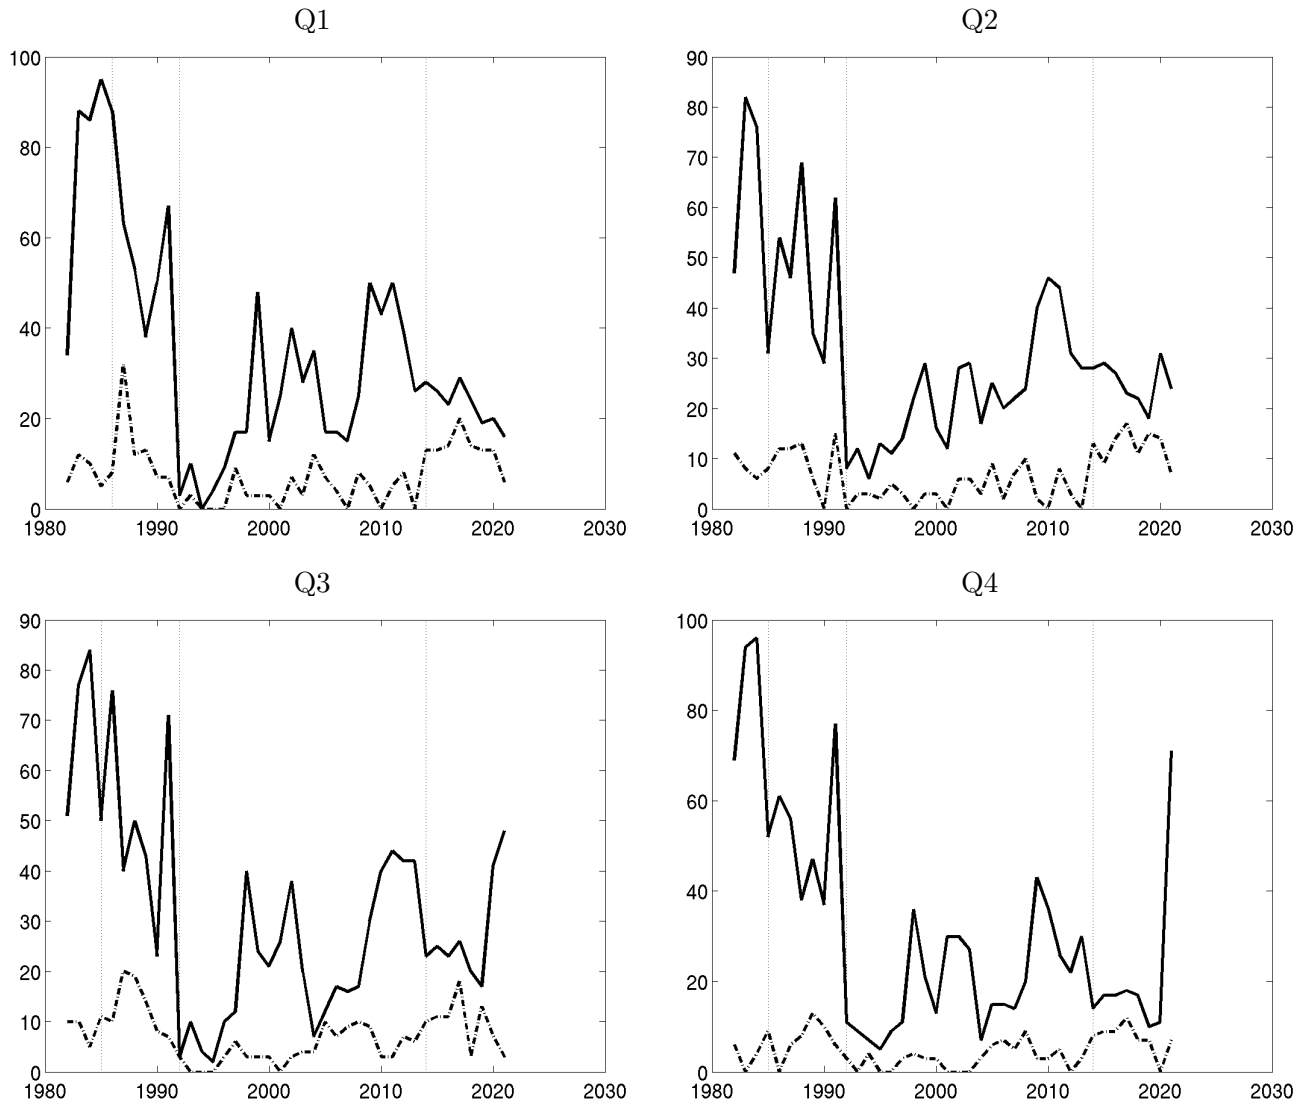
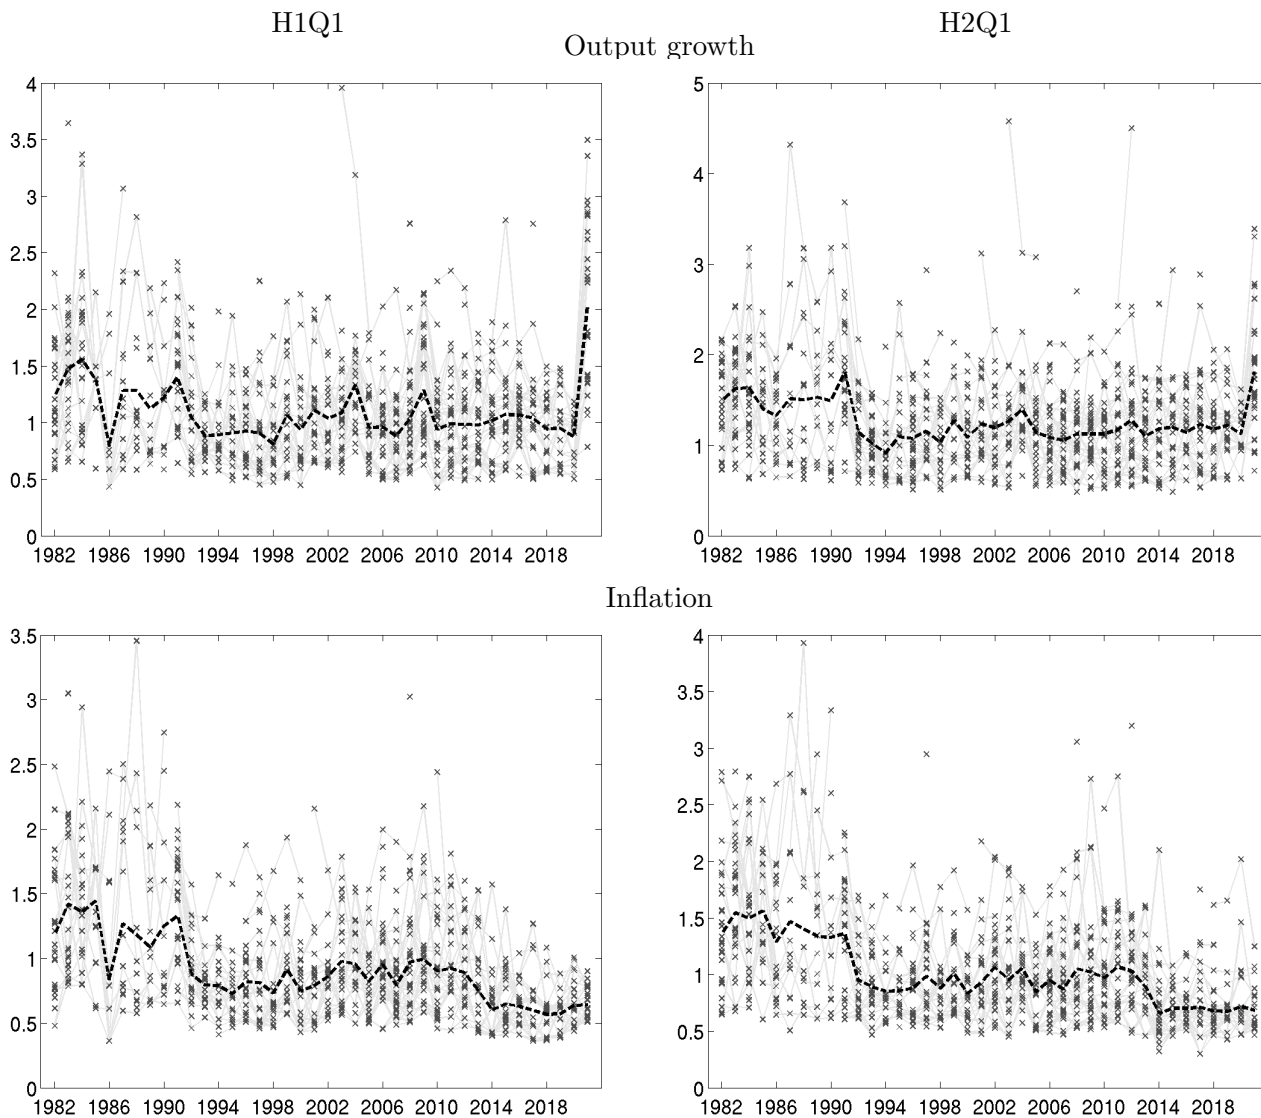


Figure A-7: Percentage of respondents for H2 inflation surveys placing positive probability on either one open bin (solid line) or both (dash-and-dotted line)



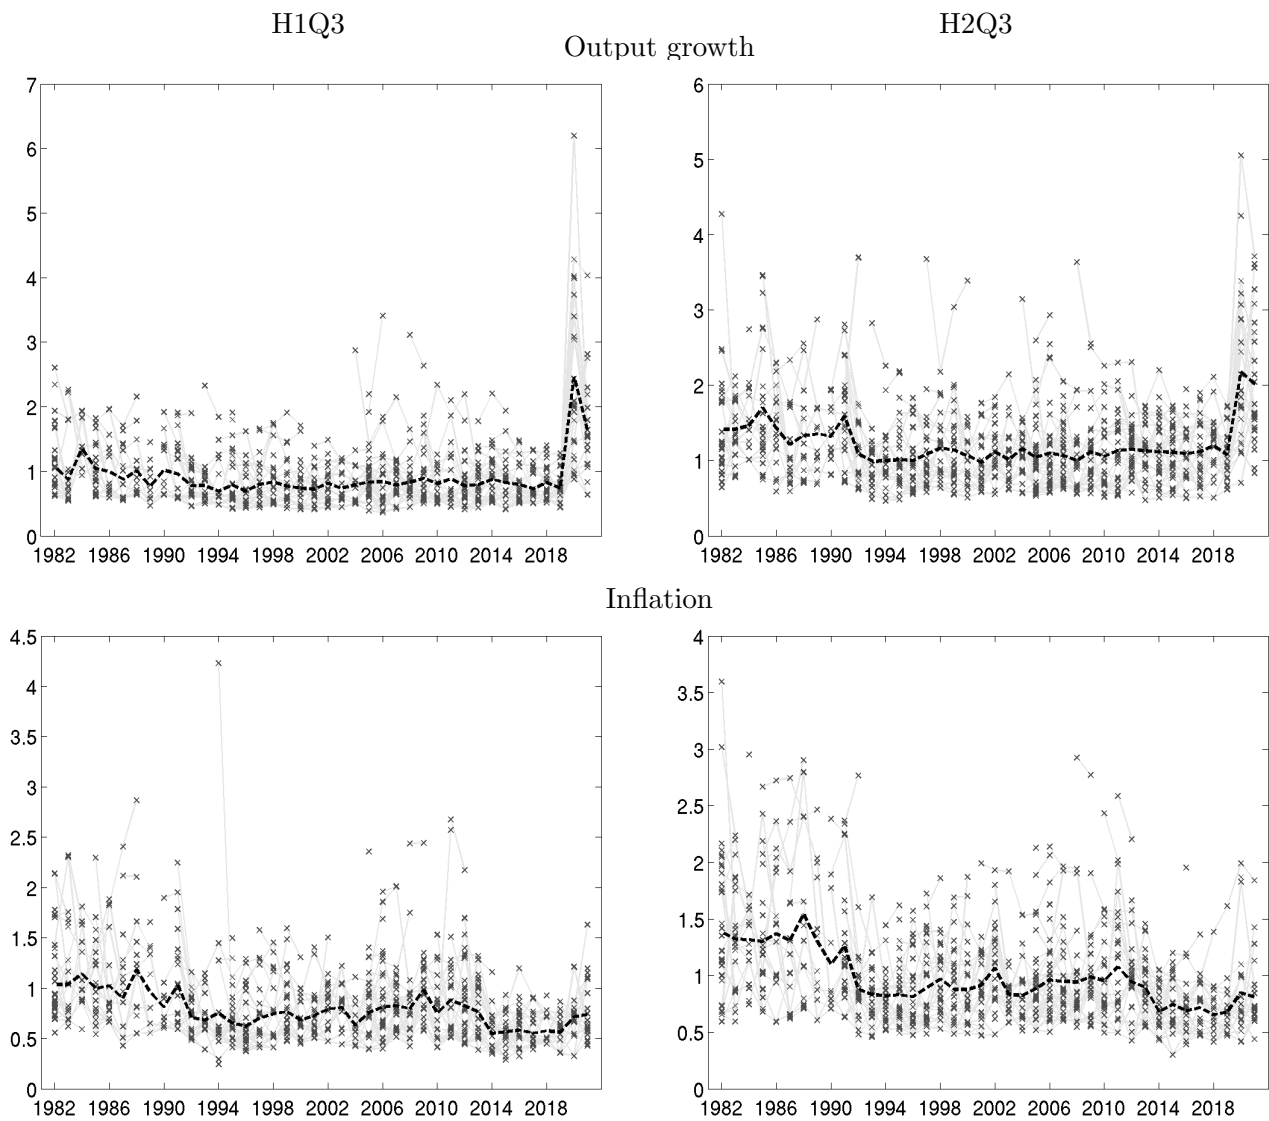
E.C Heterogeneity in subjective uncertainty

Figure A-8: Subjective uncertainty by individual respondent: Q1 survey



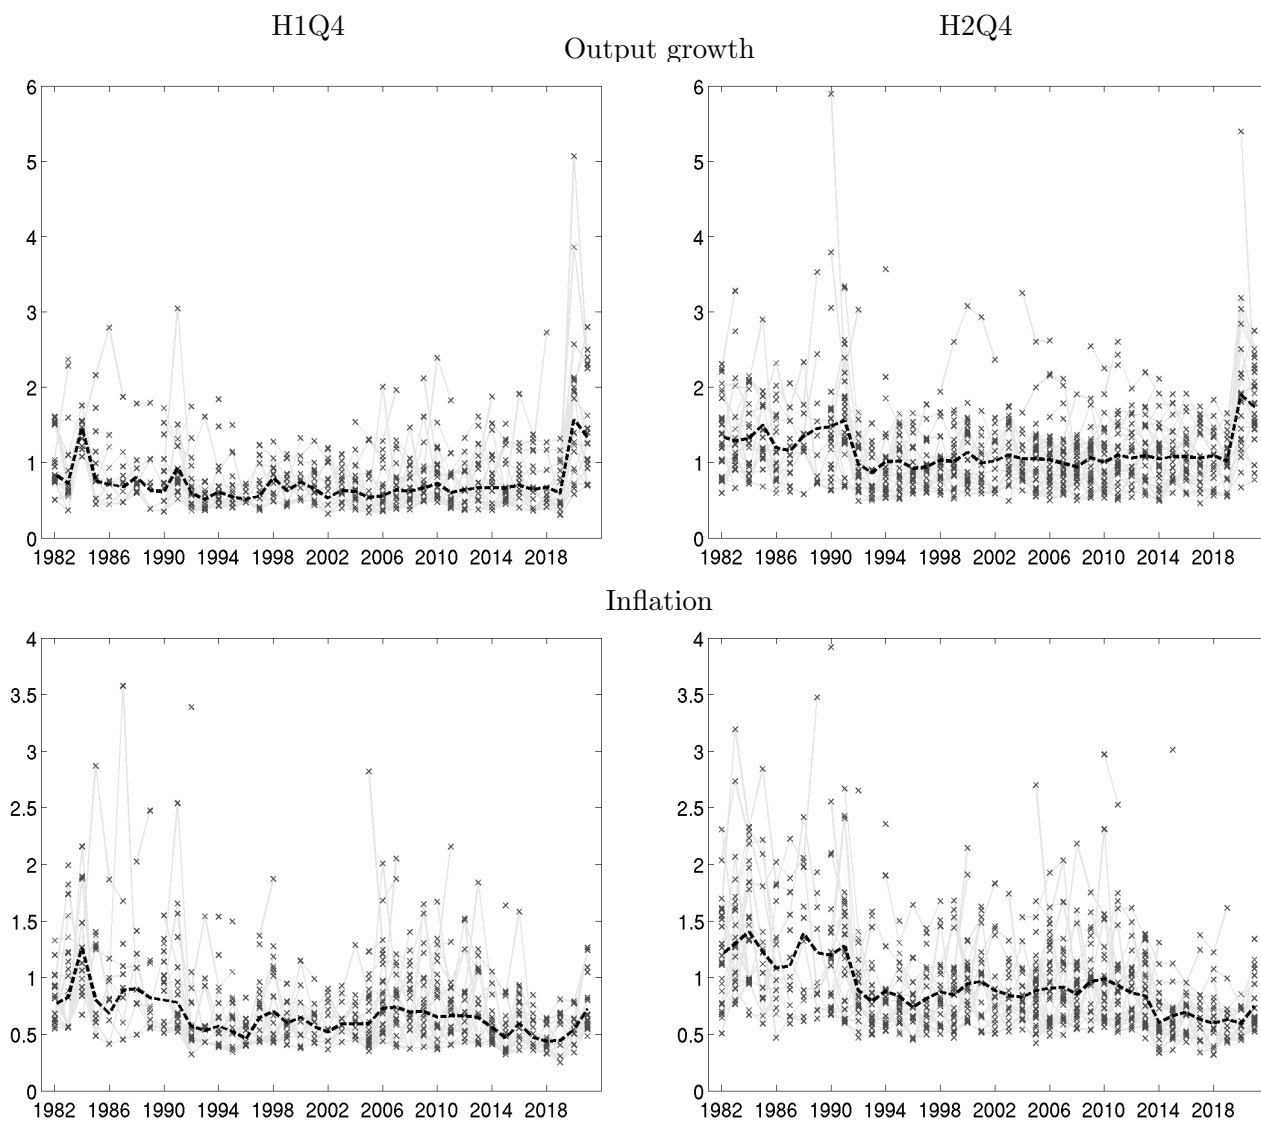
Note: Each panel displays the posterior mean of the standard deviation of the subjective predictive distribution by individual respondent (light gray crosses, connected by thin gray line whenever the respondent appears in consecutive surveys), and the cross-sectional average of the individual standard deviations (dashed black line). Top panels: Output growth projections; bottom panels: inflation projections. Left column: current year projections; right column: following year projections.

Figure A-9: Subjective uncertainty by individual respondent: Q3



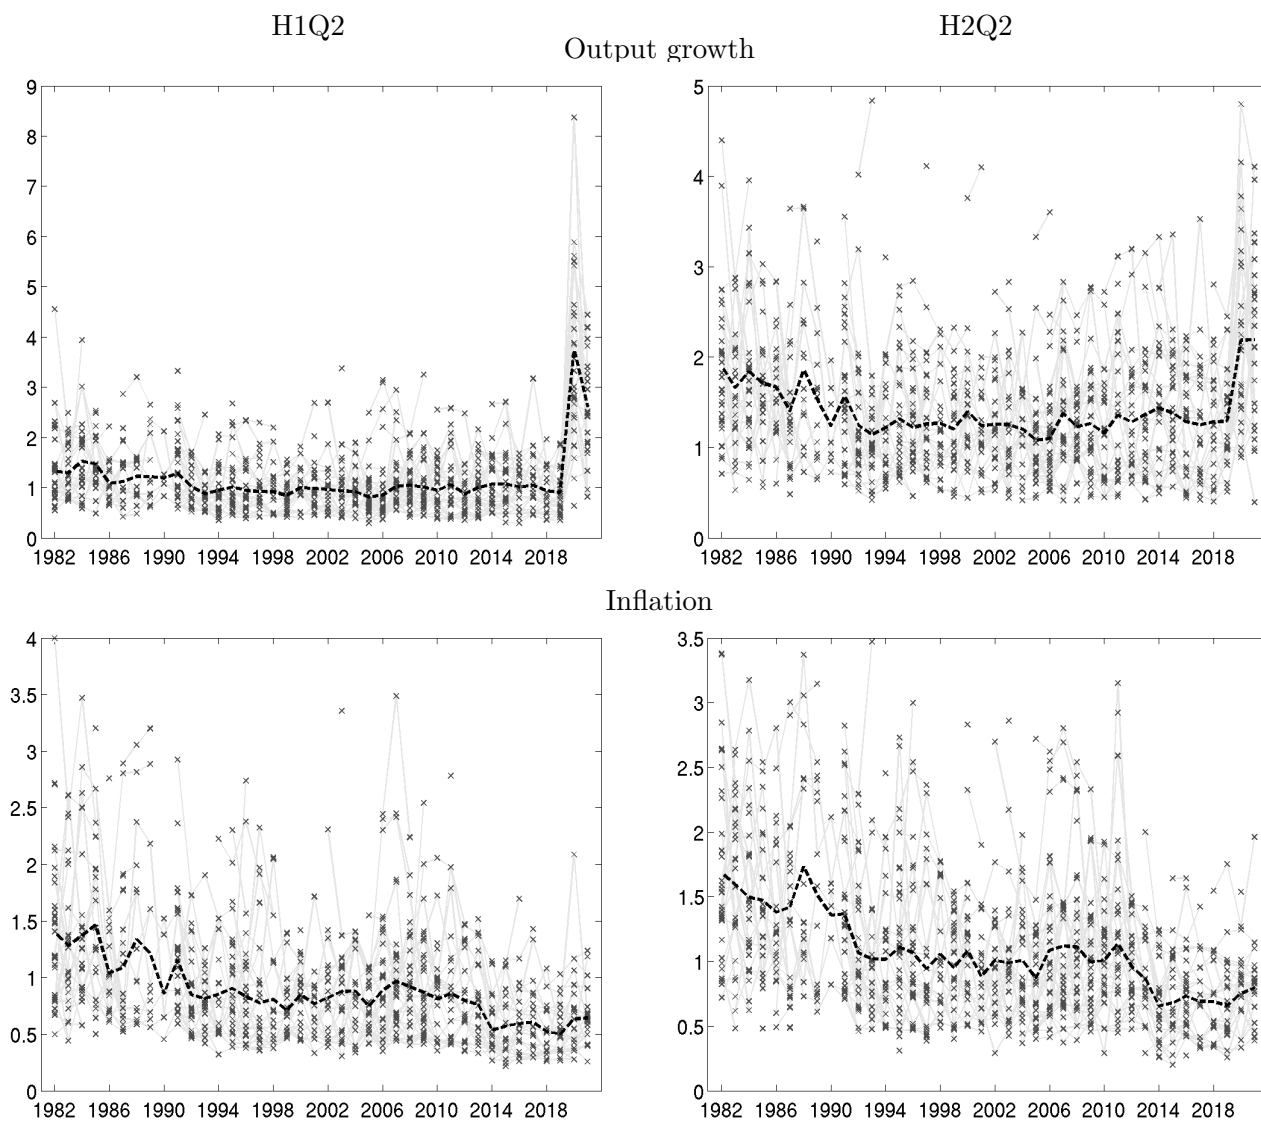
Note: Each panel displays the posterior mean of the standard deviation of the subjective predictive distribution by individual respondent (light gray crosses, connected by thin gray line whenever the respondent appears in consecutive surveys), and the cross-sectional average of the individual standard deviations (dashed black line). Top panels: Output growth projections; bottom panels: inflation projections. Left column: current year projections; right column: following year projections.

Figure A-10: Subjective uncertainty by individual respondent: Q4



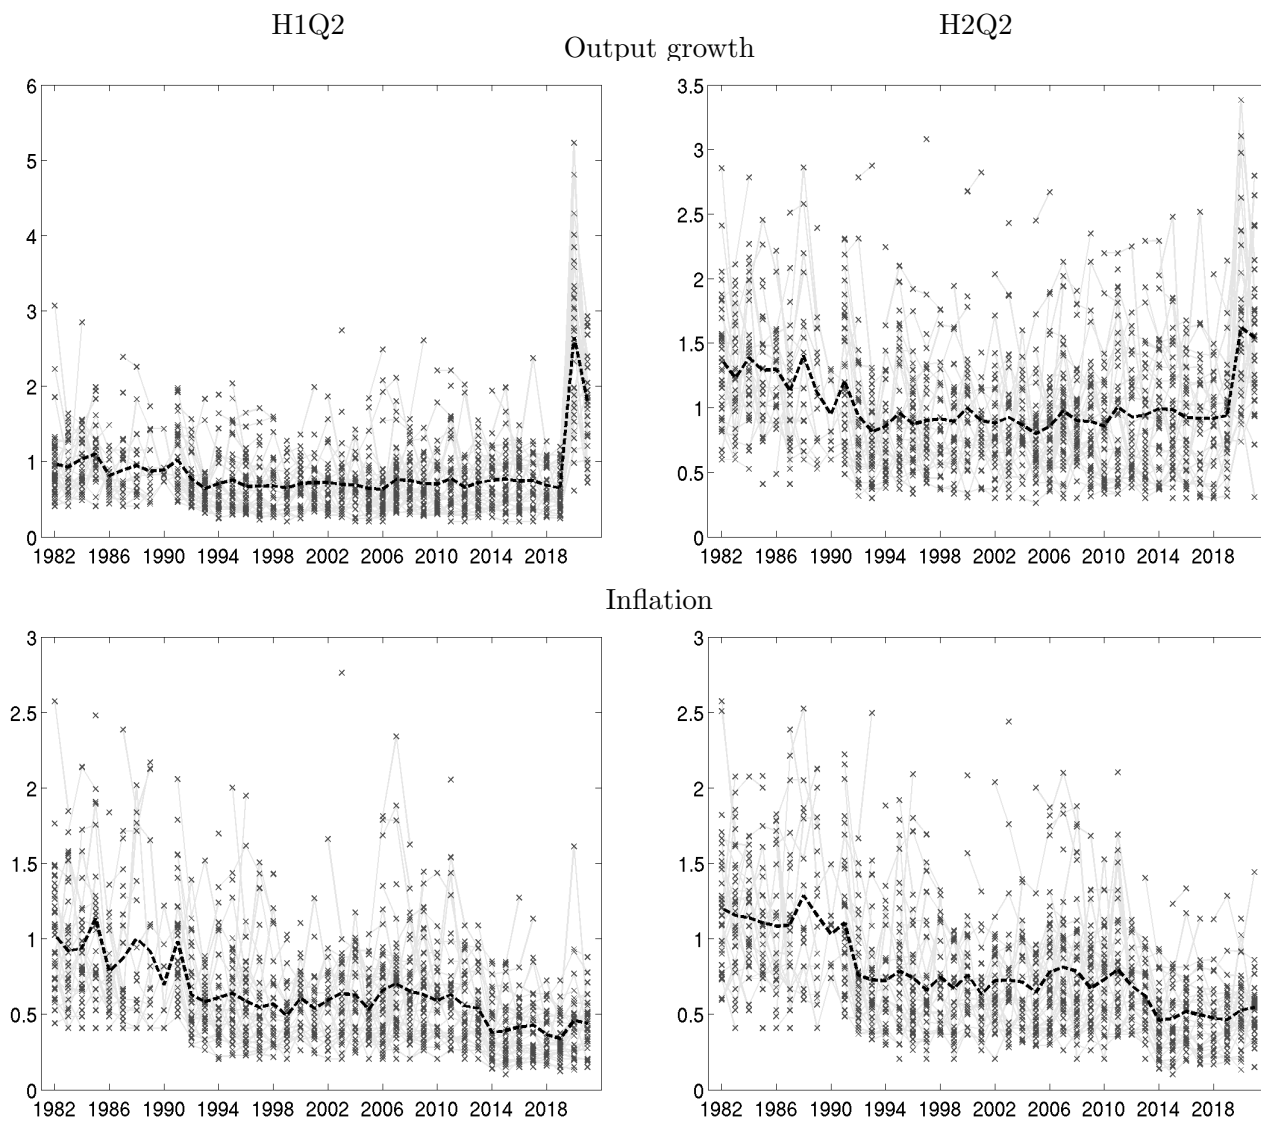
Note: Each panel displays the posterior mean of the standard deviation of the subjective predictive distribution by individual respondent (light gray crosses, connected by thin gray line whenever the respondent appears in consecutive surveys), and the cross-sectional average of the individual standard deviations (dashed black line). Top panels: Output growth projections; bottom panels: inflation projections. Left column: current year projections; right column: following year projections.

Figure A-11: Subjective uncertainty by individual respondent: IQRs, Q2 survey



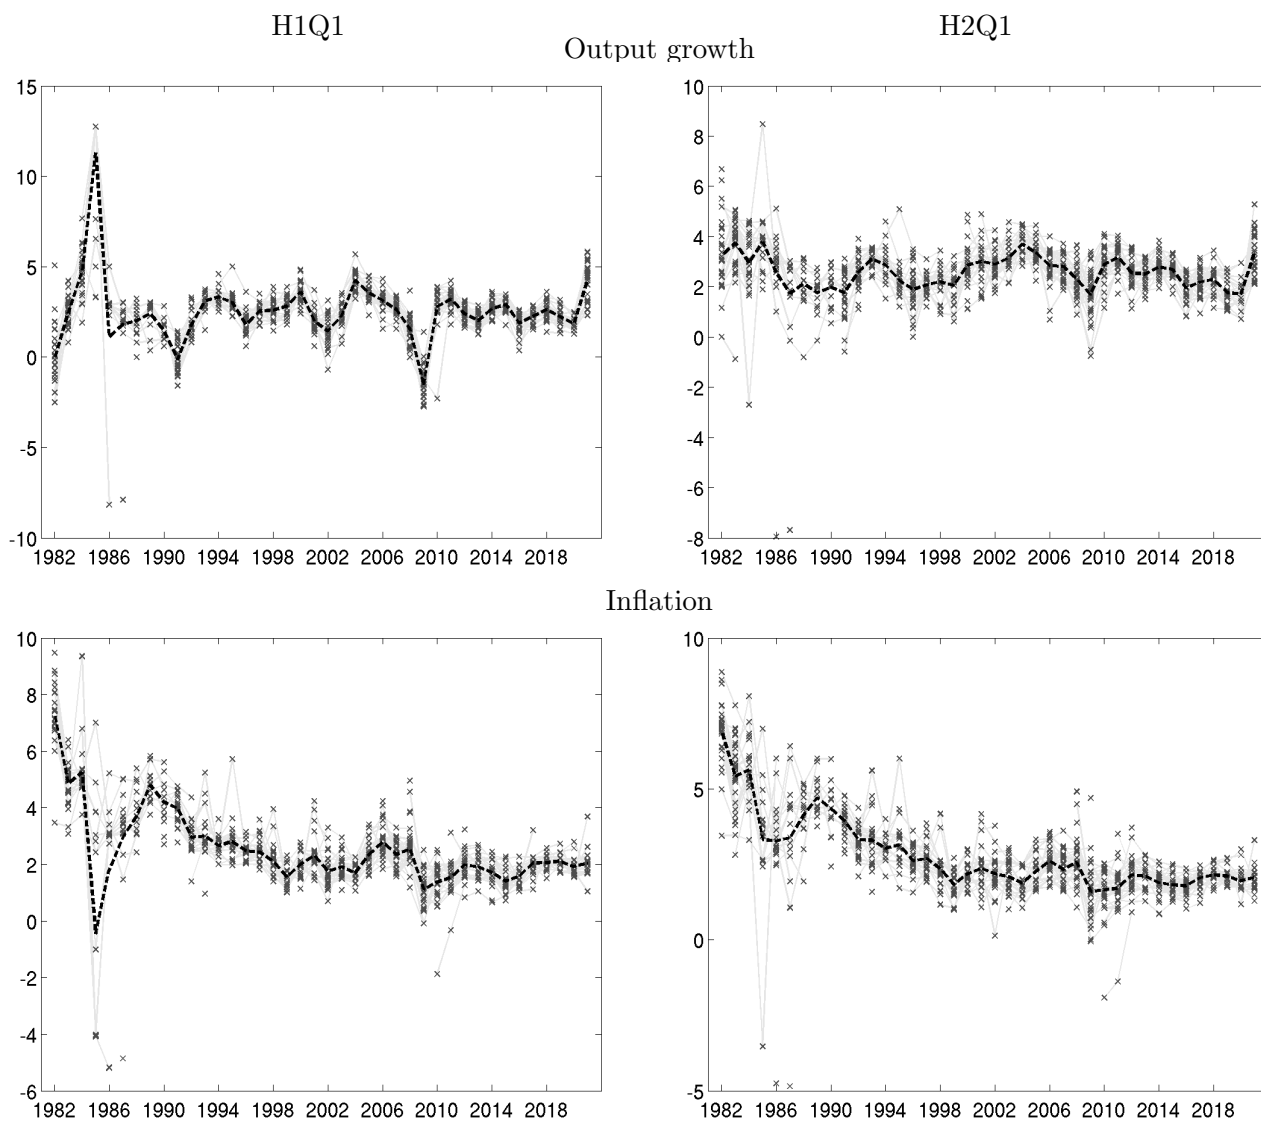
Note: Each panel displays the posterior mean of the interquartile range of the subjective predictive distribution by individual respondent (light gray crosses, connected by thin gray line whenever the respondent appears in consecutive surveys), and the cross-sectional average of individual interquartile ranges (dashed black line). Top panels: Output growth projections; bottom panels: inflation projections. Left column: current year projections; right column: following year projections.

Figure A-12: Subjective uncertainty by individual respondent: Beta, Q2 survey



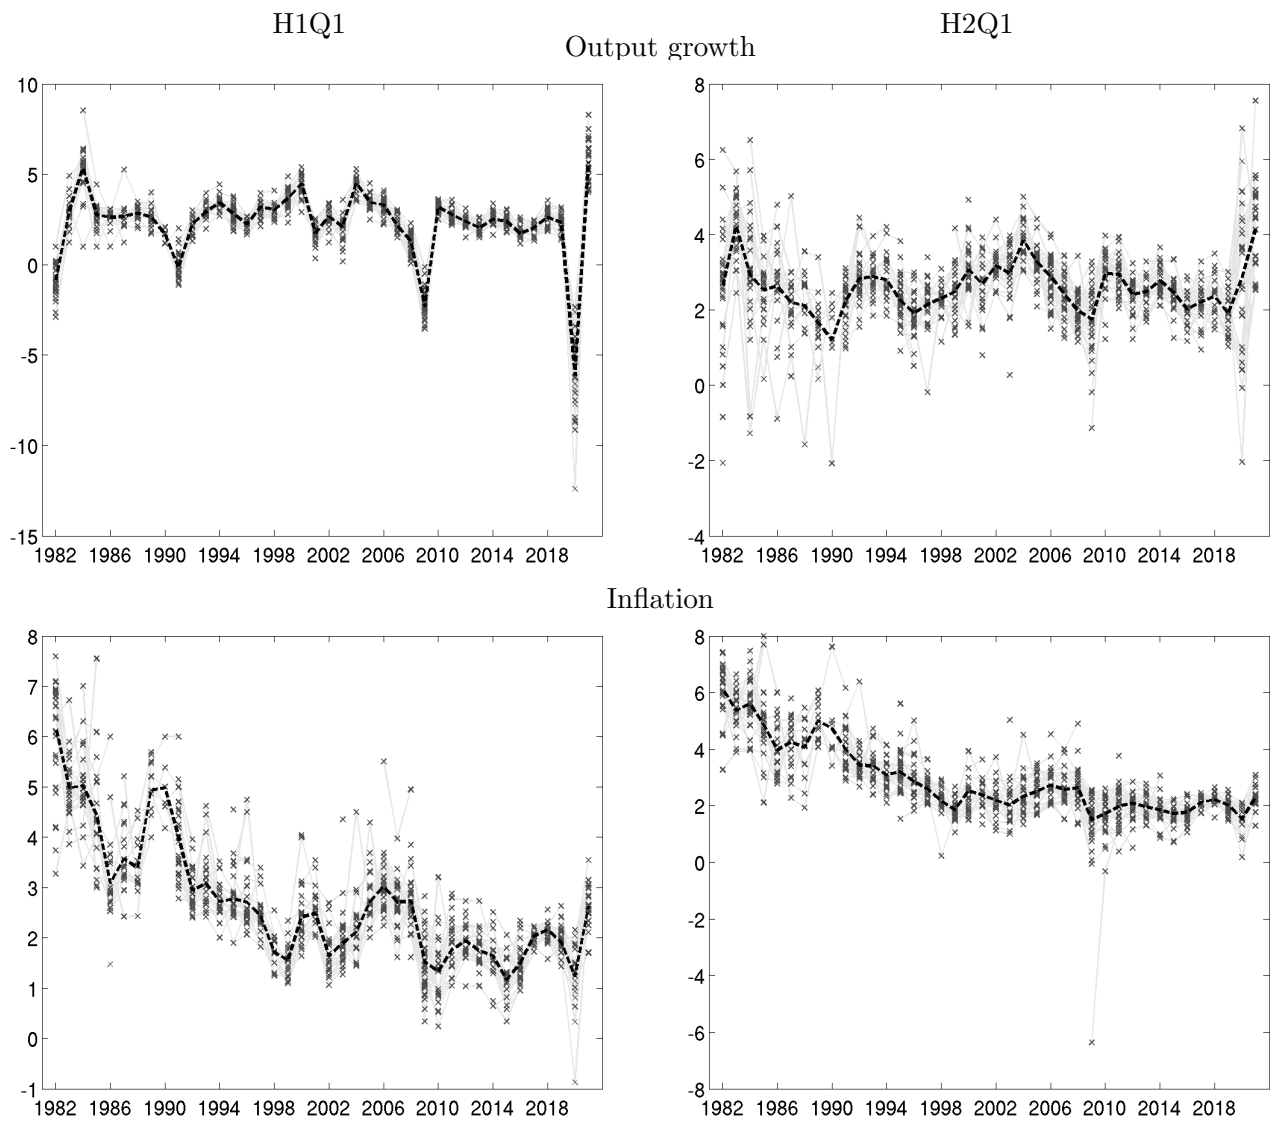
Note: Each panel displays the standard deviation of the subjective predictive distribution by individual respondent when this is estimated using least-squares under the beta parametric assumption (light gray crosses, connected by thin gray line whenever the respondent appears in consecutive surveys), and the cross-sectional average of the individual standard deviations (dashed black line). Top panels: Output growth projections; bottom panels: inflation projections. Left column: current year projections; right column: following year projections.

Figure A-13: Mean predictions by individual respondent: Q1 survey



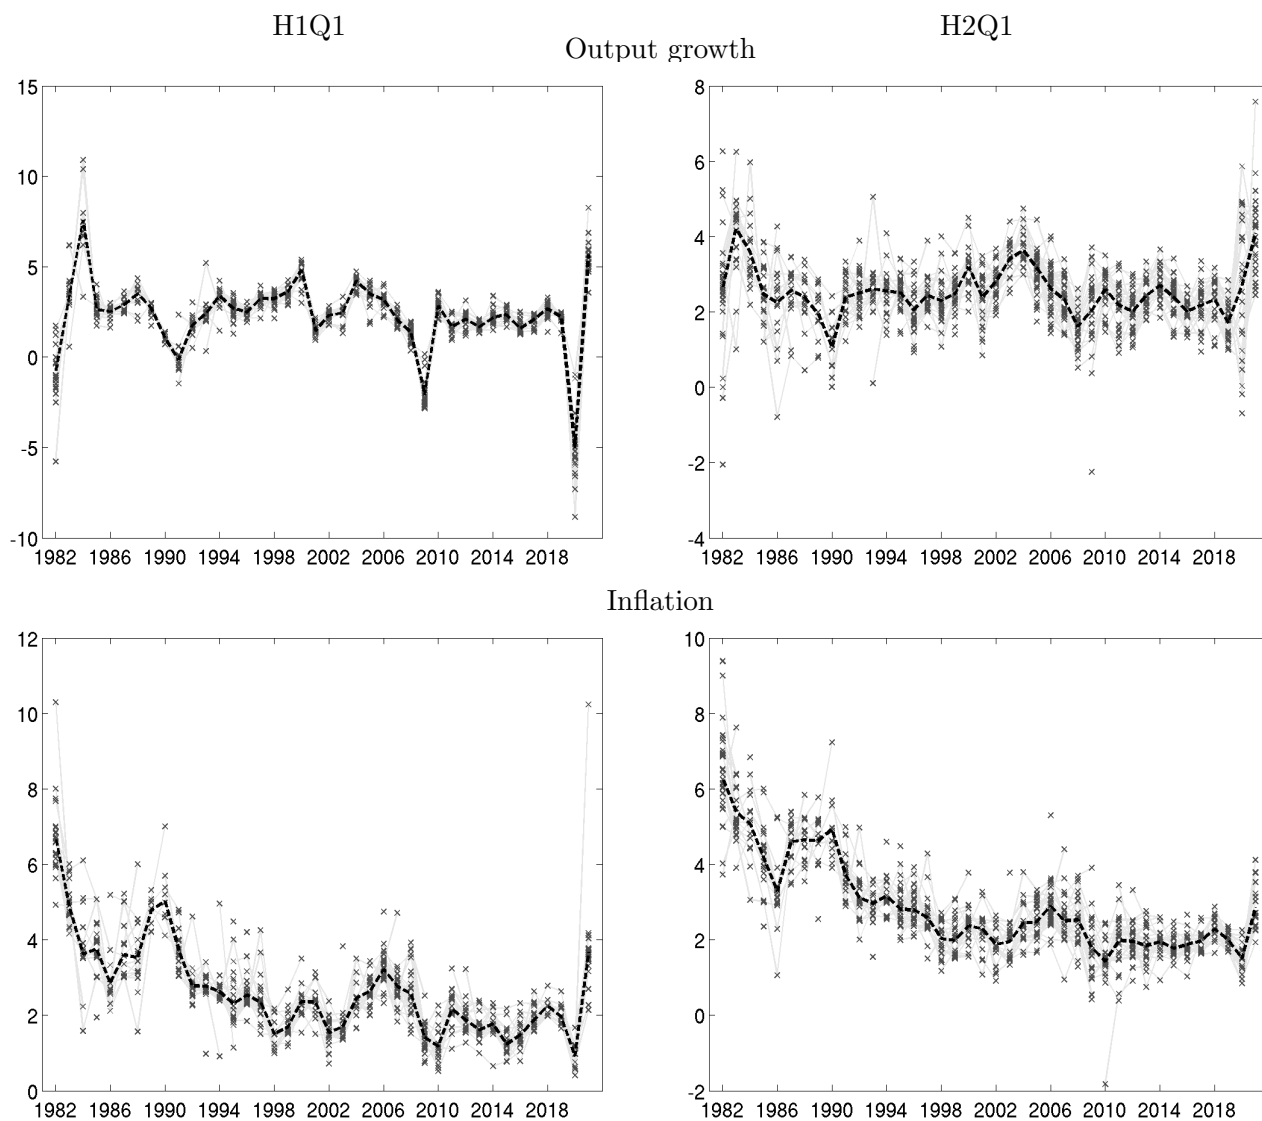
Note: Each panel displays the posterior mean of the mean of the subjective predictive distribution by individual respondent (light gray crosses, connected by thin gray line whenever the respondent appears in consecutive surveys), and the cross-sectional average of the individual means (dashed black line). Top panels: Output growth projections; bottom panels: inflation projections. Left column: current year projections; right column: following year projections.

Figure A-14: Mean predictions by individual respondent: Q2 survey



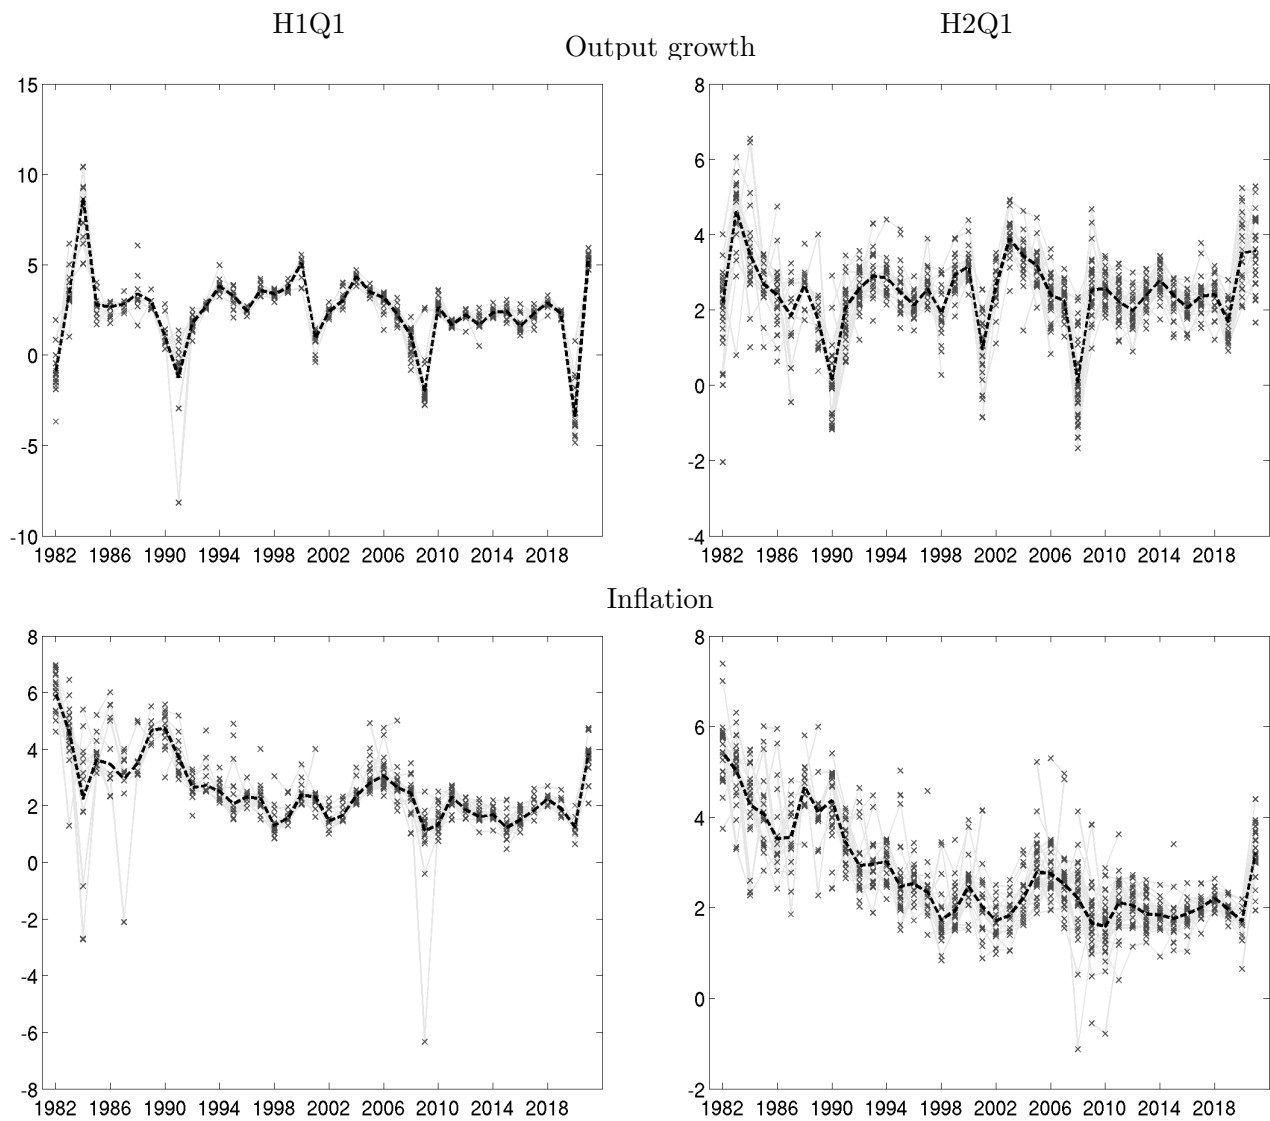
Note: Each panel displays the posterior mean of the mean of the subjective predictive distribution by individual respondent (light gray crosses, connected by thin gray line whenever the respondent appears in consecutive surveys), and the cross-sectional average of the individual means (dashed black line). Top panels: Output growth projections; bottom panels: inflation projections. Left column: current year projections; right column: following year projections.

Figure A-15: Mean predictions by individual respondent: Q3 survey



Note: Each panel displays the posterior mean of the mean of the subjective predictive distribution by individual respondent (light gray crosses, connected by thin gray line whenever the respondent appears in consecutive surveys), and the cross-sectional average of the individual means (dashed black line). Top panels: Output growth projections; bottom panels: inflation projections. Left column: current year projections; right column: following year projections.

Figure A-16: Mean predictions by individual respondent: Q4 survey




Note: Each panel displays the posterior mean of the mean of the subjective predictive distribution by individual respondent (light gray crosses, connected by thin gray line whenever the respondent appears in consecutive surveys), and the cross-sectional average of the individual means (dashed black line). Top panels: Output growth projections; bottom panels: inflation projections. Left column: current year projections; right column: following year projections.

E.D Subjective uncertainty and forecast accuracy: additional results

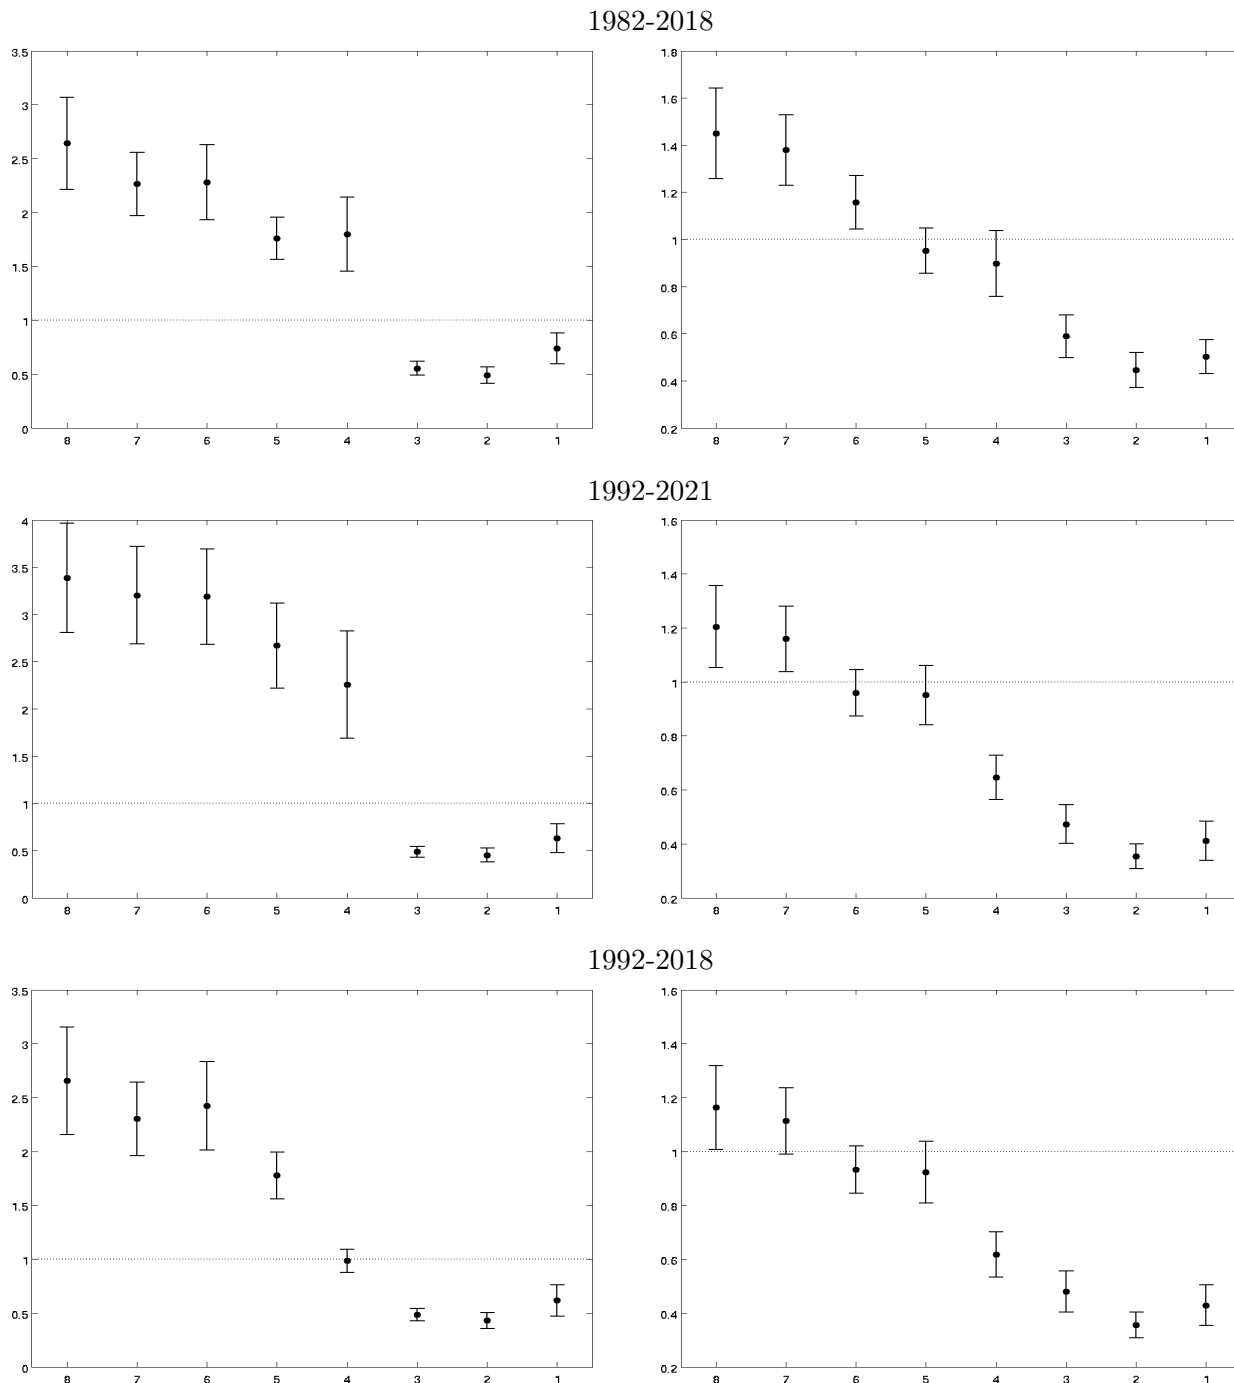
A scale test: additional results

Figure A-17: A scale test: mean vs point predictions
 Output growth Inflation



Note: Black dots correspond to OLS estimates of α_q from regression (20) for $q = 8, \dots, 1$. Gray crosses correspond to OLS estimates when the point prediction $y_{t-q,i}^{pp}$ is used in place of $E_{t-q,i}[y_t]$. Whiskers indicate 90 percent posterior coverage intervals based on robust standard errors.


Figure A-19: A scale test: different samples



Note: Black dots correspond to OLS estimates of α_q from regression (20) for $q = 8, \dots, 1$. Solid black whiskers indicate 90 percent posterior coverage intervals based on robust standard errors.

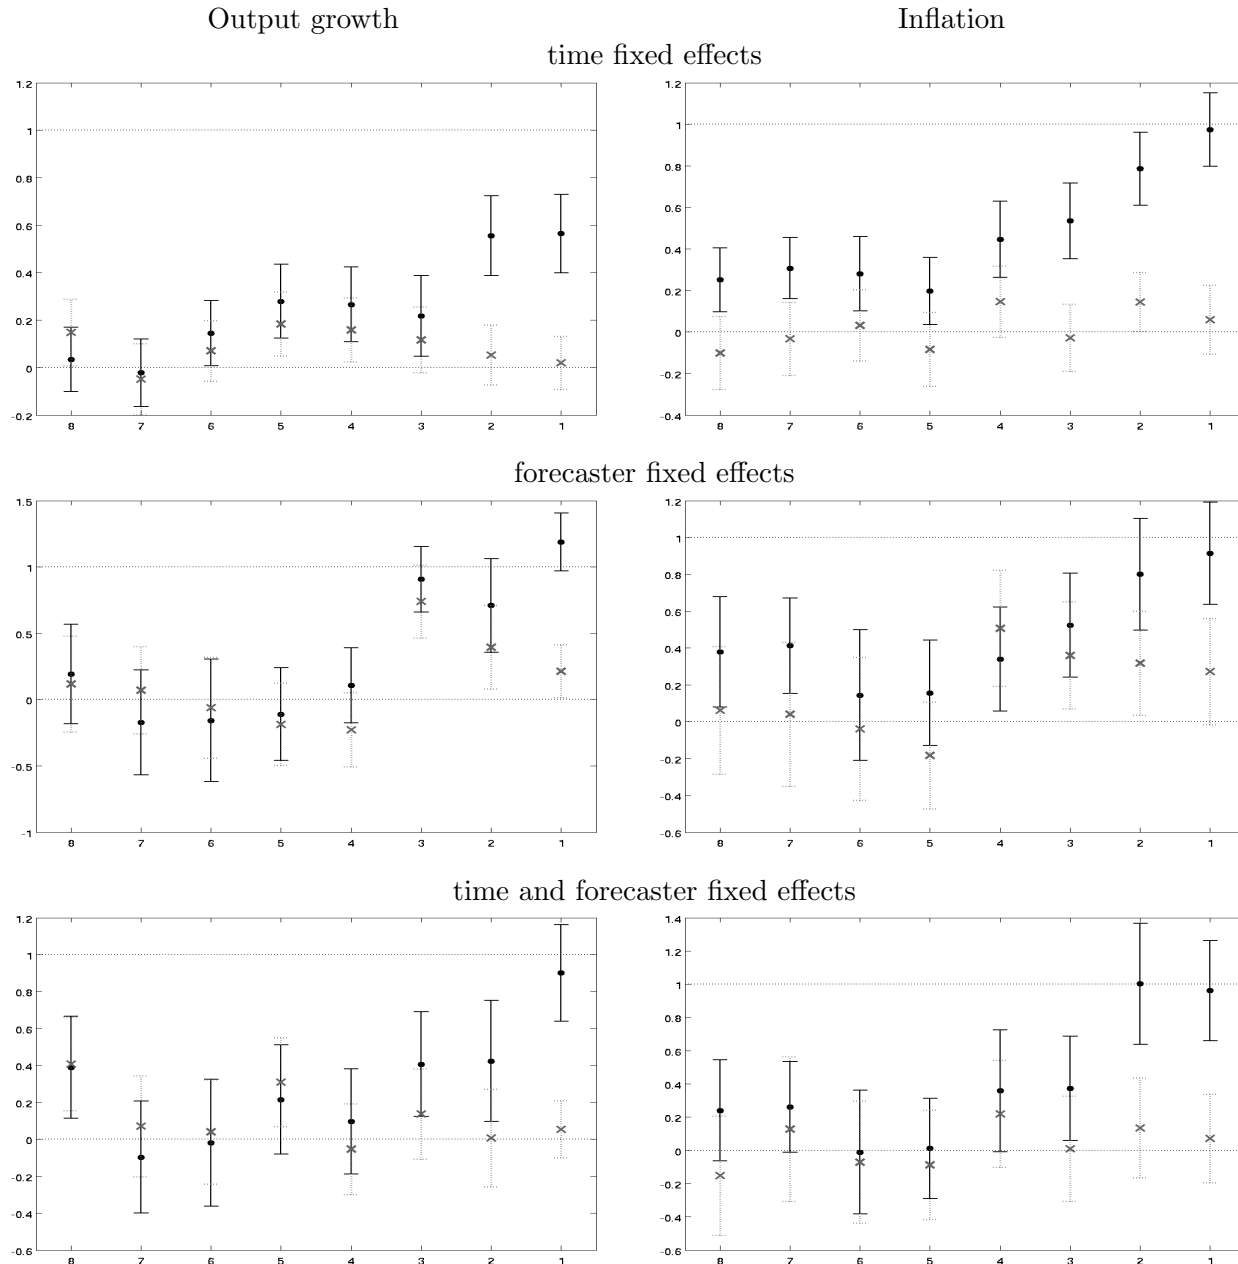
A variation test: additional results (unweighted)

Figure A-20: A variation test—mean vs point projections



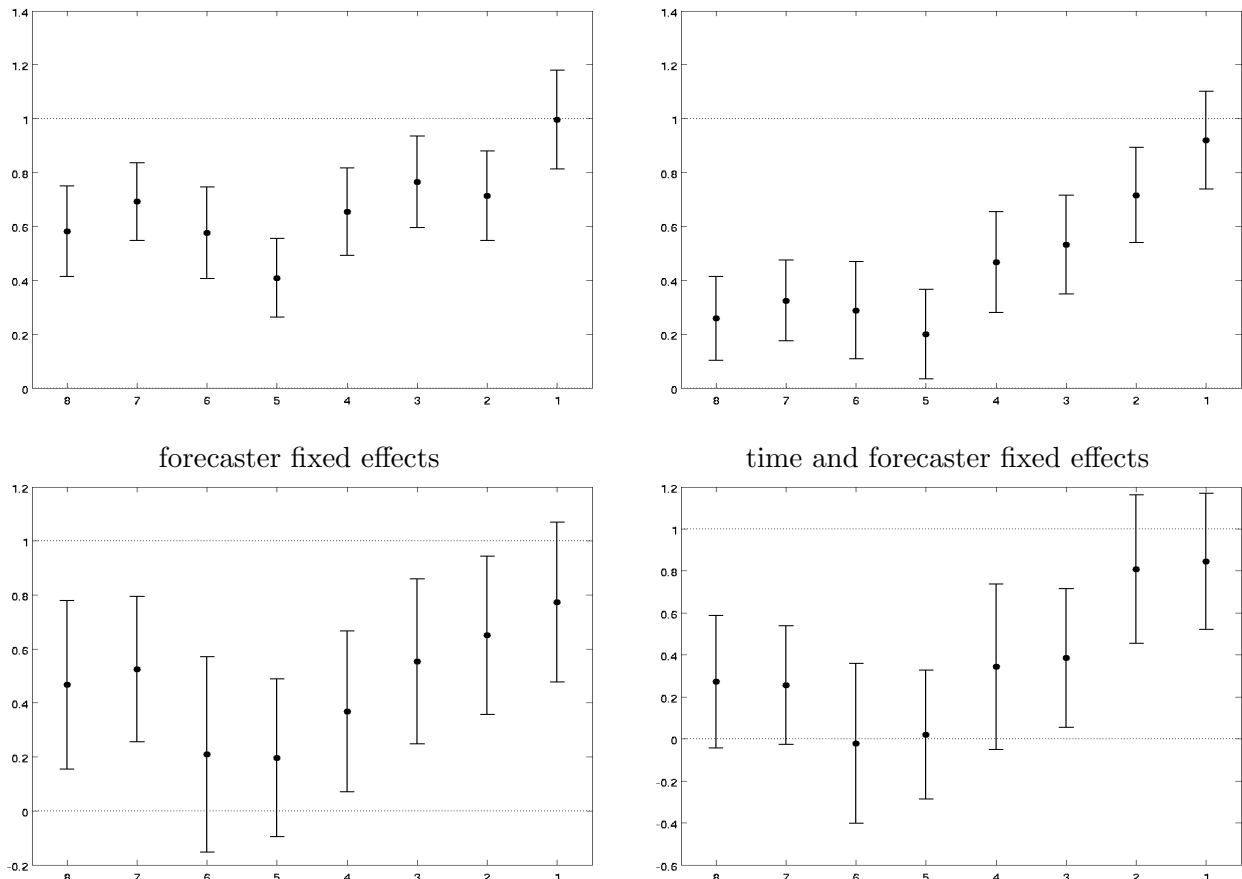
Note: Black dots correspond to OLS estimates of $\beta_{1,q}$ from regression (21) for $q = 8, \dots, 1$. Gray crosses correspond to OLS estimates when the point prediction $y_{t-q,i}^{pp}$ is used in place of $E_{t-q,i}[y_t]$. Whiskers indicate 90 percent posterior coverage intervals based on robust standard errors.

Figure A-21: A variation test: Regressions with fixed effects—mean vs point projections





Note: Black dots correspond to OLS estimates of $\beta_{1,q}$ from regression (21) using time (top panels), forecaster fixed effects (middle panels), or both (bottom panels), for $q = 8, \dots, 1$. Gray crosses correspond to OLS estimates when the point prediction $y_{t-q,i}^{pp}$ is used in place of $E_{t-q,i}[y_t]$. Dotted gray whiskers indicate 90 percent posterior coverage intervals based on robust standard errors.

Figure A-25: A variation test: inflation; 1982-2018 sample
 No fixed effects time fixed effects




Note: Black dots correspond to OLS estimates of $\beta_{1,q}$ from regression (21) for $q = 8, \dots, 1$. Solid black whiskers indicate 90 percent posterior coverage intervals based on robust standard errors.


Figure A-27: A variation test: inflation; 1992-2021 sample
 No fixed effects time fixed effects



forecaster fixed effects





time and forecaster fixed effects




Note: Black dots correspond to OLS estimates of $\beta_{1,q}$ from regression (21) for $q = 8, \dots, 1$. Solid black whiskers indicate 90 percent posterior coverage intervals based on robust standard errors.


Figure A-29: A variation test: inflation; 1992-2018 sample
 No fixed effects time fixed effects



forecaster fixed effects




time and forecaster fixed effects



Note: Black dots correspond to OLS estimates of $\beta_{1,q}$ from regression (21) for $q = 8, \dots, 1$. Solid black whiskers indicate 90 percent posterior coverage intervals based on robust standard errors.


A variation test: Additional Results (weighted)

Figure A-30: A variation test: Mean vs Point Projections-weighted output growth inflation



Note: Black dots correspond to OLS estimates of $\beta_{1,q}$ from regression (21) for $q = 8, \dots, 1$. Gray crosses correspond to OLS estimates when the point prediction $y_{t-q,i}^{pp}$ is used in place of $E_{t-q,i}[y_t]$. Whiskers indicate 90 percent posterior coverage intervals based on robust standard errors.

Figure A-31: A variation test: regressions with fixed effects for both mean and point forecasts-weighted



Note: Black dots correspond to OLS estimates of $\beta_{1,q}$ from regression (21) using time (top panels), forecaster fixed effects (middle panels), or both (bottom panels), for $q = 8, \dots, 1$. Gray crosses correspond to OLS estimates when the point prediction $y_{t-q,i}^{pp}$ is used in place of $E_{t-q,i}[y_t]$. Whiskers indicate 90 percent posterior coverage intervals based on robust standard errors.

