# 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*

April 18, 2022

**Abstract**

We propose a non-parametric 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. 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) made economists appreciate the advantages of probabilistic surveys relative 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 interquantile 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 non-parametric approach for the estimation of the survey respondents' forecast densities.[3] The approach starts by making parametric assumptions 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

[1]Indeed, a number of recent surveys, including the Federal Reserve Bank of New York Survey of Consumer Expectations, rely heavily on probabilistic questions.

[2]For a few quantities of interest, such as the median, one can compute non-parametric bounds as in Engelberg et al. (2009), which depends on how one deals with reporting "noise" (e.g., rounding).

[3]In economics, the Bayesian non-parametric approach so far has applied to the analysis of treatment effects (Chib and Hamilton, 2002), autoregressive panel data (Hirano, 2002; Gu and Koenker, 2017; Liu, 2021), 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 up to 2011. Outside of economics, these methods are widely used in biostatistics (Mitra and Müller, 2015), machine learning (Blei et al., 2010, Hannah et al., 2011), and psychology (Griffiths and Tenenbaum, 2006).

in the reporting (e.g., rounding toward zero). We then relax this parametric model by embedding it into the more general Bayesian non-parametric approach, thereby amending the potential misspecification associated with the parametric assumptions. This is because, loosely speaking, Bayesian non parametric replaces any 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 provide 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 noise). The non-parametric 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, 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 the Bayesian non-parametric 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 forecast subjective distributions. The rest of the section is devoted to the description of our Bayesian non-parametric approach to inference.

## 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, such as inflation or GDP growth, falls within different contiguous ranges, where these ranges are pre-specified by the survey designer (some 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). For each forecaster $i = 1, ..., n$ the available data consists of a vector of probabilities $\mathbf{z}_i = (z_{i,1}, \ldots, z_{i,J})$, with $z_{i,j} \geq 0$ and $\sum_{j=1}^{J} z_{i,j} = 1$, measuring the predictive likelihood that continuous variable $y$ (e.g. inflation or GDP growth) falls within the respective bin. The bins are mutually exclusive and contiguous, and generally cover the

Figure 1: Probability Forecasts for Selected Examples

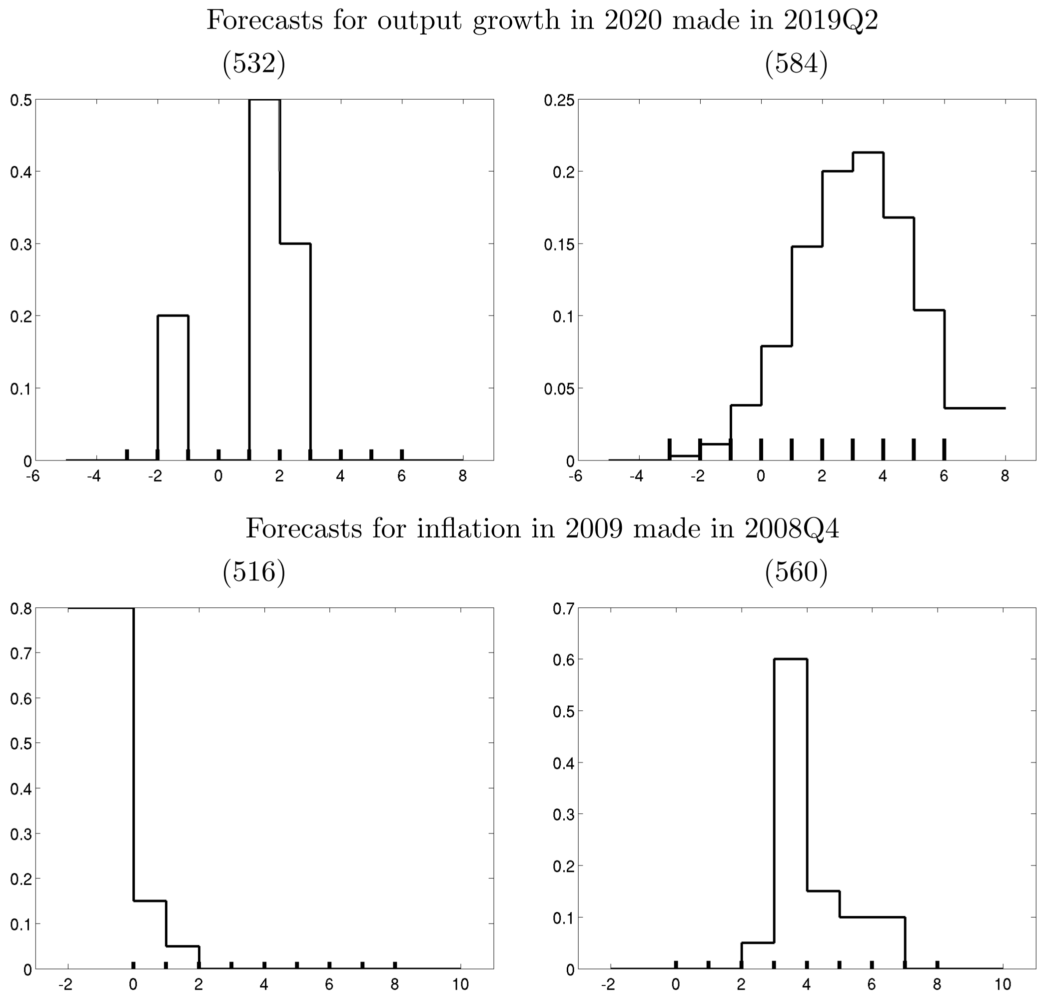


*Note*: Each panel displays the forecast probabilities $z_{i,j}$, $j = 1, \ldots, J$ (step-wise solid lines) for a given forecaster $i$ (forecaster number shown in parentheses) and the bin bounds (black ticks, horizontal axis).

entire real line. In what follows, we denote by $(y_{j-1}, y_j]$, $j = 1, \ldots, J$ the bins and assume that $y_0 < y_1 < \ldots < y_J$, where $y_0$ and $y_J$ are equal to $-\infty$ (left open bin) and $+\infty$ (right open bin), respectively. Figure D-1 in the Appendix displays the evolution of the bin ranges from the beginning of our sample, in 1982, 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 fact that the bin boundaries change over time needs to be borne in mind when comparing surveys for different years.

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) discusses at length some of the features of the SPF survey, and the Philadelphia Fed's site provides a manual for interpreting the data that includes the history up to the present of bin boundaries for the various variables being forecast.[4]

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 made 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. The probabilities $\mathbf{z}_i$'s are displayed as histograms, while the black ticks on the horizontal axis mark the boundaries of the bins.

The 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. Another feature is the fact that 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. Should the econometrician interpret this information literally, or as an indication that this respondent has a bimodal forecast distribution with some probability on a recession, and a larger probability on an expansion, with very small but not literally zero 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 D-4 and D-5 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 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

[4]Figure D-2 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). The number of respondents 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 D-3 shows survey participation by respondent, and provides a visual description of the panel's composition.

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 skew in their predictive densities? The general approach for macroeconomic surveys has been to postulate that forecasters $i = 1, \dots, n$ have in mind a given predictive probability distribution $F_i(y)$ over the variable being forecast, which they use to assign the bin probabilities $\mathbf{z}_i$. The task of the econometrician is then to infer the underlying $F_i(y)$ based on the data $\mathbf{z}_i$, and then use the estimated $F_i(y)$ to answer the questions of interest. To our knowledge, 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 $Z_{ij} = z_{i,1} + \cdots + z_{i,j}$ $j = 1, \dots, J$, $i = 1, \dots, n$ 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/piece-wise 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 skew-normal (Garcia and Manzanares, 2007), a generalized beta (Engelberg et al., 2009)[5] and a skew-t distribution (e.g., Ganics et al., 2020). The Gaussian and the generalized beta assumptions have been the most popular approaches in academic research, although in applied work at central banks the mixture of uniforms approach is often followed. 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

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

These approaches have been popular but have some limitations. A first limitation is that the assumed parametric distribution may be misspecified, in the sense that it may not fit the individual responses well. Relatedly, the width of the bins can be large, as is obviously the case when the respondent places probability on open bins (interior real output growth bins after 2020 are also very wide). This implies that even if the distributions fit the $Z_{ij}$'s, the inference results on moments and quantiles can be sensitive to the distributional assumption. A second issue is that bounded distributions such as the beta or the mixture of uniforms

[5] Whenever the number of (adjacent) bins with positive probability is two or fewer, Engelberg et al. (2009) uses a triangular distribution.

take literally the $z_{ij}$ that are zero, in that they place no probability mass on bins where the respondents place no mass. More in general, for all assumed $F(\cdot)$'s the approach outlined in expression (1) ignores the issue of rounding, in that it takes 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., 2008a, 2014; Engelberg et al., 2009; Manski and Molinari, 2010; Manski, 2011; Giustinelli et al., 2020, among others, discuss the issue of rounding; Binder, 2017 uses rounding to measure uncertainty).[6] Finally, almost all existing approaches ignore inference uncertainty, even that concerning $\boldsymbol{\theta}_i$ for a given parametric assumption, 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] These limitations are well known in the literature (see Clements et al., forthcoming). There have been attempts to address some of these issues, in particular the potential misspecification, by choosing more flexible families of distributions such as the skew-normal or the skew-Student-t distribution (e.g., Garcia and Manzanares, 2007; Ganics et al., 2020). But the possibility of misspecification remains. Most importantly, if the econometrician does not account for inference uncertainty, this flexibility comes at the price of overparamterization.

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 forecaster distribution. This model follows the literature in that it assumes that each forecasters has in mind a specific predictive distribution $F(\cdot)$ which he uses to assign probabilities $\boldsymbol{\nu}$ to the bins. It is different from the literature in that it explicitly postulates that the data $\mathbf{z}$ are noisy versions of the $\boldsymbol{\nu}$'s, where the noise is assumed to follow a parametric distribution. We then depart from this parametric framework by embedding it into a more general Bayesian *non-parametric* model, which assumes the parameters of the forecasters in the cross-section are drawn form an infinite mixture prior. The combination of different parameter draws

[6] Manski and Molinari (2010) and Giustinelli et al. (2020) propose to treat 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. It is important to note that the inferential approach based on interval data followed by these researchers is very different from the one described at the beginning of this section.

[7] Researchers recognize the emergence of an inference issue especially when the information provided by the respondent is very limited, but the proposed solution mostly amounts to either choosing less parameterized distributions or discarding the respondent. For instance, some researchers simply discard histograms with fewer than three bins Clements (2010), others (Engelberg et al., 2009; Clements, 2014b,a; Clements and Galvão, 2017) use a triangle distribution in these cases, as mentioned above. Liu and Sheng (2019), however, make an attempt to account for parameter uncertainty for given parametric assumptions. They propose maximum likelihood estimation of parametric distributions on artificial data generated from the histogram.

accomodates different shapes of the predictive distribution and is flexible enough to approximate a wide range of densities. The model flexibility potentially amends misspecification associated with the parametric assumptions; it accounts for forecaster heterogeneity; and it allows for some degree of information sharing in the cross-section when making inference.

## II.B A parametric probabilistic model

We assume that the probability vector $\mathbf{z}_i$ reported by a forecaster is a noise-ridden measurements of an unobserved vector of subjective probabilities over the $J$ bins $\boldsymbol{\nu}_i = (\nu_{i1}, \ldots, \nu_{iJ})$, with $\nu_{ij} \geq 0$ and $\nu_{i1} + \ldots + \nu_{iJ} = 1$ (Boero et al., 2008b, discuss the issue of noise). If each forecaster has a subjective probability distribution $F_i(\cdot)$ over the variable being forecast ($y \in \mathcal{Y} \subset \mathbb{R}$), which they use to compute the bin probabilities $\nu_{ij}$, then

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

where $\boldsymbol{\theta}_i \in \Theta$ includes the parameters 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

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

but the general approach accommodates many other choices for $F(\cdot|\boldsymbol{\theta})$.

The uncertainty in $\mathbf{z}_i$ is encoded into a probability distribution $h(\cdot)$, that is

$$\mathbf{z}_i = (z_{i,1}, \ldots, z_{i,J}) \overset{ind}{\sim} h(\mathbf{z}_i|\boldsymbol{\theta}_i), \tag{3}$$

which captures the noise due to approximations or to actual mistakes in reporting. In choosing the random histogram distribution $h(\cdot)$, one needs to account for the fact that $\mathbf{z}_i$ belongs to the simplex; that is, the elements of $\mathbf{z}_i$ are positive and sum up to one. A convenient choice for the distribution $h(\cdot)$ is the standard Dirichlet distribution which is defined on the simplex. A drawback of this distribution is that its PDF is null for $\mathbf{z}_i$'s that have some elements equal to zero, when in fact forecasters often assign zero probability to one or more bins. To specify $h$, we follow Zadora et al. (2010) and Scealy and Welsh (2011) and use a distribution which allows for values of the random vector on the boundary of the simplex. This distribution can be described in term of the augmented representation $(\mathbf{z}; \boldsymbol{\xi}) = (z_1, \ldots, z_J; \xi_1, \ldots, \xi_J)'$ where the indicator variables $\xi_j$ is 0 if and only if $z_j = 0$ and $\xi_j = 1$ otherwise. We impose that $\xi_1 + \ldots + \xi_J < J$ to rule out the event all reported probabilities are zero. The joint distribution of $\mathbf{z} = (z_1, \ldots, z_J)$ and $\boldsymbol{\xi} = (\xi_1, \ldots, \xi_J)'$ is a zero-augmented Dirichlet distribution with probability density function

$$h(\mathbf{z}, \boldsymbol{\xi}|\boldsymbol{\theta}) = \frac{1}{c(\boldsymbol{\theta})} \prod_{j=1}^{J} \alpha_j(\boldsymbol{\theta})^{\xi_j}(1-\alpha_j(\boldsymbol{\theta}))^{1-\xi_j} h(\mathbf{z}|\boldsymbol{\theta}, \boldsymbol{\xi}), \tag{4}$$

where $\boldsymbol{\theta}$ is a parameter, $\boldsymbol{\alpha}(\boldsymbol{\theta}) = (\alpha_1(\boldsymbol{\theta}), \ldots, \alpha_J(\boldsymbol{\theta}))$ are the probabilities that a forecaster will report a zero probability on the $J$ bins, $c(\boldsymbol{\theta}) = 1 - (\alpha_1 \cdot \ldots \cdot \alpha_J(\boldsymbol{\theta}))$ is a normalizing constant, and $h(\mathbf{z}|\boldsymbol{\theta}, \boldsymbol{\xi})$ is the standard Dirichlet distribution defined on the elements of $\mathbf{z}$ that are non zero:

$$h(\mathbf{z}|\boldsymbol{\theta}, \boldsymbol{\xi}) = \frac{\Gamma\left(\sum_{j\in\mathcal{J}^*} \phi(\boldsymbol{\theta})\nu_j(\boldsymbol{\theta})\right)}{\prod_{j\in\mathcal{J}^*} \Gamma(\phi(\boldsymbol{\theta})\nu_j(\boldsymbol{\theta}))} \prod_{j\in\mathcal{J}^*} z_j^{\phi(\boldsymbol{\theta})\nu_j(\boldsymbol{\theta})-1}, \tag{5}$$

where $\mathcal{J}^* = \mathcal{J}^*(\boldsymbol{\xi}) = \{j = 1, \ldots, J; \xi_j = 0\}$ is the set indexes of the non-zeros elements of $\mathbf{z}$, $\phi(\boldsymbol{\theta})\kappa$ is the rescaled precision, with $\kappa = \sum_{j\in\mathcal{J}^*} \nu_j(\boldsymbol{\theta})$, and $\nu_j(\boldsymbol{\theta})/\kappa$ for $j \in \mathcal{J}^*$ are the renormalized $\nu(\boldsymbol{\theta})$'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.[8] The distribution used in (3) is the marginal distribution of $\mathbf{z}$ implied by expression (5),

$$h(\mathbf{z}|\boldsymbol{\theta}) = \sum_{\boldsymbol{\xi}\in\mathcal{X}} h(\mathbf{z}, \boldsymbol{\xi}|\boldsymbol{\theta})$$

where $\mathcal{X}$ is the set of all vectors with 0-1 binary entries of length $J$ which are not all zeros, i.e. $\mathcal{X} = \{\boldsymbol{\xi} = (\xi_1, \ldots, \xi_J)' \in \{0,1\}^J,\ s.t.\ \xi_1 + \ldots + \xi_J < J\}$.

The probability $\boldsymbol{\alpha}$ of reporting zero mass is modeled as $\alpha_j(\boldsymbol{\theta}) = \alpha(\nu_j(\boldsymbol{\theta}), \epsilon(\boldsymbol{\theta}))$ where the function $\alpha(\nu, \epsilon)$ is decreasing in $\nu$ such that $\alpha \to 1$ for $\nu \to 0$ and $\alpha \to 0$ for $\nu \to 1$. $\epsilon(\boldsymbol{\theta})$ measures the sensitivity of $\alpha_j(\boldsymbol{\theta})$ to $\nu$ (that is, $\alpha \to 1$ for $\epsilon \to 0$ and $\alpha \to 0$ for $\epsilon \to 0$). In

[8] Note that the conditional Dirichlet satisfies some relevant properties of the unconditional Dirichlet, that are the elements of $\mathbf{z}$ and their marginal conditional means

$$\mathbb{E}(z_j|\boldsymbol{\xi}) = \frac{\phi\nu_j}{\sum_{j\in\mathcal{J}^*} \phi\nu_j} = \frac{\nu_j}{\sum_{j\in\mathcal{J}^*} \nu_j},\ j \in \mathcal{J}^* \tag{6}$$

sum up to one, and their marginal conditional variances

$$\mathbb{V}(z_j|\boldsymbol{\xi}) = \frac{\nu_j(\kappa - \nu_j)}{\kappa^2(\phi\sum_{j\in\mathcal{J}^*} \nu_j + 1)},\ j \in \mathcal{J}^* \tag{7}$$

go to zero with $\phi \to \infty$.

practice, the probabilities $\boldsymbol{\alpha}$ are parameterized as:

$$\alpha_j(\boldsymbol{\theta}) = \int_0^{\epsilon(\boldsymbol{\theta})} \mathcal{B}e(x|\nu_j(\boldsymbol{\theta}), r)dx \tag{8}$$

$j = 1, \ldots, J$, where $\mathcal{B}e(x|m, r)$ denotes the pdf of a beta distribution $\mathcal{B}e(m, r)$ with mean $m$ and precision $r$ parameters. We assume $r$ is fixed at 100 and $\epsilon(\boldsymbol{\theta}) = \epsilon$. The parameter vector of $h(\mathbf{z}, \boldsymbol{\xi}|\boldsymbol{\theta})$ 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 "noise" around the non-zero $z_{i,j}$'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 (e.g, the $F(\cdot|\boldsymbol{\theta})$, or the $\alpha(\cdot|\boldsymbol{\theta})$), it can still be wrong. Embedding these parametric assumptions into a mode general non parametric model arguably protects us, at least to some extent, from misspecification. We describe this approach in the next section.

## II.C A Bayesian non-parametric model

The Bayesian non-parametric hierarchical model works as follows. We assume that the distribution generating the $\mathbf{z}_i$ has a respondent-specific parameter $\boldsymbol{\theta}_i$ and the parameters $\boldsymbol{\theta}_i$, $i = 1, \ldots, N$ are 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 of distributions the parameter $\boldsymbol{\theta}_i$ of $i$-th forecaster is distributed according to

$$\boldsymbol{\theta}_i \overset{iid}{\sim} \begin{cases} \boldsymbol{\theta}_1^* & \text{with probability } w_1 \\ \vdots & \\ \boldsymbol{\theta}_K^* & \text{with probability } w_K \end{cases} \tag{9}$$

with $w_k > 0$ and $w_1 + \ldots + w_K = 1$. At the second stage of the hierarchy, it is assumed that the unknown parameter types are sampled from a common distribution $\boldsymbol{\theta}_k^* \overset{iid}{\sim} G_0(\boldsymbol{\theta})$, $k = 1, \ldots, K$, and that the type probabilities have prior distribution

$$(w_1, \ldots, w_K) \sim \text{Dir}\Big(\frac{\psi_0}{K}, \ldots, \frac{\psi_0}{K}\Big), \tag{10}$$

where $\psi_0$ is a concentration parameter and $\text{Dir}(a_1, \ldots, a_K)$ a Dirichlet distribution of parameters $(a_1, \ldots, a_K)$.

Now let the number of types $K$ go to infinity. When this happens, we obtain the discrete random measure

$$G(\boldsymbol{\theta}) = \sum_{k=1}^{\infty} w_k \delta(\boldsymbol{\theta} - \boldsymbol{\theta}_k^*) \tag{11}$$

where $\delta(x)$ denotes a point mass distribution located at 0, the so-called "atoms" $\boldsymbol{\theta}_k^*$ are i.i.d. random variables from the base measure $G_0$, and the random weights $w_k$ are generated by the stick-breaking representation $SB(\psi_0)$ given by

$$w_k = v_k \prod_{l=1}^{k-1} (1 - v_l) \tag{12}$$

where the stick-breaking components $v_l$ are i.i.d. random variables from a Beta distribution $\mathcal{B}e(1, \psi_0)$ (e.g., see Pitman, 2006). Following Sethuraman (1994), 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:

$$\boldsymbol{\theta}_i \stackrel{iid}{\sim} G, \quad G \sim \mathcal{DP}(\psi_0, G_0).$$

The base measure $G_0$ has the interpretation of mean type distribution, and the precision parameter $\psi_0$ measures the concentration of $G$ around $G_0$, so that when $\psi_0 \to +\infty$ all forecasters are assumed to be of the same type and when $\psi_0 \to 0$ the inference is done forecaster by forecaster (using the same prior). See Ghosh and Ramamoorthi (2003) for an introduction to Dirichlet process priors and Hjort et al. (2010) for a review on the state-of-the-art practice of Bayesian non-parametrics.

Sethuraman (1994)'s constructive representation, in addition to being computationally convenient as discussed in the next section, implies that our model has the infinite mixture representation

$$h_G(\mathbf{z}) = \int h(\mathbf{z}|\boldsymbol{\theta}) G(d\boldsymbol{\theta}) = \sum_{k=1}^{\infty} w_k h(\mathbf{z}|\boldsymbol{\theta}_k) \tag{13}$$

where the weights $w_k$ come from the same prior distribution (12) for all forecasters. This representation indicates that the Bayesian non-parametric model is flexible and, as such, can overcome the inherent misspecification implied by the use of a specific parametric assumptions.

In conclusion, in our Bayesian non-parametric model each forecaster is described by a prior distribution over a very rich parameter space. At the same time, Bayesian non-parametrics allows for some degree of pooling: the approach allocates forecasters whose predictive distributions are similar to one another into groups and allows the number of groups to grow naturally as more data becomes available. This pooling mitigates overfitting.

## II.D Posterior inference

The Dirichlet process $G$ generates a priori dependence among the forecaster-specific parameters $\boldsymbol{\theta}_i$'s via the formation of clusters of forecasters of the same type. The relationship between $\boldsymbol{\theta}_i$ and the mixture components $\boldsymbol{\theta}^*_k$, is encoded by the auxiliary indicator variables $d_i$'s, which are equal to $k$ if $\boldsymbol{\theta}_i$ is from the $k^{th}$ mixture component, that is $\boldsymbol{\theta}_i = \boldsymbol{\theta}^*_{d_i}$.[9] The allocation variables $d_i$ ($i = 1, \ldots, n$) are then used to construct the posterior estimate of the forecaster-specific subjective probability:

$$F_i(y) \quad = \quad \mathbb{E}\left[F(y|\boldsymbol{\theta}^*_{d_i})|\mathbf{z}_{1:n}, \boldsymbol{\xi}_{1:n}\right] = \mathbb{E}\left[\sum_{k=1}^{\infty} F(y|\boldsymbol{\theta}^*_k)\mathbb{I}\{d_i = k\}\middle|\mathbf{z}_{1:n}, \boldsymbol{\xi}_{1:n}\right]. \tag{14}$$

Monte Carlo sampling can be used to approximate the posterior distribution and the quantities of interest such as subjective probabilities, point estimates and posterior credible intervals. Building on Walker (2007) and Kalli et al. (2011), we use a slice sampling algorithm which generates random draws from the posterior distribution of $\boldsymbol{\theta}_i$ for $i = 1, \ldots, n$ (Appendix B provides the details of the Gibbs sampler). The output of the Gibbs sampler can be used to approximate Eq. 14 as follows

$$\hat{F}_i(y) \quad = \quad \frac{1}{M}\sum_{m=1}^{M}\sum_{k=1}^{\infty} F(y|\boldsymbol{\theta}^{*(m)}_k)\mathbb{I}\{d^{(m)}_i = k\} \tag{15}$$

where $d^{(m)}_i$, $\boldsymbol{\theta}^{*(m)}_k$, $m = 1, \ldots, M$ are the MCMC samples for the infinite mixture atoms and allocation variables.

## II.E Posterior consistency

In this section, we also discuss asymptotic properties of the posterior distribution as the number of forecasters goes to infinity.

[9] As shown in Escobar and West (1995), the predictive distribution of $\boldsymbol{\theta}_{n+1}$ conditional on $(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_n)$ can be represented as a Polya's urn process

$$\boldsymbol{\theta}_{n+1}|\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_n \sim \frac{\psi_0}{\psi_0 + n}G_0(\boldsymbol{\theta}_{n+1}) + \frac{1}{\psi_0 + n}\sum_{i=1}^{n}\delta(\boldsymbol{\theta}_i - \boldsymbol{\theta}_{n+1}).$$

With probability $\frac{\psi_0}{\psi_0 + n}$ the new draw $\boldsymbol{\theta}_{n+1}$ is generated from $G_0$, but it is otherwise equal to one of the previous $n$ draws. In fact, the $n$ forecasters' distribution can be characterized using $N$ different clusters, where $N$ is a random variable with prior mean $\mathbb{E}[N] \approx \psi_0 log(\frac{\psi_0 + n}{\psi_0})$. When $\psi_0 \to \infty$ 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 $G_0$.

We formalize convergence with respect the number of observations via weak consistency of the posterior distribution (see, e.g. Ghosh and Ramamoorthi, 2003), which provides a widely accepted minimal requirement for large sample behavior of Bayesian non-parametric models (see, 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 $\mathcal{H}$ is the set of all possible data generating densities (with respect to a dominating measure) on the sample space $\mathcal{Z} \subset \mathbb{R}^J$. Given a prior $\Pi$ on $\mathcal{H}$, the posterior is said to be *weakly consistent* at $h_0$ if for every i.i.d. sequence $\mathbf{z}_1, \mathbf{z}_2, \ldots$ of random variables with common density $h_0$ the posterior probability $\Pi(U|\mathbf{z}_1, \ldots, \mathbf{z}_n)$ converges a.s. to 1 as $n \to +\infty$ for every weak neighbourhood $U$ of $h_0$. For some background material on posterior consistency, see e.g. Ghosh and Ramamoorthi (2003).

In our setting, the prior $\Pi$ on $\mathcal{H}$ is a type I mixture prior (see Wu and Ghosal (2009a)) induced by the map $G \mapsto h_G(\mathbf{z}) = \int_\Theta h(\mathbf{z}|\boldsymbol{\theta}) G(d\boldsymbol{\theta})$ where $\Theta$ is the mixing parameter space, and $G$ has a Dirichlet process prior $\mathcal{DP}(\psi_0, G_0)$, where $G_0$ is a base measure on $\Theta$ and $\psi_0$ the concentration parameter.

To prove weak consistency for our model, we use Schwartz theorem (see e.g. Chapter 4 in Ghosh and Ramamoorthi (2003)). This result states that weak consistency at a "true density" $h_0$ holds if the prior assigns positive probabilities to Kullback-Leibler neighborhoods of $h_0$. We state here only one main result on consistency, all the details, proofs and some additional results are available in Section D of the Appendix.

If the probability of zero bins is positive, $\mathcal{H}$ includes also distribution $h$ with assigns positive mass to sub-symplex of lower dimension of the $J$-dimensional simplex $\Delta^J = \{(z_1, \ldots, z_{J-1}) : z_1 + \cdots + z_{J-1} \leq 1,\ 0 < z_j < 1\}$. In this case $\mathcal{Z}$ can be seen as the augmented space of the possible values of $(\mathbf{z}, \boldsymbol{\xi})$ and we need to properly generalize the definition of Kullback-Leibler divergence to mixed densities. Let $\mathbf{z}_{\boldsymbol{\xi}} = [z_j : j \in \mathcal{J}^*(\boldsymbol{\xi})]$ and observe that, given $\boldsymbol{\xi}$, $\mathbf{z}_{\boldsymbol{\xi}}$ takes values in the open $J - |\boldsymbol{\xi}|$-dimensional simplex $\Delta^{J-|\boldsymbol{\xi}|}$ where $|\boldsymbol{\xi}| = \xi_1 + \cdots + \xi_J$. On the sample space $\mathcal{Z}$, given by all the pairs $(\boldsymbol{\xi}, \mathbf{z})$, one can thus define a $\sigma$-finte measure $\lambda(d\boldsymbol{\xi} d\mathbf{z}) = c(d\boldsymbol{\xi}) \otimes \mathcal{L}_{\boldsymbol{\xi}}(d\mathbf{z}_{\boldsymbol{\xi}})$ where $c$ is the counting measure on $\mathcal{X}$ and, given $\boldsymbol{\xi}$, $\mathcal{L}_{\boldsymbol{\xi}}$ is the Lebesgue measure on $\Delta^{J-|\boldsymbol{\xi}|}$.

The set of all possible data generating densities $\mathcal{H}$ is the set of all the densities with respect to $\lambda$. These densities $g$ factorize as $g(\boldsymbol{\xi}, \mathbf{z}) = g(\boldsymbol{\xi})g(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\xi})$, where $\mathbf{z}_{\boldsymbol{\xi}} = [z_j : j \in \mathcal{J}^*(\boldsymbol{\xi})]$.

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

$$KL(h_0, g) = \int_{\mathcal{Z}} h_0(\boldsymbol{\xi}, \mathbf{z}) \log\Big(\frac{h_0(\boldsymbol{\xi}, \mathbf{z})}{g(\boldsymbol{\xi}, \mathbf{z})}\Big)\lambda(d\boldsymbol{\xi} d\mathbf{z}). \tag{16}$$

Some more details on $KL(h_0, g)$ are given in Section D of the Appendix.

We define $\mathcal{M}^*$ as the set of finite mixtures of densities (4), and denote by $\mathcal{H}_0^*$ the set of densities that can be approximated in the Kullback-Leibler sense by densities in $\mathcal{M}^*$, i.e.

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

**Theorem 1.** Assume that $\boldsymbol{\theta} \mapsto (\alpha_1(\boldsymbol{\theta}), \ldots, \alpha_J(\boldsymbol{\theta}), \phi(\boldsymbol{\theta})\nu_1(\boldsymbol{\theta}), \ldots, \phi(\boldsymbol{\theta})\nu_J(\boldsymbol{\theta}))$ is a continuous function such that $\nu_j(\boldsymbol{\theta}) > 0$ and $0 < \alpha_j(\boldsymbol{\theta}) < 1$ for every $j = 1, \ldots, J$. If $G_0$ has full support, then the posterior is weakly consistent at any density $h_0$ in $\mathcal{H}_0^*$ such that

$$\sum_{\boldsymbol{\xi} \in \mathcal{X}} h_0(\boldsymbol{\xi}) \int_{\Delta^{J-|\boldsymbol{\xi}|}} \left| \log\left( \prod_{j \in \mathcal{J}^*(\boldsymbol{\xi})} z_j \right) \right| h_0(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\xi}) d\mathbf{z}_{\boldsymbol{\xi}} < +\infty. \tag{17}$$

In the result given above, the number of bins is finite. Additional asymptotic results are obtained when the number of bins $J$ goes to infinity, the bin size goes to zero and the rounding disappears. First, for $n$ fixed, we proved that the random histogram model $\mathbf{z}_i$ converges to an infinite dimensional model where each forecaster's response is modelled by a (random) CDF with mean $F(\cdot|\boldsymbol{\theta}_i)$. This infinite dimensional prior model gives positive probability to any (weak) neighbourhood of any distribution defined on the support set of $F(\cdot|\boldsymbol{\theta}_i)$. See Theorem 2 and Corollary 1 in Appendix C. Moreover, under suitable assumptions, as the number of forecasters and the number of bins go to infinity the posterior consensus distribution converges to the true consensus CDF of the forecasters, see Proposition 5 in Appendix C.

In the applications, the choice of the probabilistic model, in particular of the distribution family $F(\cdot|\boldsymbol{\theta}_i)$, and of the prior distribution can have an impact on the results given that $n$ is far from infinity (around 30 for the SPF) and the number of bins $J$ is small (e.g., $J = 10$ in the US SPF on GDP in 2020). When the number of bins decreases and/or the bin width increases, the amount of information available to reconstruct the subjective CDF diminishes, and model assumptions can have a more significant impact on the empirical results. An advantage of the Bayesian approach is that it accounts for the lack of information returning wider credible intervals. In general, the approach provides a measure of the level of estimation uncertainty for all objects of interest. Nevertheless, a robustness check with respect to the specification of the prior distribution and the distribution family should be considered in all applications of this method.

## II.F Prior parameters specification

In the following we discuss the prior setting used in the real data application. The parametric family chosen for the subjective CDF is $F(y|\boldsymbol{\theta}) = (1-\omega)\Phi(y|\mu, \sigma_1^2) + \omega\Phi(y|\mu+\mu_\delta, \sigma_2^2)$ where $\boldsymbol{\theta} = (\mu, \mu_\delta, \sigma_1, \sigma_2, \omega, \phi, \epsilon)$. The parameters $\epsilon$ and $\phi$ are used to specify $h$ and $\alpha$ in (5) and (8), respectively.

The base measure $G_0$ of the DPP is given by the product of the following distributions. The location of the first mixture component is $\mu \sim \mathcal{N}(2, 5^2)$. Note that the standard deviation is 5, so this is a very loose prior. The scales of the mixture components follow $\sigma_j \sim \mathcal{IG}a(a_\sigma, b_\sigma)\mathcal{I}(\sigma_1)_{(0,10)}$, $j = 1, 2$ where $a_\sigma, b_\sigma$ are chosen s.t. the standard deviation has mean $E[\sigma_1] = 2$ and a variance $V[\sigma_1] = 4$. We truncate the distribution at 10 for numerical reasons. The parameters $\mu_\delta$ 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 $\omega$, the weight on the second component of the mixture, is $\omega \sim \mathcal{B}(0.5, 3)$. Its mode is zero, implying that the prior favors models with one mixture only. The prior places roughly 20 percent probability on $\{\omega \geq 0.25\}$.

As regards the prior for the rounding-off parameter $\epsilon$ we assume a $\mathcal{G}a(a_\epsilon, b_\epsilon)$ and set $a_\epsilon$, $b_\epsilon$ such that $\alpha_j$ is close to one for $\nu_j$ less than 1%, very small for any $\nu_j$ larger than 5%, and virtually zero when $\nu_j$ is larger than 10%.[10]

For the precision parameter of the random histogram, $\phi$, we assume a $\mathcal{G}a(a_\phi, b_\phi)$ where $a_\phi$ and $b_\phi$ are s.t. $E[\phi] = 100$, $V[\phi] = 100$. The left panel of Figure 2 shows the 50 and 90 percent bands for the noise associated with three different values of $\nu$: 0.1, 0.6 and 0.3. The right panel of Figure 2 shows the mean and the 90% coverage intervals of $\alpha_j(\boldsymbol{\theta})$ as a function of $\nu_j(\boldsymbol{\theta})$. The probability of reporting zero becomes non negligible only for $\nu < 0.04\%$.

[10]We 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 $\mathcal{B}e(x|\nu, r) = \frac{1}{B(\nu r, (1-\nu)r)} x^{\nu r - 1}(1-x)^{(1-\nu)r-1}$ with $x \in (0, 1)$, $m \in (0, 1)$ and precision $r > 0$.

Figure 2: Noise and zero probability

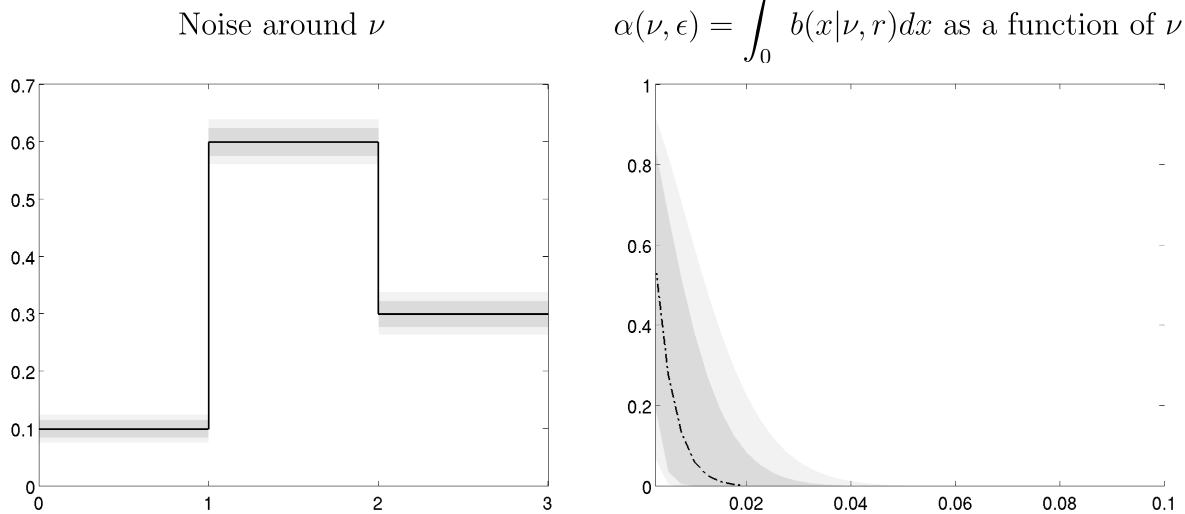


As regards the concentration hyperparameter $\psi_0$ of the Bayesian non-parametric prior, which determines the prior number of clusters, we follow the standard choice and set $\psi_0 = 1$. This implies that the expected number of clusters for a cross-section of 30 survey respondents is roughly 4.

# III Results

In this section we first discuss the application of the non-parametric 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 next 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

In this section we provide posterior estimates of the subjective predictive distributions $F(y|\theta_i)$ obtained with our approach for the examples discussed in section II.A, and compare these

Figure 3: Inference Using Bayesian Non-parametric Approach: CDFs and PDFs for Selected Examples

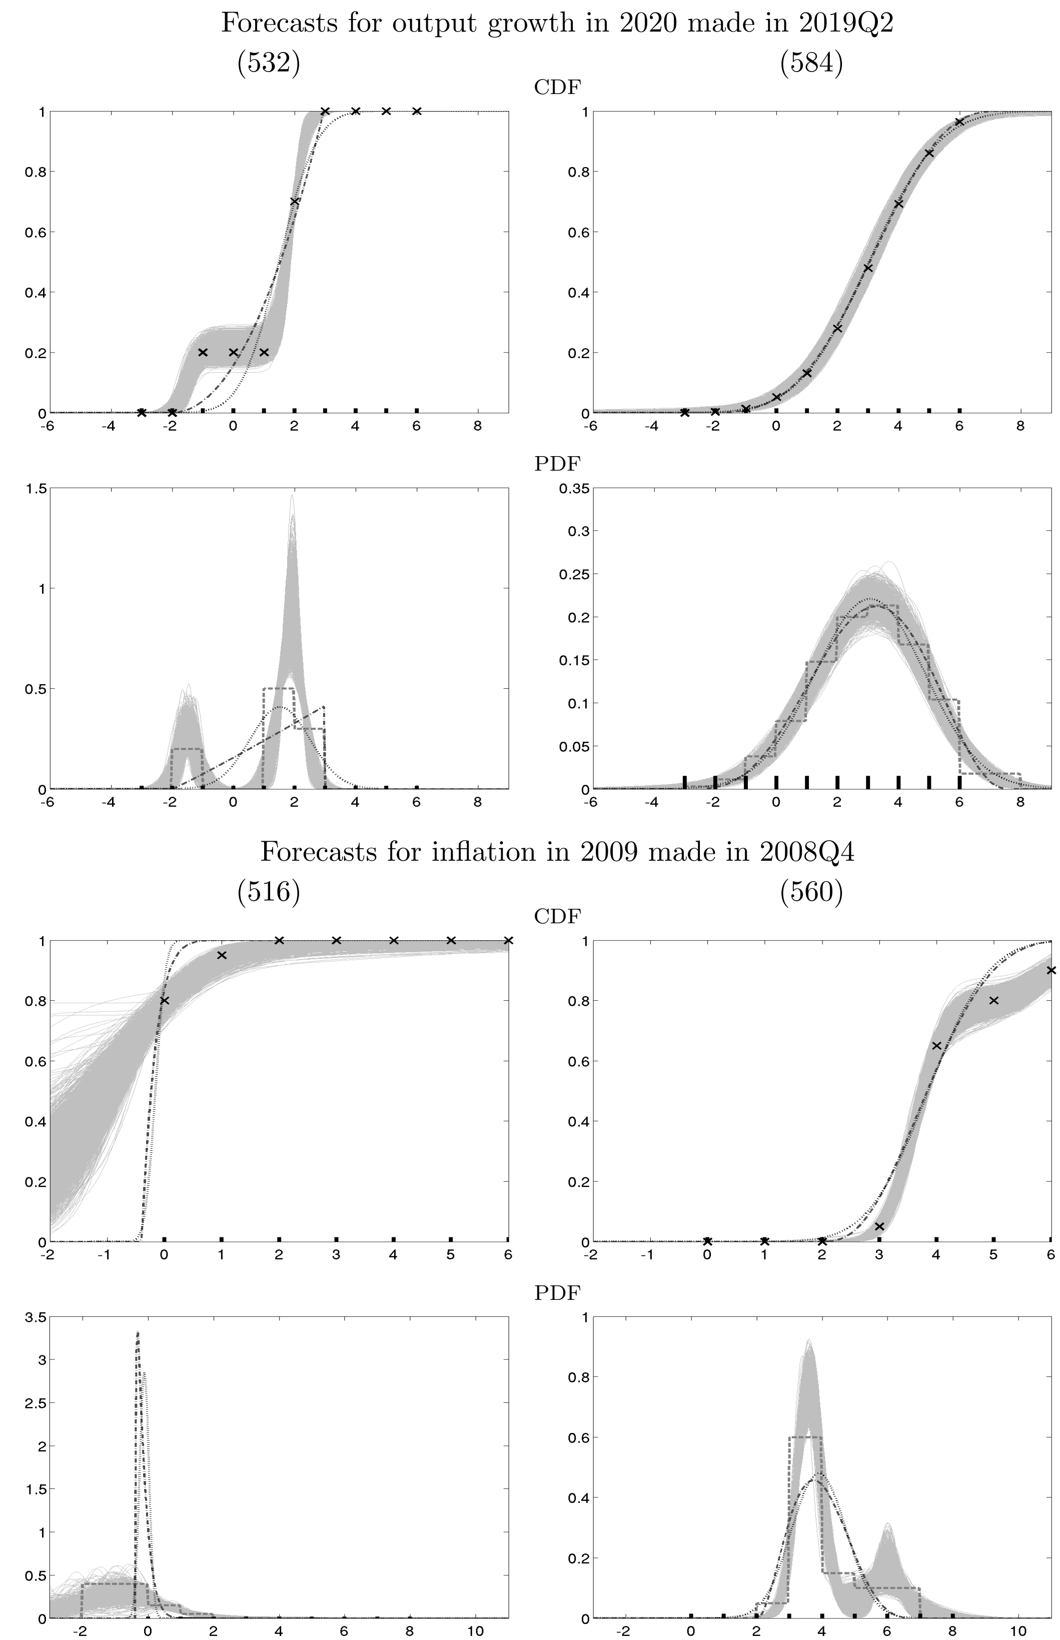


*Note*: In each panel: the subjective CDFs (top panels) and selected quantiles (bottom panels). Top panels: subjective CDF using least-squares approach with normal (gray, dashed line) or beta (black, dash-and-dotted line) assumption; subjective CDF using BNP approach (posterior random draws in light gray); and observed cumulated histogram probabilities $Z_{ij}$ $j = 1, \ldots, J$ (crosses).

estimates with the results obtained under some of the approaches currently used in the literature.

Figure 3 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 (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 dotted lines) obtained from step-wise uniform PDF (gray dashed lines) implied by the histogram probabilities $z_{ij}$.

Figure 3 is helpful in illustrating a few points about the Bayesian non-parametric approach. First, the observed cumulative probabilities (the $Z_{ij}$'s; crosses) belongs to the high posterior density region for all these respondents. This suggest that our approach is flexible enough to capture a variety of arguably challenging cases. Bassetti et al. (forthcoming) provide several other examples obtained during the recent Covid episode which confirms this impression. In contrast, the beta and the normal approaches do not fit the $Z_{ij}$'s well in these examples, with the exception of respondent (584), and their CDFs and PDFs do not belong to the high posterior density region obtained from the BNP approach. This implies that there can be important differences in the objects of interests, such as the measure of uncertainty, or quantiles, implied by the different approaches. Bassetti et al. (forthcoming) again discuss some of these differences during the Covid period.

Figure 3 also shows that whenever there is less information from the respondent, the BNP approach delivers wider posterior coverage intervals that reflect this higher degree of uncertainty. The case of respondent 516 is exemplary. This respondent places 80 percent probability on the left open bin (see Figure 1), implying that we know very little about the left-tail behavior of this forecaster. The posterior coverage intervals for both the BNP CDF and PDF reflect this uncertainty, 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

Heterogeneity in macroeconomic probabilistic forecasts was noted a long time ago. While much of the early literature focused on disagreement in point projections or central ten-

Figure 4: Subjective uncertainty by individual respondent–Output Growth

H1Q2

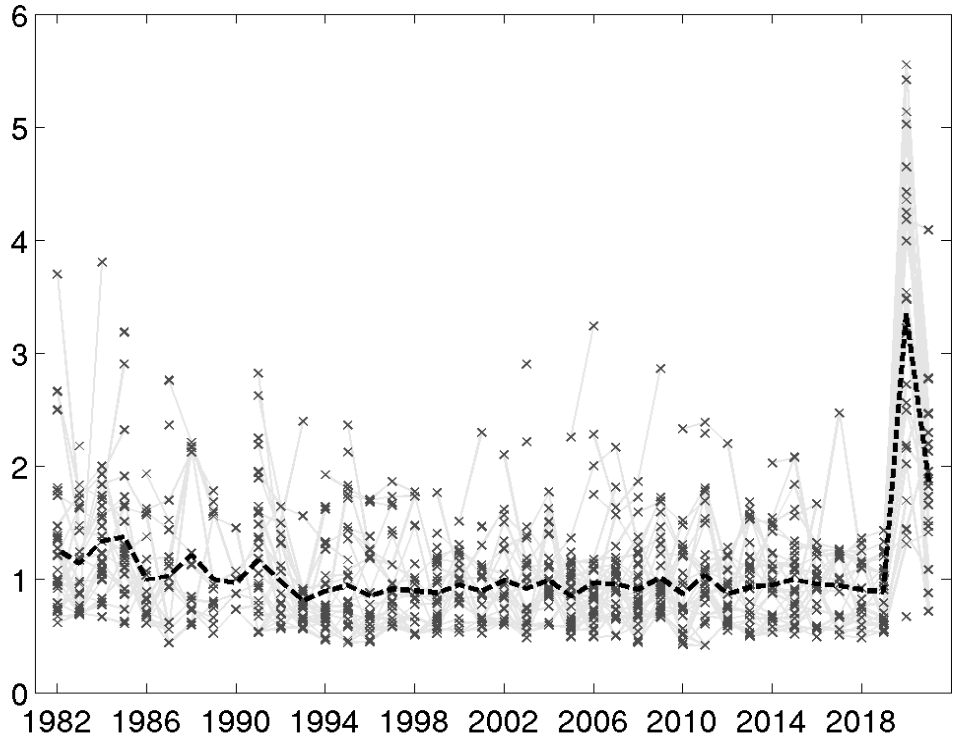


H2Q2

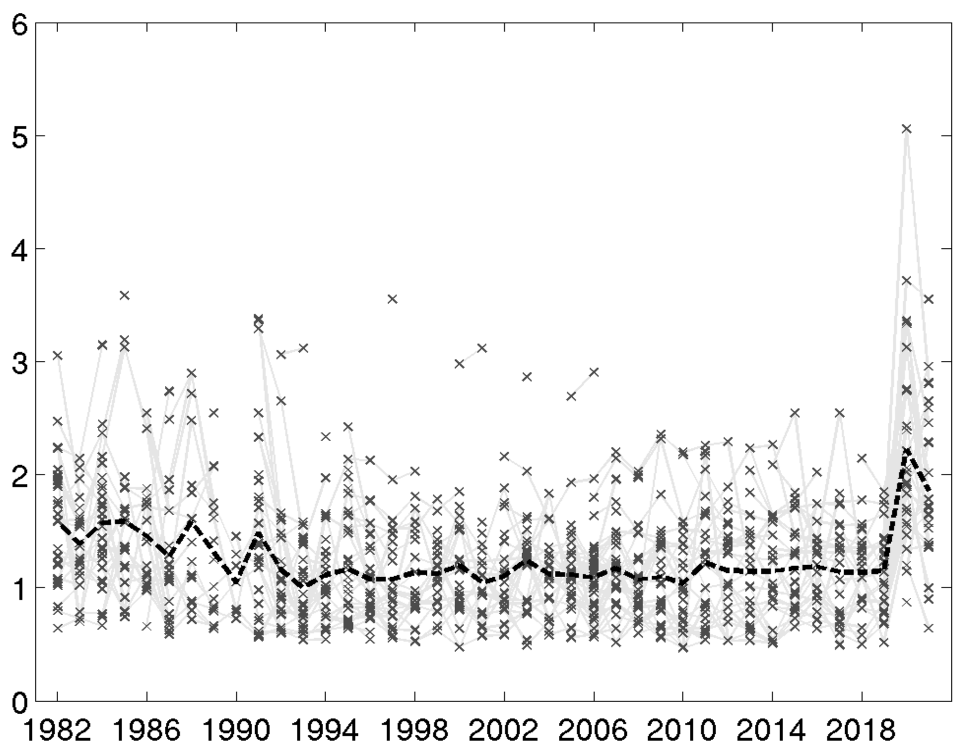


*Note*:

Figure 5: Subjective uncertainty by individual respondent–Inflation

H1Q2

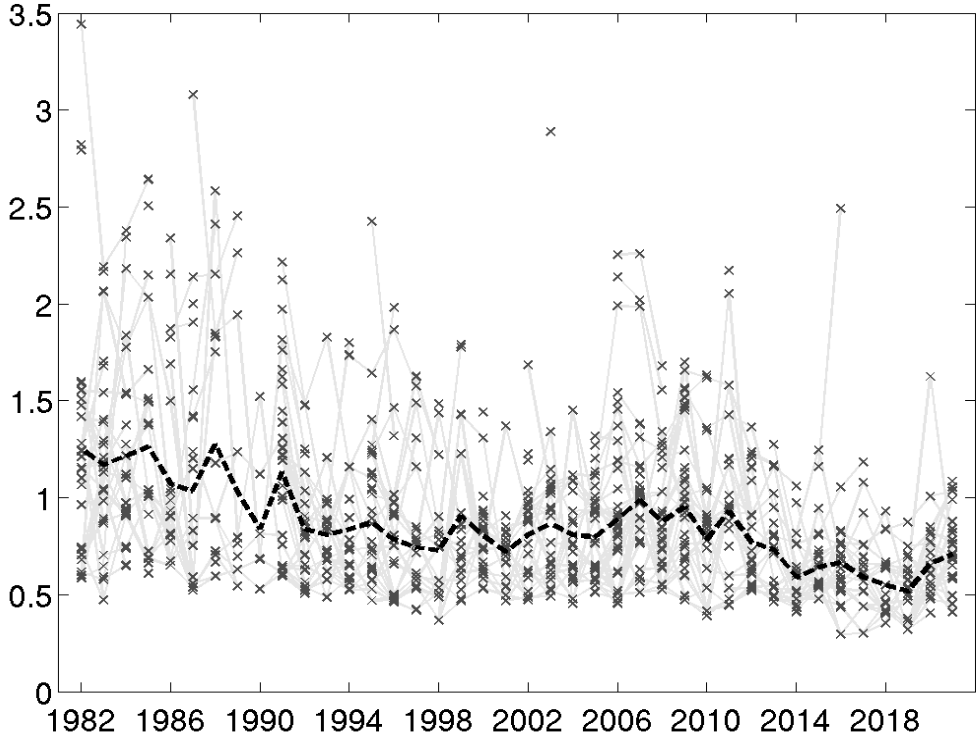


H2Q2

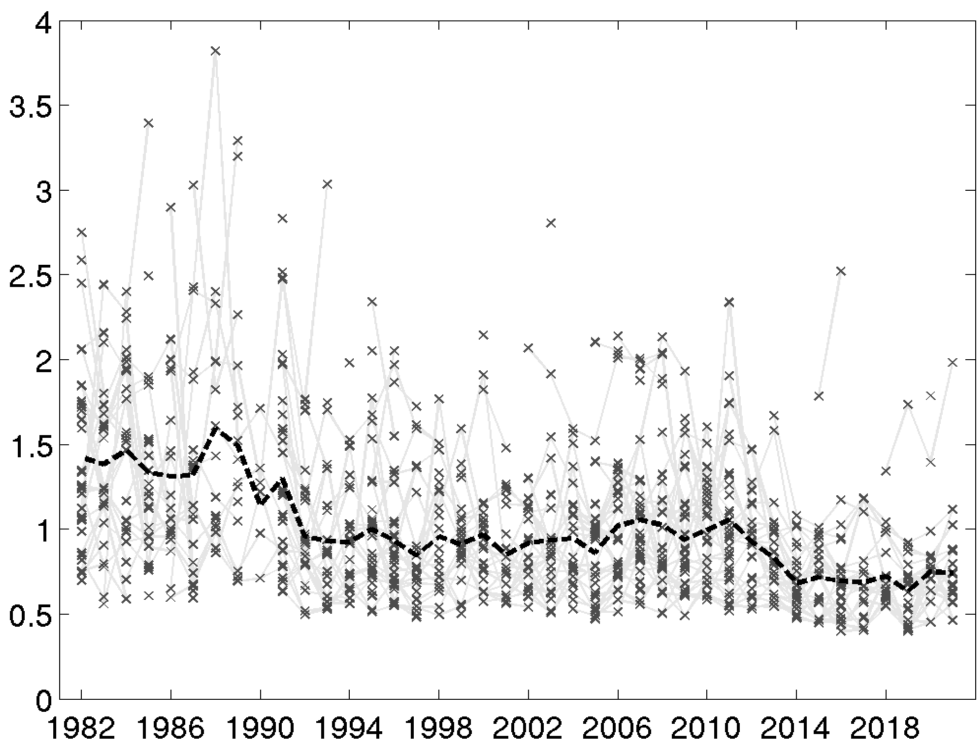


*Note*:

Figure 6: Cross-sectional standard deviation of individual uncertainty

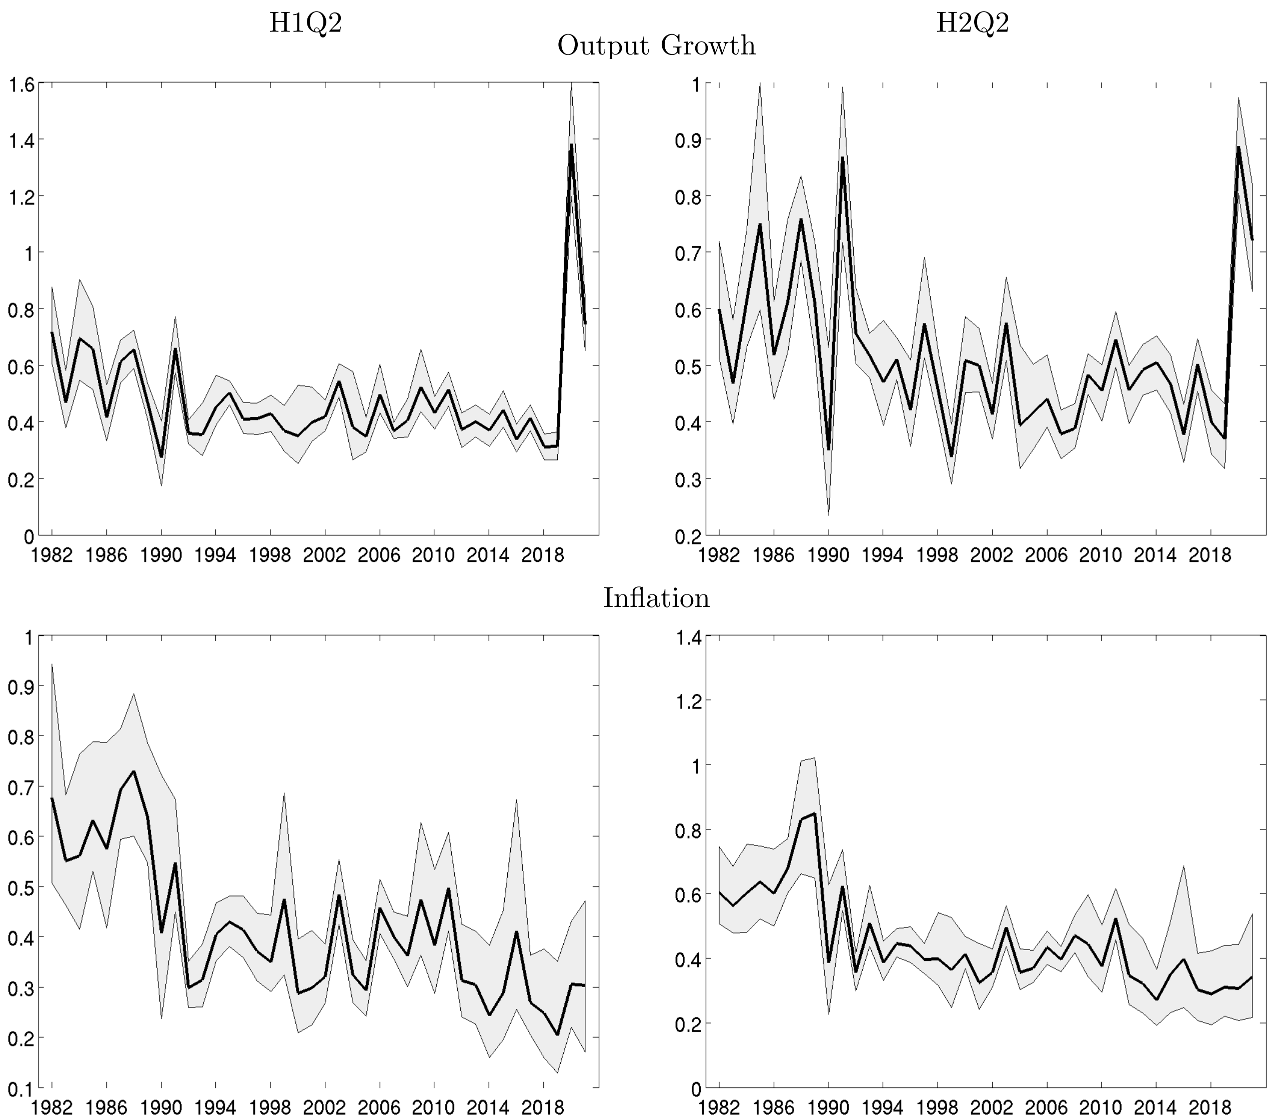


*Note*:

dencies,[11] more recent work documents the fact that forecasters disagree about uncertainty. Lahiri and Liu (2006) and D'Amico and Orphanides (2008) are to our knowledge some of the first papers to highlight heterogeneity in individual uncertainty, and do so in the context of SPF predictions for inflation.[12] A number of articles also provide evidence of persistent differences in subjective uncertainty. Bruine De Bruin et al. (2011) for instance measure individual uncertainty for consumers' density forecasts using the interquartile range obtained by fitting a beta distribution to each forecaster and find that heterogeneity in perceived uncertainty is significant and persistent, as it appears to be associated with demographic characteristics and financial literacy. Boero et al. (2014) and Rich and Tracy (2021), using the Bank of England Survey of External Forecasters and the European SPF, respectively, also find that relative differences in uncertainty are long lasting, and interpret this fact as suggesting that the degree of uncertainty is a forecaster-specific characteristic akin to the individual optimism and pessimism established in the literature on point forecasts.[13]

In this section we document the evolution of individual measures of uncertainty obtained using our approach in our 1982-2021 sample. We do 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 aforementioned literature and show that professional forecasters differ significantly in terms assessment of uncertainty,

[11]See Mankiw et al., 2003; Carroll, 2003; Capistrán and Timmermann, 2009; Patton and Timmermann, 2010, 2011; Andrade and Le Bihan, 2013; Andrade et al., 2016 and other work mentioned in the excellent recent survey by Clements et al. (forthcoming). Following Zarnowitz and Lambros (1987), a large literature has also investigated the question of whether disagreement and the average dispersion of density forecasts (average uncertainty) move together (see Giordani and Soderlind, 2003; Lahiri et al., 1988; Rich and Tracy, 2010; Lahiri and Sheng, 2010; Abel et al., 2016, and Rich and Tracy, 2021, among many others; Kozeniauskas et al., 2018 provide a clear discussion of the conceptual differences between macroeconomic uncertainty and disagreement using a model where forecasters have private information and update their beliefs using Bayes' law).

[12]Lahiri and Liu (2006) plot the evolution over time of the distribution of individual measures of uncertainty (which they obtain by fitting a Gaussian CDF to for each forecaster's histogram), and show that the persistence in uncertainty is much less than what the aggregate time series data would suggest. D'Amico and Orphanides (2008) fit a Gamma distribution to the cross-sectional CDF of individual variances, which they also obtain under the Normal parametric assumption, and then use the variance of this Gamma to measure disagreement about uncertainty and its evolution over time.

[13]Rich and Tracy (2021) propose the Wasserstein distance as a way of directly measuring heterogeneity in predictive densities, and in computing this distance assume that individual PDFs are step-wise Uniform distributions in computing the distance. Relatedly, Clements (2014b) and Manzan (2021) discuss the updating of density forecasts and in particular uncertainty in light of new information.

and that these differences vary over time. 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.

Figures 4 and 5 show the evolution of subjective uncertainty by individual respondent for output growth and inflation, respectively. The top and bottom panels display uncertainty for the current and the next year projections, respectively, made in the second quarter (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 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). 6 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 4 shows that on average uncertainty for current year (top panel) 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 (bottom panel) 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 interquantile range to measure uncertainty, as done in Figure D-9 in the Appendix, leads to very similar results. Using the generalized beta approach to fit histograms (see Figure D-10) 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 6, 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. Looking at Figure 4 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. Relatedly, while differences in subjective uncertainty are persistent, forecasters do change their mind from period to period about their subjective uncertainty, and their relative ranking varies as indicated by the fact that the thing gray lines very often cross one another.

Figure 5 shows that on average subjective uncertainty for inflation in both for current (top) and following (bottom) 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 2021 and reached a lower plateau since the mid 2010s. 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 Bassetti et al. (forthcoming). In the case of inflation changes in the bins, which took place in 1985, 1992, and 2014 (see Figure D-1 in the Appendix), may have payed 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 earlier when survey design had not yet changed.

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 6) 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 somewhat tightly estimated, and these level shifts, especially that from the 1980s to the mid-1990s, 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 using density forecasts

Is there a correspondence between forecast errors and subjective uncertainty? The answer to this question 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. It is also of interest because it represents a test of the noisy rational expectations hypothesis (see Coibion and Gorodnichenko, 2012, 2015, for instance). According to this model, forecasters receive both public and private signals about the state of the economy. In the cross-section, the quality of forecasters' private signals ought to be reflected in equal measure in their density forecasts and in their *ex-post* forecast accuracy. Similarly, in the time series changes in the precision of either public or private signals, due for instance to variations in macroeconomic uncertainty, should be equally reflected in changes in both subjective confidence and forecast errors.

### A Scale Test

If survey respondent $i$ is forecasting the variable of interest $y_t$ at time $t-q$, their subjective uncertainty is defined by

$$\sigma^2_{t|t-q,i} = E_{t-q,i}[(y_t - E_{t-q,i}[y_t])^2], \tag{18}$$

where $E_{t-q,i}[.]$ denotes expectations taken using $i$'s predictive distribution. If we construct the random variable

$$\eta_{i,t,t-q} = (y_t - E_{t-q,i}[y_t])^2/\sigma^2_{t|t-q,i}, \tag{19}$$

under rational expectations (that is, if the data generating process for $y_t$ is consistent with the predictive distribution) its unconditional expectation has to be equal to one, that is,

$$E[(y_t - E_{t-q,i}[y_t])^2/\sigma^2_{t|t-q,i}] = 1. \tag{20}$$

Figure 7: Subjective Uncertainty and Forecast Accuracy: A Scale Test

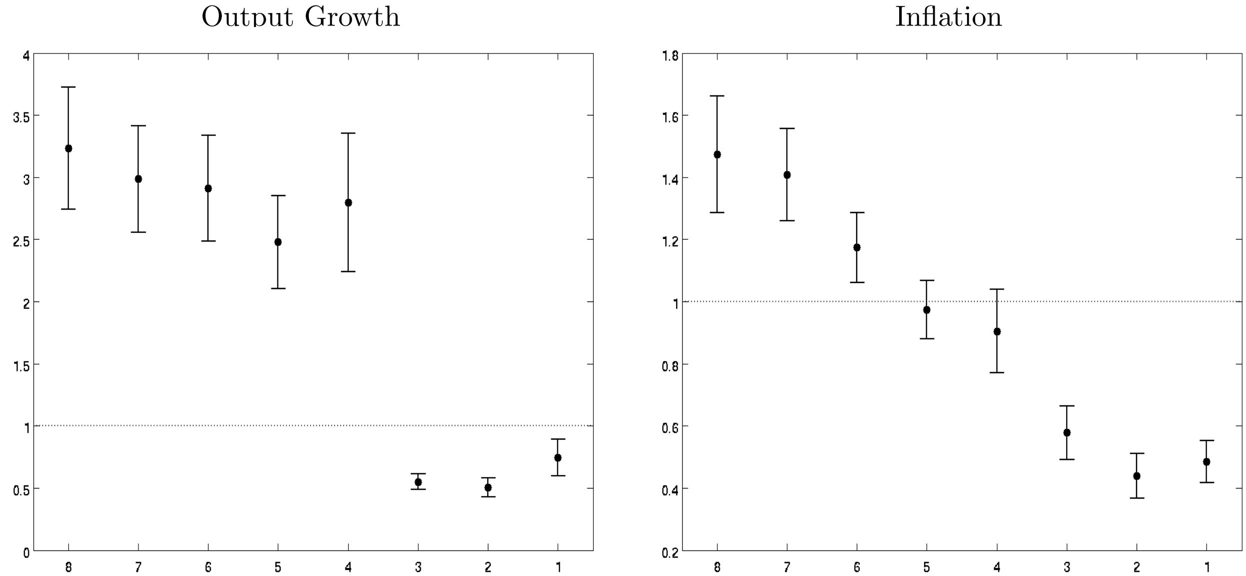


We can assess this hypothesis by testing whether $\alpha_q = 1$ in the panel regression

$$(y_t - E_{t-q,i}[y_t])^2/\sigma^2_{t|t-q,i} = \alpha_q + \epsilon_{t,i,q},\ t = 1,..,T,\ i = 1,..,N. \tag{21}$$

where we use the posterior means of $E_{t-q,i}[y_t]$ and $\sigma^2_{t|t-q,i}$ from our approach. Estimates of $\alpha_q$ that are significantly greater/lower than 1 indicate that forecasters under/over estimate uncertainty, and we hence refer to this as a "scale test," meaning that it is an assessment of whether density forecasts are appropriately scaled.

Figure 7 shows estimates of $\alpha_q$ 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 crosses indicate the OLS point estimates and the whiskers the two-standard deviations posterior intervals, which are robust to heterogeneity (Müller, 2013). The figure shows that for horizons between two and one-and-half years (e.g., $q = 6, 7$ or 8) $\alpha_q$ is significantly larger than 1 for both output growth and inflation. In fact, for output growth $\alpha_q$ 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 ($q = 5, 4$) $\alpha_q$ remains well above 1 for output, but is not significantly different from 1 for inflation. For shorter horizons ($q$ lower than 3) $\alpha_q$ 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 $\alpha_q$ is barely significantly below 1. Figure D-12 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).

Figure 8: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy?

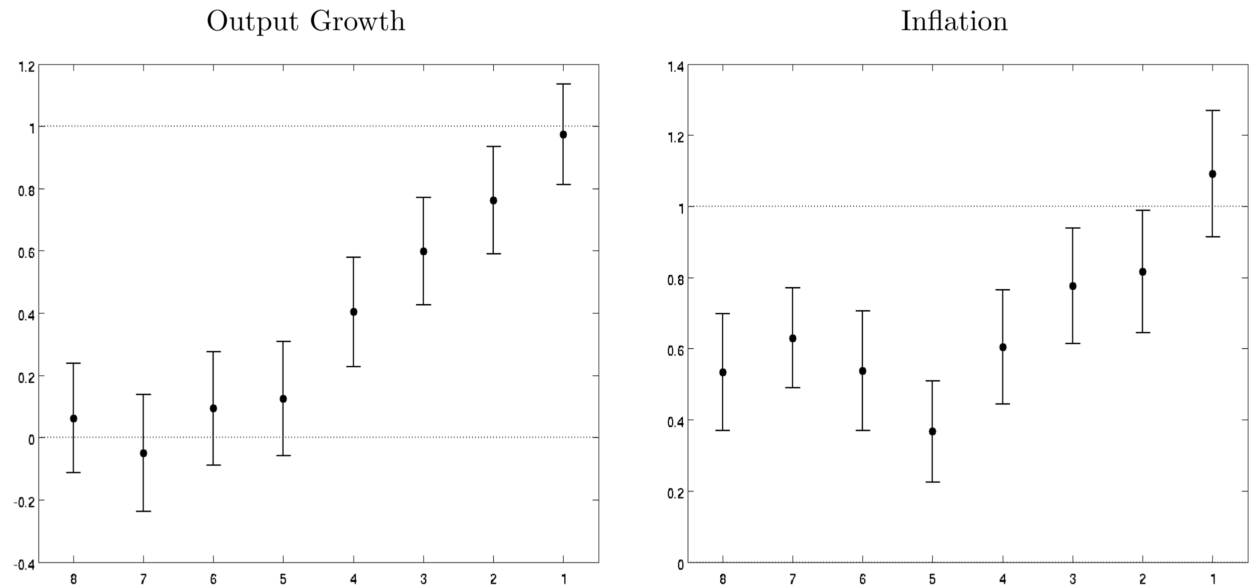


The idea behind the regression in (21) borrows heavily from existing literature. Clements (2014a) in particular computes values for $\eta_{i,t,t-q}$ using the point predictions in place of the mean $E_{t-q,i}[y_t]$, and estimates of $\sigma^2_{t|t-q,i}$ obtained from fitting a generalized Beta distribution. Clements then computes $\alpha_{iq}$ 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. Clements also plots time series averages of $\sigma_{t|t-q,i}$ against each forecaster's root mean square error (again, computed using the point forecasts). Both exercises are conducted for US SPF surveys for output growth and inflation from 1981Q3 to 2010Q4. Casey (2021) applies Clements (2014a)'s approach to Euro area, UK, and US SPF, using a sample from 1999 to 2015. The gist of Clements (2014a)'s and Casey (2021)'s findings are broadly in line with those reported above: at longer horizons forecasters generally tend to be overconfident, and this overconfidence diminishes, or becomes underconfidence, as the horizon gets shorter.

The benefit from running a panel regression as in (21) is twofold. First, we explicitly

Figure 9: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy?

Baseline vs Weighted

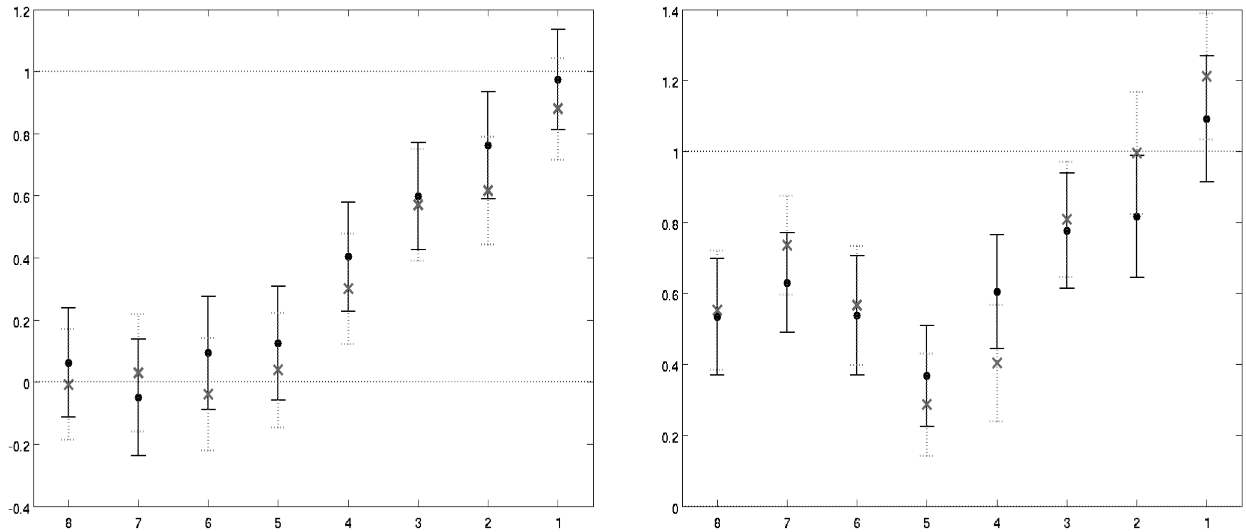


Baseline vs Beta

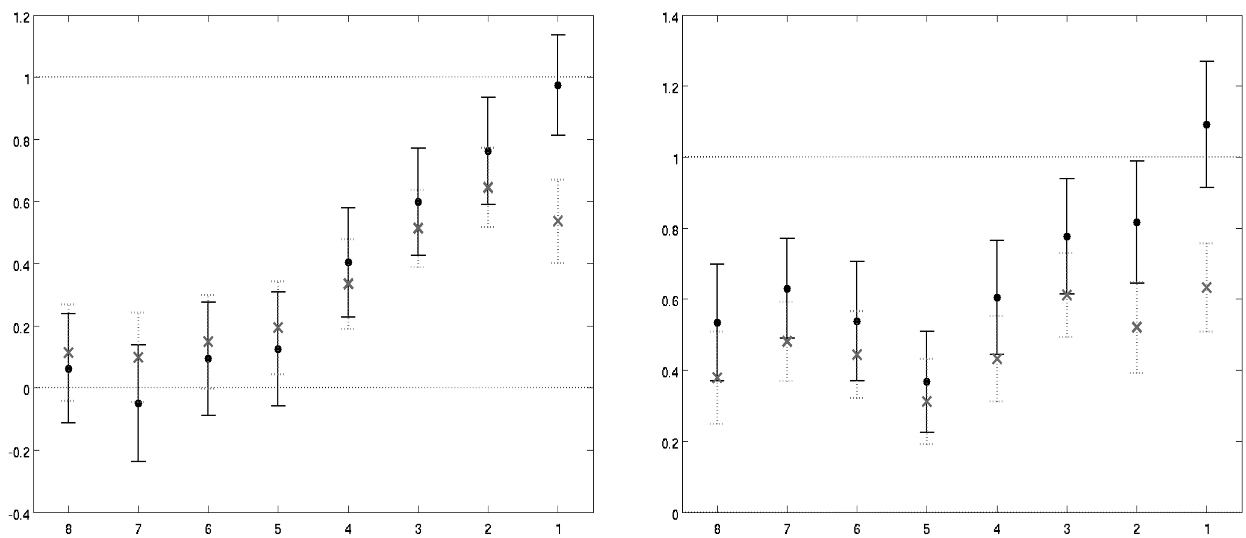


test whether predictive distributions are correctly scaled using the entire panel, rather than forecaster by forecaster, thereby getting a clear answer on whether the rational expectation hypothesis is rejected or not for the SPF. Second, we obtain quantitative estimates of the average degree of over/under-confidence 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 as confident as they should be, for instance, was not known to our knowledge. And so is the finding that at horizon of about one year we cannot reject the hypothesis that $\alpha_q = 1$ for inflation forecasts. Also, previous literature mostly used point forecasts, while of course under rational expectations equation (20) holds for the mean, but not necessarily for the point forecast if this differs from the mean (Figure D-11 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).

**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. Taking logs of both sides of equation (19) and dividing by 2 we obtain:

$$\ln |y_t - E_{t-q,i}[y_t]| - \ln \sigma_{t|t-q,i} = \frac{1}{2} \ln \eta_{i,t,t-q}, \tag{22}$$

implying that in the panel regression

$$\ln |y_t - E_{t-q,i}[y_t]| = \beta_{0,q} + \beta_{1,q} \ln \sigma_{t|t-q,i} + \epsilon_{t,i,q}, \;\; t = 1, .., T, \;\; i = 1, .., N. \tag{23}$$

the coefficient $\beta_{1,q}$ ought to be equal to 1 under rational expectations. Equation (23) is estimated via OLS where the standard deviation $\sigma_{t|t-q,i}$ is measured using the posterior mean of the standard deviation estimated using our approach, and robust standard errors are computed. Figure 8 plots the point estimates of $\beta_{1,q}$ for different horizons (crosses) and the whiskers denote the two-standard deviations posterior intervals.

It is striking that for output growth there is no significant 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 10 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 8 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 forecasters with lower subjective uncertainty have lower absolute forecast errors on average, even for short

Figure 10: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Regressions with Fixed Effects

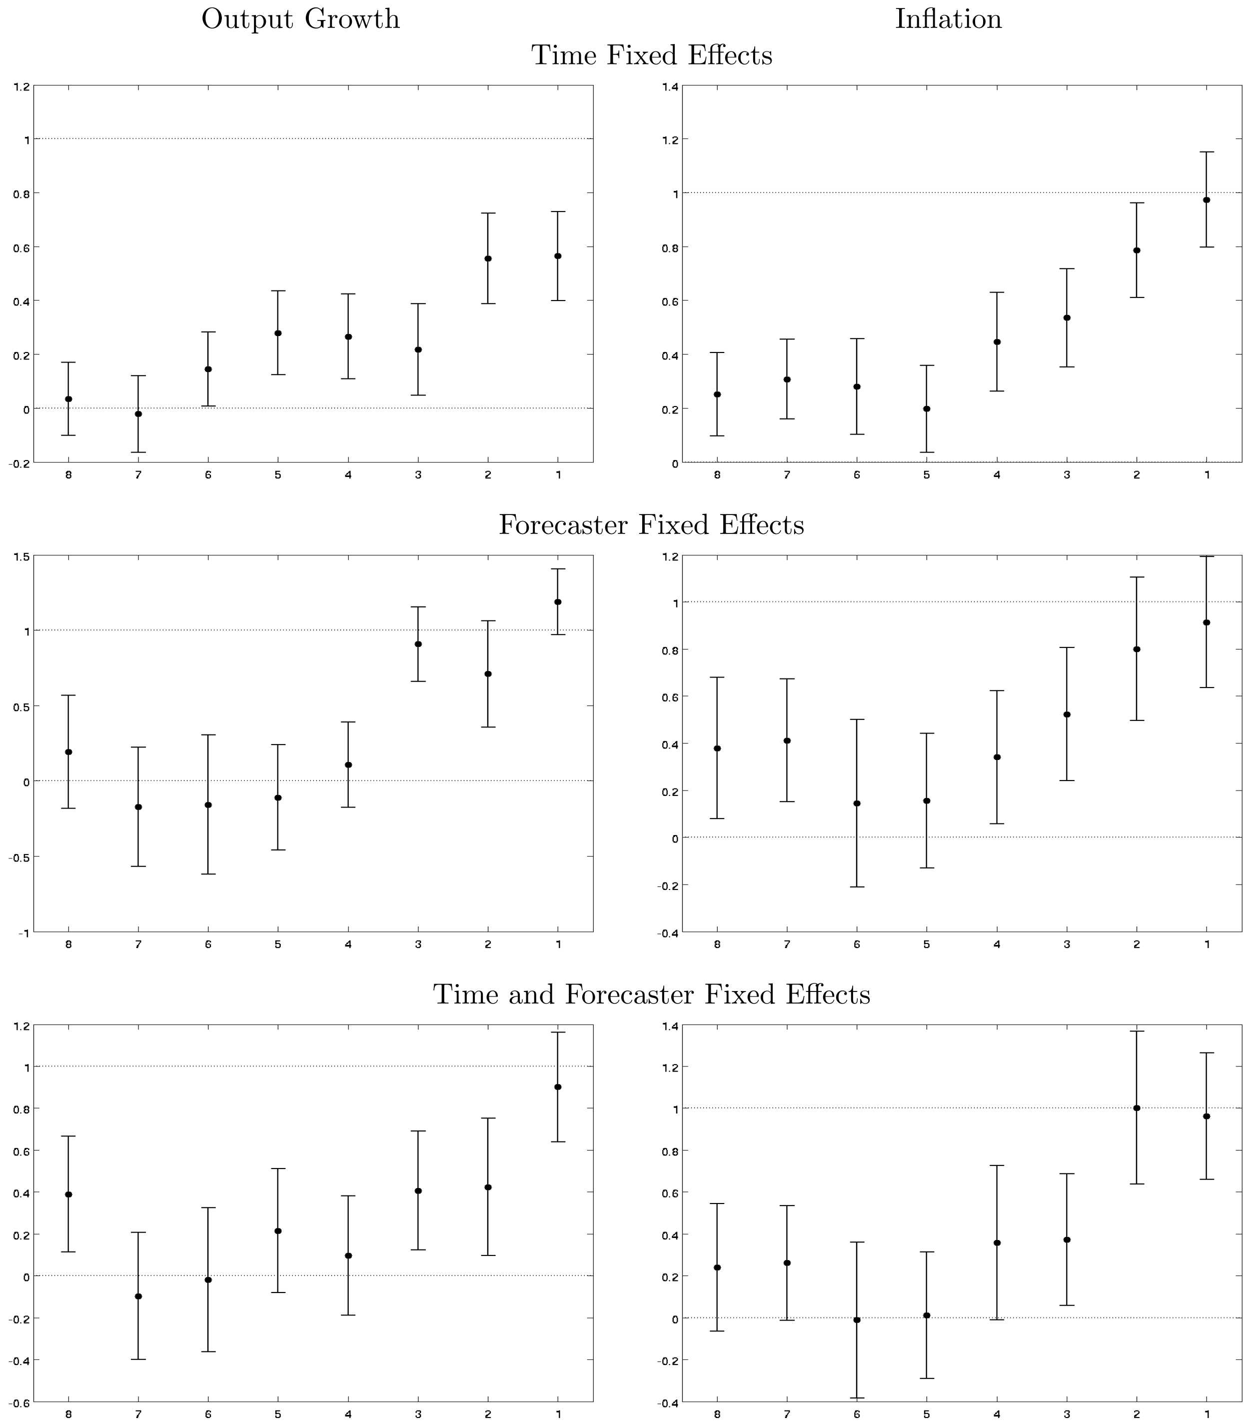


horizons. At longer horizons there is little relationship also for inflation, although for very short horizons $\beta_{1,q}$ is one or very close to one. The results with forecaster fixed effects (middle panels) suggest that changes over time in the standard deviation of the predictive density for inflation or output do not map into changes in forecast accuracy for long horizons, but they do so pretty well for horizons of two quarters or less. This is the case also when we include both forecaster and time fixed effects (bottom panels). 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 Appendix shows that all these results are broadly robust to different samples.

Figure 11: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Mean vs Point Projections

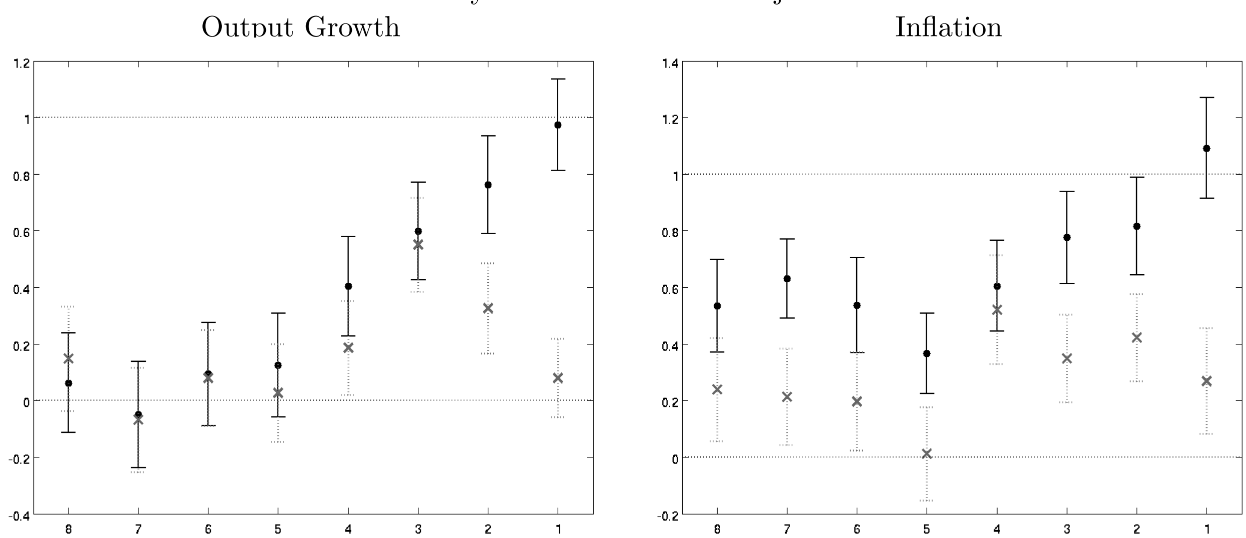


To our knowledge, this approach to testing the noisy rational expectations model and set of results are both new to the literature. As mentioned, Clements (2014a) computes time series averages of $\sigma_{t|t-q,i}$ for each forecaster and plots them against the corresponding predictive root mean square error (RMSE) computed during the same period (Clements 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). 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 RMSEs are computed using predictions other than the mean, there is no *a priori* reason why they should match the subjective standard deviation, even under rational expectations. In fact Figure 11 shows that when we use the point predictions (gray diamonds) the correspondence between subjective uncertainty and forecast error vanishes at short horizons.[14]

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. It also misses the fact that while differences in subjective uncertainty are persistent, forecasters do change their mind about their subjective uncertainty. This is a feature of the data that cross-sectional regressions do not exploit. The finding that in the time dimension the mapping between subjective uncertainty and forecast accuracy is just not there for output, and is only partial for inflation, at longer horizons, but is in line with the noisy rational expectations model for both output and inflation at short horizons, is entirely novel to our knowledge.

### A Location Test: The Relative Accuracy of Mean and Point Predictions

In light of the different results obtained for mean and point forecasts displayed in Figure 11, we now briefly discuss the relative forecasting accuracy of mean versus point forecasts. The top panels of Figure 12 shows OLS estimates of the coefficient $\gamma_q$ in the panel regression

$$\ln \frac{(y_t - E_{t-q,i}[y_t])^2}{(y_t - y^{pp}_{t,t-q,i})^2} = \gamma_q + \epsilon_{t,i,q},\ t = 1,..,T,\ i = 1,..,N. \tag{24}$$

where $y^{pp}_{t,t-q,i}$ is the point forecast for $y_t$ made by forecaster $i$ in period $t-q$.

Estimates of $\gamma_q$ significantly greater than 0 indicate that on average mean projections fare worse than point 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 $\gamma_q$ are not significantly different from 0 for output growth, and only slightly positive for inflation. This

[14] A apter comparison with Clements cross sectional results is in Figure D-14 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 12: Relative Accuracy of Mean vs Point Projections

Output Growth Inflation

Log Ratio of Squared Forecast Errors ($\gamma_q$)

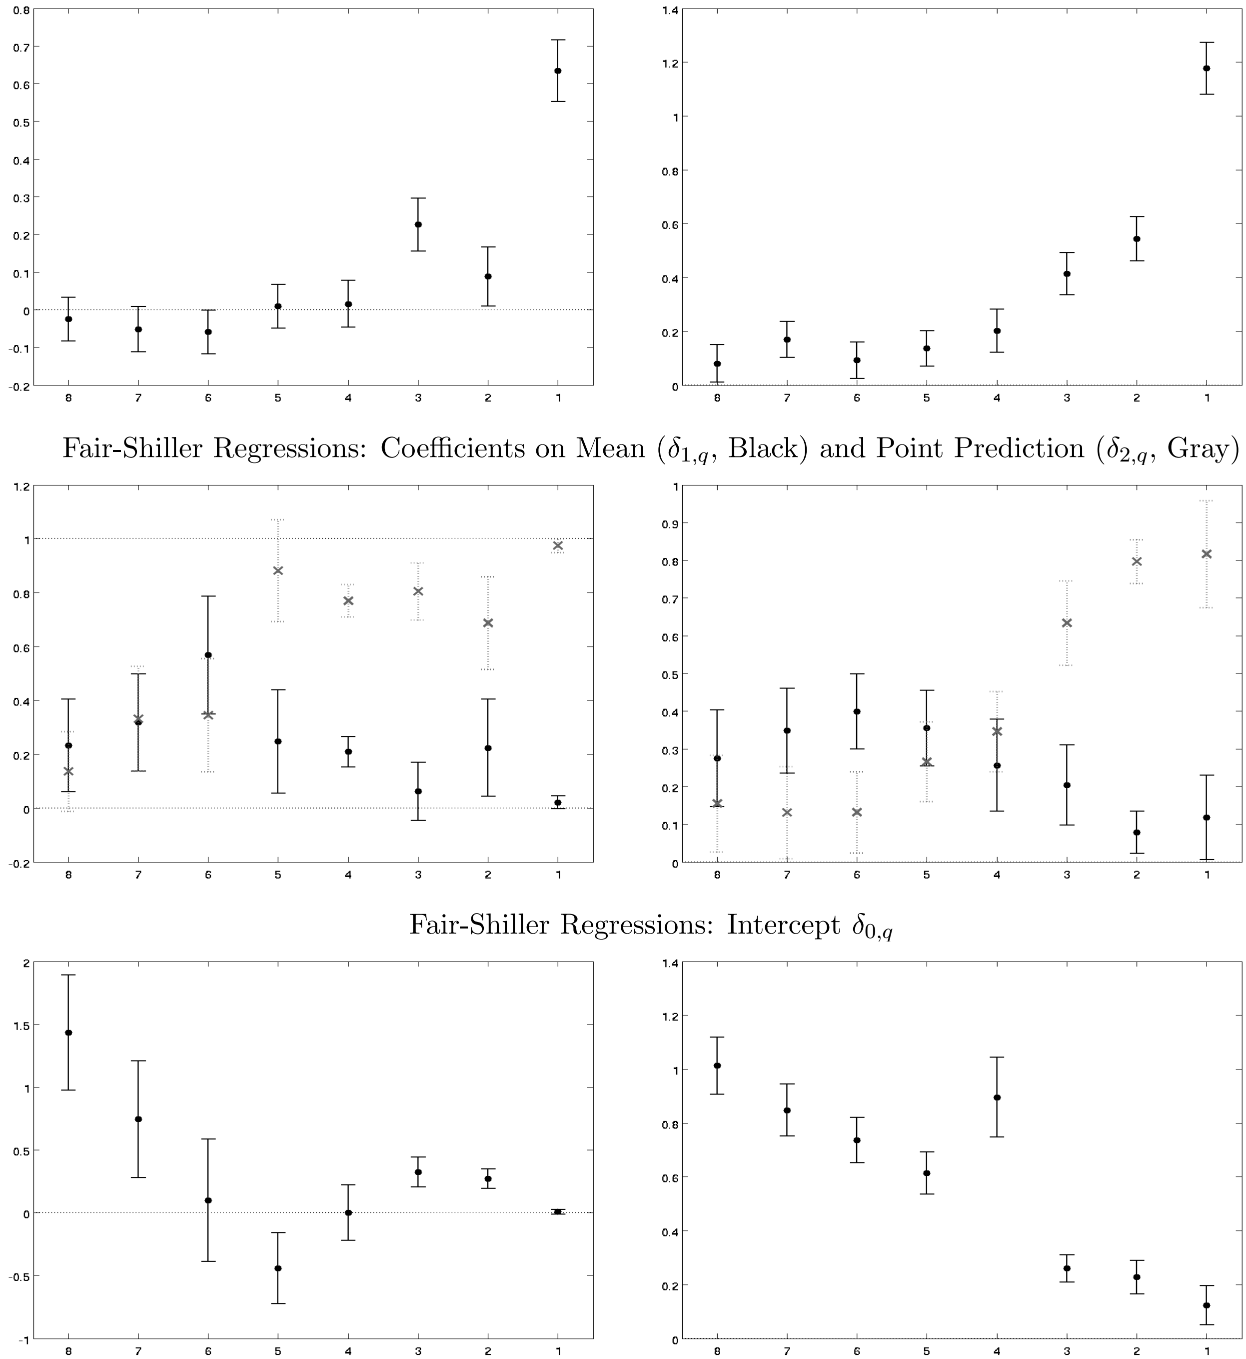


result may partly reflect the fact that for these horizons point and mean predictions are not very different (see Engelberg et al., 2009). As the horizon gets shorter the estimates tend to become much larger and significantly positive for both output growth and inflation.

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 (2009, 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

$$y_t = \delta_{0,q} + \delta_{1,q} E_{t-q,i}[y_t] + \delta_{2,q} y^{pp}_{t,t-q,i} + \epsilon_{t,i,q},\ t = 1, .., T,\ i = 1, .., N. \tag{25}$$

The rationality of density projections would imply $\delta_{0,q} = 0$, $\delta_{1,q} = 1$, and $\delta_{2,q} = 0$. If point projections $y^{pp}_{t,t-q,i}$ coincide with mean forecasts then the two regressors are multicollinear. The middle panels of Figure 12 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 0. 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 0, indicating that point predictions are much closer to actual outcomes than mean forecasts.

**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. This is the case for both output growth and inflation, although overconfidence for output growth is quite striking. 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. For both, the mapping between *ex-ante* uncertainty and ex-post forecast errors is far from one. For very short horizons, density forecasts are correctly scaled for output growth, and slightly underconfident for inflation. For both output and inflation 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. One hypothesis is that these differences stem from heterogeneity in models.

# IV Conclusions

In this paper we presented a novel approach for conducting inference using data from probabilistic surveys, and used it to investigate whether US 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 non parametric 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

**Abel, Joshua, Robert Rich, Joseph Song, and Joseph Tracy**, "The measurement and behavior of uncertainty: evidence from the ECB Survey of Professional Forecasters," *Journal of Applied Econometrics*, 2016, *31* (3), 533–550.

**Andrade, Philippe and Hervé Le Bihan**, "Inattentive professional forecasters," *Journal of Monetary Economics*, 2013, *60* (8), 967–982.

**_ , Richard K Crump, Stefano Eusepi, and Emanuel Moench**, "Fundamental disagreement," *Journal of Monetary Economics*, 2016, *83*, 106–128.

**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.

**Billio, M., R. Casarin, and L. Rossini**, "Bayesian Nonparametric Sparse VAR Models," *Journal of Econometrics*, 2019, *212*, 97–115.

**Binder, Carola C**, "Measuring uncertainty based on rounding: New method and application to inflation expectations," *Journal of Monetary Economics*, 2017, *90*, 1–12.

**Blei, David M., Thomas L. Griffiths, and Michael I. Jordan**, "The nested chinese restaurant process and bayesian nonparametric inference of topic hierarchies," *Journal of the ACM (JACM)*, 2010, *57* (2), 7.

**Boero, G., Jeremy Smith, and Kenneth F. Wallis**, "The Measurement and Characteristics of Professional Forecasts' Uncertainty," *Journal of Applied Econometrics*, 2014, *7* (30), 1029–1046.

**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 _** , "Uncertainty and disagreement in economic prediction: the Bank of England Survey of External Forecasters," *The Economic Journal*, 2008, *118* (530), 1107–1127.

**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.

**Carroll, Christopher D**, "Macroeconomic expectations of households and professional forecasters," *the Quarterly Journal of economics*, 2003, *118* (1), 269–298.

**Casey, Eddie**, "Are professional forecasters overconfident?," *International Journal of Forecasting*, 2021, *37* (2), 716–732.

**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**, "Internal consistency of survey respondents' forecasts: Evidence based on the Survey of Professional Forecasters," *The methodology and practice of econometrics. A festschrift in honour of David F. Hendry*, 2009, pp. 206–226.

_ , "Explanations of the inconsistencies in survey respondents' forecasts," *European Economic Review*, 2010, *54* (4), 536–549.

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

**Clements, Michael P**, "Probability distributions or point predictions? Survey forecasts of US output growth and inflation," *International Journal of Forecasting*, 2014, *30* (1), 99–117.

_ **and Ana Beatriz Galvão**, "Model and survey estimates of the term structure of US macroeconomic uncertainty," *International Journal of Forecasting*, 2017, *33* (3), 591–604.

**Clements, Michael P., 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.

**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.

**Escobar, M. D.**, "Estimating normal means with a Dirichlet process prior," *Journal of the American Statistical Association*, 1994, *89*, 268–277.

— **and M. West**, "Bayesian density estimation and inference using mixtures," *Journal of the American Statistical Association*, 1995, *90*, 577–588.

**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.

**Garcia, Juan A and Andrés Manzanares**, "What can probability forecasts tell us about inflation risks?," 2007.

**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.

**Griffin, J. and M. Kalli**, "Bayesian Nonparametric Vector Autoregressive Models," *Journal of Econometrics*, 2018, *203*, 267–282.

**Griffin, J. E. and M. F. J. Steel**, "Stick-breaking Autoregressive Processes," *Journal of Econometrics*, 2011, *162*, 383–396.

**Griffin, Jim E.**, "Inference in infinite superpositions of non-Gaussian Ornstein–Uhlenbeck processes using Bayesian nonparametic 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.

**Griffiths, Thomas L. and Joshua B. Tenenbaum**, "Optimal predictions in everyday cognition," *Psychological science*, 2006, *17* (9), 767–773.

**Gu, Jiaying and Roger Koenker**, "Unobserved heterogeneity in income dynamics: An empirical Bayes perspective," *Journal of Business & Economic Statistics*, 2017, *35* (1), 1–16.

**Hannah, Lauren A, David M Blei, and Warren B Powell**, "Dirichlet process mixtures of generalized linear models," *Journal of Machine Learning Research*, 2011, *12* (Jun), 1923–1953.

**Hirano, K.**, "Semiparametric Bayesian Inference in Autoregressive Panel Data Models," *Econometrica*, 2002, *70*, 781–799.

**Hjort, Nils Lid, Chris Holmes, Peter Müller, and Stephen G Walker**, *Bayesian nonparametrics*, Vol. 28, Cambridge University Press, 2010.

**Ishwaran, H. and L. F. James**, "Gibbs sampling methods for stick-breaking priors," *Journal of the American Statistical Association*, 2001, *96*, 161–173.

_ **and M. Zarepour**, "Markov chain Monte Carlo in approximate Dirichlet and beta two-parameter process hierarchical models," *Biometrika*, 2000, *87*, 371–390.

**Jensen, Mark J. and John M. Maheu**, "Bayesian semiparametric stochastic volatility modeling," *Journal of Econometrics*, 2010, *157* (2), 306–316.

**Kalli, M., J. E. Griffin, and S. G. Walker**, "Slice sampling mixture models," *Statistics and Computing*, 2011, *21*, 93–105.

**Karabatsos, George and Stephen G Walker**, "Coherent psychometric modelling with Bayesian nonparametrics," *British Journal of Mathematical and Statistical Psychology*, 2009, *62* (1), 1–20.

**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 avd Economic Statistics Section-JSM*, 2009, pp. 2396–2408.

_ **and Xuguang Sheng**, "Measuring forecast uncertainty by disagreement: The missing link," *Journal of Applied Econometrics*, 2010, *25* (4), 514–538.

_ **, Christie Teigland, and Mark Zaporowski**, "Interest rates and the subjective probability distribution of inflation forecasts," *Journal of Money, Credit and Banking*, 1988, *20* (2), 233–248.

**Li, Yuelin, Elizabeth Schofield, and Mithat Gönen**, "A tutorial on Dirichlet process mixture modeling," *Journal of mathematical psychology*, 2019, *91*, 128–144.

**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 M. 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.

**Manzan, Sebastiano**, "Are professional forecasters Bayesian?," *Journal of Economic Dynamics and Control*, 2021, *123*, 104045.

**Mitra, Riten and Peter Müller**, *Nonparametric Bayesian Inference in Biostatistics*, Springer, 2015.

**Müller, Ulrich K**, "Risk of Bayesian inference in misspecified models, and the sandwich covariance matrix," *Econometrica*, 2013, *81* (5), 1805–1849.

**Navarro, Daniel J, Thomas L Griffiths, Mark Steyvers, and Michael D Lee**, "Modeling individual differences using Dirichlet processes," *Journal of mathematical Psychology*, 2006, *50* (2), 101–122.

**Norets, A. and J. 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, J.**, "Bayesian regression with heteroscedastic error density and parametric mean function," *Journal of Econometrics*, 2014, *178*, 624–638.

**Pitman, J.**, *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**, "The relationships among expected inflation, disagreement, and uncertainty: evidence from matched point and density forecasts," *The Review of Economics and Statistics*, 2010, *92* (1), 200–207.

_ **and** _ , "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.

**Scealy, J. L. and A. H. Welsh**, "Regression for compositional data by using distributions defined on the hypersphere," *Journal of the Royal Statistical Society, Series B*, 2011, *73* (3), 351–375.

**Sethuraman, J.**, "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, S. 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, V. and L.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

For computational reasons, we take a data augmentation approach and write the Gibbs sampler using the joint distribution $h(\mathbf{z}, \boldsymbol{\xi}|\boldsymbol{\theta})$. Our infinite mixture model is then

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

Our Gibbs sampler applied to the cross section of $(\mathbf{z}_i, \boldsymbol{\xi}_i)$, $i = 1, \ldots, n$ 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

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

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 $\mathbf{z}_i$ and $\boldsymbol{\xi}_i$ given $u_i$, the weights and the atoms, is

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

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

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

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 (12).

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

   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) $v_{1:\bar{d}}|d_{1:n}, \boldsymbol{\theta}_{1:\infty}, \psi, \mathbf{z}_{1:n}, \boldsymbol{\xi}_{1:n}$.

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

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

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

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

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

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

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

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

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

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

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 $\sum_{\bar{n}_i+1}^{\infty} w_k < u_i$ 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. $\boldsymbol{\theta}_{1:\infty}|v_{1:\infty},u_{1:n},d_{1:n},\psi,\mathbf{z}_{1:n},\boldsymbol{\xi}_{1:n}$
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 $k \le \bar{n}$.

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

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

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

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

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

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

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

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

3. $d_{1:n}|v_{1:\infty}, u_{1:n}, \boldsymbol{\theta}_{1:\infty}, \psi, \mathbf{z}_{1:n}, \boldsymbol{\xi}_{1:n}$
Draws for each $d_i,\ i = 1,\dots,n$, are obtained by drawing from a multinomial with weights proportional to

$$\pi(d_i|\dots) \propto \mathbb{I}(u_i < w_{d_i}) h(\mathbf{z}_i, \boldsymbol{\xi}_i|\boldsymbol{\theta}_{d_i}) \tag{A-11}$$

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 $h(\mathbf{z}_i, \boldsymbol{\xi}_i|\boldsymbol{\theta}_k)$.

# C Further theoretical results

## C.A Model properties

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

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

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 non-parametric prior alleviates possible misspecification issues.

Introduce a latent Dirichlet process $Z_{i,\infty}(\cdot)|\tilde{\boldsymbol{\theta}}_i \sim \mathcal{DP}(\phi(\tilde{\boldsymbol{\theta}}_i), F(\cdot|\tilde{\boldsymbol{\theta}}_i))$ with parameters $\phi(\tilde{\boldsymbol{\theta}}_i)$ and $F(\cdot|\tilde{\boldsymbol{\theta}}_i)$, given $\tilde{\boldsymbol{\theta}}_i$ 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

$$Z_{i,\infty}(y) = \sum_{j=1}^{\infty} w_{ij} \mathbb{I}\{y_{ij} \leq y\} \tag{E-2}$$

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

$\Delta^J$ and $Z_{i,\infty}(y)$ on $\mathcal{Y}$.

**Proposition 1.** If $\alpha_j = 0$, $j = 1, \ldots, J$, the Bayesian model

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

where $\mathbf{z}_i = (z_{i,1}, \ldots, z_{i,J})$ admits the following stochastic representation:

$$(z_{i,1}, \ldots, z_{i,J}) := (Z_{i,\infty}(y_1), Z_{i,\infty}(y_2) - Z_{i,\infty}(y_1), \ldots, 1 - Z_{i,\infty}(y_{J-1})) \quad i = 1, \ldots, n$$
$$Z_{i,\infty} \overset{ind}{\sim} \mathcal{DP}(\phi(\tilde{\boldsymbol{\theta}}_i), F(\cdot|\tilde{\boldsymbol{\theta}}_i)) \quad i = 1, \ldots, n$$
$$\tilde{\boldsymbol{\theta}}_i \overset{i.i.d.}{\sim} G \quad i = 1, \ldots, n$$
$$G \sim \mathcal{DP}(\psi, G_0).$$

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 $(z_{i,1}, \ldots, z_{i,J})$ 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 $\phi(\boldsymbol{\theta}_i)$.

When $\alpha(\cdot|\epsilon) \neq 0$ an extra noise is set in, resulting in a random proportions of bins which are randomly set to zero with probability $\alpha(\nu_j(\boldsymbol{\varphi}_i)|\epsilon)$. After this deletion, in order to obtain the $\mathbf{z}_i$, the increments of the Dirichelet process $\tilde{F}_i$ are simply normalized to sum one.

forecaster-specific subjective distribution. Given the subjective probability distribution $F(\cdot|\tilde{\boldsymbol{\theta}}_i)$ and the level of noise $\phi(\tilde{\boldsymbol{\theta}}_i)$, the forecaster reports the weights $(z_{i,1}, \ldots, z_{i,J})$ corresponding to the increments of a "noisy" version $Z_{i,\infty}$ of $F(\cdot|\tilde{\boldsymbol{\theta}}_i)$. This "noisy" version is the CDF obtained by a forecaster-specific Dirichlet process with base measure $F(\cdot|\tilde{\boldsymbol{\theta}}_i)$ and concentration parameter $\phi(\tilde{\boldsymbol{\theta}}_i)$.

By (E-2), 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 E-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 $F(\cdot|\tilde{\boldsymbol{\theta}}_i)$. 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 $Z_{i,J}$. For any $J$ we consider the partition $\mathcal{P}_J = \{y_0^J = -\infty < y_1^J < \ldots < y_J^J = +\infty\}$ and define the following one-to-one mapping between $\mathbf{z}_i$ and the CDF $Z_{i,J}$. Without loss of generality, we assign to the middle point of each interval the bin probability mass, and

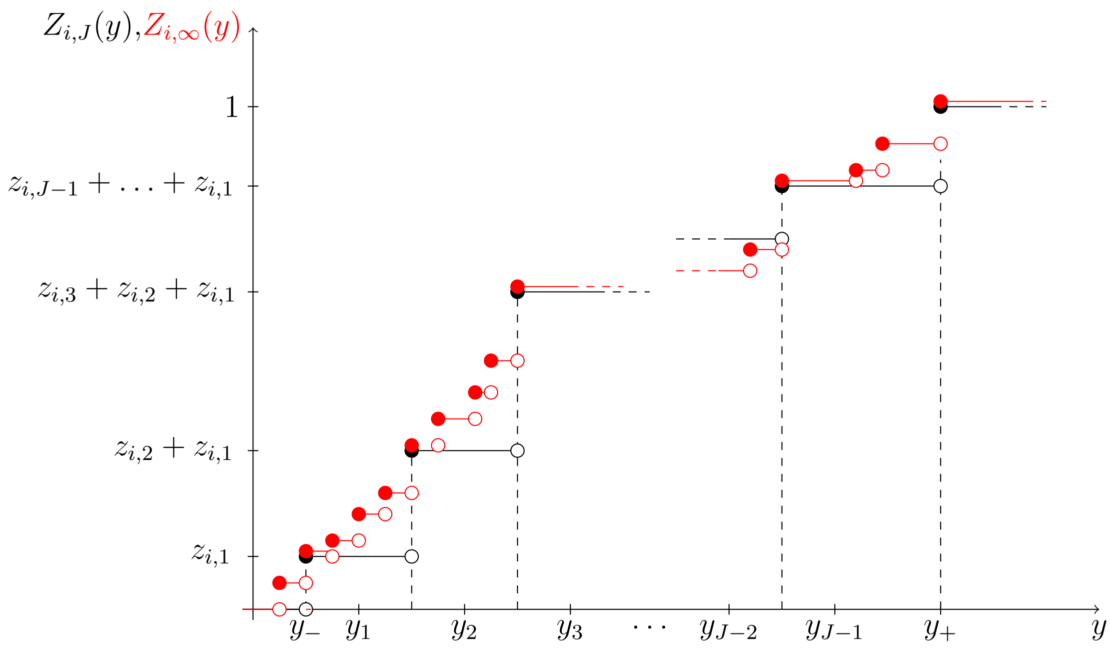


Figure E-1: Mapping between $z_{i,j}$, $j=1,\ldots,J$, $Z_{iJ}(y)$ and $Z_{i,\infty}(y)$.

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 E-1):

$$Z_{i,J}(y) = \begin{cases} 0 & \text{if } y < y_-^J \\ z_{i,1} & \text{if } y_-^J \le y < (y_1^J + y_2^J)/2 \\ z_{i,1}+\cdots+z_{i,j} & \text{if } y \in [(y_{j-1}^J + y_j^J)/2, (y_j^J + y_{j+1}^J)/2) \text{ for } 2 < j \le J-2 \\ z_{i,1}+\cdots+z_{i,J-1} & \text{if } y \in [(y_{J-2}^J + y_{J-1}^J)/2, y_+^J) \\ 1 & \text{if } y \ge y_+^J \end{cases}$$

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

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

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

If $F(\cdot|\boldsymbol{\theta}_i)$ is a continuous CDF, then

$$\lim_{J\to+\infty}\mathbb{E}[Z_{i,J}(y)|\tilde{\boldsymbol{\theta}}_i]=\mathbb{E}[Z_{i,\infty}(y)|\boldsymbol{\theta}_i]=F(y|\tilde{\boldsymbol{\theta}}_i)\qquad a.s.$$

## C.B Further asymptotics

### C.B.1 Posterior consistency

If $\alpha_j=0$ for $j=1,\dots,J$, i.e. forecasters give non-zero probability to each bin, the sample space is $\mathcal{Z}=\Delta^J$ and $\mathcal{H}$ is the set densities (with respect to the Lebesgue measure) on $\Delta^J$ . In this case, Kullback-Leibler divergence between two distribution $h_0, g$ on $\mathcal{Z}=\Delta^J$ is easily defined as

$$KL(h_0,g):=\int_{\mathcal{Z}}h_0(\mathbf{z})\log\Big(\frac{h_0(\mathbf{z})}{g(\mathbf{z})}\Big)d\mathbf{z}.$$

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

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

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

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

**Theorem 3.** Let $\Theta$ be an open subset of $\mathbb{R}^m$ for some $m$ and $\alpha_j(\boldsymbol{\theta})=0$ 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

$$\int_{\Delta^J}\left|\log\left(\prod_{j=1}^{J-1}z_j\Big(1-\sum_{j=1}^{J-1}z_j\Big)\right)\right|h_0(\mathbf{z})d\mathbf{z}<+\infty.\tag{E-3}$$

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

$$F(y|\boldsymbol{\theta})=\sum_{i=1}^M\omega_i\Phi(y|\mu_i,\sigma_i^2)\tag{E-4}$$

then the parameter vector is $\boldsymbol{\theta} = (\mu_1, \ldots, \mu_M, \sigma_1^2, \ldots, \sigma_M^2, \omega_1, \ldots, \omega_M, \phi)$. If $G_0$ has full support, then the posterior is weakly consistent at any $h_0$ in $\mathcal{H}_0$ satisfying (E-3). Indeed, in this case $(\phi\nu_1(\boldsymbol{\theta}), \ldots, \phi\nu_J(\boldsymbol{\theta}))$ is a continuous function on $\mathbb{R}_+^J$ and $\phi\nu_j(\boldsymbol{\theta}) > 0$ for every $j = 1, \ldots, J$.

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

**Proposition 2.** Assume $\alpha_j(\boldsymbol{\theta}) = 0$ for all $j = 1, \ldots, J$ and that $\boldsymbol{\theta} \mapsto (\phi(\boldsymbol{\theta})\nu_1(\boldsymbol{\theta}), \ldots, \phi(\boldsymbol{\theta})\nu_J(\boldsymbol{\theta}))$ is a continuous function on $\mathbb{R}_+^J$ such that $\phi(\boldsymbol{\theta})\nu_j(\boldsymbol{\theta}) > 0$ for every $j = 1, \ldots, J$. If for every $\boldsymbol{a} = (a_1, \ldots, a_J) \in [1, +\infty)^J$ and $\delta > 0$, there is $\boldsymbol{\theta}_\delta$ in $\Theta$ such that $\|\boldsymbol{a} - \boldsymbol{a}_\delta\|_\infty \leq \delta$ with $\boldsymbol{a}_\delta = \phi(\boldsymbol{\theta}_\delta)(\nu_1(\boldsymbol{\theta}_\delta), \ldots, \nu_J(\boldsymbol{\theta}_\delta))$, then any continuous density function on $\Delta^J$ 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 $\Delta^J$ and satisfies (E-3). An example in which all the conditions of Proposition 2 are met is the fully non-parametric case

$$F(y|\boldsymbol{\theta}) = \sum_{j=1}^{J} \varphi_j \mathbb{I}_{A_j}(y) \tag{E-5}$$

where $A_j = [y_j, +\infty)$, $j = 1, \ldots, J-1$, $A_J = [y^+, +\infty]$ and $\nu_j(\boldsymbol{\theta}) = \varphi_j$, $j = 1, \ldots, J$. Conditions in Proposition 2 are satisfied also in the Gaussian mixture case of (E-4) 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 defined in (14). In what follows, $F_{n+1}$ denotes the posterior predictive distribution of $y$, defined as

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

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,

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

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

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

As in Subsection C.A, we consider set of nested partitions $\mathcal{P}_J = \{y_0^J = -\infty < y_1^J < \ldots < y_J^J = +\infty\}$ in such a way $\mathcal{P}_{J+1}$ is a refinement of $\mathcal{P}_J$. We assume that observations $\mathbf{z}_1^J, \ldots, \mathbf{z}_n^J$ 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 $(w_{1,0}, \ldots, w_{M,0})$ and parameters $\boldsymbol{\theta}_{1,0}, \ldots, \boldsymbol{\theta}_{M,0}$ in $\Theta$. Note that with these hypotheses $\mathbf{z}_1^J, \ldots, \mathbf{z}_n^J$ 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 ($n \to +\infty$) and in the number of bins ($J \to +\infty$). 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

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

for every $y$ point of continuity of $F^*$.

# D Proofs

## D.A Details on formula (16)

Let $\mathcal{H}$ be the set of all the densities with respect to $\lambda$, i.e. the densities $g$ that factorize as $g(\boldsymbol{\xi},\mathbf{z}) = g(\boldsymbol{\xi})g(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\xi})$. This assumption is coherent with the model assumption in (4), indeed

$$h(\boldsymbol{\xi},\mathbf{z}|\boldsymbol{\theta}) = h(\boldsymbol{\xi}|\boldsymbol{\theta})h(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\theta},\boldsymbol{\xi}), \tag{B-1}$$

where

$$h(\boldsymbol{\xi}|\boldsymbol{\theta}) = \frac{1}{c(\boldsymbol{\theta})}\prod_{j=1}^{J}\alpha_j(\boldsymbol{\theta})^{\xi_j}(1-\alpha_j(\boldsymbol{\theta}))^{1-\xi_j}$$

and $h(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\theta},\boldsymbol{\xi}) = h(\mathbf{z}|\boldsymbol{\theta},\boldsymbol{\xi})$ is the Dirichlet distribution of parameters $[\phi(\boldsymbol{\theta})\nu_j(\boldsymbol{\theta}) : j \in \mathcal{J}(\boldsymbol{\xi})]$ defined on the non zero elements $\mathbf{z}_{\boldsymbol{\xi}}$.

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

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

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

$$\sum_{\boldsymbol{\xi}\in\mathcal{X}} h_0(\boldsymbol{\xi})\int_{\Delta^{J-|\boldsymbol{\xi}|}} h_0(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\xi})\log\Big(\frac{h_0(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\xi})h_0(\boldsymbol{\xi})}{g(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\xi})g(\boldsymbol{\xi})}\Big)d\mathbf{z}_{\boldsymbol{\xi}}$$
$$= \sum_{\boldsymbol{\xi}\in\mathcal{X}} h_0(\boldsymbol{\xi})\log\Big(\frac{h_0(\boldsymbol{\xi})}{g(\boldsymbol{\xi})}\Big) + \sum_{\boldsymbol{\xi}\in\mathcal{X}} h_0(\boldsymbol{\xi})\int_{\Delta^{J-|\boldsymbol{\xi}|}} h_0(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\xi})\log\Big(\frac{h_0(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\xi})}{g(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\xi})}\Big)d\mathbf{z}_{\boldsymbol{\xi}}.$$

## D.B Proofs of Theorem 1 and 3

The proof of Theorem 1 is based on an application of Theorem 1 and Lemma 3 of Wu and Ghosal (2009a,b). 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 $\mu$. We assume that $\mathcal{X}_0$ is a subset the finte set $\mathcal{X} = \{\boldsymbol{\xi} \in \{0,1\}^J : |\boldsymbol{\xi}| < J\}$. Following the notation introduced above, the sample space $\mathcal{Z}$ is the set of all the pairs $(\boldsymbol{\xi},\mathbf{z})$, where $\boldsymbol{\xi} = (\xi_1,\dots,\xi_J)$, $\mathbf{z} = (z_1,\dots,z_J)$, $\xi_i = \mathbb{I}\{z_i = 0\}$. The non-null elements of $\mathbf{z}$, denoted by $\mathbf{z}_{\boldsymbol{\xi}}$, takes values in an open subset $\mathcal{Z}_{\boldsymbol{\xi}}$ of $\mathbb{R}^{J-|\boldsymbol{\xi}|}$. In our application $\mathcal{Z}_{\boldsymbol{\xi}} = \Delta^{J-|\boldsymbol{\xi}|}$. On the sample space $\mathcal{Z}$, one can thus define a $\sigma$-finte measure $\lambda(d\boldsymbol{\xi} d\mathbf{z}) = c(d\boldsymbol{\xi}) \otimes \mathcal{L}_{\boldsymbol{\xi}}(d\mathbf{z}_{\boldsymbol{\xi}})$ where $c$ is the counting measure on $\mathcal{X}$ and, given $\boldsymbol{\xi}$, $\mathcal{L}_{\boldsymbol{\xi}}$ is the Lebesgue measure on $\mathcal{Z}_{\boldsymbol{\xi}} \subset \mathbb{R}^{J-|\boldsymbol{\xi}|}$. Let $\mathcal{H}$ be the set of all the densities with respect to $\lambda$, i.e. the densities $g$ that factorize as $g(\boldsymbol{\xi}, \mathbf{z}) = g(\boldsymbol{\xi}) g(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\xi})$. We also assume that the kernel $h(\boldsymbol{\xi}|\boldsymbol{\theta}) h(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\theta}, \boldsymbol{\xi})$ factorizes in the same way, i.e.

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

Finally, given a probability measure $G$ on $\Theta$, we write

$$h_G(\boldsymbol{\xi}, \mathbf{z}) = \int_{\Theta} h(\boldsymbol{\xi}, \mathbf{z}|\boldsymbol{\theta}) G(d\boldsymbol{\theta}). \tag{B-2}$$

and we assume that $\Pi$ is the prior on $\mathcal{H}$ induced by the map (B-2) when $G$ has prior $\hat{\Pi}$.

In our application, $h_G(\boldsymbol{\xi}, \mathbf{z})$ is given by (4) and $\hat{\Pi}$ is the Dirichlet process prior $\mathcal{DP}(\psi, G_0)$.

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

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

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

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

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

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

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

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

*Proof of Theorem 4.* One has

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

If we show that there is an open neighbourhood $V$ of $G_\varepsilon$ 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

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

and

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

Let $V_0:=\{G:G(D_\varepsilon)>1/2\}$. Since $G_\varepsilon(D_\varepsilon)=1$, by Portmanteau Theorem $V$ is an open neighbourhood of $G_\varepsilon$. 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 $V_1$,

$$\begin{aligned}&\int_{C^c_{\boldsymbol{\xi}}}\log\Big(\frac{h_{G_\varepsilon}(\boldsymbol{\xi},\mathbf{z}_{\boldsymbol{\xi}})}{h_G(\boldsymbol{\xi},\mathbf{z}_{\boldsymbol{\xi}})}\Big)h_0(\boldsymbol{\xi},\mathbf{z}_{\boldsymbol{\xi}})d\mathbf{z}_{\boldsymbol{\xi}}\\&\leq\int_{C^c_{\boldsymbol{\xi}}}\log\Big(\frac{h_{G_\varepsilon}(\boldsymbol{\xi},\mathbf{z}_{\boldsymbol{\xi}})}{\inf_{\boldsymbol{\theta}\in D_\varepsilon}h(\boldsymbol{\xi}|\boldsymbol{\theta})h(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\xi},\boldsymbol{\theta})}\Big)h_0(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\xi})d\mathbf{z}_{\boldsymbol{\xi}}+\log(2)\int_{C^c_{\boldsymbol{\xi}}}h_0(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\xi})d\mathbf{z}_{\boldsymbol{\xi}}\leq\frac{\varepsilon}{2}.\end{aligned}\tag{B-3}$$

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

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

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

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

is a weak neighbourhood of $G_\varepsilon$. Hence, for $G$ in $V_{\boldsymbol{\xi}}$

$$\Big|\int_{D_\varepsilon}h(\boldsymbol{\xi},\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\theta})G_\varepsilon(d\boldsymbol{\theta})-\int_{D_\varepsilon}h(\boldsymbol{\xi},\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\theta})G(d\boldsymbol{\theta})\Big|\leq\frac{c\varepsilon}{4}\tag{B-4}$$

Since $supp(G_\varepsilon) \subset D_\varepsilon$,

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

Hence, using $\log(x+1) \le x$ and (B-4), for $G$ in $V_0 \cap V_{\boldsymbol{\xi}}$ one obtains

$$\int_{C_{\boldsymbol{\xi}}} \log\Big(\frac{h_{G_\varepsilon}(\boldsymbol{\xi},\mathbf{z}_{\boldsymbol{\xi}})}{h_G(\boldsymbol{\xi},\mathbf{z}_{\boldsymbol{\xi}})}\Big) h_0(\mathbf{z}_{\boldsymbol{\xi}}|\boldsymbol{\xi})d\mathbf{z}_{\boldsymbol{\xi}} \le \frac{\varepsilon}{2}. \tag{B-5}$$

At this stage, combining (B-3) and (B-5), one obtains that $A_\varepsilon(G)$ for every $G$ in $V = V_0 \cap (\cap_{\boldsymbol{\xi}} V_{\boldsymbol{\xi}})$. □

We can now prove both Theorem 3 and Theorem 1.

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

$$\tilde{\boldsymbol{\nu}}(\boldsymbol{\theta}) := (\tilde\nu_1(\boldsymbol{\theta}),\dots,\tilde\nu_J(\boldsymbol{\theta})) = (\phi(\boldsymbol{\theta})\nu_1(\boldsymbol{\theta}),\dots,\phi(\boldsymbol{\theta})\nu_J(\boldsymbol{\theta})) \tag{B-6}$$

and

$$Z_{\boldsymbol{\theta}} = \frac{\prod_{j=1}^J \Gamma(\tilde\nu_j(\boldsymbol{\theta}))}{\Gamma\Big(\sum_{j=1}^J \tilde\nu_j(\boldsymbol{\theta})\Big)}.$$

*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 $KL(h_0, g_\varepsilon) \le \varepsilon$. 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 $supp(G_\varepsilon) = \cup_{i=1}^{M_\varepsilon}\{\boldsymbol{\theta}_{i,\varepsilon}\}$. To conclude recall that if $\Pi$ is $\mathcal{DP}(\psi, G_0)$ and $supp(G_\varepsilon) \subset supp(G_0)$, then $G_\varepsilon \in supp(\Pi)$; see, for instance, Theorem 3.2.4 of Ghosh and Ramamoorthi (2003).

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

Now

$$\begin{aligned} I_\varepsilon(\mathbf{z}) &:= \inf_{\boldsymbol{\theta}\in D_\varepsilon} h(\mathbf{z}|\boldsymbol{\theta}) \\ &= \inf_{\boldsymbol{\theta}\in D_\varepsilon} \frac{1}{Z_{\boldsymbol{\theta}}} \prod_{j=1}^{J-1} z_j^{\tilde\nu_j(\boldsymbol{\theta})-1}\Big(1-\sum_{j=1}^{J-1} z_j\Big)^{\tilde\nu_j(\boldsymbol{\theta})-1} \\ &\ge C_{1,\varepsilon}\prod_{j=1}^{J-1} z_j^{\mu_{j,\varepsilon}-1}\Big(1-\sum_{j=1}^{J-1} z_j\Big)^{\mu_{J,\varepsilon}-1} =: I^*_\varepsilon(\mathbf{z}) \end{aligned}$$

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 $D_\varepsilon$ 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

$$\int \log\Big(\frac{h_{G_\varepsilon}(\mathbf{z})}{I_\varepsilon(\mathbf{z})}\Big)h_0(\mathbf{z})d\mathbf{z} \leq \int \log\Big(\frac{h_{G_\varepsilon}(\mathbf{z})}{I^*_\varepsilon(\mathbf{z})}\Big)h_0(\mathbf{z})d\mathbf{z}$$

$$\leq \int \left|\log\left(\frac{g_\varepsilon(\mathbf{z})}{\prod_{j=1}^{J-1} z_j^{\mu_{j,\varepsilon}-1}\Big(1-\sum_{j=1}^{J-1} z_j\Big)^{\mu_{J,\varepsilon}-1}}\right)\right| h_0(\mathbf{z})d\mathbf{z} + |\log(C_{1,\varepsilon})|.$$

Since

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

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

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

$$\leq C_{4,\varepsilon}\left[1+\left|\log\left(\prod_{j=1}^{J-1} z_j\Big(1-\sum_{j=1}^{J-1} z_j\Big)\right)\right|\right]$$

Combining all the estimates, one gets

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

by assumption (E-3). Hence

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

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

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

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 $D_\varepsilon$.

□

*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

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

where

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

Moreover,

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

where

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

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 $D_\varepsilon$ is compact, $\nu_j(\boldsymbol{\theta})$ and $\alpha_j(\boldsymbol{\theta})$ are continuous, $0 < \alpha_j(\boldsymbol{\theta}) < 1$ and $\nu_j(\boldsymbol{\theta}) > 0$, $j = 1,\dots,J$. Finally,

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

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

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

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

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

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

## D.C Proof of Proposition 2

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

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

Note that

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

where $\tilde{\boldsymbol{\nu}}(\boldsymbol{\theta})$ is defined in (B-6).

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

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

**Lemma 2.** Let $a = (a_1, \ldots, a_J) \in [1, +\infty)^J$. 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

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

*Proof.* The Proof is left to the reader. □

**Lemma 3.** Assume that, for every $a = (a_1, \ldots, a_j) \in [1, +\infty)^J$ 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 $\Delta^J$ 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 $g_\varepsilon(\mathbf{z}) = \sum_{i=1}^{M_\varepsilon} q_{i,\varepsilon} D(\mathbf{z}; a_{i,1,\varepsilon}, \ldots, a_{i,J,\varepsilon})$ 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 $\|D(\cdot; a_{i,1,\varepsilon}, \ldots, a_{i,J,\varepsilon}) - D(\cdot; \tilde{\nu}_1(\boldsymbol{\theta}_{i,\varepsilon}), \ldots, \tilde{\nu}_J(\boldsymbol{\theta}_{i,\varepsilon}))\|_\infty \leq \varepsilon/2$. Hence, setting $\tilde{g}_\varepsilon(\mathbf{z}) := \sum_{i=1}^{M_\varepsilon} q_{i,\varepsilon} D(\mathbf{z}; \tilde{\nu}_1(\boldsymbol{\theta}_{i,\varepsilon}), \ldots, \tilde{\nu}_J(\boldsymbol{\theta}_{i,\varepsilon}))$, one gets

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

□

**Lemma 4.** For every densities $g_1$ and $g_2$ in $\Delta^J$

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

*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 $\Delta^J$, then, for every $\eta > 0$, there is a density $g_\eta$ 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 $g$

$$KL(h_0, g) \leq (2+\varepsilon)\log(1+\varepsilon) + (1+\varepsilon)[KL(h_\varepsilon, g) + \sqrt{KL(h_\varepsilon, g)}]. \tag{B-7}$$

By Lemma 3 there is a density $\tilde{g}_\varepsilon$ 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.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, $\tilde{\boldsymbol{\theta}}_i := \boldsymbol{\theta}_{d_i}$ are drawn form 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

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

see (**??**). Hence by the law of iterated expectations

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

Since

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

we obtain the result

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

□

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

## D.E Proofs of Proposition 1

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

$$\begin{aligned}(z_{i,1}, \ldots, z_{i,J}) &\sim Dir_J(\phi(\boldsymbol{\theta}_i)\nu_1(\boldsymbol{\theta}_i), \ldots, \phi(\boldsymbol{\theta}_i)\nu_J(\boldsymbol{\theta}_i))\\ \boldsymbol{\theta}_i &\overset{i.i.d.}{\sim} G\\ G &\sim \mathcal{DP}(\psi, G_0),\end{aligned}$$

and the thesis follows. □

# E Additional Results

## E.A Survey design

Figure D-1: Bin Ranges

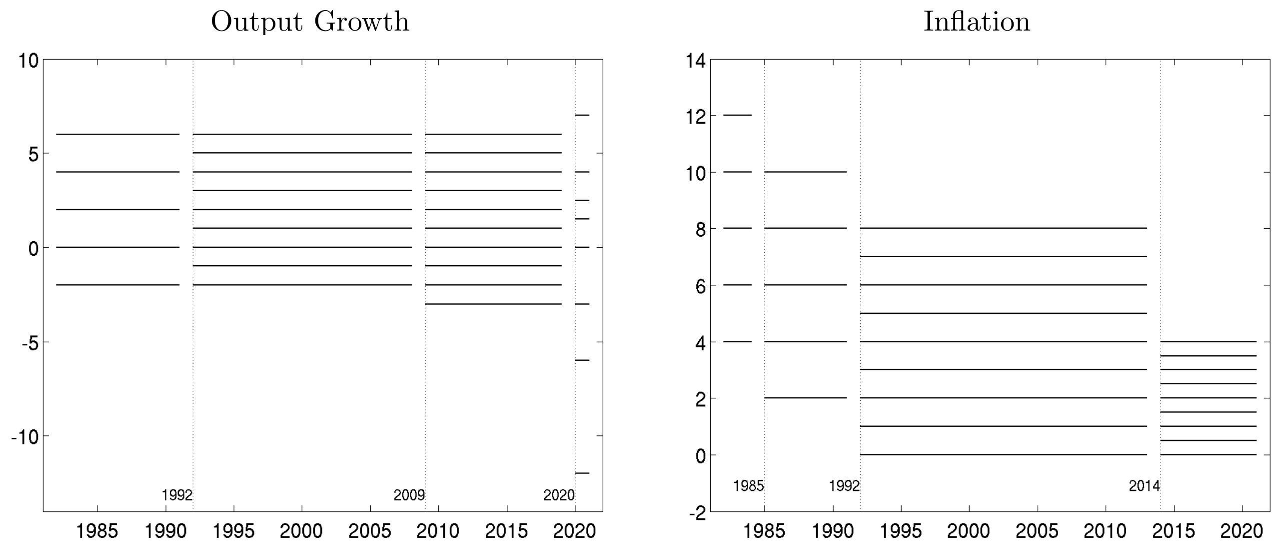

Figure D-2: Number of respondents for H1 output growth surveys

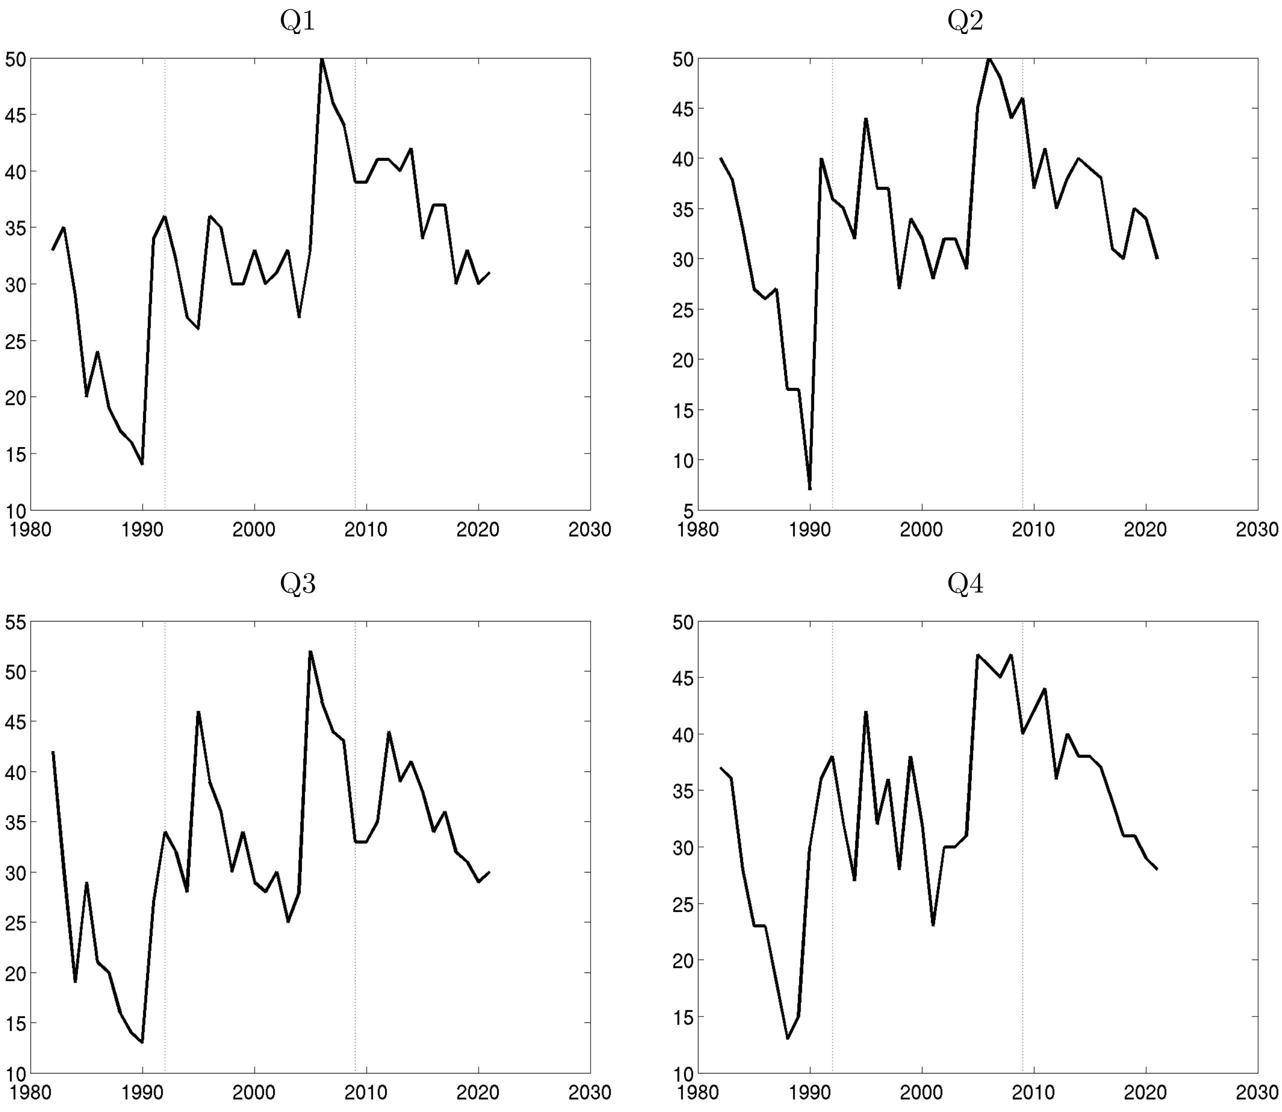

Figure D-3: SPF survey participation by respondent

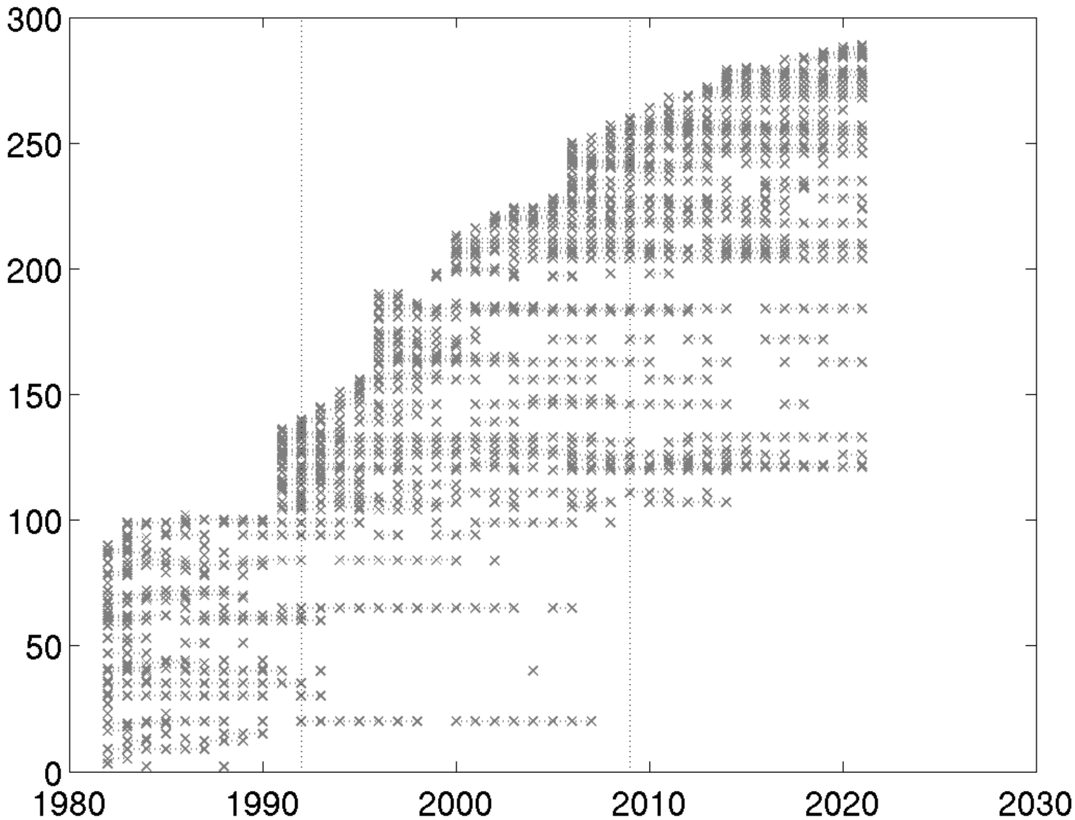

Figure D-4: Percentage of respondents for H2 output growth surveys placing positive probability on either one open bin or both

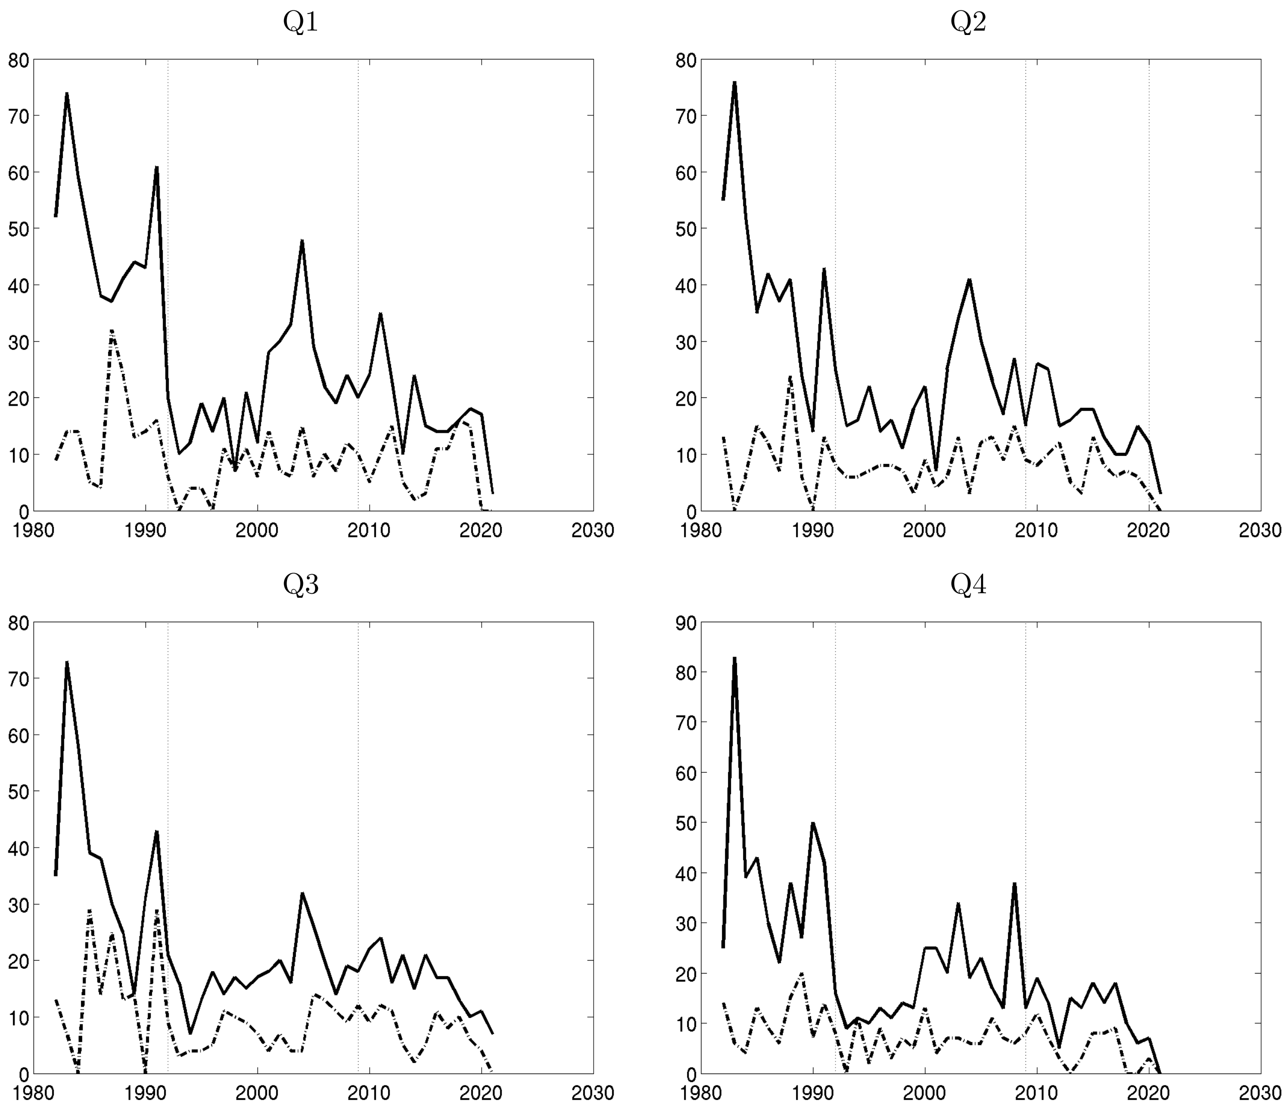

Figure D-5: Percentage of respondents for H2 inflation surveys placing positive probability on either one open bin or both

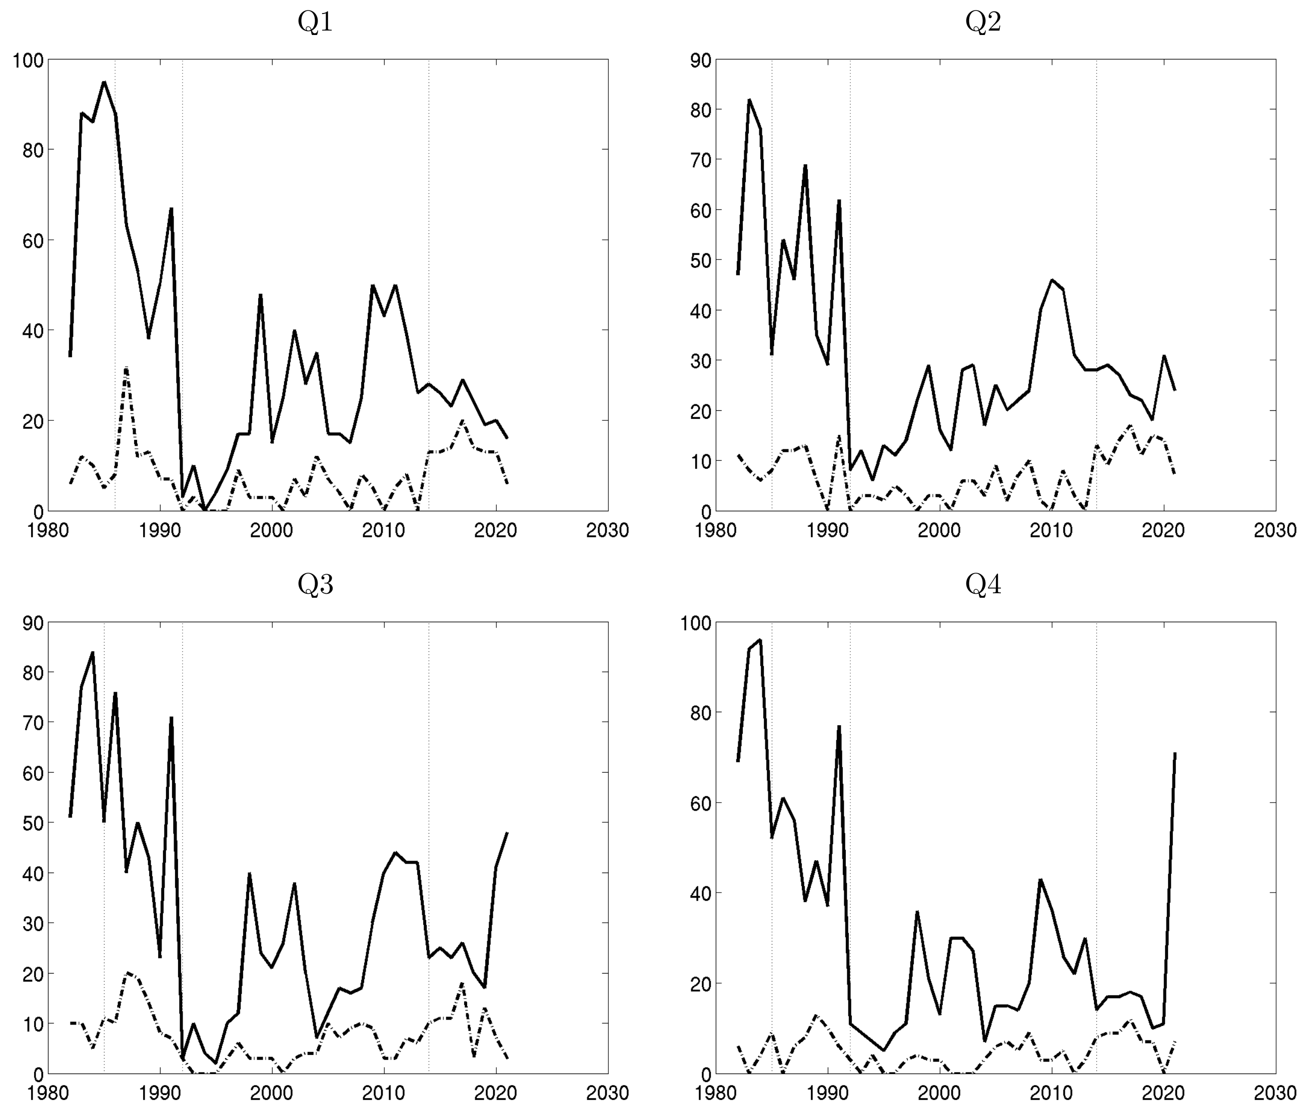

## E.B Heterogeneity in subjective uncertainty

Figure D-6: Subjective uncertainty by individual respondent: Q1

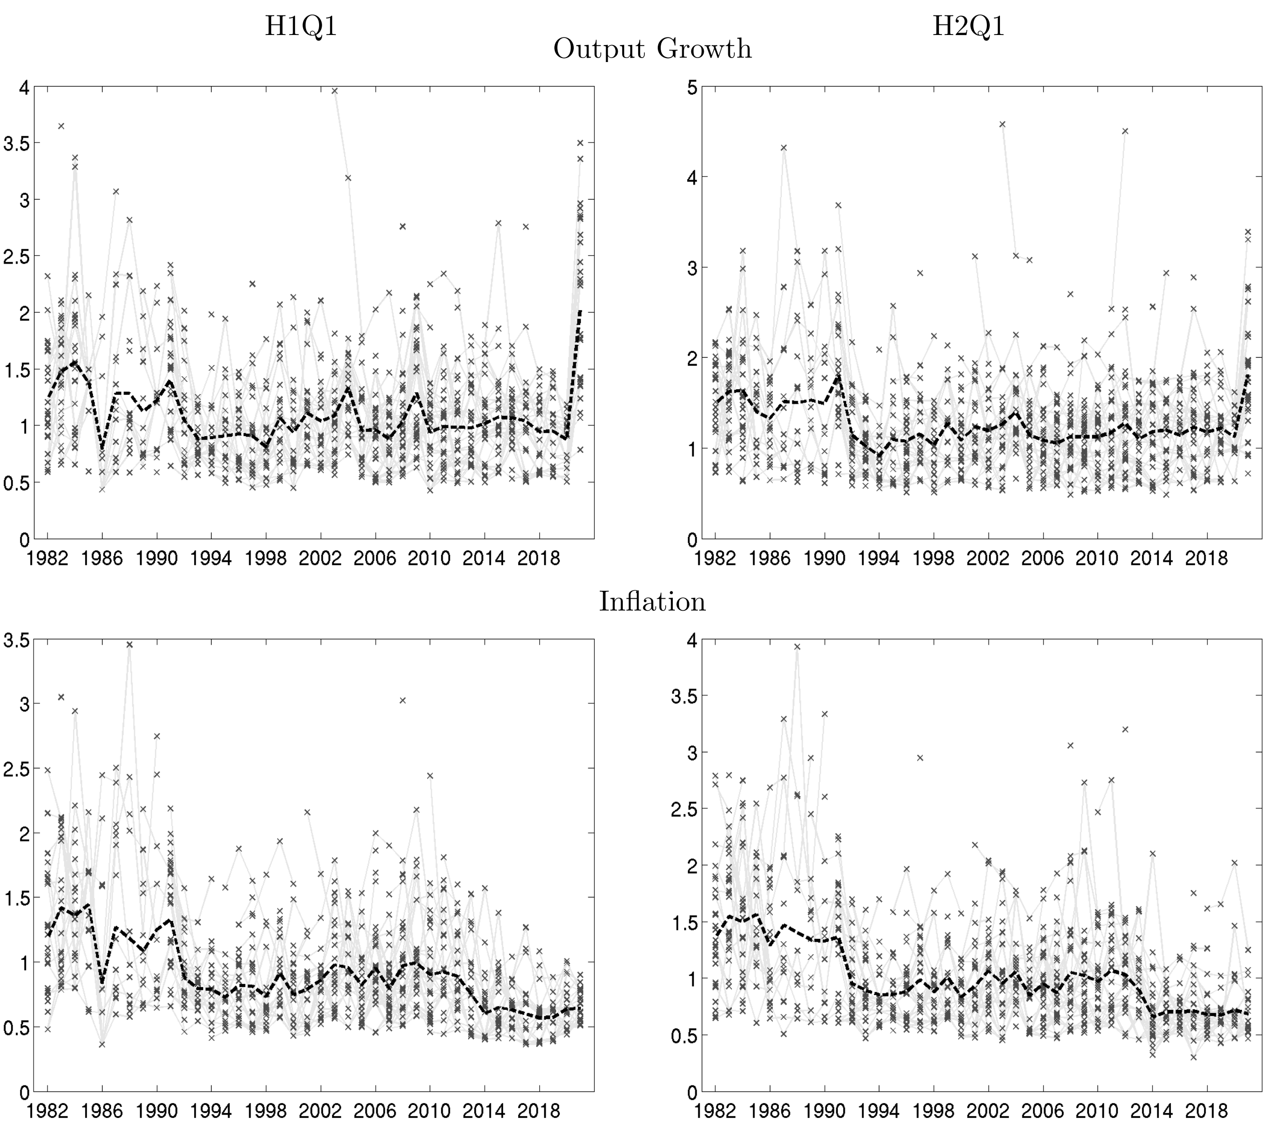


*Note*:

Figure D-7: Subjective uncertainty by individual respondent: Q3

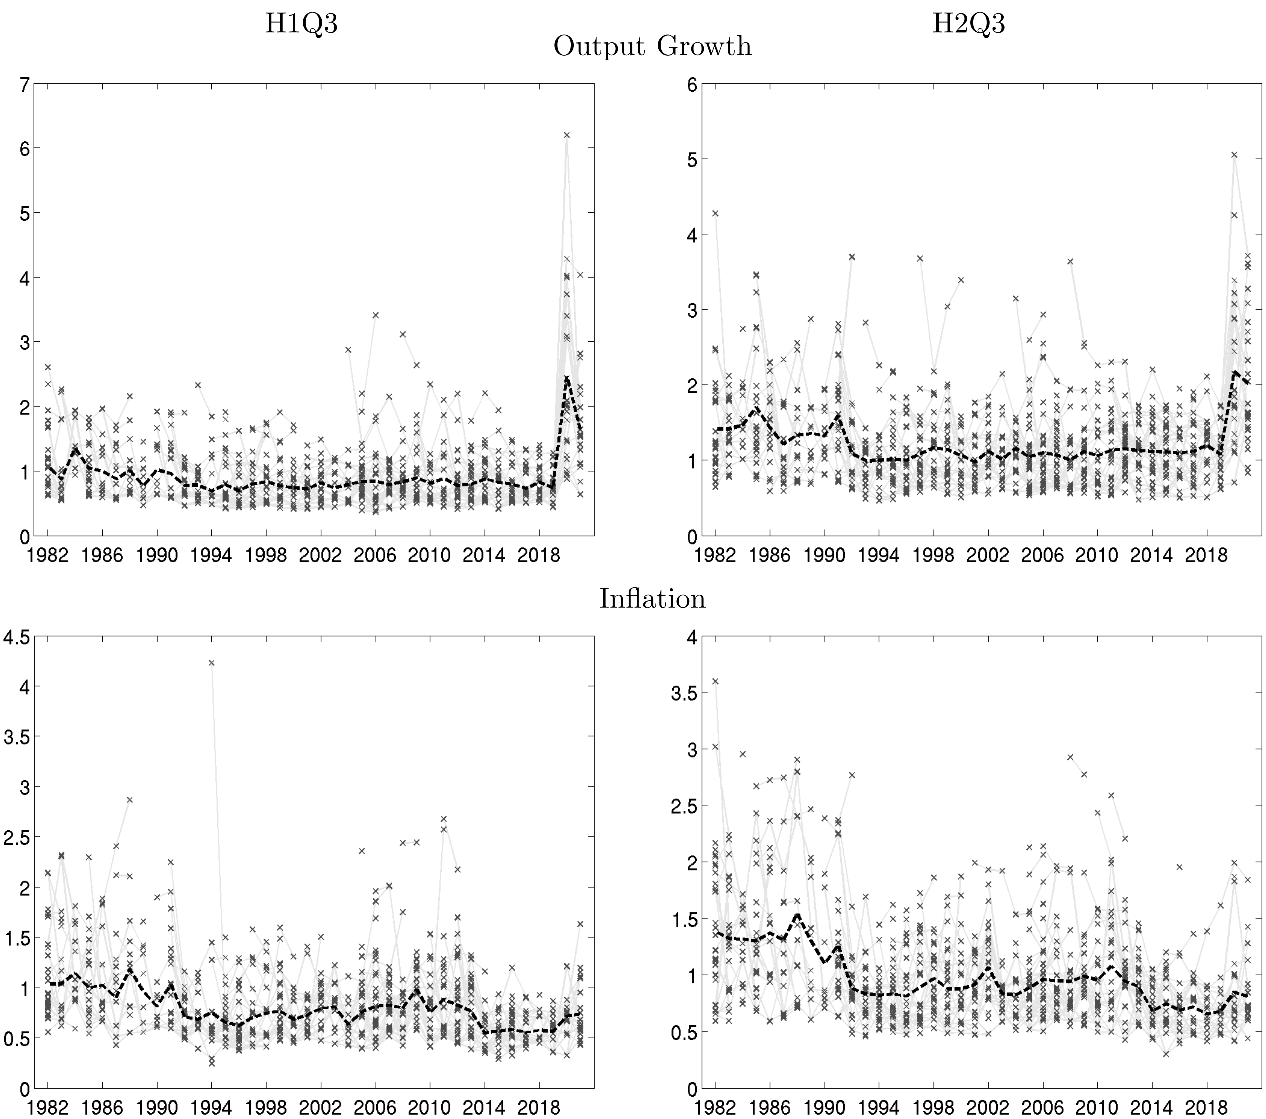


*Note*:

Figure D-8: Subjective uncertainty by individual respondent: Q4

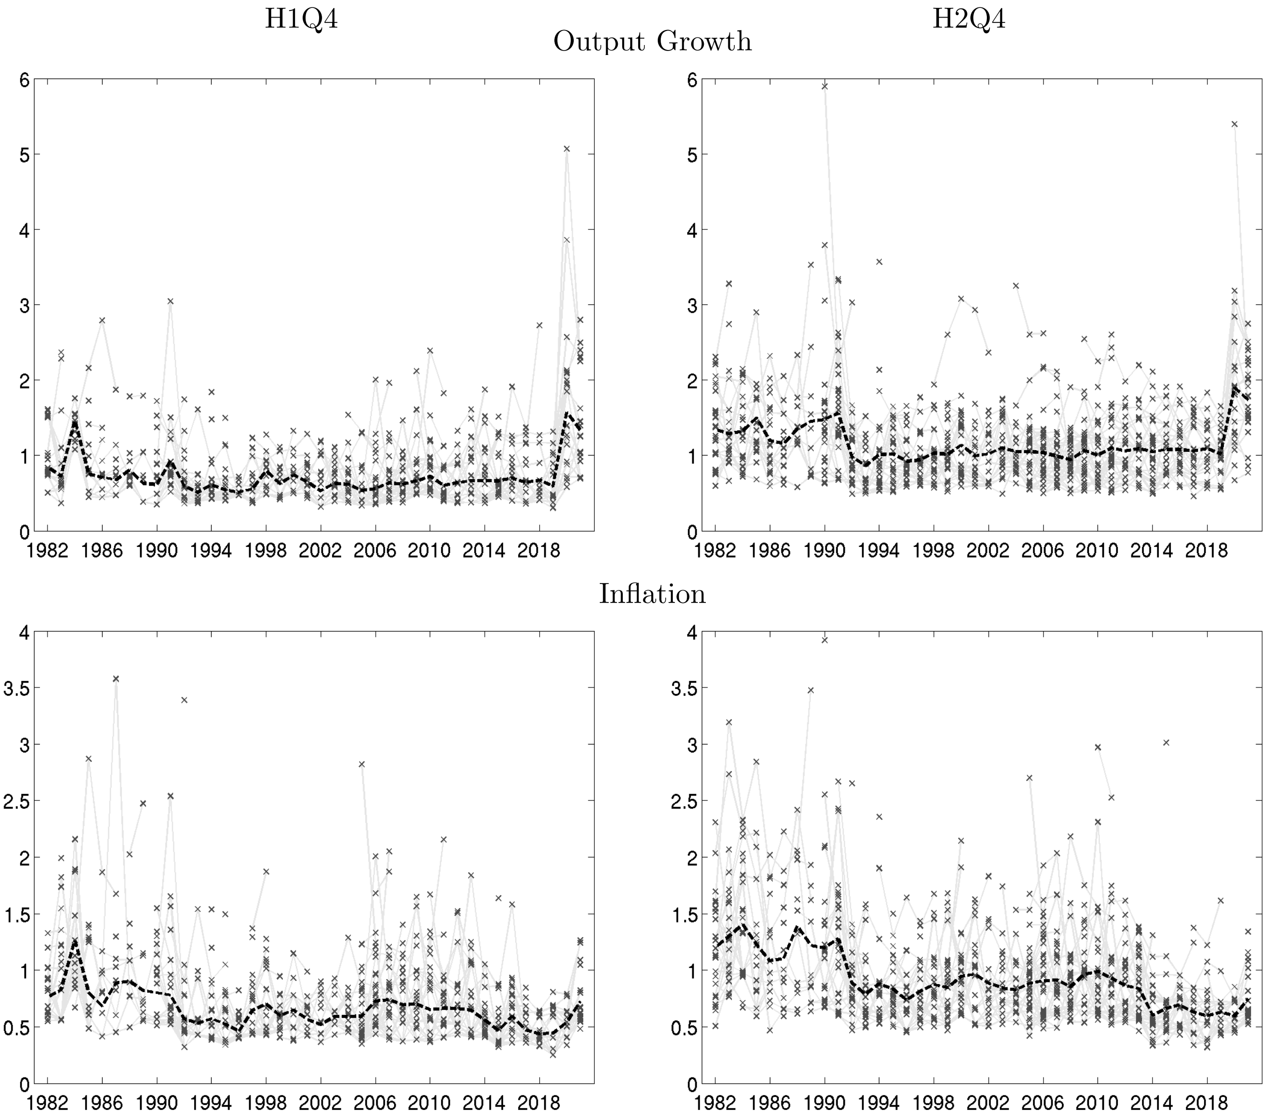


*Note*:

Figure D-9: Subjective uncertainty by individual respondent: IQRs

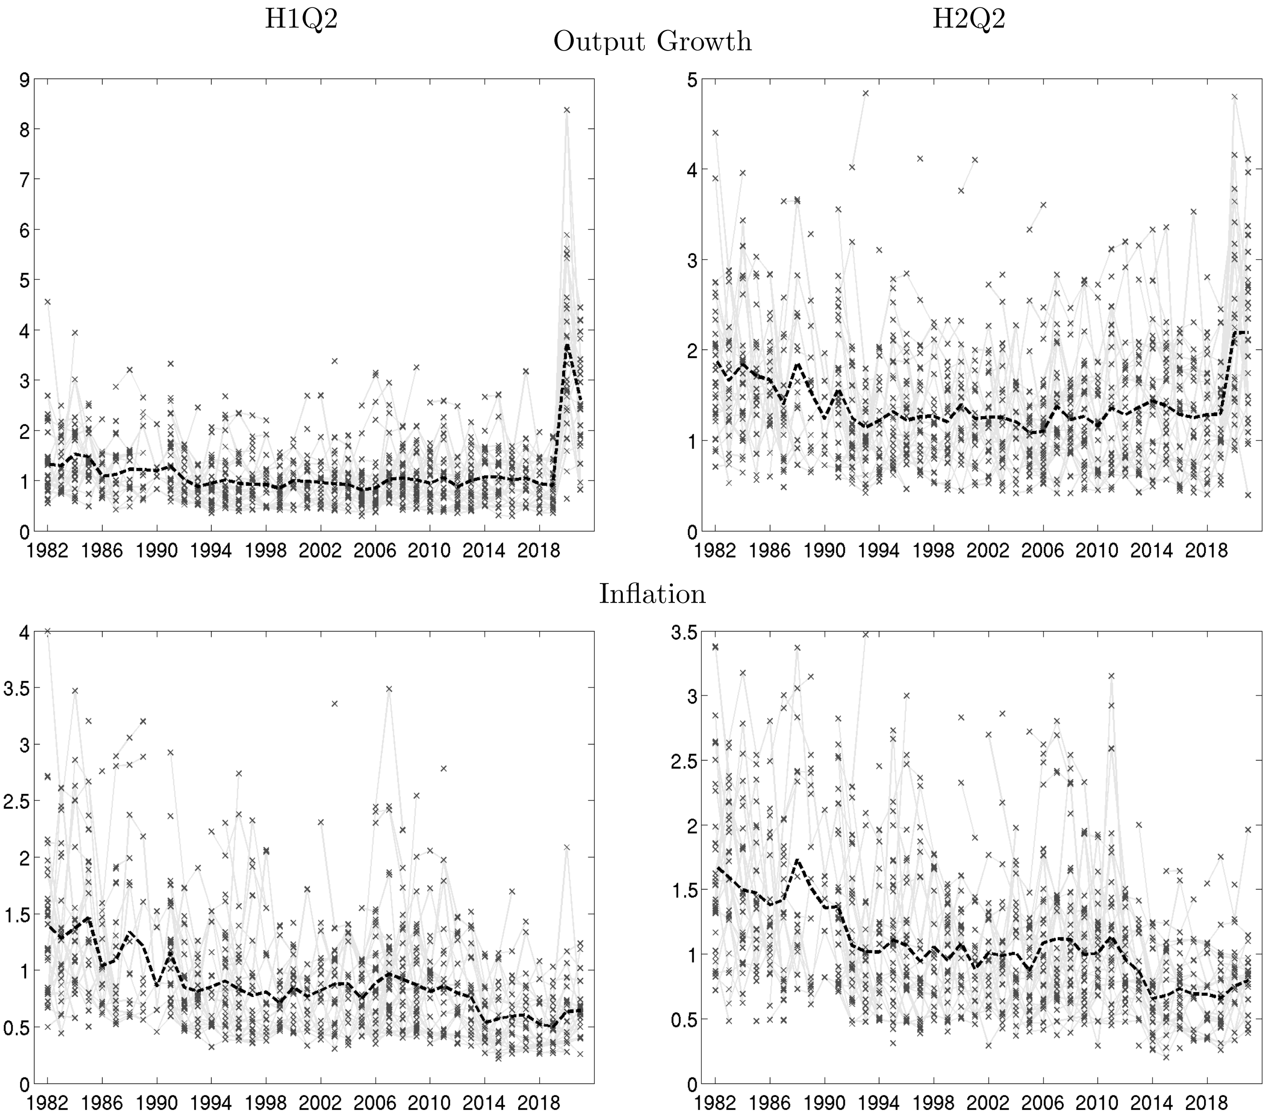


*Note*:

Figure D-10: Subjective uncertainty by individual respondent: Beta

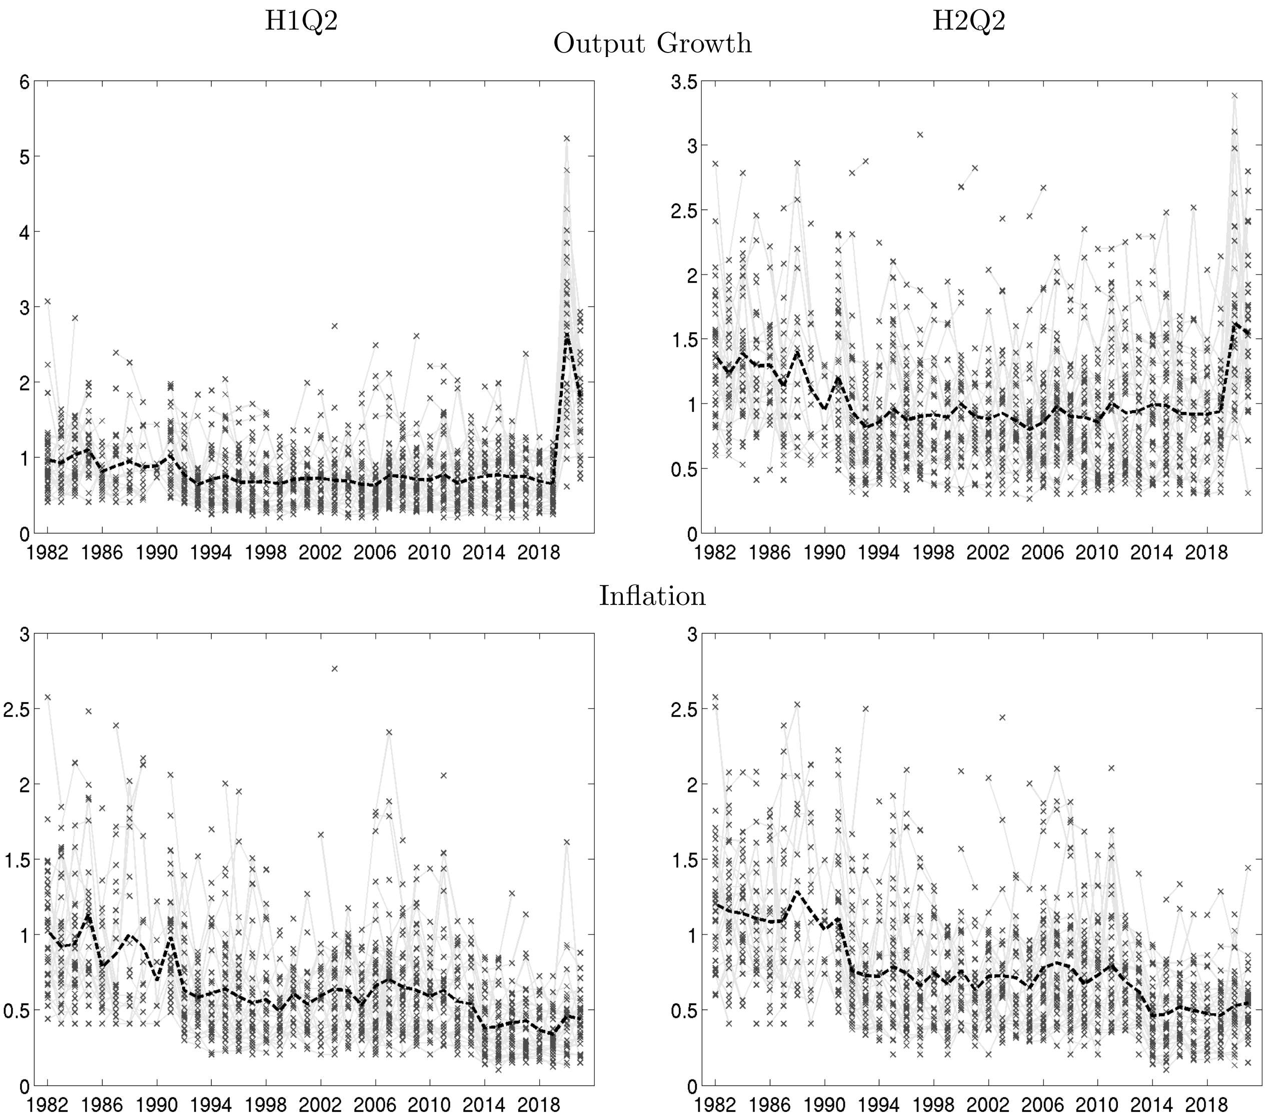


*Note*:

## E.C Subjective uncertainty and forecast accuracy

Figure D-11: Subjective Uncertainty and Forecast Accuracy: Mean vs Point Predictions

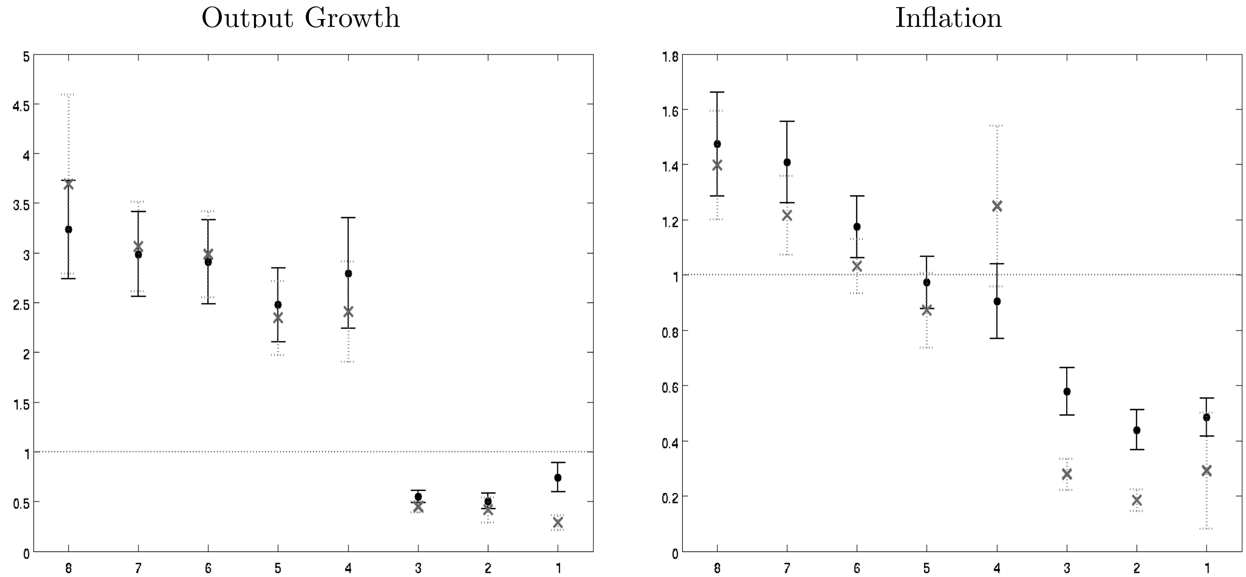

Figure D-12: Subjective Uncertainty and Forecast Accuracy: Different Samples

Output Growth Inflation

1982-2018

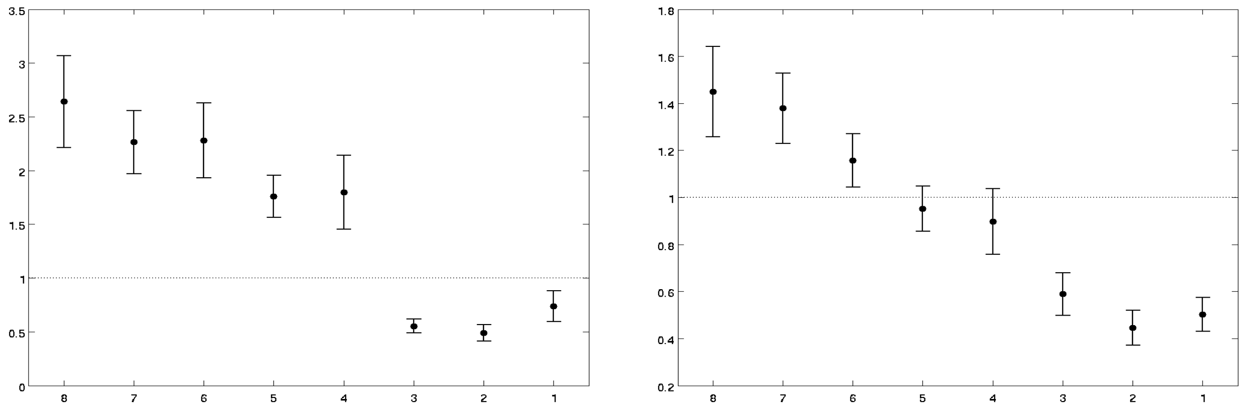

1992-2021

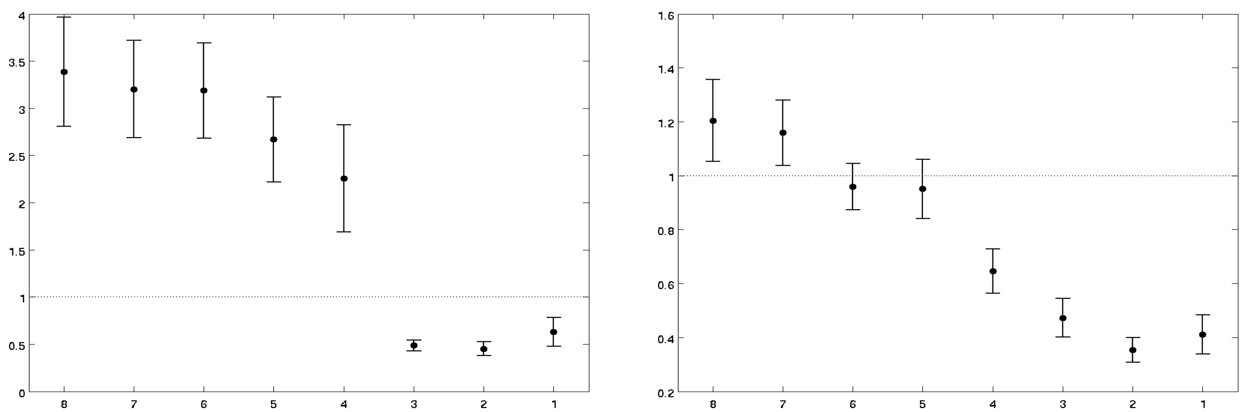

1992-2018

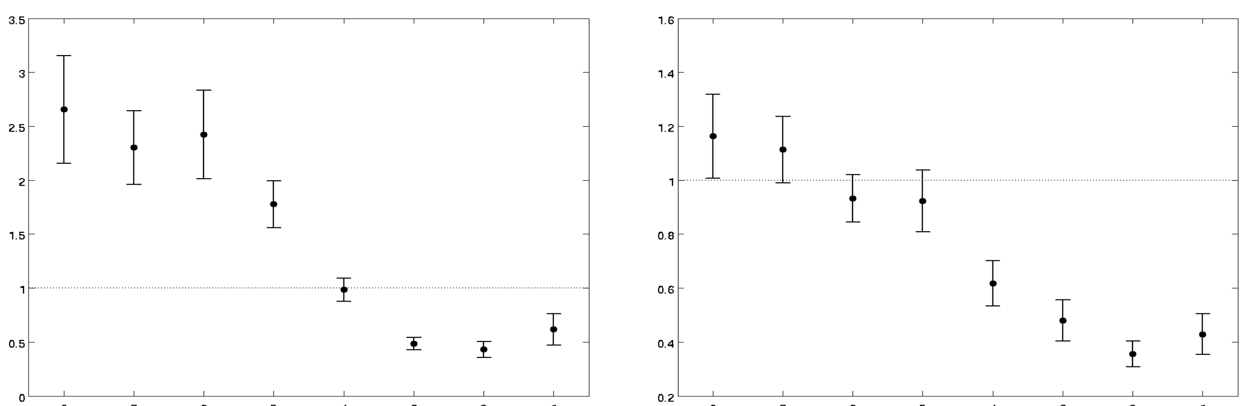

Figure D-13: Subjective Uncertainty and Forecast Accuracy: Using Std2

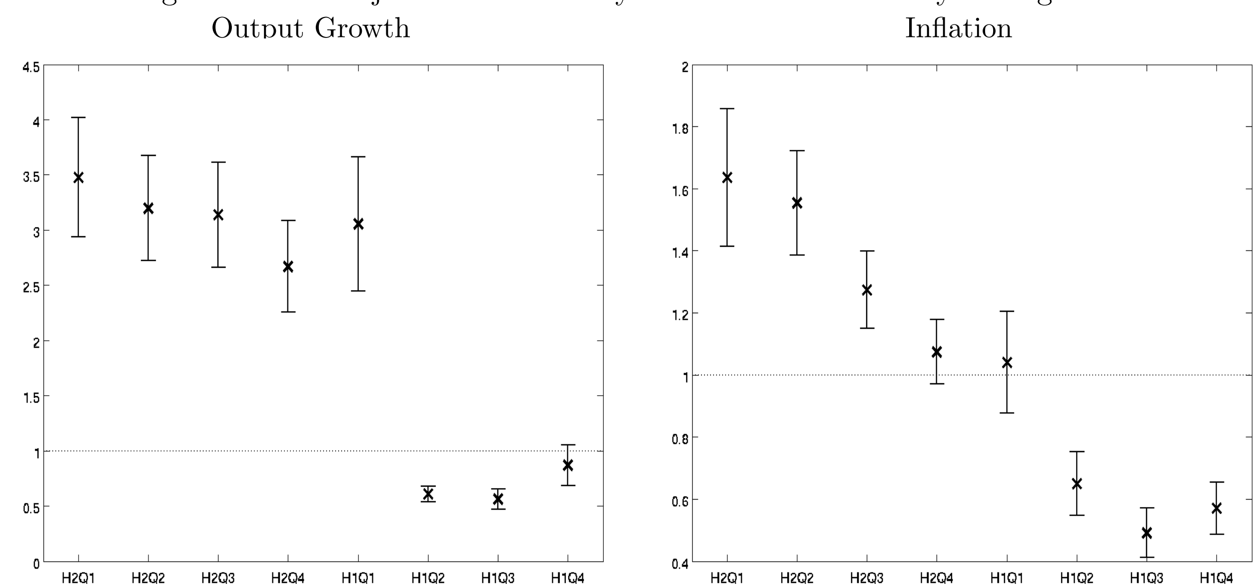

## Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Additional Results (Unweighted)

Figure D-14: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Regressions with Fixed Effects for both Mean and Point Forecasts

Output Growth | Inflation

Time Fixed Effects

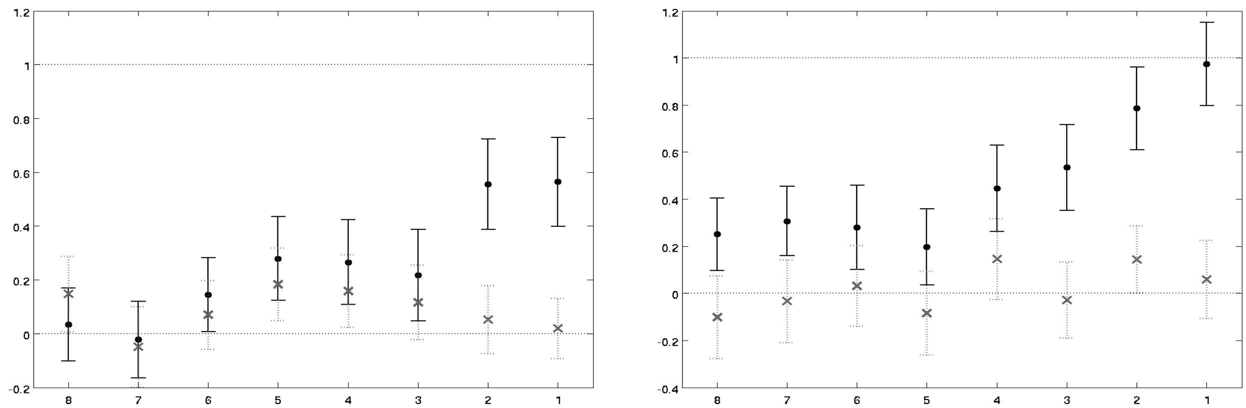


Forecaster Fixed Effects

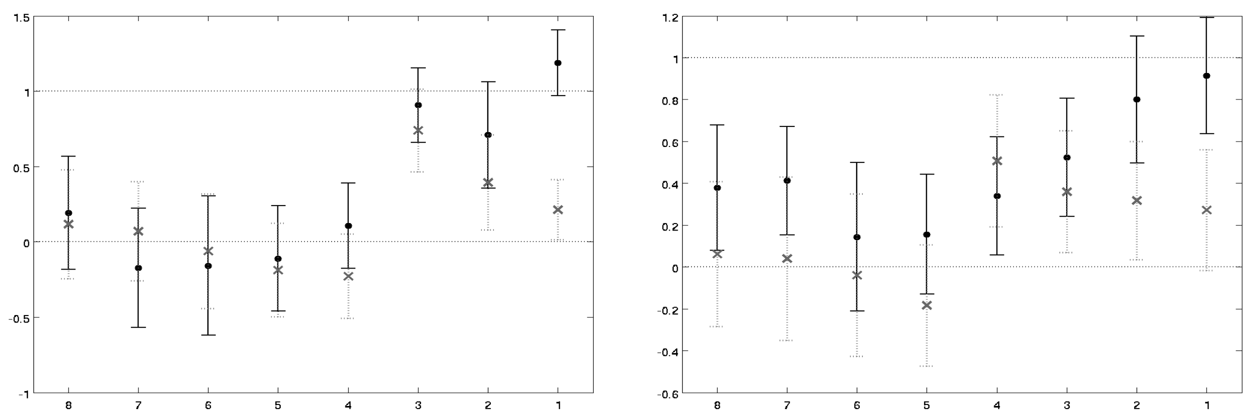


Time and Forecaster Fixed Effects

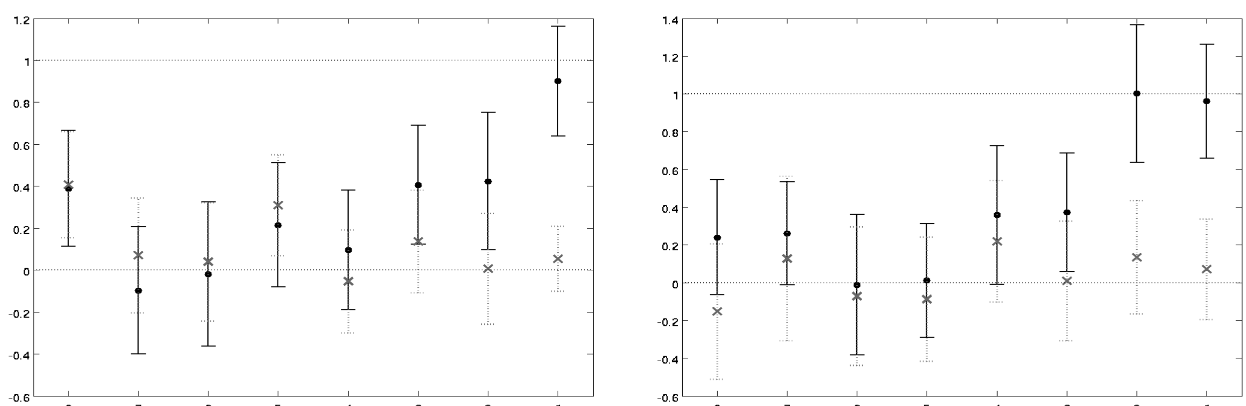

Figure D-15: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Output Growth; 1982-2018 Sample

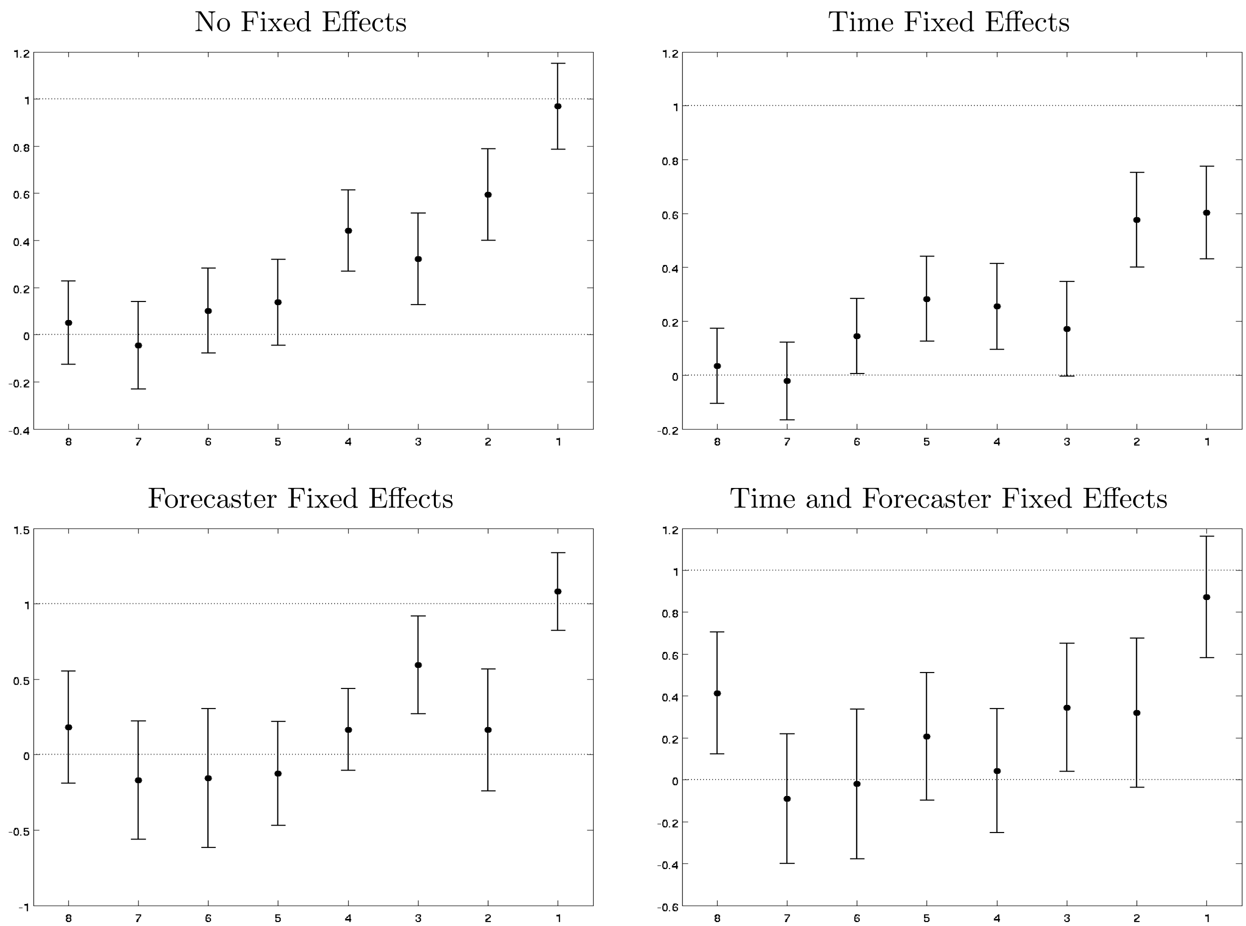

Figure D-16: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Inflation; 1982-2018 Sample

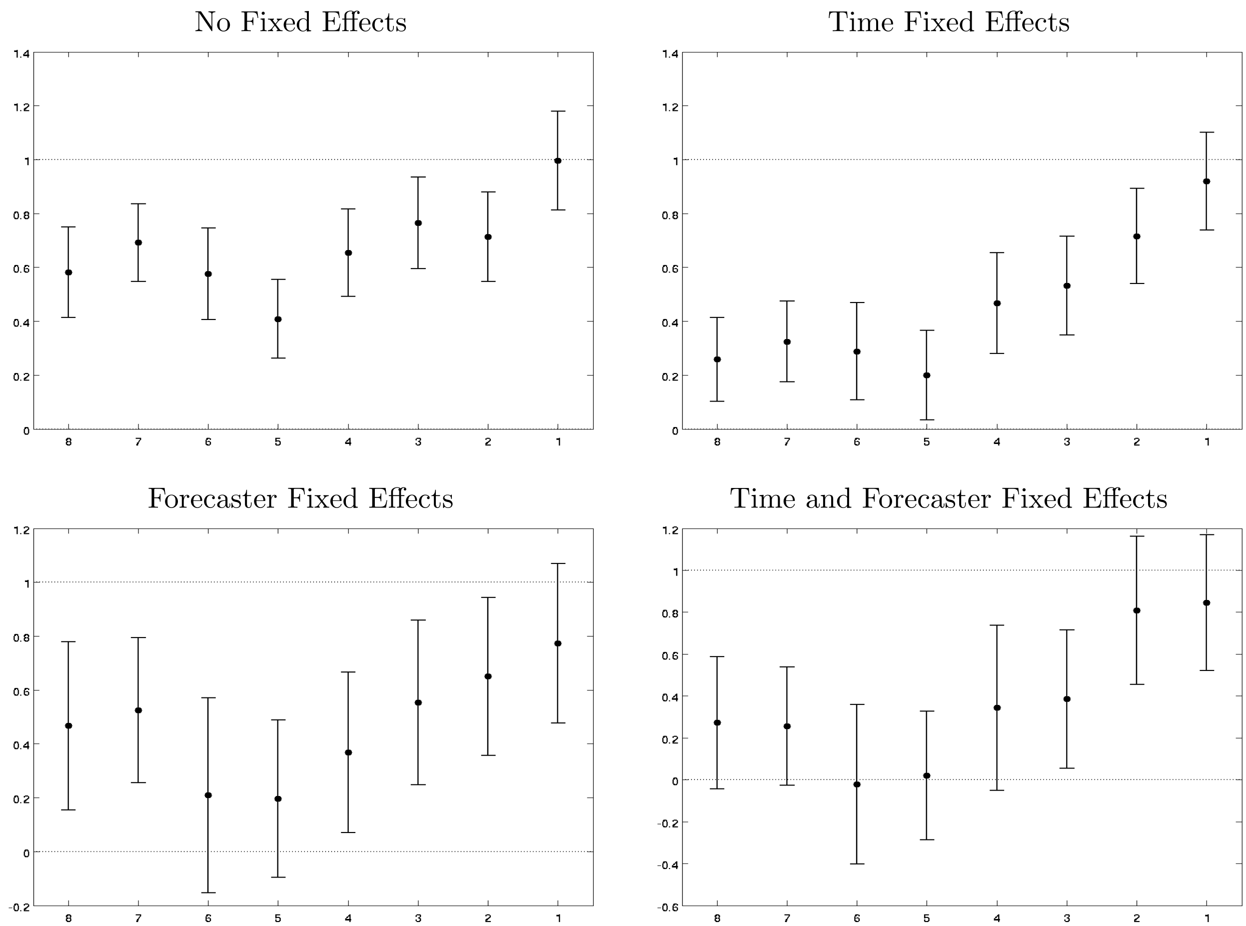

Figure D-17: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Output Growth; 1992-2021 Sample

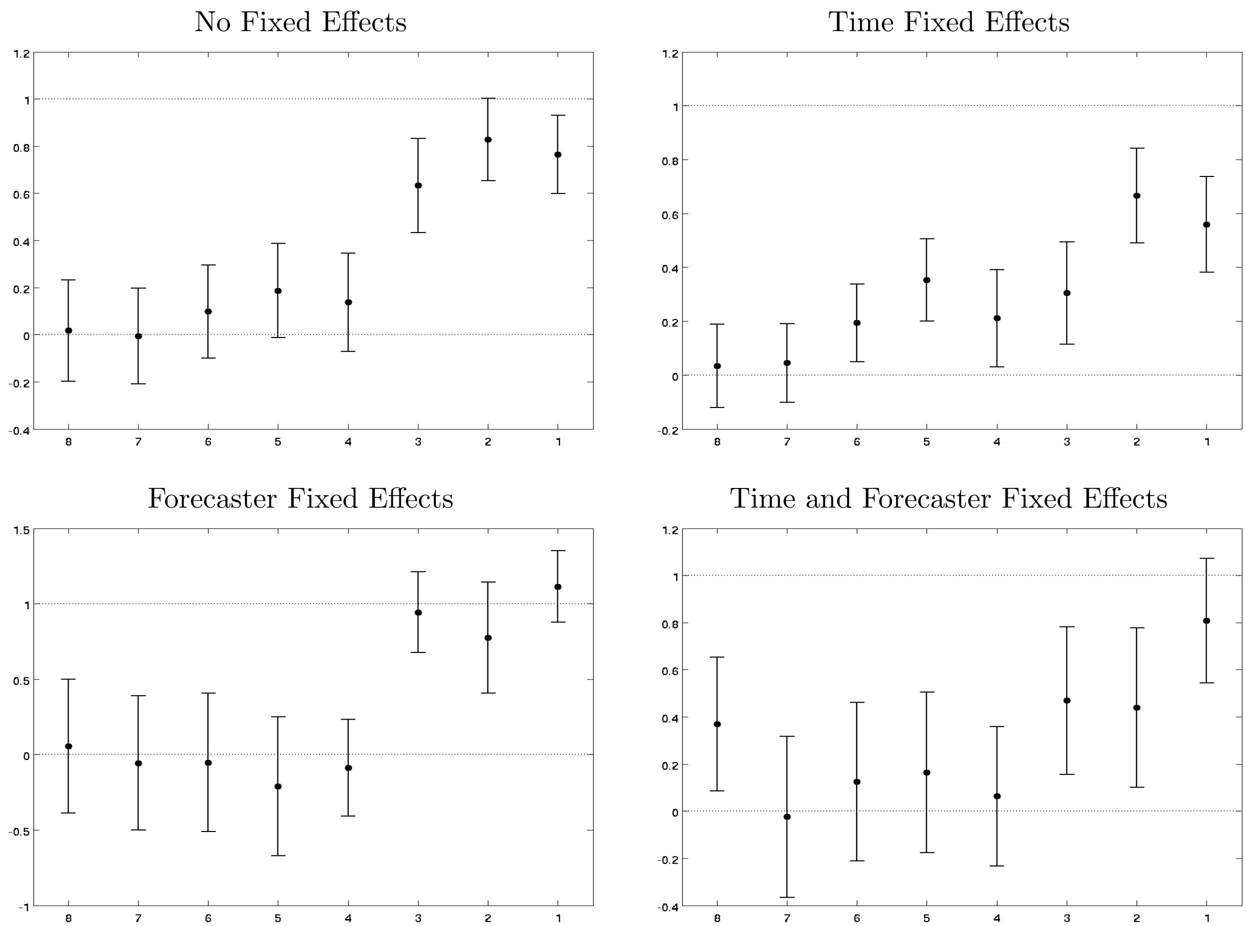

Figure D-18: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Inflation; 1992-2021 Sample

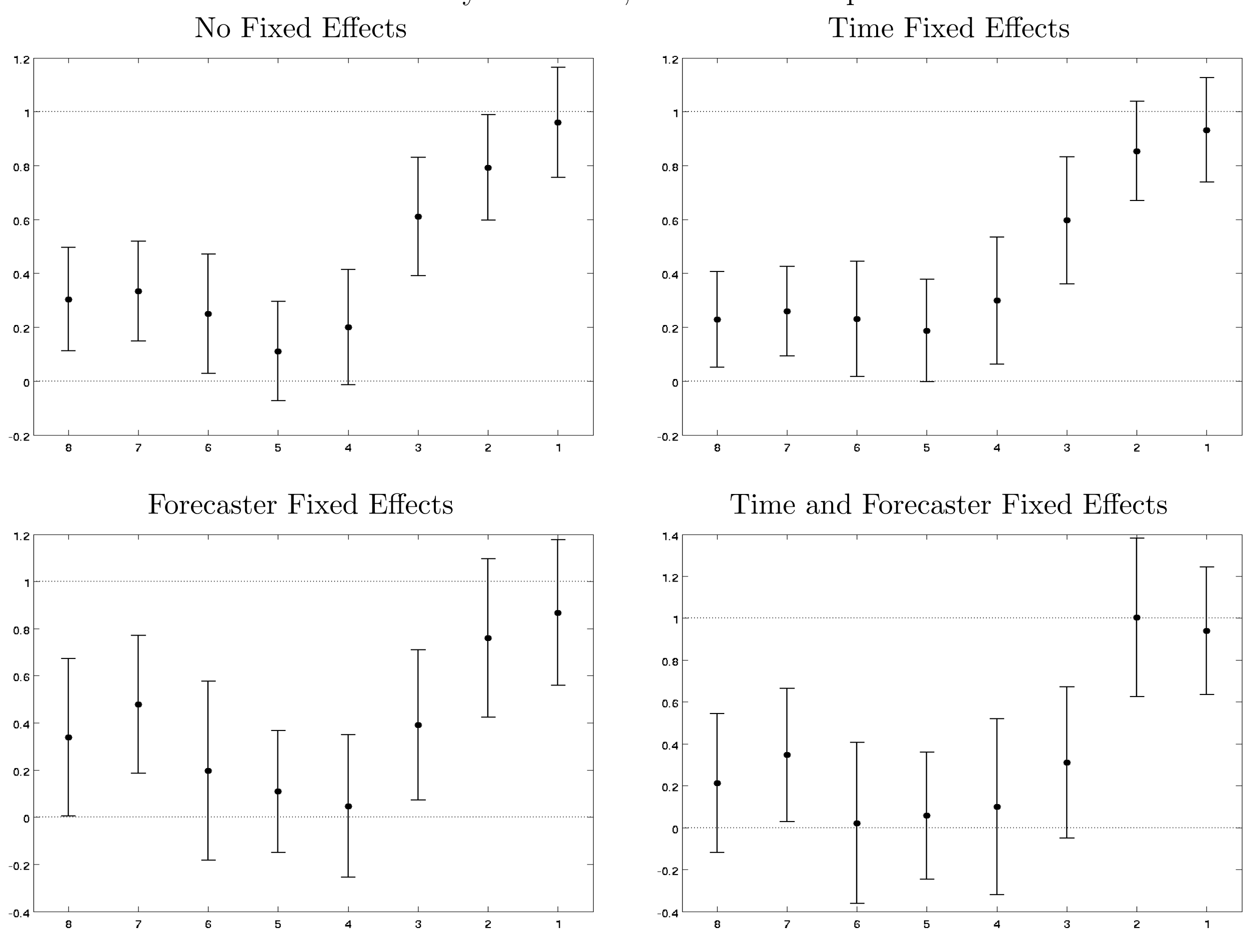

Figure D-19: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Output Growth; 1992-2018 Sample

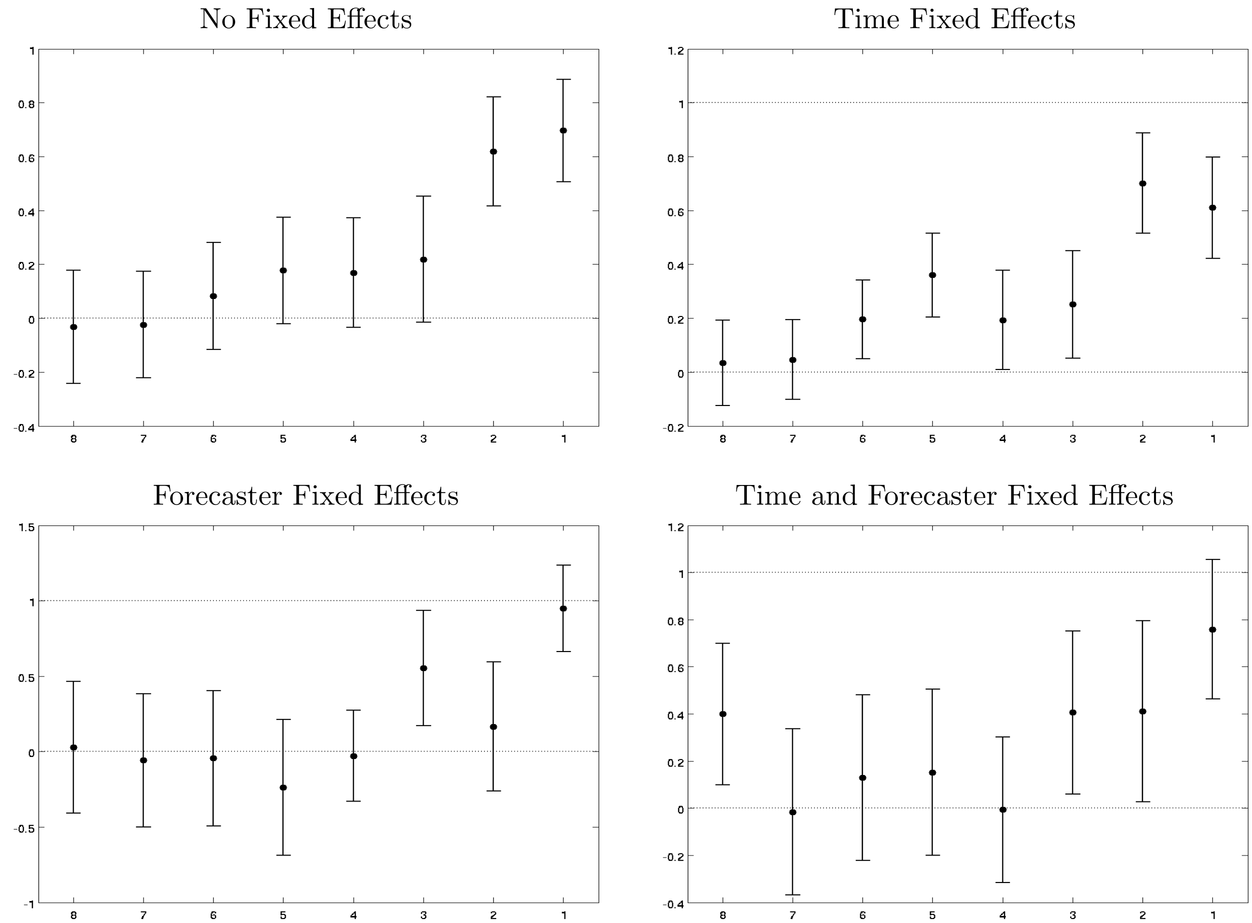

Figure D-20: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Inflation; 1992-2018 Sample

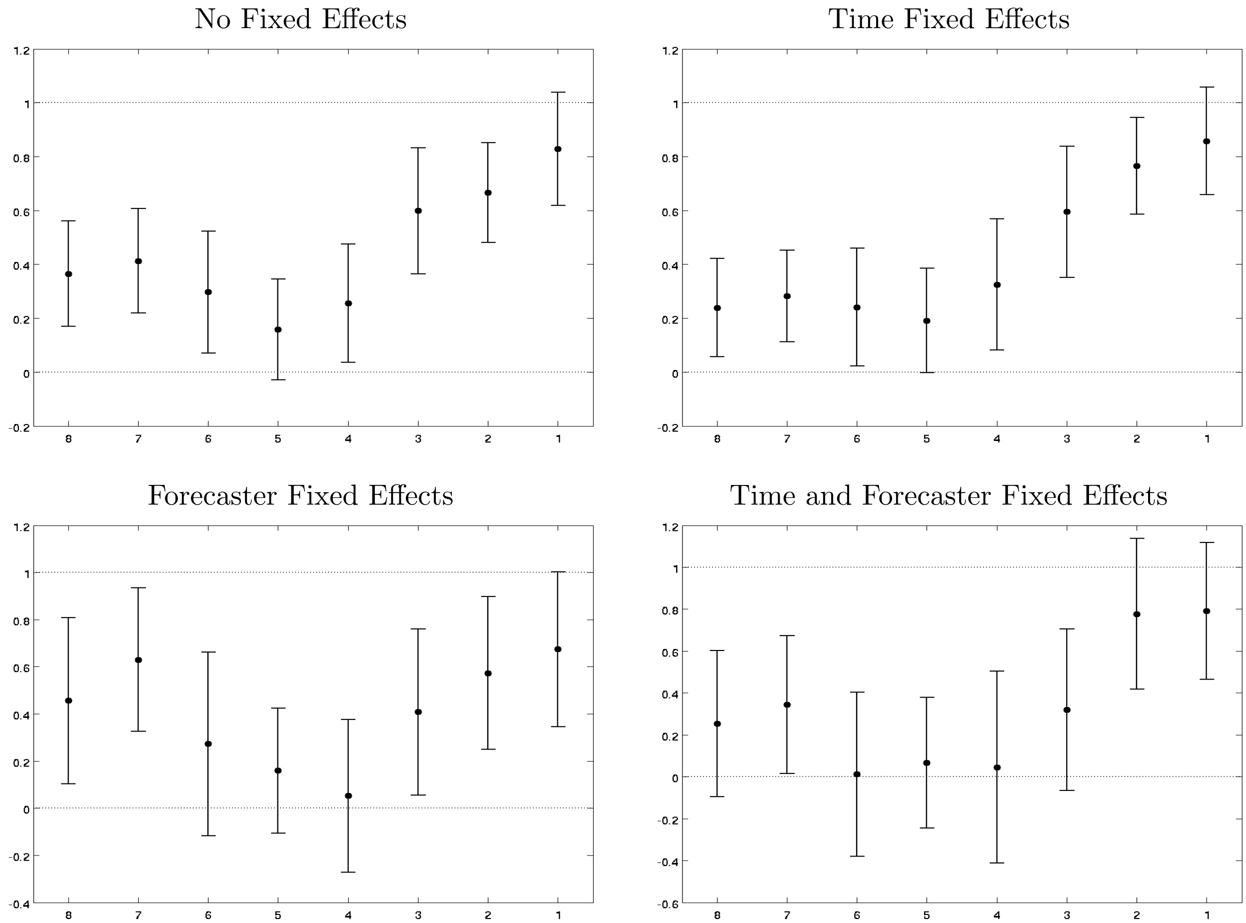

**Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Additional Results (Weighted)**

Figure D-21: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Mean vs Point Projections–Weighted

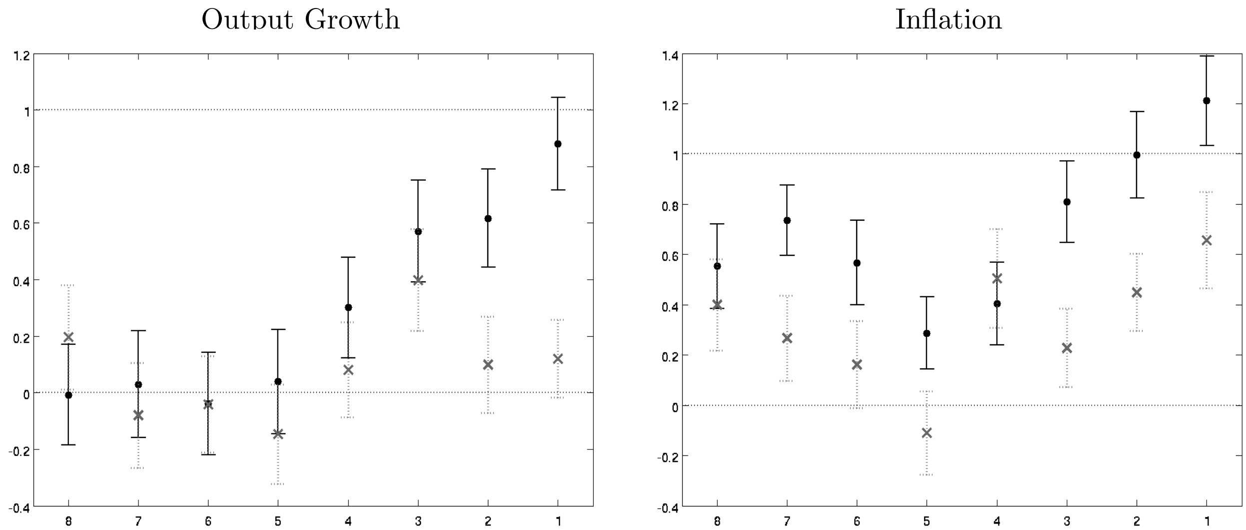

Figure D-22: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Regressions with Fixed Effects–Weighted

Output Growth | Inflation

Time Fixed Effects

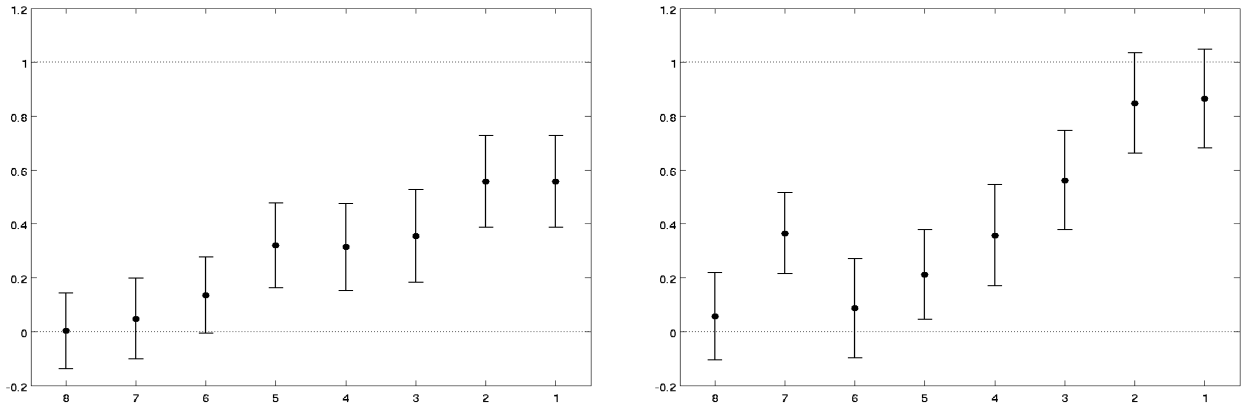


Forecaster Fixed Effects

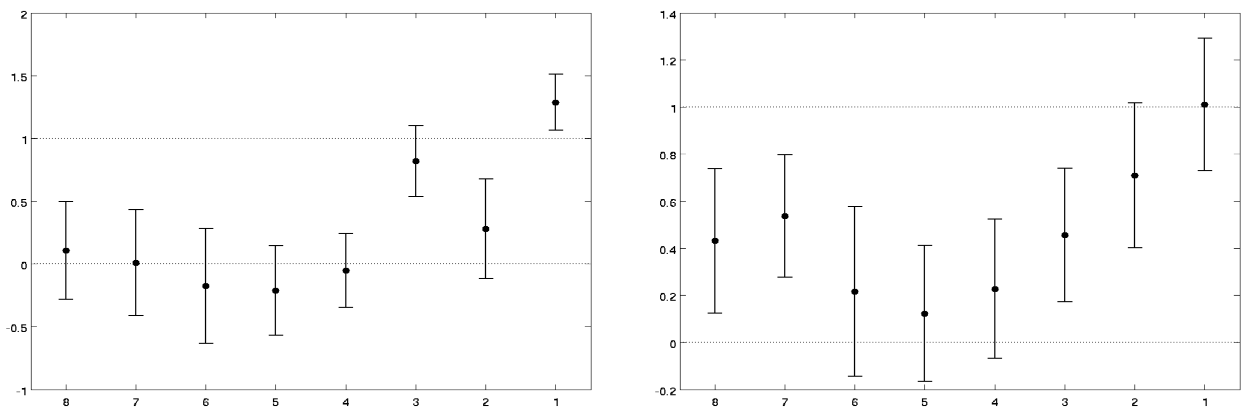


Time and Forecaster Fixed Effects

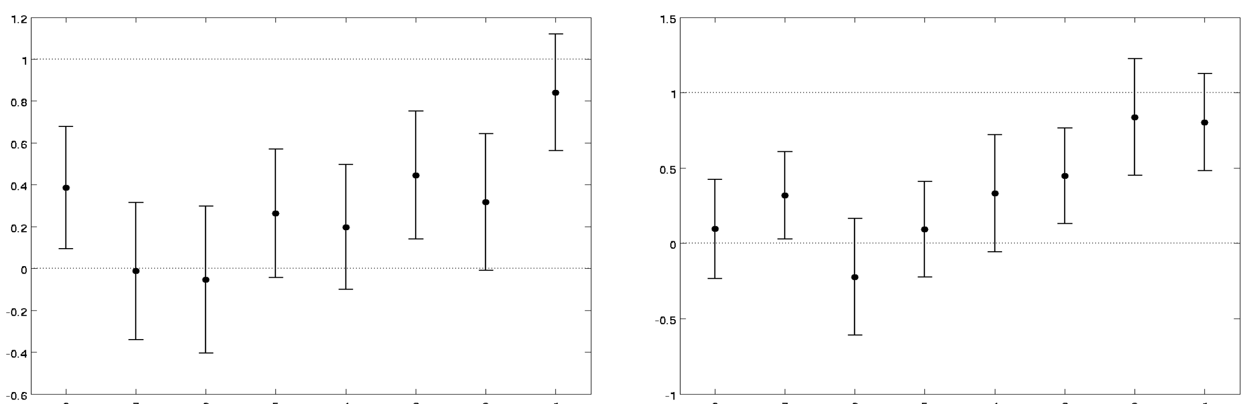

Figure D-23: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Regressions with Fixed Effects for both Mean and Point Forecasts–Weighted

Output Growth | Inflation

Time Fixed Effects

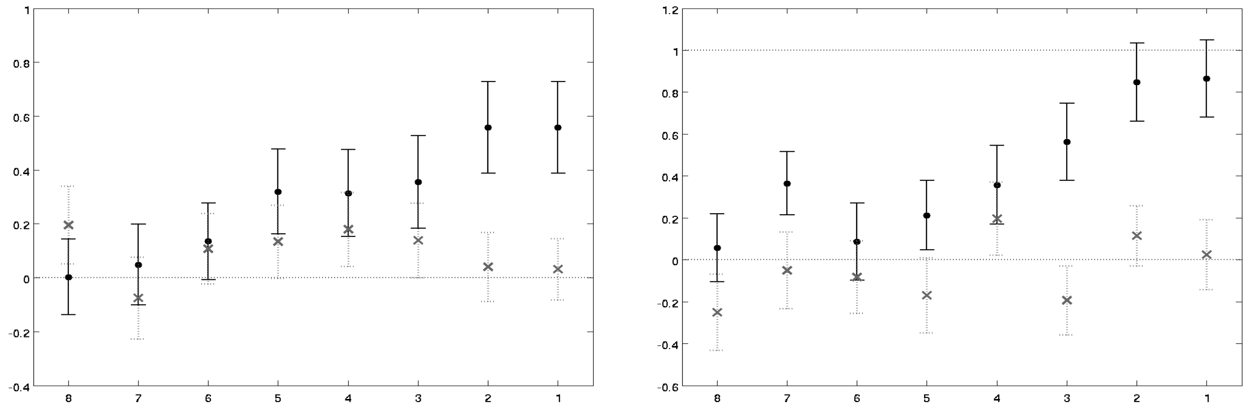


Forecaster Fixed Effects

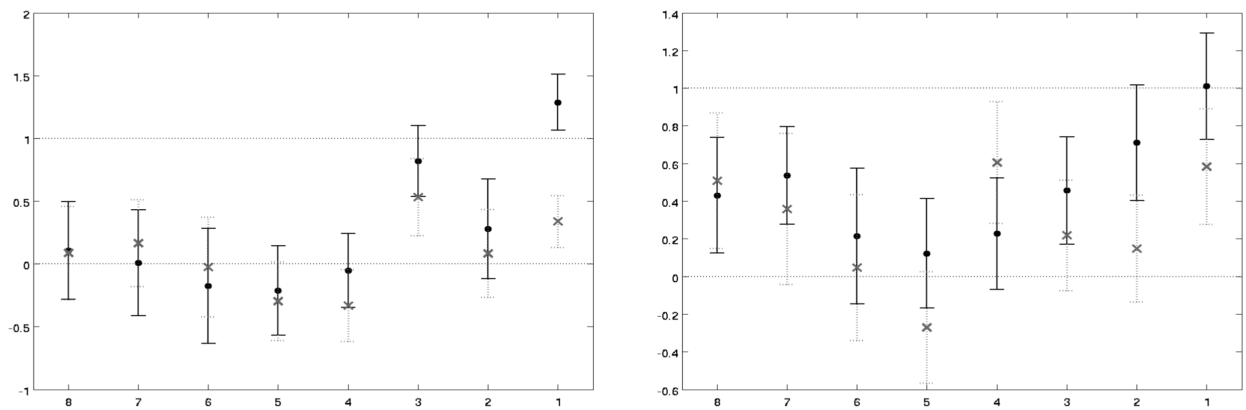


Time and Forecaster Fixed Effects

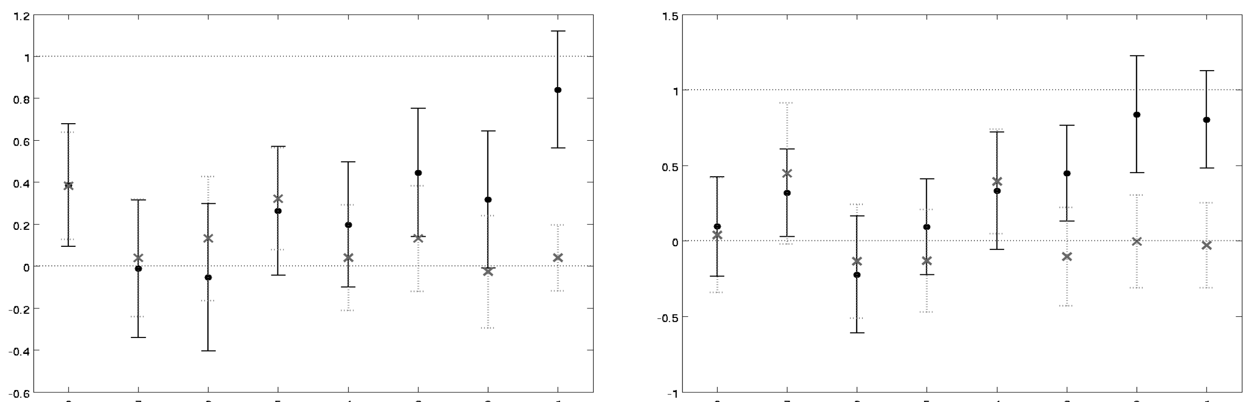

Figure D-24: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Output Growth; 1982-2018 Sample–Weighted

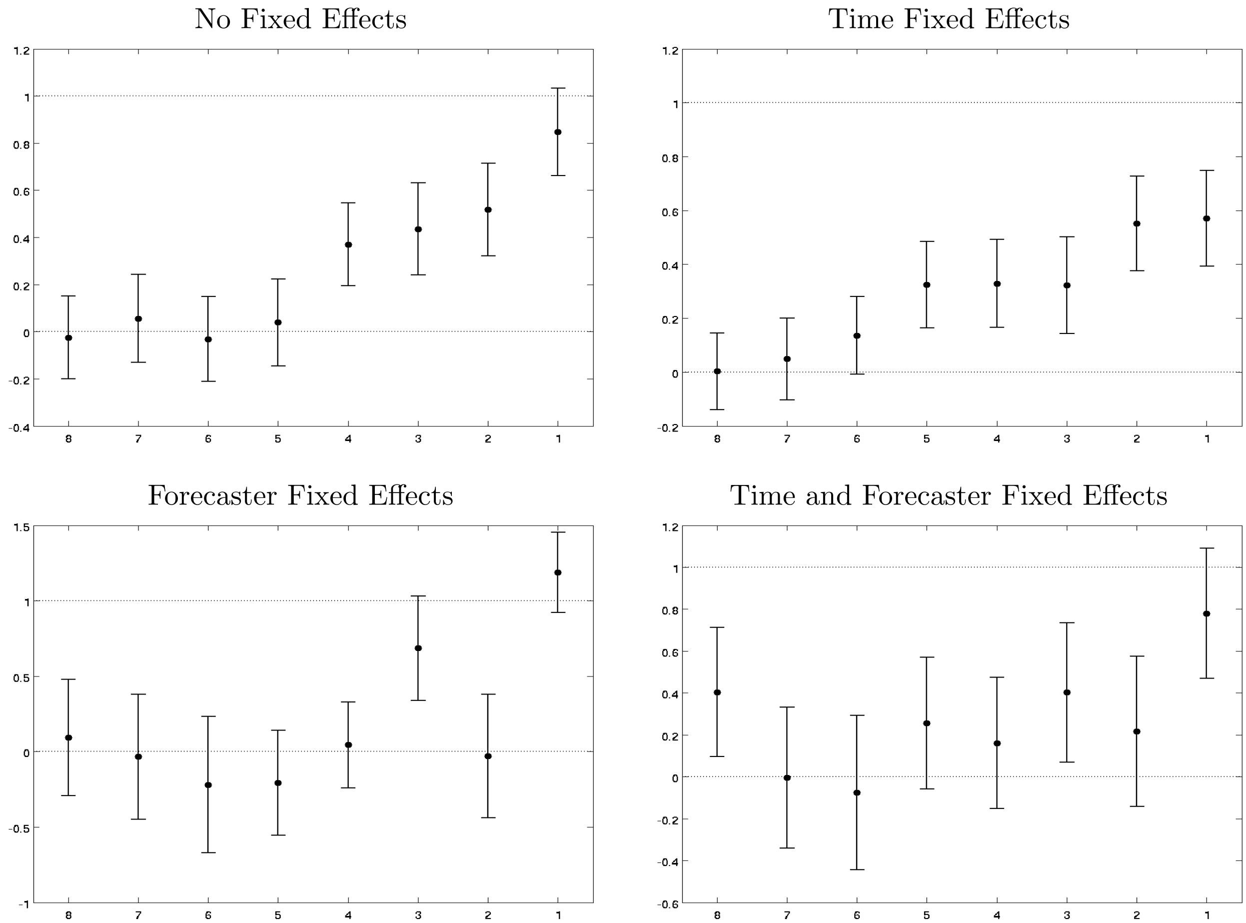

Figure D-25: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Inflation; 1982-2018 Sample–Weighted

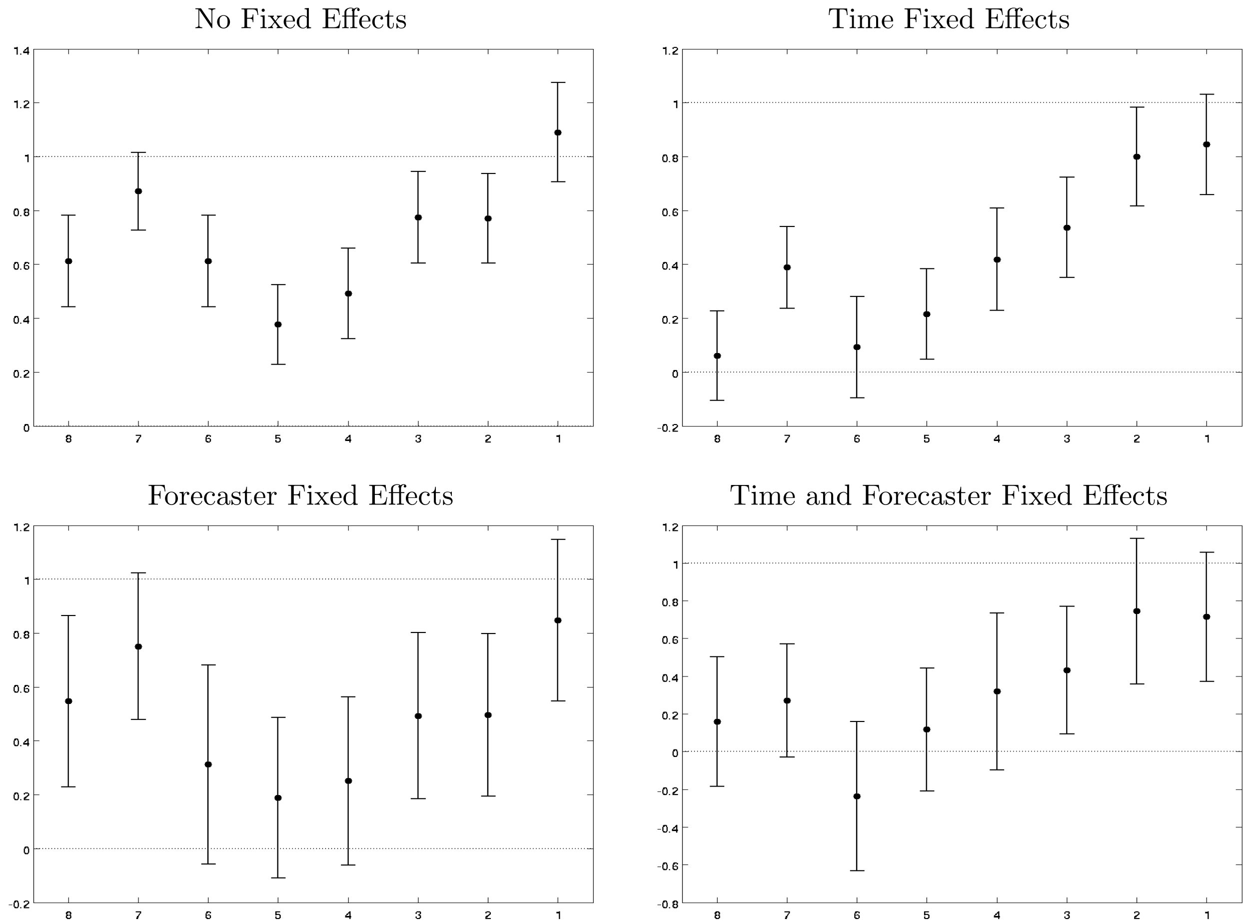

Figure D-26: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Output Growth; 1992-2021 Sample–Weighted

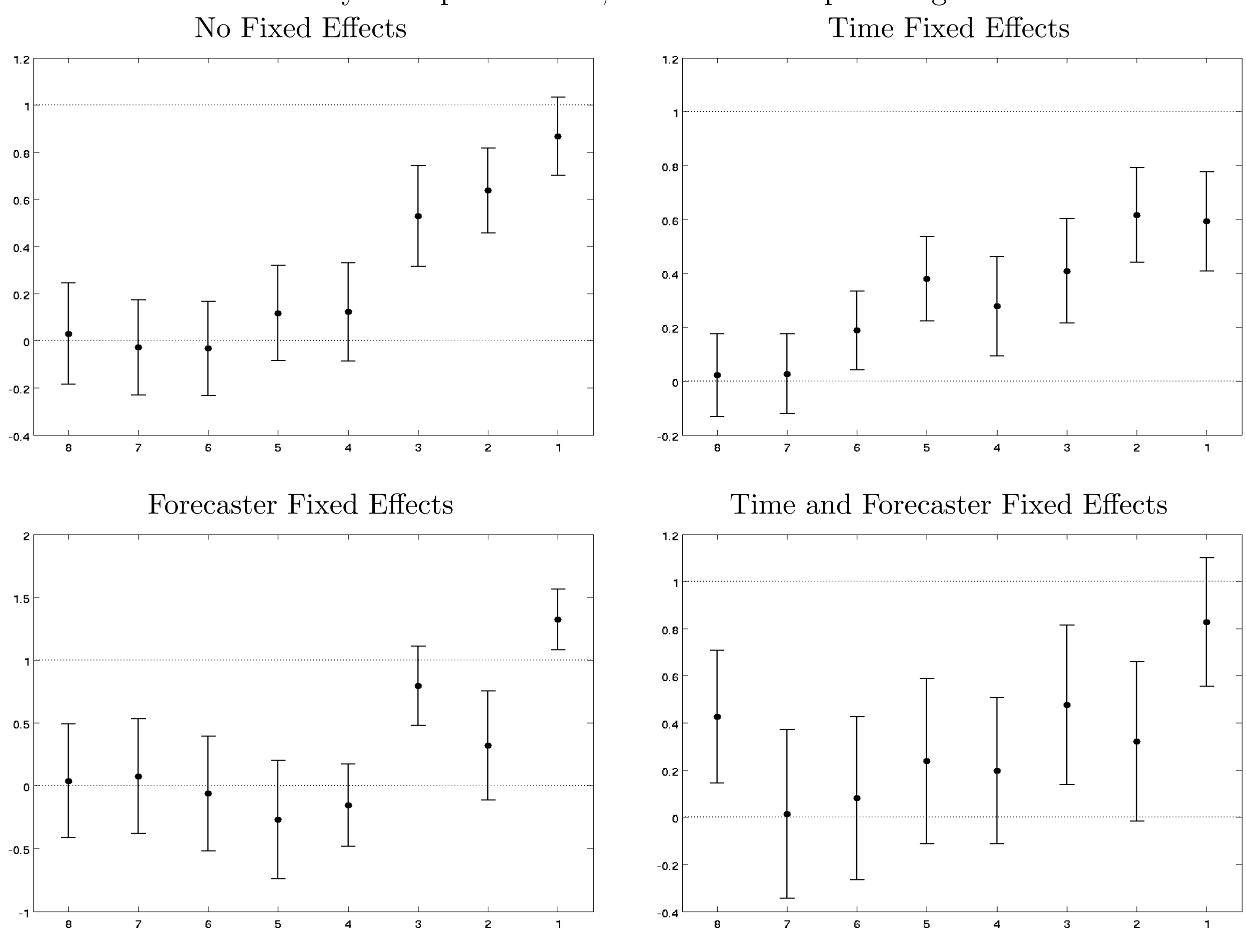

Figure D-27: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Inflation; 1992-2021 Sample–Weighted

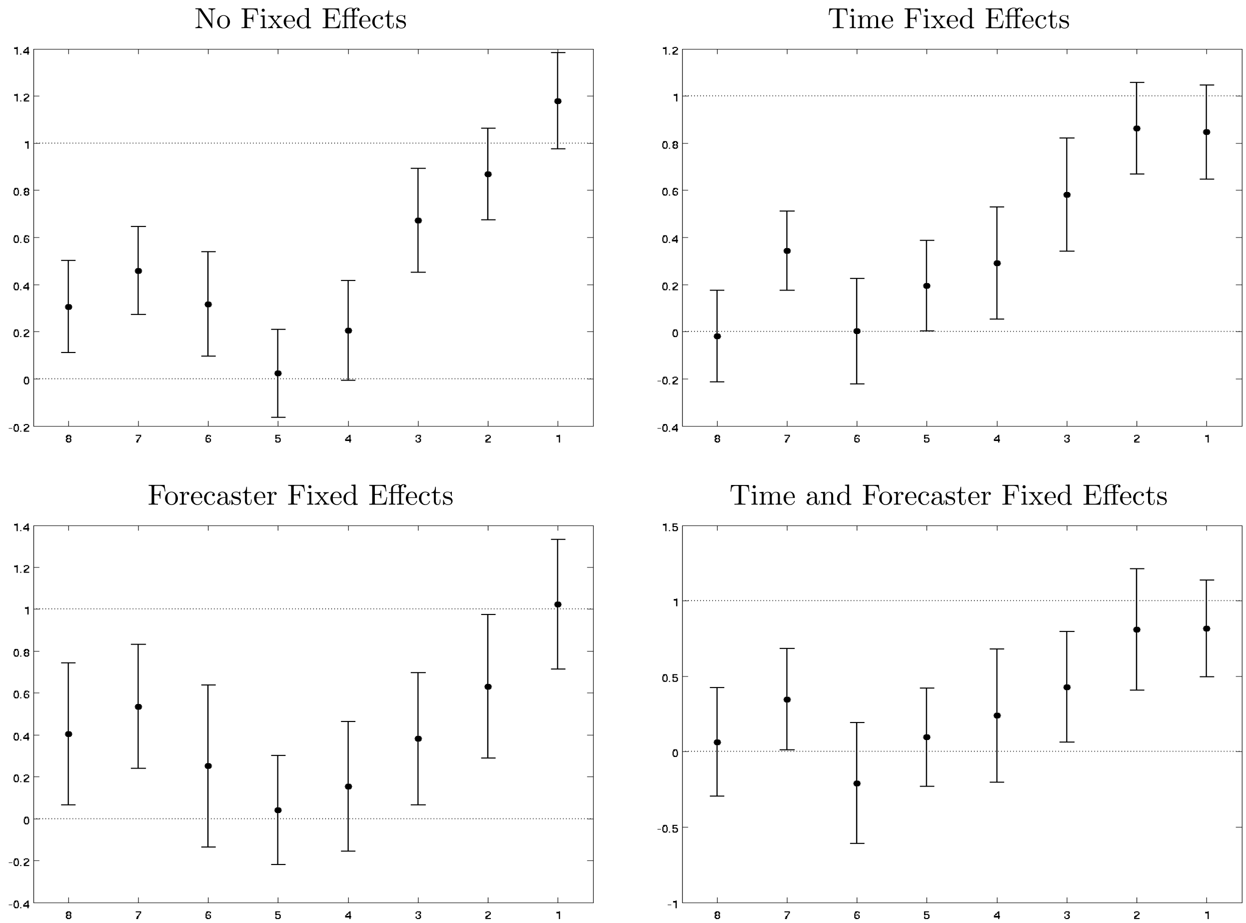

Figure D-28: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Output Growth; 1992-2018 Sample–Weighted

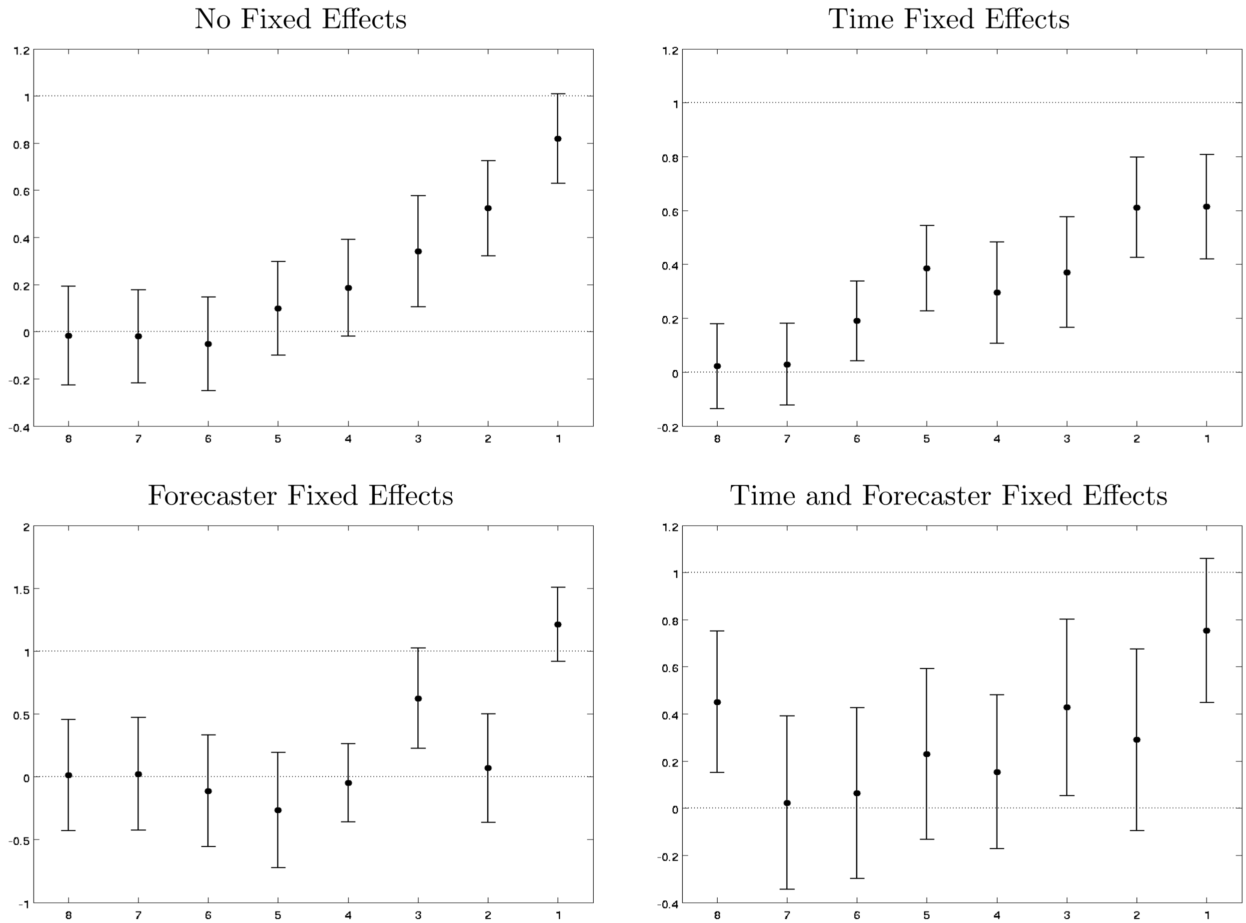

Figure D-29: Do Differences in Subjective Uncertainty Map into Differences in Forecast Accuracy? Inflation; 1992-2018 Sample–Weighted