



Contents lists available at ScienceDirect

Journal of Monetary Economics

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

Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion

Andrew J. Patton^{a,*}, Allan Timmermann^{b,1}

^a Department of Economics, Duke University, Box 90097, Durham, NC 27708-0097, United States

^b Rady School of Management, University of California San Diego, and CREATES, 9500 Gilman Drive, La Jolla, CA 92093-0553, United States

ARTICLE INFO

Article history:

Received 20 December 2008

Received in revised form

27 July 2010

Accepted 29 July 2010

ABSTRACT

Key sources of disagreement among economic forecasters are identified by using data on cross-sectional dispersion in forecasters' long- and short-run predictions of macroeconomic variables. Dispersion among forecasters is highest at long horizons where private information is of limited value and lower at short forecast horizons. Moreover, differences in views persist through time. Such differences in opinion cannot be explained by differences in information sets; our results indicate they stem from heterogeneity in priors or models. Differences in opinion move countercyclically, with heterogeneity being strongest during recessions where forecasters appear to place greater weight on their prior beliefs.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Differences in agents' beliefs play an important role in macroeconomic analysis. In models where agents observe noisy private and public information, heterogeneity in beliefs has been offered as an explanation for why monetary policy shocks can have real and persistent effects on output growth due to limited capacity for processing information (Woodford, 2003; Mackowiak and Wiederholt, 2009), infrequent updating of beliefs (Mankiw and Reis, 2002) or slow aggregate learning arising from dispersed information (Lorenzoni, 2009). Differences in beliefs also play a key role in determining the effect of public information signals in the literature on the social value of information in which agents have a coordination motive due to the strategic complementarity of their actions (Morris and Shin, 2002; Amador and Weill, 2009).²

While heterogeneity in agents' beliefs can be an important determinant of the "average opinion" about macroeconomic conditions, the reasons why agents disagree are not well understood. This is important since differences in agents' priors versus differences in their private information signals need not display the same degree of persistence and thus could influence macroeconomic dynamics very differently. Moreover, a better understanding of what determines heterogeneity in agents' beliefs and how this heterogeneity evolves over time can facilitate sharper tests of macroeconomic models for which subjective beliefs are a driver of economic activity. This point is highlighted by the sensitivity of some of the conclusions drawn from models with heterogeneous information to the type of signals observed by agents (e.g., Hellwig and Venkateswaran, 2009).

* Corresponding author. Tel.: +1 919 660 1849.

E-mail address: andrew.patton@duke.edu (A.J. Patton).

¹ Timmermann acknowledges support from CREATES, funded by the Danish National Research Foundation.

² Heterogenous beliefs and information asymmetries also play an important role in models of financial markets (Brunnermeier, 2001) and in foreign exchange rate models (Bacchetta and van Wincoop, 2006).

Hence, it is important to establish empirically why agents disagree and how this disagreement evolves over time and across different states of the economy. To this end, our paper explores survey data on differences in agents' subjective views at several forecast horizons and develops a novel approach for comparing these to model-based predictions of forecast dispersion. This allows us to address to what extent agents disagree, whether this disagreement has diminished over time, whether the primary source of disagreement is differences in models or differences in information, and how disagreements depend on the state of the economy.

A unique data set on forecasts of GDP growth and inflation for a given year recorded at different forecast horizons is used in the analysis. Fixing the time period and varying the forecast horizon allows us to identify the source of disagreement among forecasters. This holds because heterogeneity in private signals versus heterogeneity in model priors have very different effects on the cross-sectional dispersion of beliefs at long, medium and short forecast horizons. If instead the conventional approach of fixing the forecast horizon and varying the time period was used, variations in disagreement might simply reflect changes in the volatility of the underlying variable (e.g., the “Great Moderation”, [McConnell and Perez-Quiros, 2000](#)) and so the two effects would be difficult to disentangle.

Our analysis accomplishes five objectives. First, it documents empirically how the dispersion among agents' beliefs varies over time as well as across different forecast horizons, whether there is any relation between “average” beliefs and dispersion in beliefs, and how persistent differences in individual agents' beliefs tend to be.

Second, our analysis addresses the question from the title, namely the key sources of disagreement among forecasters. At the most basic level of analysis, agents can disagree either because of differences in their information signals or because of differences in their priors or models. Intuitively, in a stationary world differences among agents' information signals should matter most at short forecast horizons and less so at long horizons since variables will revert to their mean. Conversely, differences in prior beliefs about long-run inflation or output growth, or differences in their models of these quantities, should matter relatively more at long horizons where signals are weaker. If cross-sectional dispersion was only available for a single horizon it would not be possible to infer the relative magnitude of priors versus information signals underlying the cross-sectional dispersion. By studying the term-structure of dispersion in beliefs – i.e., differences in forecasts at long, medium and short horizons – the key sources of disagreement can thus be identified. Empirically, heterogeneity in information signals is found not to be a major factor in explaining the cross-sectional dispersion in forecasts of GDP growth and inflation: heterogeneity in priors or models is more important.

Third, our paper develops an approach for comparing the observed dispersion in subjective beliefs to that implied by a simple reduced-form model (whose moments are matched as closely as possible to the survey data) for how uncertainty about macroeconomic variables evolves. Our analysis uncovers evidence of “excess dispersion” in inflation forecasts at short horizons: at horizons of less than six months the observed disagreement between agents' predictions of inflation is high relative to the degree of uncertainty about inflation implied by our model. In contrast, the benchmark model does a good job at matching the empirically observed dispersion in views about GDP growth.

Fourth, our model is generalized to incorporate the effect of economic state variables on time-variation in the (conditional) cross-sectional dispersion measured at different horizons. Theoretical models such as [Van Nieuwerburgh and Veldkamp \(2006\)](#) suggest that macroeconomic uncertainty and dispersion in beliefs should be greater during recessions, where fewer information signals are received, than during expansions. Consistent with this, empirical evidence is found that differences in opinion move counter-cyclically, with disagreements being larger in recessions than in expansions. Our analysis suggests that greater differences in opinion are not due to increased heterogeneity in information signals but can be related to a shift toward agents putting more weight on model-based forecasts during recessions.

Fifth, our paper offers a variety of methodological contributions. A model is developed that incorporates heterogeneity in agents' prior beliefs and information sets while accounting for measurement errors and the overlapping nature of the forecasts for various horizons. A simulation-based method of moments (SMM) framework is employed for estimating the model parameters in a way that accounts for how agents update their beliefs as new information arrives. The shape of the cross-sectional dispersion in forecasts at different horizons is taken as the object to be fitted and SMM estimation is used to account for the complex covariance patterns arising in forecasts recorded at different (overlapping) horizons.

The plan of the paper is as follows. Section 2 takes a first look at the data. Section 3 presents our framework for modelling the evolution in the cross-sectional dispersion among forecasters across multiple forecast horizons in a way that allows for heterogeneity in agents' information and their prior beliefs. Section 4 develops our econometric approach. Empirical findings on the cross-sectional forecast dispersion are presented in Section 5 and Section 6 presents results for a model of time-varying dispersion. Section 7 concludes. Additional details on the estimation of the model are presented in a technical appendix.

2. A first look at the data

Before setting up a formal model, it is useful to take a first look at the data used in the analysis. Our data are taken from the Consensus Economics Inc. forecasts which comprise quantitative predictions of private sector forecasters. Each month survey participants are asked for their forecasts of a range of macroeconomic and financial variables for the major economies. The number of survey respondents varies between 15 and 33 during our sample, with an average of 26 respondents. Our analysis focuses on US real GDP growth and CPI inflation for the current and subsequent calendar year.

This gives 24 monthly next-year and current-year forecasts over the period 1991–2008 or a total of $24 \times 18 = 432$ monthly observations. The target dates for the predicted variable are labeled $t=1991, \dots, 2008$, while $h=1, \dots, 24$ months are the forecast horizons.

To document how the spread in individual forecasters' views around the mean depends on the forecast horizon, and to see how it evolves through time, Figs. 1 and 2 plot for each year in our sample the individual forecasts against the consensus (average) forecast at horizons of $h=1, 6, 12$ and 24 months. Movements in mean forecasts from year to year tend to be very smooth at the longest forecast horizon but are more volatile at shorter forecast horizons. Conversely, the cross-sectional spread in forecasts is highest at the 24-month horizon and is sharply reduced as the horizon shrinks, with the dispersion being particularly low at the one-month horizon. Since agents' information signals can be expected to be of less value at the long horizons where disagreement seems to be greatest, these plots provide an early indication that differences in opinion are not primarily driven by differences in information.

In the heterogeneous information approach to macroeconomics, differences in information are a key to the formation of the "average opinion" about macroeconomic conditions. It is therefore of interest to see whether there is a relation between the mean forecast and the dispersion in beliefs. To this end, Table 1 presents the correlation between the consensus forecast and the dispersion in forecasts for each horizon. For GDP growth a strong negative correlation emerges – with 23 of 24 correlation estimates being negative and 14 being significant at the 10% level – indicating higher dispersion in beliefs during years with low economic growth, i.e., countercyclical movements in disagreements about GDP growth. Conversely, for inflation, a positive relation emerges between the dispersion in beliefs and the consensus view – with 22 of 24 correlation estimates being positive and 6 being significant at the 10% level – suggesting that dispersion grows with the average expected inflation rate.

Further insight into the sources of differences in opinion can be gained from studying the extent to which individual forecasters are regularly above or below the mean forecast. Differences in prior beliefs might suggest persistent patterns in individual forecasters' optimism or pessimism relative to the average forecaster, whereas differences in private information are perhaps more suggestive of short-lived differences. As a first illustration, Fig. 3 plots for all horizons the time-series average of four individual forecasters' positions in the cross-sectional distribution of forecasts (with 0.1

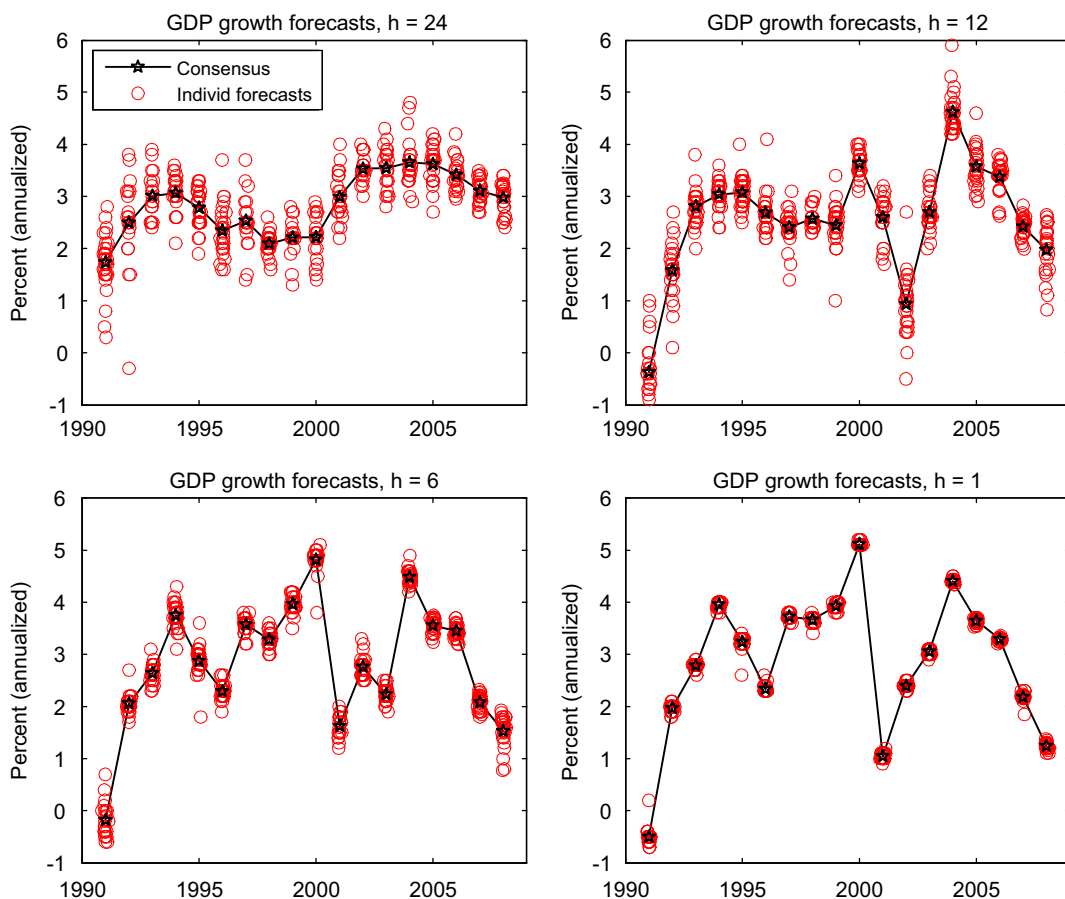


Fig. 1. Consensus and individual forecasts for GDP growth over the period 1991–2008, for four forecast horizons (24 months, 12 months, 6 months and one month).

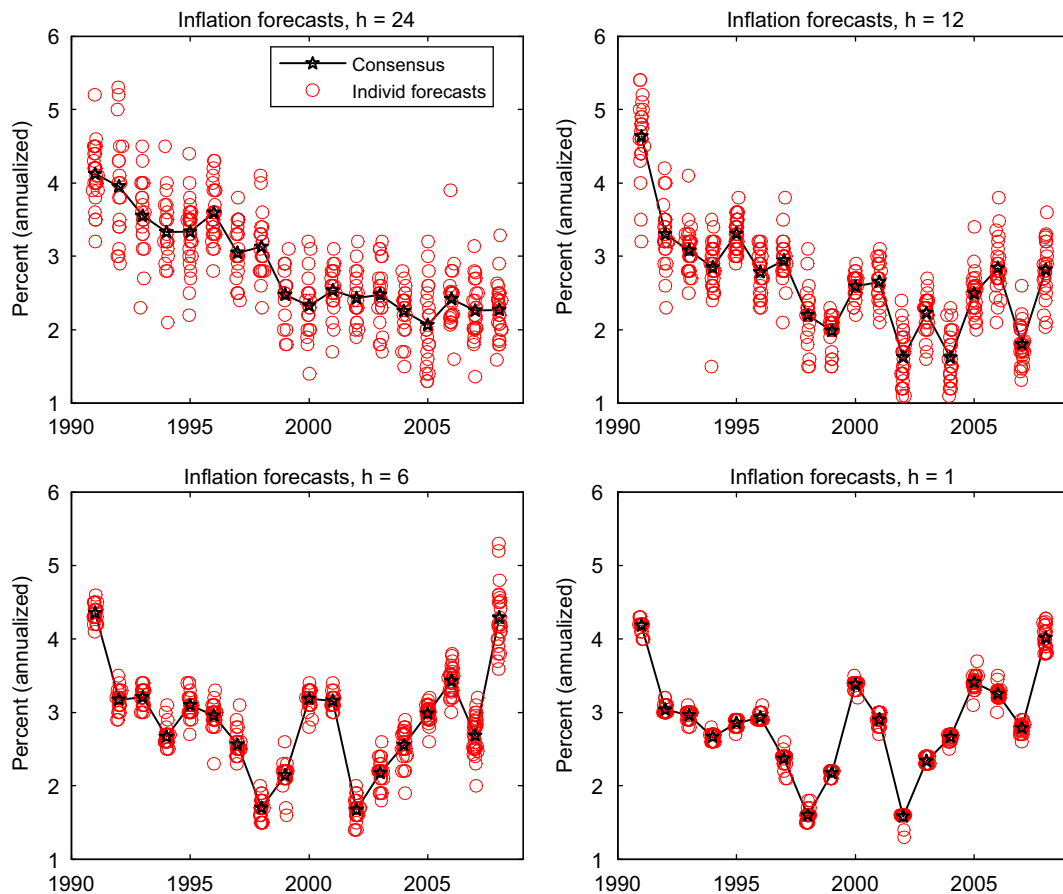


Fig. 2. Consensus and individual forecasts for inflation over the period 1991–2008, for four forecast horizons (24 months, 12 months, 6 months and one month).

meaning that a forecaster is at the 10th percentile of this distribution; 0.5 means the forecaster is at the median; 0.9 means the forecaster is at the 90th percentile, etc.) If differences in beliefs across forecasters were short-lived, the percentiles of the individual forecasters should be tightly clustered around the median. This is not what the data suggest, particularly at the longest horizons where some of the forecasters are consistently optimistic or pessimistic. However, as the forecast horizon gets shorter, views tend to become more densely clustered around the median, particularly for the inflation series.

To address persistence in (relative) views more systematically, all individual forecasters are ranked according to whether, in a given year, t , their forecast is in the bottom, middle or top tercile. This exercise is repeated for all years in the sample and used to compute transition probabilities to see whether forecasters who are in, say, the top tercile (i.e., the most optimistic forecasters in the case of GDP growth) in year t continue to be in the top tercile in year $t+1$. Results from this exercise, conducted separately at short, (1–12 months) and long (13–24 months) forecast horizons are reported in Table 2. In the absence of persistence in the relative views of individual forecasters, the entries in this table should all be approximately one-third (0.33). In contrast, if differences in forecasters' views persist, terms on the diagonal should be significantly higher than 0.33 and off-diagonal terms smaller than 0.33. There is strong evidence that disagreements among forecasters tend to persist. For GDP growth, at the short horizon, there is a 63% chance (nearly twice what is expected under no persistence) that the most optimistic forecasters continue to be relatively optimistic in the following period, while the most pessimistic forecasters repeat with a 45% probability. At the long forecast horizons there is even greater persistence in the relative ranking of forecasters by their degree of optimism or pessimism with repeat probabilities always above 50%. Similar conclusions hold for inflation. In all cases the estimated probabilities of remaining in the same tercile are significantly greater than 33%.

Forecasters enter and exit from our sample, and variations in the length of time a forecaster has been reporting to the survey could be a source of cross-sectional dispersion beyond the two channels modeled in the next section (differences in signals and differences in models or priors). However, this does not appear to be a main concern here. The probability of a forecaster remaining in the sample conditional on having reported in the previous month is 95% (i.e., there is only a 5% probability of leaving the sample), while the probability of remaining out of the sample if previously excluded is 0.90, so

Table 1
Summary statistics for the consensus forecast and the dispersion of forecasts, across horizons.

Forecast horizon	GDP growth			Inflation		
	Mean forecast	Mean dispersion	Correlation	Mean forecast	Mean dispersion	Correlation
1	2.861	0.071	−0.602*	2.839	0.075	0.448
2	2.824	0.088	−0.414	2.865	0.083	0.454
3	2.788	0.127	−0.410*	2.890	0.108	0.257
4	2.810	0.140	−0.524*	2.907	0.125	0.359
5	2.785	0.153	−0.375*	2.909	0.147	0.383
6	2.824	0.197	−0.482*	2.888	0.177	0.264
7	2.827	0.222	−0.534*	2.881	0.198	0.176
8	2.817	0.252	−0.462*	2.814	0.212	−0.084
9	2.748	0.306	−0.694*	2.713	0.259	0.563*
10	2.677	0.341	−0.640*	2.648	0.287	0.522*
11	2.606	0.355	−0.799*	2.591	0.286	0.261
12	2.559	0.380	−0.688*	2.653	0.323	0.478*
13	2.569	0.388	−0.646*	2.720	0.327	0.248
14	2.558	0.407	−0.698*	2.772	0.347	0.424*
15	2.626	0.419	−0.516*	2.814	0.355	0.388
16	2.739	0.414	−0.231	2.855	0.369	0.294
17	2.806	0.391	0.066	2.841	0.385	0.034
18	2.848	0.374	−0.102	2.847	0.405	0.129
19	2.852	0.391	−0.120	2.870	0.433	0.512*
20	2.856	0.387	−0.148	2.872	0.411	−0.020
21	2.849	0.396	−0.092	2.828	0.416	0.325
22	2.869	0.396	−0.297	2.825	0.435	0.228
23	2.865	0.400	−0.305	2.836	0.430	0.461*
24	2.853	0.420	−0.300	2.865	0.436	0.386

Note: This table presents the average consensus forecast, average cross-sectional dispersion in forecasts, and the time-series correlation between the consensus forecast and the dispersion in forecasts, for each horizon between one month and 24 months, computed across all years in the sample period (1991–2008). Correlation coefficients that are significantly different from zero at the 10% level (using Newey and West (1987) standard errors) are marked with an asterisk.

there is a 10% chance of re-entering the following month. Moreover, the cross-sectional dispersion is very similar whether calculated with or without new entrants in the sample.

To get an early indication of whether learning effects are important in the sample, the number of reported forecasts can be used as a crude indicator of experience. This is admittedly an imperfect measure of experience since a forecaster could have produced predictions long before being included in the Consensus Economics survey. At each point in time our forecasters are sorted into two groups according to whether the number of their reported forecasts is higher or lower than the median number of reports filed up to that point. Then separate measures of cross-sectional dispersion are computed for the most experienced and least experienced forecasters. Unreported results (available upon request) show that the cross-sectional dispersion in the two groups is almost identical, with only mild evidence of slightly higher dispersion among the most experienced group of inflation forecasters.

Three conclusions can be drawn from this brief look at the data. First, differences in opinions among forecasters tend to be much greater at long forecast horizons than at short forecast horizons. Second, there is a systematic relationship between the cross-sectional dispersion and average beliefs, with differences in opinion about GDP growth varying countercyclically. Third, there is considerable persistence through time in individual forecasters' views relative to that of the median forecaster and persistence tends to be higher at the longer forecast horizons.

3. The term structure of cross-sectional dispersion

Survey data on economic forecasts have been the subject of a large literature – see Pesaran and Weale (2006) for a recent review – and many studies have found this type of data to be of high quality, e.g., Romer and Romer (2000) and Ang et al. (2007). The focus of this literature has, however, mainly been on measuring the precision of average survey expectations as opposed to understanding why and by how much forecasters disagree.

Dispersion in beliefs observed at different forecast horizons turns out to provide important clues on why forecasters disagree. In fact, the importance of heterogeneity in priors can be identified primarily from the long end of the term structure of cross-sectional dispersion, while the importance of heterogeneity in signals is primarily identified from the short end of the term structure. Intuition for this comes from considering a simple AR(1) example. For this case, the h -period forecast is simply the present state times the AR(1) coefficient raised to the appropriate power, ϕ^h . Using parameter values similar to those obtained in our empirical analysis, less than one-third of the current signal carries over after 24 months. Hence, any difference between agents' signals about the current state is not going to be very important for




Fig. 3. Average position in the cross-sectional distribution of forecasters of four selected forecasters, for GDP growth and inflation, for each forecast horizon.

the long-horizon forecasts, and so disagreement in long-term forecasts must largely reflect different beliefs about the long-run mean. The next section introduces a model that formalizes this intuition in a more general setting.

3.1. A model for disagreement between forecasters

Consider how the disagreement among forecasters about an “event” measured at a fixed time period, t , (e.g., GDP growth in 2011) changes as the forecast horizon, h , is reduced. These so-called fixed-event forecasts (Nordhaus, 1987; Clements, 1997) with a time-varying forecast horizon match the focus in some theoretical models (e.g., Amador and Weill, 2009), that study how heterogeneity among agents evolves leading up to the revelation of the true value of a predicted variable.

Our analysis addresses how agents update their forecasts of some variable measured, e.g., at the annual frequency, when they receive news on this variable more frequently, e.g., on a monthly basis. To this end, let y_t denote the single-period variable (e.g., monthly log-first differences of GDP or a price index tracking inflation), while the rolling sum of the 12

Table 2
Transition matrices for low, middle and high ranked forecasters.

	GDP growth			Inflation		
	Low	Mid	High	Low	Mid	High
<i>Short-horizon forecasts</i>						
Low	0.451 [*]	0.323	0.225 [†]	0.491 [*]	0.290	0.218 [†]
Mid	0.258 [†]	0.433 [*]	0.309	0.204 [†]	0.415 [*]	0.381 [*]
High	0.106 [†]	0.266 [†]	0.629 [*]	0.082 [†]	0.246 [†]	0.672 [*]
<i>Long-horizon forecasts</i>						
Low	0.583 [*]	0.333	0.085 [†]	0.603 [*]	0.285 [†]	0.112 [†]
Mid	0.201 [†]	0.519 [*]	0.280 [†]	0.230 [†]	0.554 [*]	0.216 [†]
High	0.073 [†]	0.264 [†]	0.663 [*]	0.091 [†]	0.219 [†]	0.700 [*]

Note: This table presents the probability of a forecaster transitioning from a given tercile of the cross-sectional distribution of forecasts (low, middle, and high) to another tercile in the following year. The upper panel presents the average of these probabilities for all forecasts with horizon up to 12 months ("short horizon") and the lower panel presents the corresponding results for forecasts with horizon greater than 12 months ("long horizon"). Estimated probabilities that are significantly greater (lower) than 0.33 at the 10% level are marked with an asterisk (dagger).

most recent single-period observations of y is denoted z_t :

$$z_t = \sum_{j=0}^{11} y_{t-j}. \quad (1)$$

That is, y_t is the monthly variable (e.g., monthly GDP growth) and z_t is the corresponding annual variable. Our use of a variable tracking monthly changes in GDP (y_t) is simply a modelling device. US GDP figures are currently only available quarterly, but economic forecasters can be assumed to employ higher frequency data when constructing their monthly forecasts of GDP. Giannone et al. (2008), for example, propose methods to incorporate into macroeconomic forecasts news about the economy between formal announcement dates. When taking our model to data, the focus is naturally on those aspects of the model that have empirical counterparts. Because the analysis is concerned with flow variables that forecasters gradually learn about as new information arrives prior to and during the period of their measurement, the fact that part of the outcome could be known prior to the end of the measurement period (the "event date") means that the timing of the forecasts has to be carefully considered.

Agents are assumed to choose their forecasts to minimize the expected value of the squared forecast error, $e_{t|t-h} \equiv z_t - \hat{z}_{t|t-h}$, where z_t is the predicted variable, $\hat{z}_{t|t-h}$ is the forecast computed at time $t-h$, t is again the event date and h is the forecast horizon. Under this loss function, the optimal h -period forecast is simply the conditional expectation of z_t given information at time $t-h$, \mathcal{F}_{t-h} :

$$\hat{z}_{t|t-h}^* = E[z_t | \mathcal{F}_{t-h}]. \quad (2)$$

To track the evolution in the predicted variable, our analysis follows Patton and Timmermann (forthcoming) and uses a simple reduced-form model that, in common with popular macroeconomic models, decomposes y_t into a persistent first-order autoregressive component, x_t , and a temporary component, u_t :

$$\begin{aligned} y_t &= x_t + u_t \\ x_t &= \phi x_{t-1} + \varepsilon_t, \quad -1 < \phi < 1 \\ u_t &\sim iid(0, \sigma_u^2), \varepsilon_t \sim iid(0, \sigma_\varepsilon^2), \quad E[u_t \varepsilon_s] = 0 \quad \text{for all } t, s. \end{aligned} \quad (3)$$

Here ϕ measures the persistence of x_t , while u_t and ε_t are innovations that are both serially uncorrelated and mutually uncorrelated. Without loss of generality, the unconditional mean of x_t , and thus y_t and z_t , is assumed to be zero.

The advantage of using this highly parsimonious model is that it picks up the stylized fact that variables such as GDP growth and inflation clearly contain a persistent component. Unlike more structural approaches, it avoids having to take a stand on which particular variables agents use to compute their forecasts, a decision that in practice can be very complicated, see Stock and Watson (2002, 2006). The model can easily be extended to account for higher order dynamics, although given the relatively short time series under consideration, this is unlikely to be feasible in our empirical application.³

The model in Eq. (3) represents the data generating process for the macroeconomic variable being forecasted. To understand the cross-sectional dispersion in beliefs, heterogeneity across forecasters is next introduced. Disagreement between forecasters is modeled as arising from two possible sources: differences in information signals observed by

³ For similar reasons, heteroskedasticity in the underlying data generating process is also ignored, although this is unlikely to be important over the sample period studied here.

individual forecasters, or differences in their prior beliefs about, or econometric models for, long-run average levels. Define the cross-sectional dispersion among forecasters as

$$d_{t|t-h}^2 \equiv \frac{1}{N_{t-h}} \sum_{i=1}^{N_{t-h}} (\hat{z}_{it|t-h} - \bar{z}_{t|t-h})^2, \quad (4)$$

where $\bar{z}_{t|t-h} \equiv (1/N_{t-h}) \sum_{i=1}^{N_{t-h}} \hat{z}_{it|t-h}$ is the consensus forecast of z_t at time $t-h$, $\hat{z}_{it|t-h}$ is forecaster i 's prediction of z_t at time $t-h$ and N_{t-h} is the number of forecasters at time $t-h$. Notice that $d_{t|t-h}^2$ is a measure of subjective uncertainty reflected in agents' perceptions as distinct from objective measures of risk derived, e.g., from structural or time-series forecasting models.

To capture heterogeneity in forecasters' information, each forecaster is assumed to observe a different signal of the current value of y_t , denoted \tilde{y}_{it} . This assumption replicates the fact that different agents employ slightly different high-frequency variables for forming their forecast of GDP growth and inflation. Of course, many of the variables they examine will be common to all forecasters, such as government announcements of GDP growth, inflation and other key macroeconomic series, and so the signals the forecasters observe will, potentially, be highly correlated. The assumed structure is

$$\begin{aligned} \tilde{y}_{it} &= y_t + \eta_t + v_{it} \\ \eta_t &\sim iid(0, \sigma_\eta^2) \quad \forall t \\ v_{it} &\sim iid(0, \sigma_v^2) \quad \forall t, i \\ E[v_{it}\eta_s] &= 0 \quad \text{for all } t, s, i. \end{aligned} \quad (5)$$

Individual forecasters' measurements of y_t are contaminated with a common source of noise, denoted η_t , representing factors such as measurement errors, and independent idiosyncratic noise, denoted v_{it} . Participants in the survey are not formally able to observe each others' forecasts for the current period but they do observe previous survey forecasts.⁴ For this reason, we include a second measurement variable, \tilde{y}_{t-1} , which is the measured value of y_{t-1} contaminated with only the common noise:

$$\tilde{y}_{t-1} = y_{t-1} + \eta_{t-1}. \quad (6)$$

From this the individual forecaster is able to compute an optimal forecast from the variables observable to him:

$$\hat{z}_{it|t-h}^* \equiv E[z_t | \mathcal{F}_{it|t-h}], \quad \mathcal{F}_{it|t-h} = \{\tilde{y}_{i|t-h-j}, \tilde{y}_{t-h-1-j}\}_{j=0}^{t-h}. \quad (7)$$

Differences in signals about the predicted variable alone are unlikely to explain the observed degree of dispersion in the forecasts. The simplest way to verify this is to consider dispersion for very long horizons: as $h \rightarrow \infty$ the optimal forecasts converge towards the unconditional mean of the predicted variable. Because all forecasters are assumed to use the same (true) model to update their expectations about z , the dispersion should asymptote to zero as $h \rightarrow \infty$. As Figs. 1 and 2 reveal, this implication is in stark contrast with our data, which suggests instead that the cross-sectional dispersion converges to a constant but non-zero level as the forecast horizon grows. Thus there must be a second source of dispersion beyond differences in signals.⁵

The second source of dispersion is introduced by assuming that each forecaster comes with prior beliefs about the unconditional mean of z_t , denoted μ_i . Forecaster i is assumed to shrink the optimal forecast based on his information set $\mathcal{F}_{it|t-h}$ towards his prior belief about the unconditional mean of z_t . The degree of shrinkage is governed by a parameter, $\kappa^2 \geq 0$, with low values of κ^2 implying a small weight on the data-based forecast $\hat{z}_{it|t-h}^*$ (i.e., a large degree of shrinkage towards the prior belief) and large values of κ^2 implying a big weight on $\hat{z}_{it|t-h}^*$. As $\kappa^2 \rightarrow 0$, the forecaster places all weight on his prior beliefs and none on the data; as $\kappa^2 \rightarrow \infty$ the forecaster places no weight on his prior beliefs:

$$\begin{aligned} \hat{z}_{it|t-h} &= \omega_h \mu_i + (1 - \omega_h) E[z_t | \mathcal{F}_{it|t-h}], \\ \omega_h &= \frac{E[e_{it|t-h}^2]}{\kappa^2 + E[e_{it|t-h}^2]} \\ e_{it|t-h} &\equiv z_t - E[z_t | \mathcal{F}_{it|t-h}]. \end{aligned} \quad (8)$$

The weights placed on the prior and the conditional expectation, $E[z_t | \mathcal{F}_{it|t-h}]$, are allowed to vary across the forecast horizons in a manner consistent with standard forecast combinations. As $\hat{z}_{it|t-h}^* \equiv E[z_t | \mathcal{F}_{it|t-h}]$ becomes more accurate (i.e., as $E[e_{it|t-h}^2]$ decreases) the weight attached to that forecast increases. This weighting scheme lets agents put more weight on

⁴ As the participants in our survey are professional forecasters they may be able to observe each others' current forecasts through published versions of their forecasts such as investment bank newsletters or recommendations. If this is possible, then σ_v is expected to be close to zero.

⁵ While 24 months may not seem like a long forecast horizon, Lahiri and Sheng (2010) report evidence that the 24-month and 10-year survey forecasts of real GDP growth and inflation are in fact very similar.

the more precise signals in their short-term forecasts and less weight on signals at longer horizons. As pointed out by Lahiri and Sheng (2008, 2010), the “anchoring” of long-run forecasts is a consequence of Bayesian updating.⁶ Also, note that

$$\omega_h \rightarrow \frac{V[z_t]}{\kappa^2 + V[z_t]} \quad \text{as } h \rightarrow \infty.$$

Hence the weight on the prior in the long-run forecast can be quite large if κ^2 is small relative to $V[z_t]$. For analytical tractability, and for better finite sample identification of this parameter, κ^2 is restricted to be identical across all forecasters.

Our analysis assumes that forecasters know both the form and the parameters of the data generating process for z_t but do not observe this variable. Instead they only observe \tilde{y}_{it} and \tilde{y}_{t-1} which are noisy estimates of $[y_t, y_{t-1}]$. In common with many macroeconomic studies (e.g., Woodford, 2003), it is further assumed that agents use the Kalman filter to optimally predict (forecast, “nowcast” and “backcast”) the values of y_t needed for the forecast of $z_t \equiv \sum_{j=0}^{11} y_{t-j}$.⁷ Thus the learning problem faced by forecasters in our model relates to the latent state of the economy (measured by x_t and y_t), but not to the parameters of the model. This simplification is necessitated by our relatively short time series of data. Details on the state space representation of the model and the forecasters’ updating equations are provided in a technical appendix.

A possible interpretation of the heterogeneity in beliefs represented above by μ_i is that it captures differences in econometric models for long-run growth or inflation (for example, models with or without cointegrating relationships imposed), or it captures differences in sample periods used for the computation of agents’ forecasts (due to, for example, differences in beliefs about the dates of structural breaks). The short time-series dimension of our data does not allow us to distinguish between these competing interpretations.

The shrinkage of agents’ forecasts towards time-invariant long-run levels, μ_i , can alternatively be motivated by uncertainty about the value of the information signals received by agents. If agents know the interpretation of signals, under very mild conditions they will eventually hold identical beliefs. A standard Bayesian model would therefore require all disagreement to eventually be driven by differences in private signals. However, as shown by Acemoglu et al. (2007), if agents are uncertain about the interpretation of the signals, they need not agree even after observing an infinite sequence of identical signals. This is important since Figs. 1–3 show no evidence that agents’ beliefs converge even after 18 years of observations in our sample.⁸

4. Estimation of the model

The cross-sectional dispersion implied by our model is defined by

$$\delta_h^2 \equiv \frac{1}{N} \sum_{i=1}^N E[(\hat{z}_{it|t-h} - \bar{z}_{t|t-h})^2]. \quad (9)$$

Simulated method of moments (SMM, Gourieroux and Monfort, 1996a; Hall, 2005) is used to match the cross-sectional dispersion implied by our model, δ_h^2 , with its sample equivalent in the data given in Eq. (4). Unfortunately, a closed-form expression for δ_h^2 is not available and so simulations are used to evaluate δ_h^2 . In brief, this is done by simulating the state variables for T observations and then generating a different \tilde{y}_{it} series for each of the N forecasters. For each forecaster the optimal Kalman filter forecast is computed and then combined with the forecaster’s prior to obtain the final forecast using Eq. (8). Finally, the cross-sectional variance of the individual forecasts is used to compute $d_{t|t-h}^2$ which is averaged across time to obtain δ_h^2 .

Our model also yields predictions for the root mean-squared error (RMSE) of the consensus forecast, which is matched to the data to pin down the parameters of the data generating process, $(\sigma_u^2, \sigma_v^2, \phi)$. Details on these moments are presented in a technical appendix. Given our model for the term structure of dispersion in beliefs and the RMSE of the consensus forecast, all that remains is to specify a residual term for the model. Since the dispersion is measured by the cross-sectional variance, it is sensible to allow for an (additive) innovation term with variance related to the level of dispersion. One way to do this is to introduce a multiplicative innovation term of the following form:

$$d_{t|t-h}^2 = \delta_h^2 \cdot \lambda_{t|t-h},$$

$$E[\lambda_{t|t-h}] = 1, \quad V[\lambda_{t|t-h}] = \sigma_\lambda^2, \quad (10)$$

⁶ A formal Bayesian framework for the individual forecasters is not adopted because individual forecasters frequently enter and exit during our sample. This makes it impossible to capture how a single forecaster updates his views using Bayesian updating rules. The weighting scheme employed here has an intuitive Bayesian interpretation as a combination of the prior and the data to obtain the posterior.

⁷ The assumption that forecasters make efficient use of the most recent information is most appropriate for professional forecasters such as those considered in our empirical analysis, but is less likely to hold for households that only update their views infrequently, see Carroll (2003).

⁸ Agents’ beliefs could also fail to converge because of non-stationarities, see Kurz (1994). Another source of dispersion that is not considered here is differences in forecasters’ objectives (loss function). Capistran and Timmermann (2009) consider this possibility to explain differences among agents’ forecasts of US inflation measured at a given horizon and find that this can explain some of the dispersion in forecasts.

where $d_{t|t-h}^2$ is the observed value of the cross-sectional dispersion. The residual, $\lambda_{t|t-h}$, is assumed to be log-normally distributed with unit mean:

$$\lambda_{t|t-h} \sim iid \log N(-\frac{1}{2}\sigma_\lambda^2, \sigma_\lambda^2).$$

In addition to the term structures of consensus MSE-values and cross-sectional dispersion (each yielding up to 24 moment conditions), moments implied by the term structure of dispersion variances are also included to help estimate σ_λ^2 . The parameters of our model are obtained by solving the following expression:

$$\hat{\theta}_T \equiv \operatorname{argmin}_\theta \mathbf{g}_T(\theta)' \mathbf{g}_T(\theta), \tag{11}$$

where $\theta \equiv [\sigma_u^2, \sigma_v^2, \phi, \sigma_\eta^2, \sigma_v^2, \kappa^2, \sigma_\mu^2, \sigma_\lambda^2]'$, and, for $h = 1, 2, \dots, 24$,

$$\mathbf{g}_{T,h}(\theta) \equiv \frac{1}{T} \sum_{t=1}^T \begin{bmatrix} e_{t|t-h}^2 - MSE_h(\theta) \\ d_{t|t-h}^2 - \delta_h^2(\theta) \\ (d_{t|t-h}^2 - \delta_h^2(\theta))^2 - \delta_h^4(\theta)(\exp(\sigma_\lambda^2) - 1) \end{bmatrix}. \tag{12}$$

In total our model generates 72 moment conditions and contains eight unknown parameters. In practice only six forecast horizons ($h=1,3,6,12,18,24$) are used in the estimation, rather than the full set of 24, in response to studies of the finite-sample properties of GMM estimates (Tauchen, 1986) which find that using many more moment conditions than required for identification leads to poor approximations from the asymptotic theory, particularly when the moments are highly correlated, as in our application.⁹ The identity matrix is used as the weighting matrix in our initial SMM estimation, while the efficient weight matrix is used for the final parameter estimates and tests.

To compute standard errors and the test of over-identifying restrictions, the covariance matrix of the moments in Eq. (12) is needed. For this purpose, the model-implied covariance matrix of the moments based on the estimated parameters is used. This matrix is not available in closed-form and so 50 non-overlapping years of data are simulated to estimate it while imposing that the innovations to these processes are normally distributed and using the expressions given in the Appendix to obtain Kalman filter forecasts.¹⁰ As noted above, a closed-form expression for δ_h^2 is not available and so simulations are used to obtain an estimate of it. For each evaluation of the objective function, 50 non-overlapping years of data are simulated for 30 forecasters to estimate δ_h^2 .¹¹ The priors for each of the 30 forecasters, μ_i , are simulated as $iidN(0, \sigma_\mu^2)$.¹² The estimated δ_h^2 series are multiplied by $\lambda_{t|t-h}$, defined in Eq. (10), from which ‘measured’ values of dispersion, $d_{t|t-h}^2 = \delta_h^2 \cdot \lambda_{t|t-h}$, and the squared dispersion residual, $\lambda_{t|t-h}^2$, are obtained. These appear in the second and third set of moment conditions in Eq. (12) and, when combined with the MSE-values, are used to compute the sample covariance matrix of the moments.

5. Empirical results on forecast disagreement

We next turn to our empirical results from the econometric analysis of the cross-sectional dispersion in the survey forecasts of GDP growth and inflation. Revised data are used to measure the realized value of the target variable (GDP growth or inflation), but note that these are strongly correlated (correlation of 0.90) with the first release of the real-time series, the data recommended by Corradi et al. (2009). Our model in Section 3 assumes that the target variable is the December-on-December change in real GDP or the consumer price index, which can conveniently be written as the sum of the month-on-month changes in the log-levels of these series, as in Eq. (1). The Consensus Economics survey formally defines the target variable slightly differently to this but the impact of this difference on the model fit is negligible.¹³

To gain intuition for how the parameters of our constant-dispersion model are identified, notice that of those parameters, three ($\phi, \sigma_u^2, \sigma_v^2$) characterize the data generating process in Eq. (3). These parameters are mostly, though not solely, identified by the moments pertaining to the RMSE-values of the average forecast. In contrast, $\sigma_\eta^2, \sigma_\mu^2, \sigma_v^2$ and κ^2 are primarily determined by the moments capturing the term structure of cross-sectional dispersion. Fig. 4 shows how the model-implied cross-sectional dispersion in beliefs varies across different horizons as each of the four parameters take on low, medium and high values.¹⁴ The plots suggest that the parameters have very different effects on the term structure. Higher values of σ_η^2 increase the dispersion at short horizons, but have little effect on long-horizon dispersion. Increases in

⁹ The models presented in this paper were also estimated using the full set of 24 moment conditions and the results were qualitatively similar.

¹⁰ The sensitivity of this estimate to changes in the size of the simulation and to re-simulating the model was examined. When 50 non-overlapping years of data are used, changes in the estimated covariance matrix were found to be negligible.

¹¹ The actual number of forecasters in each survey exhibited some variation across t and h , but the simulations set $N = 30$ for all t, h for simplicity. Simulation variability for this choice of N and T is small, particularly relative to the values of the time-series variation in $d_{t|t-h}^2$ observed in the data.

¹² As a normalization it is assumed that $N^{-1} \sum_{i=1}^N \mu_i = 0$ since $N^{-1} \sum_{i=1}^N \mu_i$ and $\sigma_\mu^2 \equiv N^{-1} \sum_{i=1}^N \mu_i^2$ cannot be separately identified from our data on forecast dispersions. This normalization is reasonable if the number of optimistic forecasters is approximately equal to the number of pessimistic forecasters.

¹³ Generalizing the model to accommodate the exact definition of the target variable in the Consensus Economics survey involves lengthy algebra and complicates the description of the model, see Patton and Timmermann (forthcoming) for details.

¹⁴ The medium value of each of these parameters (except for σ_v) corresponds to the fitted value for GDP growth (reported in Table 3), and the high/low values are the fitted values ± 2 times the standard errors. The fitted value for σ_v was very near to zero, and so for that parameter we use that as the “low” value and obtain medium and high values as the fitted value +2 and +4 standard errors.

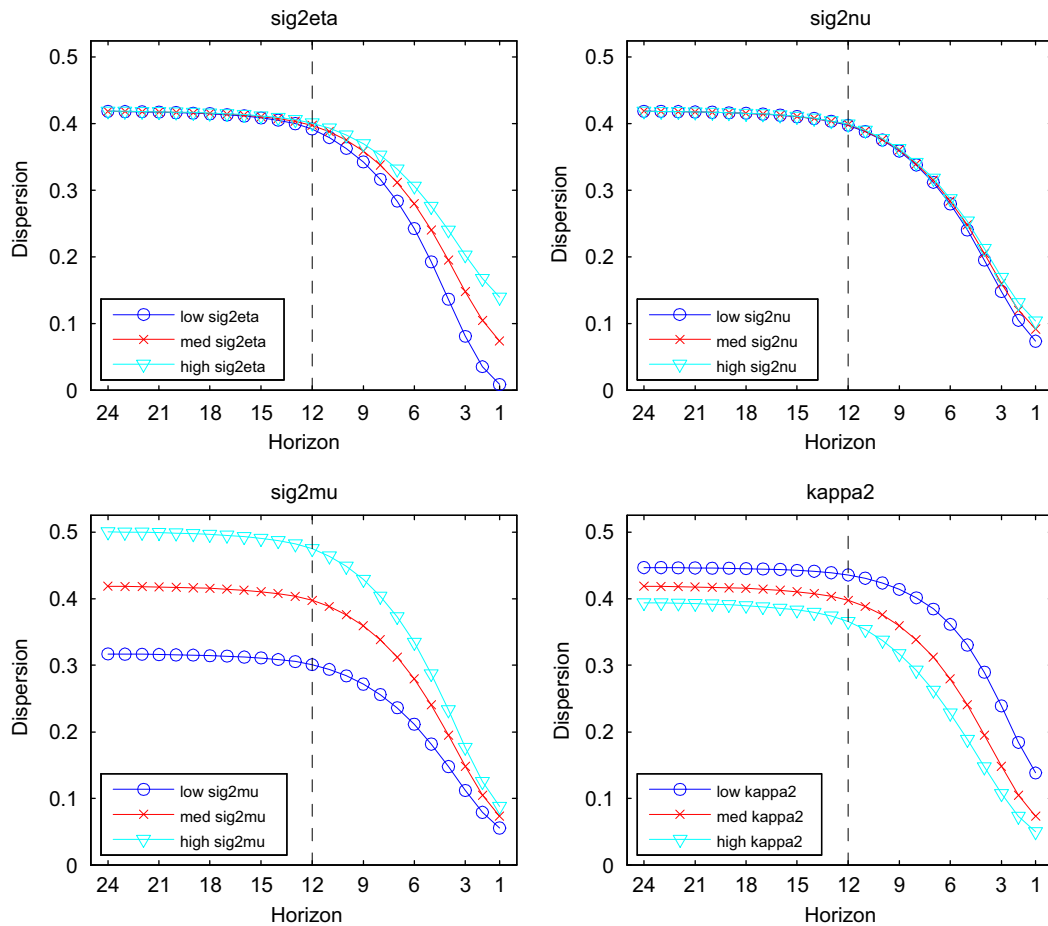


Fig. 4. Term structure of forecast dispersion when varying the variance of the common noise component (sig2eta), the variance of the idiosyncratic noise (sig2nu), the variance of the differences in prior beliefs, and the variance of perceived accuracy of prior beliefs (kappa2).

σ_v^2 have a smaller but similar effect. Variations in σ_μ^2 lead to big shifts in the long-run dispersion in beliefs, but have little effect on short-run disagreements, while conversely variations in κ^2 lead to small variations in the long-run dispersion but imply large changes in the short-run dispersion. Thus the parameters generally have very different effects on different portions of the term structure, which helps identify their values.

Fig. 5 plots the term structure of cross-sectional dispersion for GDP growth and inflation, i.e., the cross-sectional standard deviation of forecasts, averaged across the full sample, 1991–2008, listed against the forecast horizon. The cross-sectional dispersion of output growth is only slowly reduced for horizons in excess of 12 months, but declines rapidly for $h < 12$ months from 0.38 at the 12-month horizon to 0.07 at the one-month horizon. For inflation, again there is a systematic reduction in the dispersion as the forecast horizon shrinks. The cross-sectional dispersion declines from 0.44 at the 24-month horizon to 0.32 at the 12-month horizon and 0.08 at the one-month horizon.

It is interesting to contrast the pattern in Fig. 5 with the cross-sectional dispersion implied by the analysis in Amador and Weill (2009). In their model, more precise public information leads agents to rely less on private information and so slows down learning, crowding out valuable private information. Provided that prior beliefs are sufficiently dispersed, initially agents put increasingly more weight on their private information, leading to a convex segment of the aggregate learning curve. Subsequently, the learning curve becomes concave due to the fact that the true state is eventually revealed. Hence information diffuses over time along an S-shaped curve, while the cross-sectional dispersion in beliefs converges towards zero along a hump-shaped curve, i.e., it starts low, then increases monotonically, reaches a peak before decreasing towards zero. While our data are consistent with the reduced dispersion found at short forecast horizons, the very high dispersion observed at the longest horizons is clearly at odds with this type of learning model.

5.1. Parameter estimates and hypothesis tests

Table 3 reports parameter estimates for the model based on the moments in Eq. (12). For both GDP growth and inflation the estimates of σ_μ and κ suggest considerable heterogeneity across forecasters in our panel. Conversely, the estimates of

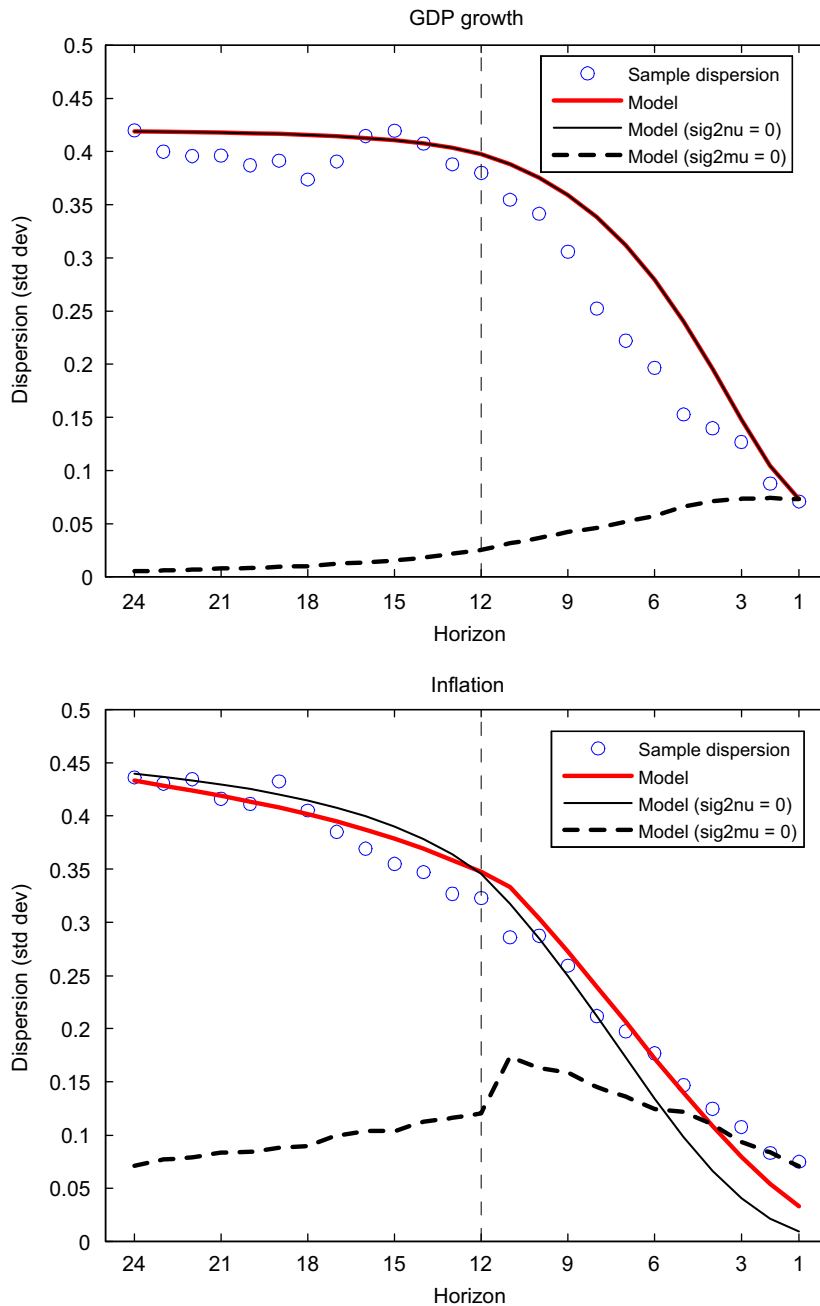


Fig. 5. Term structure of cross-sectional dispersion (standard deviation) of forecasts of GDP growth and Inflation in the U.S.

σ_v suggest that differences in individual signals are not important.¹⁵ Indeed, for GDP growth the test statistics for σ_v and σ_μ are 0 and 12.98, respectively, while for inflation the test statistics are 0.001 and 2.97. Thus for both series, the null that $\sigma_v = 0$ is not rejected while the null that $\sigma_\mu = 0$ is rejected at the 5% level. Heterogeneity in signals about GDP growth and inflation therefore does *not* appear to be a significant source of disagreement among professional forecasters, whereas heterogeneity in beliefs about the long-run levels of GDP growth and inflation is strongly significant.

¹⁵ Testing the null that σ_v (or σ_μ) is zero against it being strictly positive is complicated by the fact that zero is the boundary of the support for this parameter, which means that standard t -tests are not applicable. In such cases the squared t -statistic no longer has an asymptotic χ^2_1 distribution under the null, rather it will be distributed as a mixture of a χ^2_1 and a χ^2_0 , see, e.g., Gouriou and Monfort (1996b, Chapter 21), and the 90% and 95% critical values for this distribution are 1.64 and 2.71.

Table 3

Parameter estimates of the joint consensus forecast and constant dispersion model.

	GDP growth	Inflation
σ_u	0.000 (-)	0.074 (0.062)
σ_ε	0.073 (0.040)	0.018 (0.013)
ϕ	0.890 (0.039)	0.981 (0.031)
σ_η	0.083 (0.069)	0.000 (-)
σ_v	0.000 (2.287)	0.486 (14.264)
σ_μ	0.380 (0.081)	0.531 (0.260)
κ	0.552 (0.153)	0.593 (0.344)
<i>J</i> p-val	0.766	0.000
$H_0 : \sigma_v = 0$	0.000 (0.500)	0.001 (0.486)
$H_0 : \sigma_\mu = 0$	12.982 (0.000)	2.970 (0.042)

Note: This table reports simulated method of moments (SMM) parameter estimates of the Kalman filter model for the consensus forecasts and forecast dispersions, with standard errors in parentheses. The model is estimated using six moments each from the MSE term structure for the consensus forecast and from the cross-sectional term structure of dispersion for each variable. *p*-values from the test of over-identifying restrictions are given in the row titled "*J* p-val". The final two rows present the test statistics, with *p*-values in parentheses, of the tests for no heterogeneity in signals ($H_0 : \sigma_v = 0$) and no heterogeneity in beliefs ($H_0 : \sigma_\mu = 0$).

Our tests of the over-identifying restrictions indicate that the model provides a good fit to the GDP growth consensus forecast and forecast dispersion, with the *J*-test *p*-value being 0.77. Moreover, the top panel of Fig. 5 confirms that the model provides a close fit to the empirical term structure of forecast dispersions. This panel also shows that the model with σ_v set to zero provides an almost identical fit to the model with this parameter freely estimated, consistent with the test results for this hypothesis. In contrast, the model with σ_μ set to zero provides a poor fit to the term structure of dispersion for all horizons but the shortest. Without heterogeneity in priors, this model can only generate dispersion from differences in signals and these have limited impact at long horizons. Consistent with this, differences in individual information about GDP growth, modeled by v_{it} , do not appear important for explaining forecast dispersion. The most important features are the differences in prior beliefs about long-run GDP growth and the accuracy of Kalman filter-based forecasts as they affect the weight given to the prior relative to the Kalman filter forecast.

Unlike the model for GDP forecasts, the model for inflation forecasts and dispersions is rejected by the test of over-identifying restrictions (see the row labeled '*J* p-val' in Table 3). The model fits dispersion well for horizons greater than 12 months, but for horizons less than four months the observed dispersion is above what is predicted by our model. Given the functional form specified for the weight attached to the prior belief about long-run inflation versus the Kalman filter-based forecast, the model predicts that each forecaster will place 95.0% and 99.1% weight on the Kalman filter-based forecast for horizons of $h=3$ and 1 month, and since the Kalman filter forecasts are very similar across forecasters at short horizons our model predicts that dispersion will be low.

The observed dispersion in inflation forecasts is high relative both to the predictions of our model, and to observed forecast errors: observed dispersion (in standard deviations) for horizons $h=3$ and 1 month are 0.11 and 0.07, compared with the RMSE of the consensus forecast at these horizons of 0.19 and 0.08. Contrast this with the corresponding figures for the GDP forecasts, with dispersions of 0.13 and 0.07 and RMSE-values of 0.34 and 0.30. Thus, the dispersion of inflation forecasts is around 69% as large as the RMSE of the consensus forecast for short horizons ($1 \leq h \leq 6$), whereas the dispersion of GDP growth forecasts is around 35% as large as the RMSE of the consensus forecast.

To further illustrate this point, Fig. 6 plots the observed ratio of dispersion to RMSE, along with the predicted ratios, for horizons ranging from 24 months to one month, for both GDP growth and inflation. The upper panel of this plot reveals that our model is able to capture the basic shape of this function for GDP growth, while the lower panel shows how this ratio diverges for short horizons, especially the one-month horizon, and is not described well by our model. Patton and Timmermann (forthcoming) show that this model fits the RMSE term structure well, and so the divergence of the observed data from our model is not due to a poor model for the RMSE. The upward sloping function for the dispersion-to-RMSE ratio is difficult to explain within the confines of our model, or indeed any model assuming a quadratic penalty for forecast errors and efficient use of information, and thus poses a puzzle.

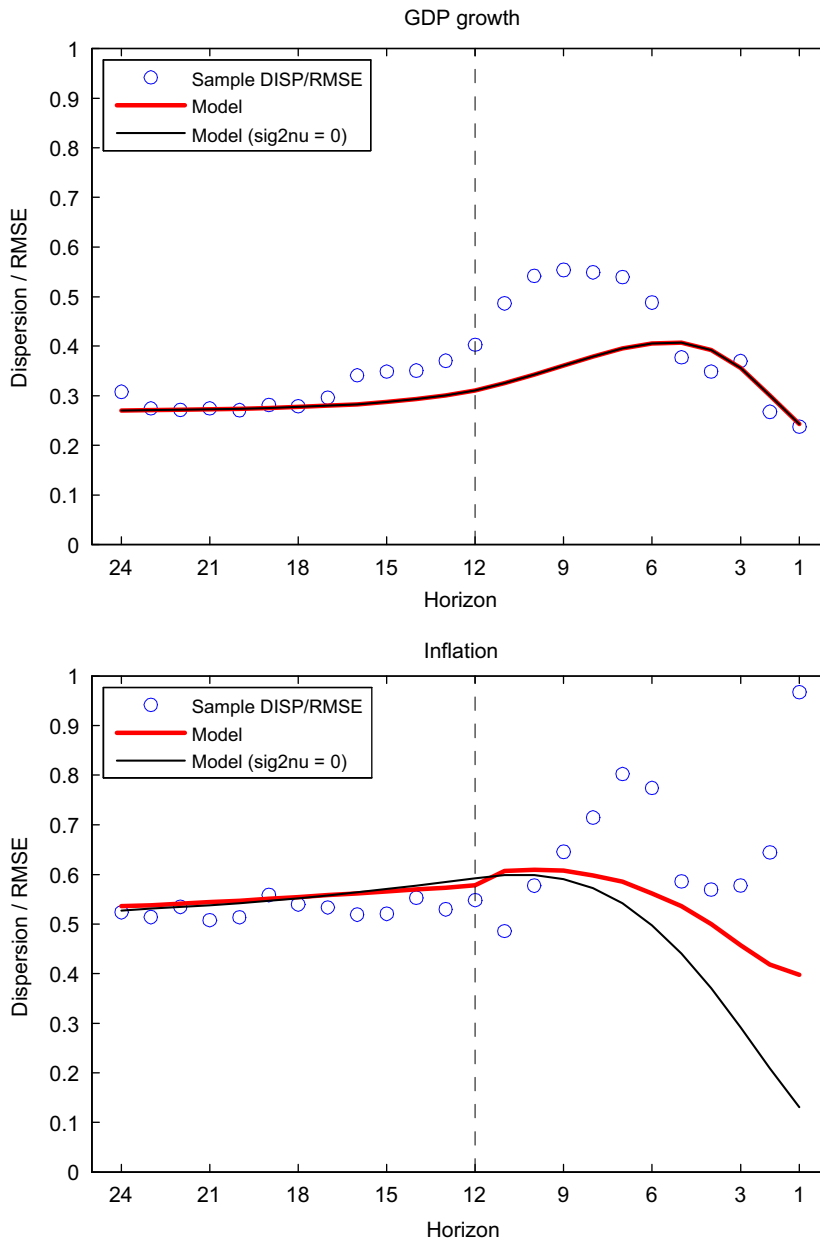


Fig. 6. Ratio of cross-sectional dispersion to root mean squared forecast errors for US GDP growth and Inflation as a function of the forecast horizon.

6. Time-varying dispersion

There is a growing amount of theoretical and empirical work on the relationship between the uncertainty facing economic agents and the economic environment. Van Nieuwerburgh and Veldkamp (2006), Veldkamp (2006) and Veldkamp and Wolfers (2007) propose endogenous information models where agents' participation in economic activity leads to more precise information about unobserved economic state variables such as (aggregate) technology shocks. In these models the number of signals observed by agents is proportional to the economy's activity level so more information is gathered in a good state of the economy than in a bad state. Recessions are therefore times of greater uncertainty which in turn means that dispersion among agents' forecasts can be expected to be wider during such periods. Similarly, Mackowiak and Wiederholt (2009) show that an increase in the variance of nominal aggregate demand leads firms to pay more attention to aggregate activity and less to idiosyncratic conditions. This could lead to a decrease in the cross-sectional dispersion in beliefs about the aggregate nominal demand. Thus, changing volatility in the variance of nominal aggregate demand – e.g., around turning points of the business cycle – can lead to time-varying cross-sectional dispersion.

Table 4
Parameter estimates for two models of time-varying dispersion.

	GDP growth	Inflation
<i>Panel A: OLS estimates of a panel model for dispersion</i>		
Fixed effects?	Yes	Yes
β_{SPR}	0.425 (0.115)	0.779 (0.082)
<i>Panel B: SMM parameter estimates of a Kalman filter model</i>		
σ_u	0.000 (-)	0.033 (0.054)
σ_ε	0.042 (0.022)	0.023 (0.014)
ϕ	0.954 (0.026)	0.957 (0.038)
σ_η	0.073 (0.066)	0.000 (-)
σ_v	0.046 (0.108)	0.486 (8.359)
σ_μ	0.682 (0.378)	0.530 (0.322)
β_0^K	4.452 (1.819)	-2.179 (3.253)
β_1^K	-4.451 (2.284)	1.190 (4.558)
J p-val	0.983	0.000

Note: The first two rows report the results from the estimation of a panel model for log-dispersion, with horizon-specific fixed effects, as a function of the log default spread. In the interests of brevity, the individual fixed effect parameters are not reported. The remainder of the table reports simulated method of moments (SMM) parameter estimates of the Kalman filter model for the consensus forecasts and forecast dispersions, with standard errors in parentheses. *p*-values from the test of over-identifying restrictions are given in the row titled “J p-val”. The model is estimated using six moments each from the MSE term structure for the consensus forecast and from the cross-sectional term structure of dispersion for each variable.

To address such issues, our model can be generalized to allow forecast dispersion to vary over the business cycle. There are of course many variables that vary with the business cycle and could be used in our empirical model for time-varying dispersion. Our choice is to simply employ the default spread (the difference in average yields of corporate bonds rated by Moody’s as BAA vs. AAA), which is known to be strongly counter-cyclical, increasing during economic downturns. Over our sample period the default spread ranges from 55 basis points in 1995, 1997 and 2000, to 338 basis points in December 2008.

6.1. Time-varying differences in beliefs

As a simple, robust way to explore the relation between disagreements among forecasters and the state of the economy, consider a pooled regression of the logarithm of the cross-sectional dispersion on the logarithm of the default spread and separate horizon fixed effects, α_h , i.e.,

$$\log(d_{t|t-h}^2) = \alpha_h + \beta_{SPR} \log(S_{t-h}) + \varepsilon_{t-h}, \quad t = 1, 2, \dots, T; \quad h = 1, 2, \dots, 24, \quad (13)$$

where S_{t-h} is the default spread in month $t-h$. This regression is robust in the sense that it does not impose our model, but conversely also does not reveal the source of any cyclical variations in belief dispersion. Panel A in Table 4 reports the estimated coefficients from Eq. (13). For GDP growth the estimate of β_{SPR} is 0.43 with a t-statistic around four. For inflation the estimate of β_{SPR} is 0.78 and the t-statistic exceeds nine. Thus dispersion in output growth and inflation expectations is significantly higher during economic recessions than during upturns.¹⁶

Next our analysis explores how counter-cyclical movements in the cross-sectional dispersion can be introduced into our model. The most natural way to allow the default spread to influence dispersion in our model is through the variance of the individual signals received by the forecasters, σ_v^2 , or through κ^2 which determines how much weight forecasters put on their data-based forecast relative to their long-run model forecast. Given that the former variable explained very little of the (unconditional) dispersion term structure, our focus is on the latter channel. This leads to the following model

$$\log(\kappa_t^2) = \beta_0^K + \beta_1^K \log(S_t). \quad (14)$$

In this model, if $\beta_1^K < 0$, then increases in the default spread coincide with periods where forecasters put less emphasis on their signals and more weight on their long-run models. Since the main source of differences in beliefs is found to be

¹⁶ These findings are consistent with the work of Döpke and Fritsche (2006) for a panel of forecasters in Germany over a different sample period.

attributable to differences in models or priors, a negative value of β_1^{κ} would indicate that periods with increased default spreads coincide with periods with greater dispersion.

Leaving the rest of the model unchanged, the model with time-varying dispersion is estimated in a similar way to the model with constant dispersion, with the following modifications. The stationary bootstrap of Politis and Romano (1994) with average block length of 12 months is used to “stretch” the default spread time series, S_t , to be 50 years in length for the simulation. This maintains the properties of the default spread process and allows us to simulate longer time series than those in our data set. Following this step the remainder of the simulation is the same as for the constant dispersion case above, noting that the combination weights applied to the Kalman-filter forecast and the “prior” are now time-varying as κ_t^2 is time-varying. The value of $\delta_h^2(\kappa_t^2)$ is needed in the estimation stage so the dispersion residual can be computed. In the constant dispersion model, this is simply the mean of $d_{t|t-h}^2$, but in the time-varying dispersion model it depends on κ_t^2 . It is not computationally feasible to simulate $\delta_h^2(\kappa_t^2)$ for each unique value of κ_t^2 in our sample, so this is estimated by setting κ_t^2 equal to its sample minimum, maximum and its [0.25,0.5,0.75] sample quantiles, and then using a cubic spline to interpolate this function, obtaining $\tilde{\delta}_h^2(\kappa_t^2)$. The accuracy of this approximation is checked for values in between these nodes and the errors are very small. Finally, dispersion residuals are computed from $\tilde{\delta}_h^2(\kappa_t^2)$ and the data and these are used in the SMM estimation of the model parameters.

Empirical results for this model are presented in Panel B of Table 4. Consistent with the work of Veldkamp (2006) and Van Nieuwerburgh and Veldkamp (2006), for GDP growth the negative sign of β_1^{κ} implies that when spreads are high, forecasters rely less on (common) information and disagree more. Moreover, the estimate β_1^{κ} is significant at the 10% level and, as indicated by the last row, the model is not rejected. For inflation forecasts, in contrast, this parameter estimate is positive but not significantly different from zero and the model is rejected.

To see the implications of our analysis for the time series of cross-sectional dispersion, Fig. 7 plots the actual dispersion versus the fitted (model-implied) dispersion at horizons of $h = 24$ and 12 months. For GDP growth the model implies considerable variation in disagreement among forecasters with dispersions increasing markedly during economic downturns as tracked by the default spread. Moreover, the model nicely tracks time-variations in the cross-sectional dispersion of GDP growth with large positive correlations between the model-implied and actual dispersion. Unsurprisingly, given the poor fit of the inflation model, for the inflation series the model fails to similarly match

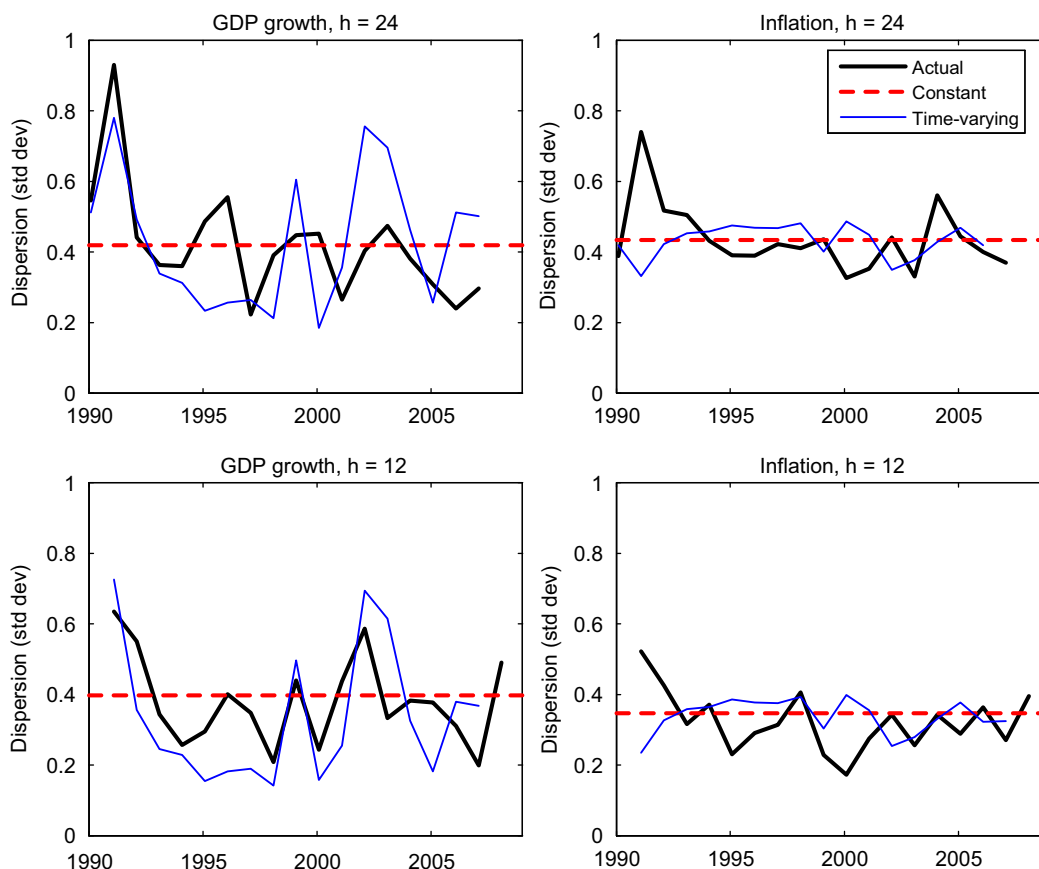


Fig. 7. Time series of observed cross-sectional dispersion of GDP growth and inflation forecasts, and model-implied time series of forecast dispersions (with constant and time-varying dispersions, respectively), for forecast horizons of 12 and 24 months.

time-variations in the actual dispersion. The contrast between the (constant and time-varying dispersion) model's ability to capture the term structure of cross-sectional dispersion for GDP growth versus its poor performance for inflation suggests that there could be fundamental differences in how agents update their beliefs about these two variables.

7. Conclusion

The degree of heterogeneity in forecasters' opinions, the nature and source of differences in opinions, and how such differences evolve over the economic cycle are important inputs to many macroeconomic models. Our empirical analysis suggests that differences in agents' forecasts of macroeconomic variables such as GDP growth and inflation tend to be much greater at long forecast horizons of up to two years compared with short horizons of a few months. Moreover, such differences in opinion tend to persist through time. To understand these patterns, our paper presents a simple and parsimonious model for the cross-sectional dispersion among forecasters that allows for heterogeneity in forecasters' information signals and in their prior beliefs or models.

Our analysis reveals several puzzling features that are difficult to explain with simple and popular forecasting models of the type used by macroeconomists. First, our empirical results suggest that heterogeneity in forecasters' information signals is not a major factor in explaining the cross-sectional dispersion in forecasts of GDP growth and inflation: heterogeneity in priors or models is more important. Second, the dispersion of forecasts does not appear to fall through time, suggesting that beliefs do not converge. Third, forecasters' views of inflation at short horizons appear to display "excess dispersion" that cannot be matched by our model and seems far greater than one would expect from differences in either prior beliefs or information signals. Fourth, and finally, our analysis shows that differences in opinions about GDP growth or inflation move strongly counter-cyclically, increasing during bad states of the world, although such variations again do not appear to be driven by heterogeneity in signals.

Acknowledgments

We thank the Editor and an anonymous referee for many constructive comments. Roy Batchelor, Steve Cecchetti, Jonas Dovern, Mike McCracken, Hashem Pesaran, Shaun Vahey, Michela Verardo, Mark Watson and seminar participants at the Board of Governors of the Federal Reserve, Cambridge, City University London, Duke, European Central Bank, London School of Economics, NBER Summer Institute, ESAM08 Meetings in Wellington, Oxford, Universite Libre Bruxelles (ECARES), Tilburg and Stanford (SITE workshop) provided helpful comments and suggestions.

Appendix A. Supplementary data

Supplementary data associated with this article can be found in the online version at doi:[10.1016/j.jmoneco.2010.07.001](https://doi.org/10.1016/j.jmoneco.2010.07.001).

References

- Acemoglu, D., Chernozhukov, V., Wernz, M., 2007. Learning and Disagreement in an Uncertain World. Mimeo, MIT.
- Amador, M., Weill, P.O., 2009. Learning from Private and Public Observations of Others' Actions. Mimeo, Stanford University.
- Ang, A., Bekaert, G., Wei, M., 2007. Do macro variables, asset markets, or surveys forecast inflation better? *Journal of Monetary Economics* 54, 1163–1212.
- Bacchetta, P., van Wincoop, E., 2006. Can information heterogeneity explain the exchange rate determination puzzle? *American Economic Review* 96 (3), 552–576.
- Brunnermeier, M., 2001. Asset Pricing under Asymmetric Information: Bubbles, Crashes, Technical Analysis and Herding. Oxford University Press.
- Capistran, C., Timmermann, A., 2009. Disagreement and biases in inflation expectations. *Journal of Money, Credit and Banking* 41, 365–396.
- Carroll, C., 2003. Macroeconomic expectations of household and professional forecasters. *Quarterly Journal of Economics* 118 (1), 269–298.
- Clements, M.P., 1997. Evaluating the rationality of fixed-event forecasts. *Journal of Forecasting* 16, 225–239.
- Corradi, V., Fernandez, A., Swanson, N.R., 2009. Information in the revision process of real-time data. *Journal of Business and Economic Statistics* 27, 455–467.
- Döpke, J., Fritsche, U., 2006. When do forecasters disagree? An assessment of German growth and inflation forecast dispersion. *International Journal of Forecasting* 22, 125–135.
- Giannone, D., Reichlin, L., Small, D., 2008. Nowcasting GDP: the real time informational content of macroeconomic data releases. *Journal of Monetary Economics* 55, 665–676.
- Gourieroux, C., Monfort, A., 1996a. Simulation-Based Econometric Methods. Oxford University Press, Oxford, United Kingdom.
- Gourieroux, C., Monfort, A., 1996b. *Statistics and Econometric Models*, vol. 2, translated from the French by Q. Vuong, Cambridge University Press, Great Britain.
- Hall, A.R., 2005. Generalized Method of Moments. Oxford University Press, USA.
- Hellwig, C., Venkateswaran, V., 2009. Setting the right prices for the wrong reasons. *Journal of Monetary Economics* 56, 557–577.
- Kurz, M., 1994. On the structure and diversity of rational beliefs. *Economic Theory* 4, 877–900.
- Lahiri, K., Sheng, X., 2008. Evolution of forecast disagreement in a Bayesian learning model. *Journal of Econometrics* 144, 325–340.
- Lahiri, K., Sheng, X., 2010. Learning and heterogeneity in GDP and inflation forecasts. *International Journal of Forecasting* 26, 265–292.
- Lorenzoni, G., 2009. A theory of demand shocks. *American Economic Review* 99 (5), 2050–2084.
- Mackowiak, B., Wiederholt, M., 2009. Optimal sticky prices under rational inattention. *American Economic Review* 99 (3), 769–803.
- Mankiw, G., Reis, R., 2002. Sticky information versus sticky prices: a proposal to replace the new Keynesian Phillips curve. *Quarterly Journal of Economics* 117, 1295–1328.

- McConnell, M.M., Perez-Quiros, G., 2000. Output fluctuations in the United States: what has changed since the early 1980s? *American Economic Review* 90, 1464–1476.
- Morris, S., Shin, H.S., 2002. Social value of public information. *American Economic Review* 92 (5), 1521–1534.
- Newey, W.K., West, K.D., 1987. A simple, positive semidefinite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, 703–708.
- Nordhaus, W.D., 1987. Forecasting efficiency: concepts and applications. *Review of Economics and Statistics* 69, 667–674.
- Patton, A.J., Timmermann, A., Forthcoming. Predictability of output growth and inflation: a multi-horizon survey approach. *Journal of Business and Economic Statistics*.
- Pesaran, M.H., Weale, M., 2006. Survey expectations. In: Elliott, G., Granger, C.W.J., Timmermann, A. (Eds.), *Handbook of Economic Forecasting*. North Holland, Amsterdam, pp. 715–776.
- Politis, D.N., Romano, J.P., 1994. The stationary bootstrap. *Journal of the American Statistical Association* 89, 1303–1313.
- Romer, C.D., Romer, D.H., 2000. Federal reserve information and the behavior of interest rates. *American Economic Review* 90, 429–457.
- Stock, J.H., Watson, M.W., 2002. Macroeconomic forecasting using diffusion indexes. *Journal of Business and Economic Statistics* 20, 147–162.
- Stock, J.H., Watson, M.W., 2006. Forecasting with many predictors. In: Elliott, G., Granger, C.W.J., Timmermann, A. (Eds.), *Handbook of Economic Forecasting*. North Holland, Amsterdam, pp. 515–554.
- Tauchert, G., 1986. Statistical properties of generalized method of moments estimators of structural parameters obtained from financial market data. *Journal of Business and Economic Statistics* 4, 397–416.
- Van Nieuwerburgh, S., Veldkamp, L., 2006. Learning asymmetries in real business cycles. *Journal of Monetary Economics* 53, 753–772.
- Veldkamp, L., 2006. Slow boom, sudden crash. *Journal of Economic Theory* 124 (2), 230–257.
- Veldkamp, L., Wolfers, J., 2007. Aggregate shocks or aggregate information? Costly information and business cycle comovement. *Journal of Monetary Economics* 54, 37–55.
- Woodford, M., 2003. Imperfect common knowledge and the effects of monetary policy. In: Aghion, P., Frydman, R., Stiglitz, J., Woodford, M. (Eds.), *Knowledge Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*. Princeton University Press.