

Why do Forecasters Disagree? Lessons from the Term Structure of Cross-Sectional Dispersion*

Andrew J. Patton

Allan Timmermann

University of Oxford

University of California San Diego and CREATES

12 December 2008

Abstract

Using data on cross-sectional dispersion in professional forecasters' long- and short-run predictions of macroeconomic variables, we identify key sources of disagreement in agents' forecasts. *We find that observed disagreement cannot be explained by differences in information sets; our results indicate it stems from heterogeneity in beliefs or models.* We develop methods for comparing measures of subjective cross-sectional dispersion in beliefs to the objective, model-implied measures of uncertainty. A simple reduced-form model is able to replicate the cross-sectional dispersion observed in forecasts of GDP growth but not for inflation - the latter appearing to be too high in the data at short horizons.

Keywords: fixed-event forecasts, Kalman filtering, optimal updating, dispersion in beliefs.

J.E.L. Codes: E37, C53, C32.

*We thank 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) for helpful comments and suggestions. Patton: Department of Economics, University of Oxford, Manor Road, Oxford OX1 3UQ, UK. Email: andrew.patton@economics.ox.ac.uk. Timmermann: Rady School, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0553, USA. Email: atimmerm@ucsd.edu. Timmermann is also affiliated with CREATES at the University of Aarhus, a research center funded by the Danish National Research Foundation.

1 Introduction

Why do agents disagree in their forecasts of macroeconomic variables such as output growth and inflation? The answer to this question goes to the heart of our understanding of how agents process information and to what extent they rely on firmly held “prior” beliefs, and also has important consequences for macroeconomic modeling. Theoretical models such as Lucas (1973) and Townsend (1983) suggest that heterogeneity in agents’ beliefs—as captured by, for example, the degree of cross-sectional dispersion in agents’ beliefs—is a key component of uncertainty about the state of the economy. Mankiw, Reis and Wolfers (2003, p.2) venture that “.. disagreement may be a key to macroeconomic dynamics.”

Modern macroeconomic analysis points towards the importance of measures of subjective uncertainty as reflected in agents’ perceptions, and as distinct from objective measures of risk derived, e.g., from structural or time-series forecasting models. Subjective views are ultimately what determine agents’ actions. In this paper we make use of survey data on differences in forecasters’ subjective views and develop a novel approach for comparing these to model-based objective measures of forecast dispersion. This allows us to ask whether agents’ perception of uncertainty is “rational” in the sense that it matches up with the objective degree of uncertainty surrounding economic variables.

In particular, we make use of a unique data set on forecasts of GDP growth and inflation for a given year, $t = 1, \dots, T$, recorded at different horizons, $h = 1, \dots, H$. Previous approaches (“fixed H and large T ”) to analyze survey data on expectations would mostly fix the forecast horizon, h , and consider forecasts for different time periods by varying t . Instead, we hold the time-period fixed and vary the forecast horizon. Adopting a “fixed T and large H ” framework allows us to analyze the source of disagreement among agents since time-variations in forecast error variance represent time-variation in agents’ uncertainty. This interpretation is not possible in the conventional framework where variation in forecast error variances might simply reflect unexpected changes in the volatility of the underlying variable (e.g. the “great moderation”) and so the two effects become difficult to disentangle.

Our analysis accomplishes four objectives. First, we address the question from the title, namely the key sources of disagreement among forecasters. At the most basic level of analysis, agents may 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 the variables will revert to their mean. Conversely, differences in prior beliefs about long-run inflation or output growth, *or differences in their econometric models of these quantities*, matters 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—we can therefore identify the key sources of disagreement. Empirically, we find that heterogeneity in their *information signals is not a major factor in explaining cross-sectional dispersion in forecasts of GDP growth and inflation: heterogeneity in their priors or models is more important*.

Second, we develop 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. We find evidence of “excess dispersion” in inflation forecasts at short horizons: at horizons of less than nine months the disagreement between agents’ predictions of inflation is high relative both to the prediction of our model, and relative to the objective degree of uncertainty about inflation.

Third, we generalize our model 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. Empirically, we indeed find that dispersion in forecasts of GDP growth has a strong and significant counter-cyclical component, whereas inflation forecast dispersion appears only weakly counter-cyclical.

Fourth, our analysis offers a variety of methodological contributions. We develop a model 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. We employ a simulation-based method of moments (SMM) framework for estimating the parameters of our model in a way that accounts for how agents update their beliefs as new information arrives. We view the shape of the cross-sectional dispersion in forecasts at different horizons as the object to be fitted and use SMM estimation 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 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 3 develops our econometric approach. Empirical results on the cross-sectional forecast dispersion are presented in Section 4 and Section 5 presents results for a model of time-varying dispersion. Section 6 concludes. Additional technical details on estimation of the model are presented in the Appendix.

2 The Term Structure of Cross-sectional Dispersion

Survey data on economic forecasts has 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, Bekaert and Wei (2007). The focus of the literature has, however, mainly been on testing the rationality of survey expectations as opposed to understanding how the precision of the forecasts evolves over time. This is related to the fact that survey data usually takes the form of “rolling event” forecasts of variables measured at different points in time (using a fixed forecast horizon but a varying date) such as a sequence of year-ahead forecasts of growth in GDP.

While it may be of economic interest to ask if the variance of the forecast error is the same across different subsamples, forecast efficiency implies no particular ranking of the error variances across different subsamples since the variance of the predicted variable need not be constant. For example, the forecast error associated with US GDP growth may have declined over time, but this need not imply that forecasters are getting any better if, as is widely believed, the volatility of US output growth has also come down (Kim and Nelson (1999) and McConnell and Perez-Quiros (2000)). *In contrast, if the same variable, e.g. GDP growth in 2000, is measured at different forecast horizons, one would expect the variance of the forecast error to decline as the forecast horizon shrinks. One might also expect the cross-sectional dispersion forecasts to decline with the forecast horizon, although this need not be the case.* In this Section we present a model that formalizes this intuition and allows us to quantify the size of these effects.

2.1 A first look at the data

Before setting up the model, it is useful to take a first look at the data we will be analyzing. To this end, Figures 1 and 2 plot the panel of cross-sectional dispersions for GDP growth and inflation as a function of the forecast horizon, h , and the time period, t . Further details on the data are provided in Section 4.

The average term structure (averaged across sample periods) is the object we are interested in modeling. There is a clear tendency for the dispersion to fall as the forecast horizon is reduced. Looking across years, there is also a tendency for the dispersion to change over time, a point we address in Section 5.

2.2 A model for disagreement between forecasters

We are interested in how the disagreement among forecasters about an “event” measured at a fixed time period, t , changes as the forecast horizon, h , is reduced, a so-called fixed-event forecast, see Nordhaus (1987) and Clements (1997). We study 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 most recent single-period observations of y is denoted z_t :

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

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 we take our model to data, we focus, naturally, on those aspects of the model that have empirical counterparts.* Since we shall be 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 may 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.

We assume that agents 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 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 variable, we follow Patton and Timmermann (2008) and use a simple reduced-form model which decomposes y_t into a persistent first-order autoregressive component, x_t , and a temporary component, u_t :

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

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, we assume that the unconditional mean of x_t , and thus y_t and z_t , is 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 which in practice can be very complicated, see Stock and Watson (2002, 2006). Moreover, the model can easily be extended to account for higher order dynamics, although given the relatively short time series we will consider, this is unlikely to be feasible in our empirical application.

2.2.1 Heterogeneity amongst forecasters

The model in (3) represents the *data generating process for the macroeconomic variable being forecasted*; to understand cross-sectional dispersion in beliefs, we next incorporate heterogeneity across forecasters. We shall model disagreement between forecasters as arising from two possible sources: differences in the information signals observed by individual forecasters, or differences in their prior beliefs *about, or econometric models for, long-run average levels*. We discuss other

possible sources of disagreement in Section 2.3. We define the cross-sectional dispersion among forecasters as

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

where $\bar{z}_{t,t-h} \equiv N^{-1} \sum_{i=1}^N \hat{z}_{i,t,t-h}$ is the consensus forecast of z_t at time $t-h$, $\hat{z}_{i,t,t-h}$ is forecaster i 's prediction of z_t at time $t-h$ and N is the number of forecasters.

To capture heterogeneity in the forecasters' information, we assume that each forecaster observes a different signal of the current value of y_t , denoted $\tilde{y}_{i,t}$. *This framework is designed to replicate the fact that different forecasters employ slightly different higher-frequency variables for forming their nowcast of GDP growth and inflation, which can lead them to different forecasts. 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 will, potentially, be highly correlated. The structure we assume is:*

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

Individual forecasters' measurements of y_t are contaminated with a common source of noise, denoted η_t and representing factors such as measurement errors, and independent idiosyncratic noise, denoted $\nu_{i,t}$. The participants in the survey we use are not formally able to observe each others' forecasts for the current period but they do observe previous survey forecasts.¹ 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 the optimal forecast from the variables observable to him:

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

¹As the participants in our survey are professional forecasters they may be able to observe each others' current forecasts through published versions of their forecasts, for example investment bank newsletters or recommendations. If this is possible, then we would expect to find σ_ν close to zero.

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. Since we assume that all forecasters use the same (true) model to update their expectations about z this implies that dispersion should asymptote to zero as $h \rightarrow \infty$. As Figures 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 source of dispersion beyond that deriving from differences in signals.

We therefore consider a second source of dispersion by assuming that each forecaster comes with prior beliefs about the unconditional mean of z_t , denoted μ_i . We assume that forecaster i shrinks the optimal forecast based on his information set $\mathcal{F}_{i,t-h}$ 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 $\hat{z}_{i,t,t-h}^*$ (i.e., a large degree of shrinkage towards the prior belief) and large values of κ^2 implying a high weight on $\hat{z}_{i,t,t-h}^*$. 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.

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

Notice that we allow the weights placed on the prior and the optimal expectation $E[z_t | \mathcal{F}_{i,t-h}]$ to vary across the forecast horizons in a manner consistent with standard forecast combinations: as $\hat{z}_{i,t,t-h}^* \equiv E[z_t | \mathcal{F}_{i,t-h}]$ becomes more accurate (i.e., as $E[e_{i,t,t-h}^2]$ decreases) the weight attached to that forecast increases. This weighting scheme lets agents put more weight on the more precise signals in their short-term forecasts and less weight on these at longer horizons. As pointed out by Lahiri and Sheng (2008b), the ‘‘anchoring’’ of long-run forecasts is a consequence of Bayesian

updating.² Also, note that

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

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 κ^2 , we impose that κ^2 is constant across all forecasters. Figure 3 plots ω_h as a function of the forecast horizon, h , for high and low values of κ , setting the other parameters to resemble those obtained for US GDP growth. This figure shows how the weight on priors versus signals varies with the forecast horizon, and thus by comparing long- and short-horizon forecast dispersions we gain some insight into this weight.

In Figure 4, we plot the theoretical term structures of dispersion for various values of σ_μ , setting the other parameters to resemble those obtained for US GDP growth. This figure shows that for small values of σ_μ , the dispersion term structure is increasing or roughly flat, while for larger values dispersion is quite high for long horizons and declines sharply as the forecast horizon shrinks towards zero³.

Our model identifies the importance of heterogeneity in priors 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 shorter horizons of the term structure.⁴ This latter point can be made by considering a simple AR(1) example: in such a case, the h -period forecast is simply the present state times the AR(1) coefficient ϕ^h . Using parameter values similar to those obtained in our empirical analysis, only between 10% and 20% of the current signal carries over after 24 periods. Hence, any difference between agents' signals is not going to be very important for the long-horizon forecasts, and so disagreement in long-term forecasts must largely reflect different beliefs about the long-run mean, μ_i .

²We refrain from adopting a formal Bayesian framework at the individual forecaster level as individual forecasters frequently enter and exit during our sample. This makes it impossible to capture how a single forecaster updates his/her views using Bayesian updating rules. The weighting scheme we employ has an clear intuitive Bayesian interpretation as combination of the prior and the data to obtain the posterior.

³Lahiri and Sheng (2008a) also propose a parametric model for the cross-sectional dispersion of macroeconomic forecasts as a function of the forecast horizon. However, they model the dispersion term structure directly, rather than through a combined model of the data generating process and the individual forecasters' prediction process as above.

⁴While 24 months may not seem like a long forecast horizon, Lahiri and Sheng (2008b) report evidence that the 24-month and 10-year survey forecasts of real GDP growth and inflation are in fact very similar.

2.3 Discussion of the model

Our analysis assumes that our 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 $[y_t, y_{t-1}]'$. We shall further assume that they use the Kalman filter to optimally predict (forecast, “nowcast” and “backcast”) the values of y_t needed for the forecast of $z_t \equiv \sum_{j=0}^{11} y_{t-j}$.⁵ Thus the learning problem faced by the 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 short time series of data. Technical details on the state-space representation of the model and the forecasters’ updating equations are provided in Appendix A.

Another 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 their forecasts (due to, for example, differences in beliefs about the dates of structural breaks). Given the short time series dimension of our data we are unable 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 the 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 here since Figures 1 and 2 show that there is no evidence that agents’ beliefs converge even after a decade and a half of observations in our sample.⁶

Another source of dispersion in agents’ beliefs which we do not consider here is differences in the forecasters’ objectives (loss function). Capistran and Timmermann (2008) 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.

⁵The assumption that forecasters make efficient use of the most recent information is most appropriate for professional forecasters such as those we shall consider in our empirical analysis, but is less likely to hold for households which may only update their views infrequently, see Carroll (2003).

⁶Agents’ beliefs may also fail to converge because of non-stationarities, cf. Kurz (1994).

3 Estimation of the Model

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

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

We use the simulated method of moments (SMM), see Gouriou and Monfort (1996a) or Hall (2005) for example, to match the cross-sectional dispersion implied by our model, δ_h^2 , with its sample equivalent in the data. Unfortunately, a closed-form expression for δ_h^2 is not available and so we resort to simulations to evaluate δ_h^2 . In brief, we do this 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 we obtain the optimal Kalman filter forecast and then combine this with the forecaster's prior to obtain his final forecast using equation (8). We then compute the cross-sectional variance of the individual forecasts to obtain $d_{t,t-h}^2$ and average these across time to obtain δ_h^2 .

Our model also yields predictions for the root mean-squared error (RMSE) of the consensus forecast, which we match to the data to help pin down the parameters of the DGP, $(\sigma_u^2, \sigma_\varepsilon^2, \phi)$. Details on these moments are presented in Appendix A. 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 the innovation term to be heteroskedastic, with variance related to the level of the dispersion. This form of heteroskedasticity, where the cross-sectional dispersion increases with the level of the predicted variable, has been documented empirically for inflation data by, e.g., Grier and Perry (1998) and Capistran and Timmermann (2008). We use the following model:

$$\begin{aligned} d_{t,t-h}^2 &= \delta_h^2 \cdot \lambda_{t,t-h} \\ E[\lambda_{t,t-h}] &= 1 \\ V[\lambda_{t,t-h}] &= \sigma_\lambda^2, \end{aligned} \quad (10)$$

where $d_{t,t-h}^2$ is the observed value of the cross-sectional dispersion. In particular, we assume that the residual, $\lambda_{t,t-h}$, is log-normally distributed with unit mean:

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

In addition to the term structures of consensus MSE-values and cross-sectional dispersion (each yielding up to 24 moment conditions) we also include moments implied by the term structure of

dispersion variances to help estimate σ_λ^2 . The parameters of our model are obtained by solving the following expression:

$$\hat{\boldsymbol{\theta}}_T \equiv \arg \min_{\boldsymbol{\theta}} \mathbf{g}_T(\boldsymbol{\theta})' \mathbf{g}_T(\boldsymbol{\theta}), \quad (11)$$

where $\boldsymbol{\theta} \equiv [\sigma_u^2, \sigma_\varepsilon^2, \phi, \sigma_\eta^2, \sigma_\nu^2, \kappa^2, \sigma_\mu^2, \sigma_\lambda^2]'$, and

$$\mathbf{g}_T(\boldsymbol{\theta}) \equiv \frac{1}{T} \sum_{t=1}^T \begin{bmatrix} e_{t,t-1}^2 - MSE_1(\boldsymbol{\theta}) \\ \vdots \\ e_{t,t-24}^2 - MSE_{24}(\boldsymbol{\theta}) \\ d_{t,t-1}^2 - \delta_1^2(\boldsymbol{\theta}) \\ \vdots \\ d_{t,t-24}^2 - \delta_{24}^2(\boldsymbol{\theta}) \\ (d_{t,t-1}^2 - \delta_1^2(\boldsymbol{\theta}))^2 - \delta_1^4(\boldsymbol{\theta}) (\exp(\sigma_\lambda^2) - 1) \\ \vdots \\ (d_{t,t-24}^2 - \delta_{24}^2(\boldsymbol{\theta}))^2 - \delta_{24}^4(\boldsymbol{\theta}) (\exp(\sigma_\lambda^2) - 1) \end{bmatrix}. \quad (12)$$

In total our model generates 72 moment conditions and contains 8 unknown parameters. In practice we use only six forecast horizons ($h = 1, 3, 6, 12, 18, 24$) 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.⁷ We use the identity matrix as the weighting matrix in our SMM estimation so that all horizons get equal weight in the estimation procedure; this is not fully efficient, but is motivated by our focus on modeling the entire term structure of forecast dispersions.

Of the eight parameters, three, $(\phi, \sigma_u^2, \sigma_\varepsilon^2)$ characterize the data generating process in (3), while σ_η^2 is a measurement error component. These parameters are mostly, though not solely, identified by the moments pertaining to the RMSE values of the average forecast⁸. In contrast, σ_μ^2 , σ_ν^2 and κ^2 are mostly determined by the moments capturing the term structure of cross-sectional dispersion and its degree of variability from year to year (σ_λ^2). Since the covariance matrix of the parameter

⁷We have also estimated the models presented in this paper using the full set of 24 moment conditions and the results were qualitatively similar.

⁸Even though both σ_η^2 and σ_u^2 are well-identified in theory, in practice they are difficult to estimate separately. We therefore set σ_η to be proportional to σ_u : $\sigma_\eta = k \cdot \sigma_u$. The goodness-of-fit of the model (as measured by Hansen's (1982) J -test of over-identifying restrictions) is generally robust for $1 \leq k \leq 4$ and we set $k = 2$ in the estimation.

estimates is not block diagonal, this holds only as an approximation, and all moments generally affect all parameters.

To obtain the covariance matrix of the moments in (12), used to compute standard errors and the test of over-identifying restrictions, we use the model-implied covariance matrix of the moments, based on the parameter estimate from the first-stage GMM parameter estimate. This matrix is not available in closed-form and so we simulate 50 non-overlapping years of data to estimate it, imposing that the innovations to these processes are Normally distributed, and using the expressions given in Appendix A to obtain the Kalman filter forecasts.⁹ As noted above, a closed-form expression for δ_h^2 is not available and so we use simulations to obtain an estimate of it. For each evaluation of the objective function, we simulated 50 non-overlapping years of data for 30 forecasters to estimate δ_h^2 .¹⁰ The priors for each of the 30 forecasters, μ_i , were simulated as *iid* $N(0, \sigma_\mu^2)$.¹¹ We multiply the estimated δ_h^2 series by $\lambda_{t,t-h}$, defined in equation (10) and from this we obtain ‘measured’ values of dispersion, $d_{t,t-h}^2 = \delta_h^2 \cdot \lambda_{t,t-h}$, and the squared dispersion residual, $\lambda_{t,t-h}^2$, which are used in the second and third set of moment conditions in (12), respectively. From these, combined with the MSEs, we compute the sample covariance matrix of the moments.

4 Empirical Results on Forecast Disagreement

4.1 Data used in the analysis

We next turn to our analysis of the cross-sectional dispersion in the survey forecasts of GDP growth and inflation. This measure matches our theoretical model in Section 2 and has the important advantage that it is not affected by incomplete data records due to the entry, exit and re-entry of individual forecasters.

⁹We examined the sensitivity of this estimate to changes in the size of the simulation and to re-simulating the model, and found that when 50 non-overlapping years of data are used the changes in the estimated covariance matrix are negligible.

¹⁰The actual number of forecasters in each survey exhibited some variation across t and h , with values between 22 and 32. In the simulations we set $N = 30$ for all t, h for simplicity. Simulation variability for this choice of N and T was small, particularly relative to the values of the time-series variation in $d_{t,t-h}^2$ that we observed in the data.

¹¹As a normalization we assume that $N^{-1} \sum_{i=1}^N \mu_i = 0$ since we cannot separately identify $N^{-1} \sum_{i=1}^N \mu_i$ and $\sigma_\mu^2 \equiv N^{-1} \sum_{i=1}^N \mu_i^2$ from our data on forecast dispersions. This normalization is reasonable if we think that the number of “optimistic” forecasters is approximately equal to the number of “pessimistic” forecasters.

Our data is taken from the Consensus Economics Inc. forecasts, in which the *quantitative predictions of private sector forecasters are reported*. Each month participants are asked for their forecasts of a range of macroeconomic and financial variables for the major economies. *The number of survey respondents for the variables we study varies between 22 and 32 in our sample period*. Our analysis focuses on US real GDP growth and CPI inflation for the current and subsequent calendar year. This gives us 24 monthly next-year and current-year forecasts over the period 1991-2004 or a total of $24 \times 14 = 336$ monthly observations. Naturally our observations are not independent draws but are subject to a set of tight restrictions across horizons, as revealed by the analysis in Section 2.

To measure the realized value of the target variable (GDP growth or inflation), we use second release data, following Romer and Romer (2000). Results are very similar when the first release is used instead as recommended by Corradi, Fernandez and Swanson (2007). *Our model in Section 2 assumed 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 equation (1). The Consensus Economics survey formally defines the target variable slightly differently to this but the impact of this difference on our results below is negligible*¹².

Figure 5 shows the cross-sectional dispersion (in standard deviation format) in output growth and inflation forecasts as a function of the forecast horizon. The cross-sectional dispersion of output growth declines only slowly for horizons in excess of 12 months, but declines rapidly for $h < 12$ months from a level near 0.4 at the 12-month horizon to around 0.1 at the 1-month horizon. For inflation, again there is a systematic reduction in the dispersion as the forecast horizon shrinks. The cross-sectional dispersion declines from around 0.45 at the 24-month horizon to 0.3 at the 12-month horizon and 0.1 at the 1-month horizon.

Table 1 reports parameter estimates for the model based on the moments in (12). The estimates of σ_μ suggest considerable heterogeneity across forecasters in our panel, whereas the estimates of σ_ν indicate that differences in individual signals may not be important, consistent with the possibility that the individual forecasters in our panel are able to observe each others' contemporaneous forecasts, rather than with a one-period lag. Testing the null that σ_ν (or σ_μ) is zero against it

¹²Generalizing the model to accommodate the exact definition of the target variable in the Consensus Economics survey involves lengthy but simple algebra, and makes the description of the model much more complicated.

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_1^2 distribution under the null, rather it will be distributed as a mixture of a χ_1^2 and a χ_0^2 , see, e.g., Gouriéroux and Monford (1996b, Chapter 21), and the 95% critical value for this distribution is 2.71. Table 2 shows the test statistics for σ_ν and σ_μ for GDP growth are 0.48 and 2.91 respectively, while for inflation the test statistics are 0.07 and 16.06. Thus for both of these series we fail to reject the null that $\sigma_\nu = 0$, while we are able to reject the null that $\sigma_\mu = 0$ at the 5% level. That is, heterogeneity in signals about GDP growth and inflation do *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.

Our tests of the over-identifying restrictions for each model indicate that the model provides a good fit to the GDP growth consensus forecast and forecast dispersion, with the p -value for that test being 0.86. Moreover, the top panel of Figure 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_ν set to zero provides almost as good a fit as the model with this parameter freely estimated, consistent with the results of tests of this hypothesis reported in Table 2. Differences in individual information about GDP growth, modelled by ν_{it} , thus 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).

In sharp contrast, the model for inflation forecasts and dispersions is rejected by the test of over-identifying restrictions (see the last column of Table 1). The model fits dispersion well for horizons greater than 12 months, but for horizons less than 9 months the observed dispersion is systematically 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 $h = 3$ and 1, and since the Kalman filter forecasts are very similar across forecasters at short horizons our model thus predicts that dispersion will be low.

Observed dispersion across forecasters is high both relative to the predictions of our model, and relative to observed forecast errors: observed dispersion (in standard deviations) for horizons 3 and

1 are 0.11 and 0.07, compared with the RMSE of the consensus forecast at these horizons of 0.08 and 0.05. Contrast this with the corresponding figures for the GDP forecasts, with dispersions of 0.14 and 0.08 and RMSE of 0.61 and 0.56. Thus, the dispersion of inflation forecasts is around 25% greater than the RMSE of the consensus forecast for short horizons, whereas the dispersion of GDP growth forecasts is around 75% *smaller* than the RMSE of the consensus forecast. Figure 6 plots the observed ratio of dispersion to RMSE, along with the predicted ratios, for horizons ranging from 24 months to 1 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, and is not described well by our model. Patton and Timmermann (2008) 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.

5 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. Veldkamp (2006) and van Nieuwerburgh and Veldkamp (2006) 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 periods¹³. This idea fits naturally with our model to the extent that information signals contain a common component and so lead to more similar beliefs during periods where more information is available.

To address such issues, we generalize our model to allow for time-varying dispersion in the forecasts. *There are of course many variables that vary with the business cycle that we could use*

¹³Using a panel of forecasters for Germany, Döpke and Fritsche (2006) indeed find that forecast disagreement is higher around recessions.

in our model for time-varying dispersion. We 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 and increases during economic downturns. Over our sample period, for example, the default spread ranges from 55 basis points in September and November 1997 to 141 basis points in January 1991 and January 2002.

5.1 Time-varying differences in beliefs

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_ν^2 , or through the variance of the prior beliefs about the long-run values of the series, σ_μ^2 . Given that the former variable explained very little of the (unconditional) dispersion term structure, we focus on the latter channel. We specify our model as

$$\log \sigma_{\mu,t}^2 = \beta_0^\mu + \beta_1^\mu \log S_t, \tag{13}$$

where S_t is the default spread in month t . In this model, if $\beta_1^\mu > 0$, then increases in the default spread coincide with increased differences in beliefs about the long-run value of the series, which in turn lead to an increase in the observed dispersion of forecasts.

Leaving the rest of the model unchanged, the model with time-varying dispersion was estimated in a similar way to the model with constant dispersion, with the following modifications. We used the stationary bootstrap of Politis and Romano (1994), with average block length of 12 months, to “stretch” the default spread time series, S_t , to be 50 years in length for the simulation. This maintains, asymptotically, the properties of this process and allows us to simulate longer time series than we have in our data set. The “standardized priors” for each of the 30 forecasters, μ_i^* , were simulated as *iid* $N(0, 1)$, and then the actual “prior” for each time period, $\mu_{i,t}$, was set as $\mu_i^* \times \sigma_{\mu,t}$, where $\sigma_{\mu,t} = \exp\{(\beta_0^\mu + \beta_1^\mu \log S_t)/2\}$. Following this step the remainder of the simulation was the same as for the constant dispersion case above. In the estimation stage we need to compute the value of $\delta_h^2(\sigma_{\mu,t})$, so that we can compute the dispersion residual. In the constant dispersion model, this is simply the mean of $d_{t,t-h}^2$, but in the time-varying dispersion model this also depends on $\sigma_{\mu,t}$. It was not computationally feasible to simulate $\delta_h^2(\sigma_{\mu,t})$ for each unique value of $\sigma_{\mu,t}$ in our sample, and so we estimated it for $\sigma_{\mu,t}$ equal to its sample minimum, maximum and its [0.25, 0.5, 0.75] sample quantiles, and then used a cubic spline to interpolate this function, obtaining $\tilde{\delta}_h^2(\sigma_{\mu,t})$. We

checked the accuracy of this approximation for values in between these nodes and the errors were very small. We then use $\tilde{\delta}_h^2(\sigma_{\mu,t})$, and the data, to compute the dispersion residuals and used these in the SMM estimation of the parameters of the model.

Empirical results for this model are presented in Table 3. Consistent with the work of Veldkamp (2006) and van Nieuwerburgh and Veldkamp (2006), our results reveal a positive relationship between default spreads and σ_μ , as evidenced by the signs of $\hat{\beta}_1^\mu$. This parameter is not significantly different from zero for the inflation forecast model, but is significant at the 10% level for the GDP growth forecast model. *These findings are consistent with the work of Döpke and Fritsche (2006) for a panel of forecasters for Germany, over a different sample period.*

In Figure 7 we plot the estimated dispersions as a function of the level of default spreads. When the default spread is equal to its sample 95th percentile (131 basis points), GDP growth forecast dispersion is approximately double what it is when the default spread is equal to its sample average (83 basis points). Similarly, when the default spread is equal to its 5th percentile (58 basis points) GDP growth forecast dispersion is approximately one half of the average figure. In contrast, the dispersion of inflation forecasts is only weakly affected by the default spread, with changes of approximately no more than 10% when the default spread moves from its average value to its 5th or 95th percentile value. We conclude that U.S. GDP growth forecast dispersion has a strong and significant counter-cyclical component, whereas U.S. inflation forecast dispersion appears only weakly counter-cyclical.

6 Conclusion

This paper developed a simple model for the cross-sectional dispersion among forecasters that allows for heterogeneity in forecasters' information signals and in their prior beliefs. Though highly parsimonious, our model sheds light on these important sources of disagreement between forecasters. Our empirical findings *suggest that heterogeneity in forecasters' information signals is not a major factor in explaining cross-sectional dispersion in forecasts of GDP growth and inflation: heterogeneity in their priors or models is more important.*

Differences in beliefs about GDP growth appear to be strongly counter-cyclical (increasing during bad states of the world) whereas differences in agents' inflation forecasts are less state dependent. Moreover, while our model can match the dispersion observed among survey participants'

forecasts of GDP growth, it fails to match the high dispersion in inflation forecasts observed at short horizons. Why professional forecasters' views of inflation at short horizons displays such "excess dispersion" is difficult to understand and poses a puzzle to any model based on agents' efficient use of information.

Appendix: Technical Details

This appendix provides details of how we derive the moments used in the empirical estimation in Section 3. We first introduce the state and measurement equations underlying the model from Section 2 cast in state space form and then show how the forecasters' updating equations can be solved.

A.1. State and Measurement Equations

Our model involves unobserved variables and so we cast it in state space form, using notation similar to that in Hamilton (1994). To account for the way the target variable is constructed, $z_t \equiv \sum_{j=0}^{11} y_{t-j}$, we augment the state equation with eleven lags of y_t so the target variable can be written as a linear combination of the state variable. The state equation is

$$\begin{bmatrix} x_t \\ y_t \\ y_{t-1} \\ \vdots \\ y_{t-11} \end{bmatrix} = \begin{bmatrix} \phi & 0 & 0 & \cdots & 0 \\ \phi & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 0 \end{bmatrix} \begin{bmatrix} x_{t-1} \\ y_{t-1} \\ y_{t-2} \\ \vdots \\ y_{t-12} \end{bmatrix} + \begin{bmatrix} \varepsilon_t \\ \varepsilon_t + u_t \\ 0 \\ \vdots \\ 0 \end{bmatrix} \quad (14)$$

which for short we write as

$$\boldsymbol{\xi}_t = \mathbf{F}\boldsymbol{\xi}_{t-1} + \mathbf{v}_t. \quad (15)$$

The measurement equation involves two variables: the estimate of y_t incorporating both common and idiosyncratic measurement error, and the estimate of y_{t-1} incorporating just common measurement error. In a minor abuse of notation relative to our discussion of this model in Section 2,

we will call the former \tilde{y}_{it}^* and the latter $\tilde{y}_{c,t-1}$, so that we may stack them into a vector $\tilde{\mathbf{y}}_{it}$:

$$\begin{bmatrix} \tilde{y}_{it}^* \\ \tilde{y}_{c,t-1} \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \end{bmatrix} \begin{bmatrix} x_t \\ y_t \\ y_{t-1} \\ \vdots \\ y_{t-11} \end{bmatrix} + \begin{bmatrix} \eta_t + \nu_{it} \\ \varphi_{t-1} \end{bmatrix} \quad (16)$$

which we write as

$$\tilde{\mathbf{y}}_{it} = \mathbf{H}'\boldsymbol{\xi}_t + \mathbf{w}_{it},$$

We introduce the measurement error φ_{t-1} , distinct from η_t but with the same distribution, so that the vector \mathbf{w}_{it} remains serially uncorrelated which simplifies the model.

The various shocks in the state and measurement equations are distributed as:

$$\left[u_t \ \varepsilon_t \ \eta_t \ \varphi_t \ \nu_{1t} \ \cdots \ \nu_{Nt} \right]' \sim iid N \left(\mathbf{0}, diag \left\{ \left[\sigma_u^2 \ \sigma_\varepsilon^2 \ \sigma_\eta^2 \ \sigma_\eta^2 \ \sigma_\nu^2 \ \cdots \ \sigma_\nu^2 \right] \right\} \right)$$

where $diag \{ \mathbf{a} \}$ is a square diagonal matrix with the vector \mathbf{a} on the main diagonal. Then $\mathbf{v}_t \sim iid N(0, \mathbf{Q})$, with

$$\mathbf{Q} = \begin{bmatrix} \sigma_\varepsilon^2 & \sigma_\varepsilon^2 & 0 & \cdots & 0 \\ \sigma_\varepsilon^2 & \sigma_\varepsilon^2 + \sigma_u^2 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 0 \end{bmatrix}.$$

And finally $\mathbf{w}_{it} \sim iid N(0, \mathbf{R})$, with

$$\mathbf{R} = \begin{bmatrix} \sigma_\eta^2 + \sigma_\nu^2 & 0 \\ 0 & \sigma_\eta^2 \end{bmatrix}.$$

Notice that by extending the state variable to include lags of \mathbf{y}_t we do not need to treat forecasts, nowcasts and backcasts separately; they can all be treated simultaneously as “forecasts” of the state vector $\boldsymbol{\xi}_t$. This simplifies the algebra considerably.

A.2. The Forecasters’ Updating Process

Our empirical data provides us with estimates of forecast uncertainty at different forecast horizons measured both in the form of the root mean squared forecast error (RMSE) of the “average”

or consensus forecast or in the form of the cross-sectional standard deviation of the forecasts (i.e., the dispersion). In this section we characterize how the forecasters update their beliefs and derive the model-implied counterparts of these two measures of uncertainty and disagreement.

Let

$$\begin{aligned}\mathcal{F}_{it} &= \sigma(\tilde{\mathbf{y}}_{it}, \tilde{\mathbf{y}}_{i,t-1}, \dots, \tilde{\mathbf{y}}_{i,1}) \\ \hat{\boldsymbol{\xi}}_{t|t-h,i} &\equiv E[\boldsymbol{\xi}_t | \mathcal{F}_{t-h,i}], \quad h \geq 0,\end{aligned}$$

where the expectation is obtained using standard Kalman filtering methods.

We assume that the forecasters have been using the Kalman filter long enough that all updating matrices, defined below, are at their steady-state values. This is done simply to remove any “start of sample” effects that may or may not be present in our actual data. Following Hamilton (1994):

$$\begin{aligned}\mathbf{P}_{t+1|t,i} &\equiv E\left[\left(\boldsymbol{\xi}_{t+1} - \hat{\boldsymbol{\xi}}_{t+1|t,i}\right)\left(\boldsymbol{\xi}_{t+1} - \hat{\boldsymbol{\xi}}_{t+1|t,i}\right)'\right] \\ &= (\mathbf{F} - \mathbf{K}_{t,i})\mathbf{P}_{t|t-1,i}(\mathbf{F}' - \mathbf{K}'_{t,i}) + \mathbf{K}_{t,i}\mathbf{R}\mathbf{K}'_{t,i} + \mathbf{Q} \\ &\rightarrow \mathbf{P}_1^*.\end{aligned}\tag{17}$$

Note that although the individual forecasters receive different signals, and thus generate different forecasts $\hat{\boldsymbol{\xi}}_{t+1|t,i}$, all signals have the same covariance structure and so will converge to the same matrix, \mathbf{P}_1^* . Similarly,¹⁴

$$\begin{aligned}\mathbf{K}_{t,i} &\equiv \mathbf{F}\mathbf{P}_{t|t-1,i}(\mathbf{P}_{t|t-1,i} + \mathbf{R})^{-1} \rightarrow \mathbf{K}^*, \\ \mathbf{P}_{t|t,i} &\equiv E\left[\left(\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t,i}\right)\left(\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t,i}\right)'\right] \\ &= \mathbf{P}_{t|t-1,i} - \mathbf{P}_{t|t-1,i}(\mathbf{P}_{t|t-1,i} + \mathbf{R})^{-1}\mathbf{P}_{t|t-1,i} \\ &\rightarrow \mathbf{P}_1^* - \mathbf{P}_1^*(\mathbf{P}_1^* + \mathbf{R})^{-1}\mathbf{P}_1^* \equiv \mathbf{P}_0^*.\end{aligned}\tag{18}$$

To estimate the matrices \mathbf{P}_1^* , \mathbf{P}_0^* , and \mathbf{K}^* , we simulate 100 non-overlapping years of data and update $\mathbf{P}_{t|t-1,i}$, $\mathbf{P}_{t|t,i}$ and $\mathbf{K}_{t,i}$ using the above equations. We use these matrices at the end of the

¹⁴The convergence of $\mathbf{P}_{t|t-1,i}$, $\mathbf{P}_{t|t,i}$ and $\mathbf{K}_{t,i}$ to their steady-state values relies on $|\phi| < 1$, see Hamilton (1994), Proposition 13.1, and we impose this in the estimation.

100^{th} year as estimates of \mathbf{P}_1^* , \mathbf{P}_0^* , and \mathbf{K}^* . Multi-step prediction error uses

$$\begin{aligned}
\hat{\boldsymbol{\xi}}_{t+h|t,i} &= \mathbf{F}^h \hat{\boldsymbol{\xi}}_{t|t,i}, \\
\text{so } \mathbf{P}_{t+h|t,i} &\equiv E \left[\left(\boldsymbol{\xi}_{t+h} - \hat{\boldsymbol{\xi}}_{t+h|t,i} \right) \left(\boldsymbol{\xi}_{t+h} - \hat{\boldsymbol{\xi}}_{t+h|t,i} \right)' \right] \\
&= \mathbf{F}^h \mathbf{P}_{t|t,i} (\mathbf{F}')^h + \sum_{j=0}^{h-1} \mathbf{F}^j \mathbf{Q} (\mathbf{F}')^j \rightarrow \mathbf{P}_h^*, \text{ for } h \geq 1.
\end{aligned} \tag{19}$$

The matrices \mathbf{P}_h^* for $h = 1, 2, \dots, 24$ are sufficient for us to obtain the term structure of RMSE, (that is, the RMSE values across different horizons, $h = 1, \dots, H$), for an individual forecaster, but the moments we include in the estimation are from the consensus forecasts, and so we need the RMSE term structure for the consensus, which requires slightly more work.¹⁵ Let

$$\bar{\boldsymbol{\xi}}_{t|t-h} \equiv \frac{1}{N} \sum_{i=1}^N \hat{\boldsymbol{\xi}}_{t|t-h,i} \tag{20}$$

be the consensus forecast of the state vector. We now derive the term structure of RMSE for this forecast, but first it is useful to derive the RMSE of the consensus “nowcast”:

$$\begin{aligned}
\bar{\mathbf{P}}_0^* &\equiv V \left[\boldsymbol{\xi}_t - \bar{\boldsymbol{\xi}}_{t|t} \right] \\
&= V \left[\frac{1}{N} \sum_{i=1}^N \left(\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t,i} \right) \right] \\
&= \frac{1}{N^2} \sum_{i=1}^N V \left[\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t,i} \right] + \frac{2}{N^2} \sum_{i=1}^{N-1} \sum_{k=i+1}^N Cov \left[\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t,i}, \boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t,k} \right] \\
&= \frac{1}{N} \mathbf{P}_0^* + \frac{N-1}{N} E \left[\left(\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t,i} \right) \left(\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t,k} \right)' \right],
\end{aligned} \tag{21}$$

using the assumption that all of our forecasters receive signals with identical distributions. It is possible to show that the current nowcast error is the following function of the previous period’s nowcast error and the intervening innovations:

$$\begin{aligned}
\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t,i} &= \left(\mathbf{I} - \mathbf{P}_1^* \mathbf{H} (\mathbf{H}' \mathbf{P}_1^* \mathbf{H} + \mathbf{R})^{-1} \mathbf{H}' \right) \mathbf{F} \left(\boldsymbol{\xi}_{t-1} - \hat{\boldsymbol{\xi}}_{t-1|t-1,i} \right) \\
&\quad + \left(\mathbf{I} - \mathbf{P}_1^* \mathbf{H} (\mathbf{H}' \mathbf{P}_1^* \mathbf{H} + \mathbf{R})^{-1} \mathbf{H}' \right) \mathbf{v}_t \\
&\quad - \mathbf{P}_1^* \mathbf{H} (\mathbf{H}' \mathbf{P}_1^* \mathbf{H} + \mathbf{R})^{-1} \mathbf{w}_{it} \\
&\equiv \mathbf{A} \left(\boldsymbol{\xi}_{t-1} - \hat{\boldsymbol{\xi}}_{t-1|t-1,i} \right) + \mathbf{B} \mathbf{v}_t + \mathbf{C} \mathbf{w}_{it},
\end{aligned} \tag{22}$$

¹⁵Patton and Timmermann (2008) also consider the behavior of the consensus forecast error but do not analyze cross-sectional dispersion in forecasts.

where \mathbf{v}_t and \mathbf{w}_{it} are defined above. We use this result to derive the covariance between nowcast errors across different forecasters:

$$\begin{aligned}
\mathbf{P}_{0ik}^* &\equiv E \left[\left(\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|i} \right) \left(\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|k} \right)' \right] \\
&= E \left[\left(\mathbf{A} \left(\boldsymbol{\xi}_{t-1} - \hat{\boldsymbol{\xi}}_{t-1|i} \right) + \mathbf{B}\mathbf{v}_t + \mathbf{C}\mathbf{w}_{it} \right) \left(\mathbf{A} \left(\boldsymbol{\xi}_{t-1} - \hat{\boldsymbol{\xi}}_{t-1|k} \right) + \mathbf{B}\mathbf{v}_t + \mathbf{C}\mathbf{w}_{kt} \right)' \right] \\
&= \mathbf{A}E \left[\left(\boldsymbol{\xi}_{t-1} - \hat{\boldsymbol{\xi}}_{t-1|i} \right) \left(\boldsymbol{\xi}_{t-1} - \hat{\boldsymbol{\xi}}_{t-1|k} \right)' \right] \mathbf{A}' + \mathbf{B}\mathbf{Q}\mathbf{B}' + \mathbf{C}E \left[\mathbf{w}_{it}\mathbf{w}_{kt}' \right] \mathbf{C}',
\end{aligned} \tag{23}$$

with all other terms in the two nowcast errors having zero covariance. Letting

$$E \left[\mathbf{w}_{it}\mathbf{w}_{kt}' \right] = \begin{bmatrix} \sigma_\eta^2 & 0 \\ 0 & \sigma_\eta^2 \end{bmatrix} \equiv \mathbf{R}_{ik},$$

we then have

$$\mathbf{P}_{0ik}^* = \mathbf{A}\mathbf{P}_{0ik}^*\mathbf{A}' + \mathbf{B}\mathbf{Q}\mathbf{B}' + \mathbf{C}\mathbf{R}_{ik}\mathbf{C}',$$

which exploits the stationarity of this process, and yields an implicit solution for the covariance of nowcast errors across forecasters, \mathbf{P}_{0ik}^* .¹⁶ Thus the variance of the error of the consensus nowcast of the state vector is:

$$\bar{\mathbf{P}}_0^* \equiv V \left[\boldsymbol{\xi}_t - \bar{\boldsymbol{\xi}}_{t|t} \right] = \frac{1}{N}\mathbf{P}_0^* + \frac{N-1}{N}\mathbf{P}_{0ik}^*. \tag{24}$$

The variance of the consensus forecast of the state vector for $h \geq 1$ can be similarly obtained. Using the following expression for forecast errors as a function of a previous nowcast error and the intervening innovations:

$$\boldsymbol{\xi}_t - \boldsymbol{\xi}_{t|t-h,i} = \mathbf{F}^h \left(\boldsymbol{\xi}_{t-h} - \boldsymbol{\xi}_{t-h|t-h,i} \right) + \sum_{j=0}^{h-1} \mathbf{F}^j \mathbf{v}_{t-j}, \quad h \geq 1, \tag{25}$$

we obtain

$$\begin{aligned}
\bar{\mathbf{P}}_h^* &\equiv V \left[\boldsymbol{\xi}_t - \bar{\boldsymbol{\xi}}_{t|t-h} \right] \\
&= V \left[\frac{1}{N} \sum_{i=1}^N \left(\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t-h,i} \right) \right] \\
&= \frac{1}{N^2} \sum_{i=1}^N V \left[\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t-h,i} \right] + \frac{2}{N^2} \sum_{i=1}^{N-1} \sum_{k=i+1}^N Cov \left[\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t-h,i}, \boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t-h,k} \right] \\
&= \frac{1}{N}\mathbf{P}_h^* + \frac{N-1}{N}E \left[\left(\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t-h,i} \right) \left(\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t-h,k} \right)' \right]
\end{aligned} \tag{26}$$

¹⁶Like other covariance matrices that appear in more standard Kalman filtering applications, see Hamilton (1994), Proposition 13.1 for example, it is not possible to obtain an explicit expression for \mathbf{P}_{0ik}^* .

To evaluate this expression requires knowledge of the covariance between the individual forecast errors measured at different horizons:

$$\begin{aligned}
\mathbf{P}_{hik}^* &\equiv E \left[\left(\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t-h,i} \right) \left(\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t-h,k} \right)' \right] \\
&= E \left[\left(\mathbf{F}^h \left(\boldsymbol{\xi}_{t-h} - \boldsymbol{\xi}_{t-h|t-h,i} \right) + \sum_{j=0}^{h-1} \mathbf{F}^j \mathbf{v}_{t-j} \right) \left(\mathbf{F}^h \left(\boldsymbol{\xi}_{t-h} - \boldsymbol{\xi}_{t-h|t-h,k} \right) + \sum_{j=0}^{h-1} \mathbf{F}^j \mathbf{v}_{t-j} \right)' \right] \\
&= \mathbf{F}^h E \left[\left(\boldsymbol{\xi}_{t-h} - \boldsymbol{\xi}_{t-h|t-h,i} \right) \left(\boldsymbol{\xi}_{t-h} - \boldsymbol{\xi}_{t-h|t-h,k} \right)' \right] \left(\mathbf{F}^h \right)' + E \left[\left(\sum_{j=0}^{h-1} \mathbf{F}^j \mathbf{v}_{t-j} \right) \left(\sum_{j=0}^{h-1} \mathbf{F}^j \mathbf{v}_{t-j} \right)' \right] \\
&= \mathbf{F}^h \mathbf{P}_{0ik}^* \left(\mathbf{F}^h \right)' + \sum_{j=0}^{h-1} \mathbf{F}^j \mathbf{Q} \left(\mathbf{F}^j \right)', \quad h \geq 1. \tag{27}
\end{aligned}$$

With these moment matrices in place it is simple to obtain the term structure of MSE-values for the consensus forecast of the target variable. Let $\boldsymbol{\omega} \equiv [0, \boldsymbol{\iota}'_{12}]'$, where $\boldsymbol{\iota}_k$ is a $k \times 1$ vector of ones, then:

$$V [z_t - \bar{z}_{t|t-h}] = V \left[\boldsymbol{\omega}' \left(\boldsymbol{\xi}_t - \bar{\boldsymbol{\xi}}_{t|t-h} \right) \right] = \boldsymbol{\omega}' \bar{\mathbf{P}}_h^* \boldsymbol{\omega}, \quad \text{for } h \geq 0. \tag{28}$$

The above expression yields 24 moments (the mean squared errors for the 24 forecast horizons) that can be used to help estimate the parameters of the model that govern the dynamics of GDP growth and inflation.

References

- [1] Acemoglu, D., V. Chernozhukov and M. Woldz, 2007, Learning and Disagreement in an Uncertain World. Mimeo, MIT.
- [2] Ang, Andrew, G. Bekaert and M. Wei, 2007, Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better? *Journal of Monetary Economics* 54, 1163-1212.
- [3] Aruoba, B., 2007, Data Revisions are not Well-Behaved, *Journal of Money, Credit and Banking*, forthcoming.
- [4] Bartlett, M.S., 1946, On the theoretical specification of sampling properties of autocorrelated time series, *Journal of the Royal Statistical Society*, Supplement 8, 27-41.
- [5] Capistran, C. and A. Timmermann, 2008, Disagreement and Biases in Inflation Expectations, *Journal of Money, Credit and Banking*, forthcoming.
- [6] Carroll, C., 2003, Macroeconomic Expectations of Household and Professional Forecasters, *Quarterly Journal of Economics*, 118(1), 269-298.
- [7] Clements, M.P., 1997, Evaluating the Rationality of Fixed-Event Forecasts, *Journal of Forecasting* 16, 225-239.
- [8] Corradi, V., A. Fernandez and N.R. Swanson, 2007, Information in the Revision Process of Real-time Data. Mimeo, Rutgers.
- [9] Döpke, J. and U. Fritsche, 2006, When do Forecasters Disagree? An Assessment of German Growth and Inflation Forecast Dispersion, *International Journal of Forecasting*, 22, 125-135.
- [10] Giannone, D., L. Reichlin and D. Small, 2008, Nowcasting GDP: The Real Time Informational Content of Macroeconomic Data Releases, *Journal of Monetary Economics*, 55, 665-676.
- [11] Gourieroux, C. and A. Monfort, 1996a, *Simulation-Based Econometric Methods*, Oxford University Press, Oxford, United Kingdom.
- [12] Gourieroux, C. and A. Monfort, 1996b, *Statistics and Econometric Models, Volume 2*, translated from the French by Q. Vuong, Cambridge University Press, Great Britain.
- [13] Grier, K. and M.J. Perry, 1998, On Inflation and Inflation Uncertainty in the G-7 Countries, *Journal of International Money and Finance* 17.
- [14] Hall, A.R., 2005, *Generalized Method of Moments*, Oxford University Press, U.S.A.
- [15] Hamilton, J.D., 1994, *Time Series Analysis*, Princeton University Press, Princeton, New Jersey.
- [16] Kim, C.-J. and C.R. Nelson, 1999, Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle, *Review of Economics and Statistics*, 81, 608-616.
- [17] Kurz, M., 1994, On the Structure and Diversity of Rational Beliefs. *Economic Theory* 4, 877-900.

- [18] Lahiri, K. and X. Sheng, 2008a, Evolution of Forecast Disagreement in a Bayesian Learning Model. *Journal of Econometrics* 144, 325-340.
- [19] Lahiri, K. and X. Sheng, 2008b, Learning and Heterogeneity in GDP and Inflation Forecasts. Mimeo, University of Albany.
- [20] Lucas, R.E., Jr. 1973, Some International Evidence on Output-Inflation Tradeoffs. *American Economic Review*, 63, 326-334.
- [21] Mankiw, G.N., R. Reis, J. Wolfers, J., 2003. Disagreement about inflation expectations. NBER Macroeconomics Annual.
- [22] McConnell, M.M. and G. Perez-Quiros, 2000, Output Fluctuations in the United States: What Has Changed Since the Early 1980s?, *American Economic Review*, 90, 1464-1476.
- [23] Nordhaus, W.D., 1987, Forecasting Efficiency: Concepts and Applications. *Review of Economics and Statistics* 69, 667-674.
- [24] Patton, A.J. and A. Timmermann, 2008, The Resolution of Macroeconomic Uncertainty: Evidence from Survey Forecasts. Unpublished Manuscript.
- [25] Pesaran, M.H. and M. Weale, 2006, Survey Expectations. Pages 715-776 in G. Elliott, C. Granger and A. Timmermann (eds.) *Handbook of Economic Forecasting*, North Holland: Amsterdam.
- [26] Politis, D.N., and J.P. Romano, 1994, The Stationary Bootstrap, *Journal of the American Statistical Association*, 89, 1303-1313.
- [27] Romer, C.D. and D.H. Romer, 2000, Federal Reserve Information and the Behavior of Interest Rates, *American Economic Review*, 90, 429-457.
- [28] Stock, J.H. and M.W. Watson, 2002, Macroeconomic Forecasting using Diffusion Indexes. *Journal of Business and Economic Statistics* 20, 147-162.
- [29] Stock, J.H. and M.W. Watson, 2006, Forecasting with Many Predictors. Pages 515-554 in Graham Elliott, Clive W.J. Granger and Allan Timmermann (eds.), *Handbook of Economic Forecasting*, Elsevier.
- [30] Tauchen, 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.
- [31] Townsend, R.M., 1983. Forecasting the forecasts of others. *The Journal of Political Economy* 91, 546-588.
- [32] Van Nieuwerburgh, S. and L. Veldkamp, 2006, Learning Asymmetries in Real Business Cycles. *Journal of Monetary Economics* 53, 753-772.
- [33] Veldkamp, L., 2006, Slow Boom, Sudden Crash. *Journal of Economic Theory* 124(2), 230-257.

Table 1: SMM parameter estimates of the joint consensus forecast and constant dispersion model

	σ_u	σ_ε	ϕ	σ_η	σ_ν	κ	σ_μ	J p-val
GDP growth	0.063 (0.012)	0.054 (0.013)	0.936 (0.034)	0.126 (—)	0.692 (1.001)	1.414 (0.941)	0.672 (0.394)	0.857
Inflation	0.000 (—)	0.023 (0.007)	0.953 (0.046)	0.000 (—)	0.045 (0.168)	0.493 (0.167)	0.509 (0.127)	0.000

Notes: This table reports SMM parameter estimates of the Kalman filter model of 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. The parameter σ_η was fixed at $2\sigma_u$ and is reported here for reference only.

Table 2: Testing the significance of differences in signals and differences in prior beliefs

	$H_0 : \sigma_\nu = 0$	$H_0 : \sigma_\mu = 0$
GDP growth	0.478 (0.245)	2.909 (0.044)
Inflation	0.072 (0.394)	16.063 (0.000)

Notes: This table presents the test statistics, with corresponding p -values in parentheses, of the tests for no heterogeneity in signals ($H_0 : \sigma_\nu = 0$) and no heterogeneity in beliefs ($H_0 : \sigma_\mu = 0$). The asymptotic distribution of these test statistics is an equally-weighted mixture of a χ_1^2 and a χ_0^2 variable.

Table 3: SMM parameter estimates of the joint consensus forecast and time-varying dispersion model

	σ_u	σ_ε	ϕ	σ_η	σ_ν	κ	β_0^μ	β_1^μ	J p-val
GDP growth	0.063 (0.012)	0.054 (0.013)	0.936 (0.034)	0.126 (—)	0.692 (0.869)	1.413 (0.738)	-0.560 (1.10)	3.075 (1.832)	0.906
Inflation	0.000 (—)	0.023 (0.007)	0.953 (0.046)	0.000 (—)	0.044 (0.151)	0.493 (0.156)	-1.318 (0.546)	0.178 (2.371)	0.000

Notes: This table reports SMM parameter estimates of the Kalman filter model of 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. The parameter σ_η was fixed at $2\sigma_u$ and is reported here for reference only.

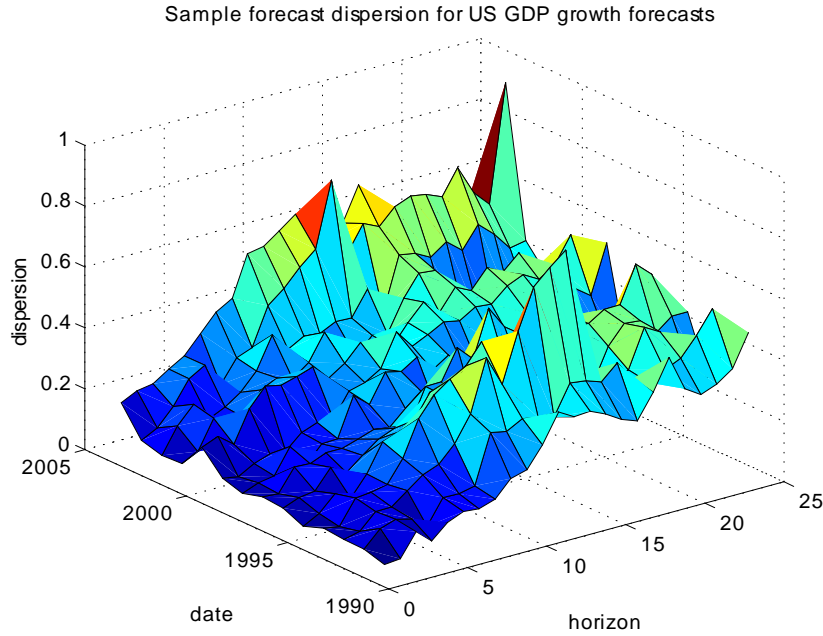


Figure 1: *Cross-sectional dispersion in forecasts of US real GDP growth for different target years, 1991 - 2004, and forecast horizons, $h = 1 - 24$ months.*




Figure 2: *Cross-sectional dispersion in forecasts of US inflation for different target years, 1991 - 2004, and forecast horizons, $h = 1 - 24$ months.*

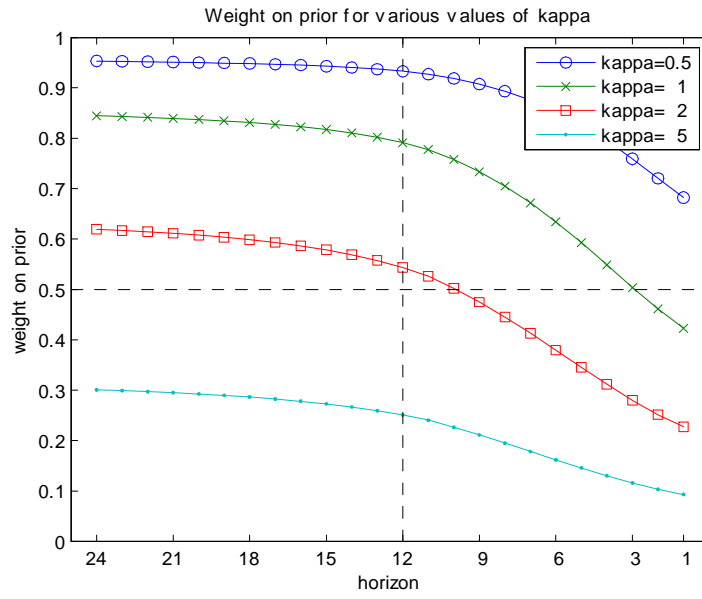


Figure 3: *Weights on prior beliefs about the long-run value of the target variable as a function of the forecast horizon and the parameter κ .*




Figure 4: *Term structure of forecast dispersion for various levels of disagreement in beliefs about the long-run value of the target variable, measured by σ_μ .*





Figure 5: *Cross-sectional dispersion (standard deviation) of forecasts of GDP growth and Inflation in the U.S.*

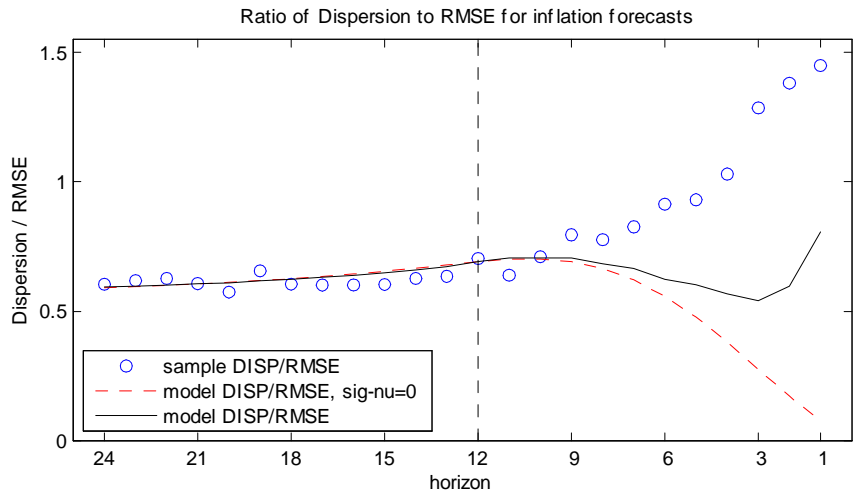
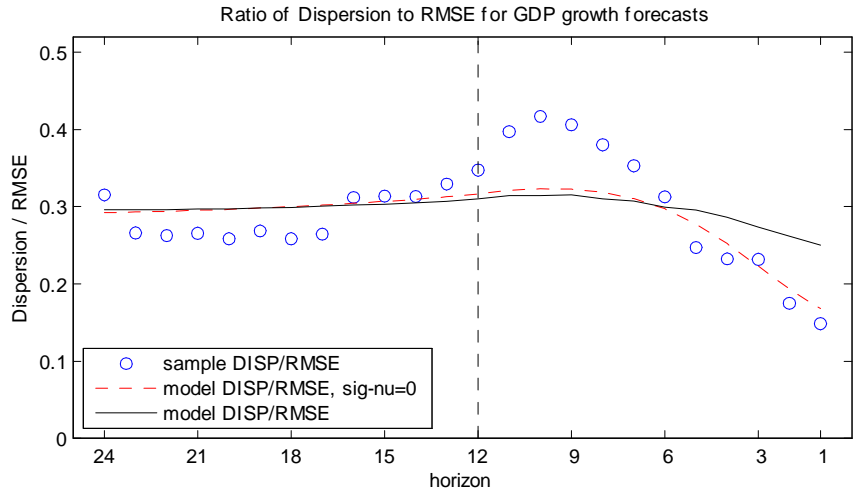


Figure 6: *Ratio of cross-sectional dispersion to root mean squared forecast errors for US GDP growth and Inflation as a function of the forecast horizon.*




Figure 7: *Cross-sectional dispersion (standard deviation) of forecasts of GDP growth and Inflation in the U.S, when the default spread is equal to its sample average, its 95th percentile or its 5th percentile.*