

Andrade, Philippe; Crump, Richard K.; Eusepi, Stefano; Moench, Emanuel

Working Paper

Noisy information and fundamental disagreement

Staff Report, No. 655

Provided in Cooperation with:

Federal Reserve Bank of New York

Suggested Citation: Andrade, Philippe; Crump, Richard K.; Eusepi, Stefano; Moench, Emanuel (2013) : Noisy information and fundamental disagreement, Staff Report, No. 655, Federal Reserve Bank of New York, New York, NY

This Version is available at:

<https://hdl.handle.net/10419/93655>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Federal Reserve Bank of New York
Staff Reports

Noisy Information and Fundamental Disagreement

Philippe Andrade
Richard K. Crump
Stefano Eusepi
Emanuel Moench

Staff Report No. 655
December 2013



This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Noisy Information and Fundamental Disagreement

Philippe Andrade, Richard K. Crump, Stefano Eusepi, and Emanuel Moench

Federal Reserve Bank of New York Staff Reports, no. 655

December 2013

JEL classification: D83, D84, E37

Abstract

We study the term structure of disagreement of professional forecasters for key macroeconomic variables. We document a novel set of facts: 1) forecasters disagree at all horizons, including the very long run; 2) the shape of the term structure of disagreement differs markedly across variables: the term structure is downward-sloping for real output growth, relatively flat for CPI inflation, and upward-sloping for the federal funds rate; 3) disagreement is time varying at all horizons, including the very long run. We suggest a model with noisy information and shifting long-run beliefs that is consistent with these stylized facts. Notably, our model does not rely on the heterogeneity of prior beliefs, bounded rationality, or differences in the precision of signals across agents.

Key words: expectations, survey forecasts, imperfect information, term structure of disagreement

Andrade: Banque de France (e-mail: philippe.andrade@banque-france.fr). Crump, Eusepi, Moench: Federal Reserve Bank of New York (e-mail: richard.crump@ny.frb.org, stefano.eusepi@ny.frb.org, emanuel.monech@ny.frb.org). The authors thank Tobias Adrian for helpful comments and discussions. Felipe Goncalves, Kirby Fears and Matthew Yeaton provided excellent research assistance. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Banque de France, the Federal Reserve Bank of New York, or the Federal Reserve System.

1 Introduction

People, even informed specialists, disagree about unknown economic outcomes. Every survey of expectations, whether it is of consumers, firms, professional forecasters, financial analysts or FOMC members shows that agents have heterogeneous beliefs about the same economic variable. Recent research incorporating heterogeneous beliefs in macroeconomic and finance theories has shown promise in answering empirical questions that have proven to be challenging for a representative agent framework. In particular, sources of disagreement can lead to inertia in price dynamics (Woodford, 2003; Mankiw and Reis, 2002; Maćkowiak and Wiederholt, 2009), non-fundamental driven business cycle fluctuations (Lorenzoni, 2009; Angeletos and La'O, 2013; Rondina and Walker, 2012), speculative bubbles and booms and busts dynamics in asset prices (Scheinkman and Xiong, 2003; Burnside, Eichenbaum, and Rebelo, 2013), or deviations from the expectations hypothesis of the term structure of interest rates (Nimark, 2012).¹ However, while the heterogeneity of beliefs is a key ingredient in all of these theoretical contributions, relatively little is known about the empirical properties of disagreement.

The first contribution of this paper is to document some new facts about disagreement among professional forecasters that should help to discipline the design of imperfect information models with heterogeneous beliefs. More precisely, we use the Blue Chip Financial Forecast (BCFF) survey to characterize the properties of disagreement for US real output growth, consumer price index (CPI) inflation, and the federal funds rate from 1986 to 2013. Our second contribution is to propose a model of expectation formation which is broadly consistent with the new facts.

The stylized facts about disagreement can be illustrated by the two graphs below. The left panel of Figure 1 shows our measure of average disagreement across time for a set of different forecasting horizons ranging from one quarter to 6-to-11 years ahead. Throughout the paper, we define disagreement as the average forecast of the highest ten responses minus that of the lowest ten responses of survey participants for a given variable and forecast horizon.² A first regularity that stands out from this figure is that, for each of the three variables we consider, the disagreement is non-zero even for long-run horizons. We refer to this persistent disagreement among forecasters as *fundamental disagreement*, since it likely

¹See, for example, Hansen (2007), Sargent (2008), and Mankiw and Reis (2010) for general discussions.

²As discussed in Section 2.1, this measure is strongly correlated with alternative measures of disagreement commonly used in the literature.




Figure 1: This figure shows selected statistics for forecast dispersion from the Blue Chip Financial Forecasts survey. Disagreement is defined as the average forecast of the highest ten responses minus that of the lowest ten responses of survey participants. The left panel shows the term structure of disagreement averaged across time for real output growth, CPI inflation, and the federal funds rate for various forecast horizons. The longest horizon captures the average forecast for horizons from 6-to-11 years ahead. The right panel displays the time series of the 6-to-11 years ahead forecast disagreement for the three variables.

captures different views about low-frequency changes in the fundamentals of the economy, such as shifts in productivity growth or monetary and fiscal regime changes. A second striking fact is that the shape of the average *term structure of disagreement* varies across variables. It is downward sloping for real output growth, almost flat for CPI inflation, and upward sloping for the federal funds rate. Hence, the dispersion of opinions about the long term can differ substantially from the disagreement in the short term. Finally, a third fact is shown in the right panel of Figure 1 which reports the time series of the long-run forecasts for the three variables from 1986 through 2013. It underlines that in addition to being non-zero, fundamental disagreement is not constant over time and that it covaries between variables.

These facts suggest that convincing theories in macroeconomics and finance should be able to generate persistent and time-varying disagreement. In this paper, we develop a parsimonious, multivariate model of expectation formation which can broadly match these facts. The model captures two important challenges that economic agents face. The first one is that they do not know the true state of the economy, but instead observe signals which may be noisy. For example, most macroeconomic variables are only published with a delay and are often

subject to large subsequent revisions. More generally, data releases, no matter how accurately measured, offer only a partial snapshot of the economy’s fundamentals. For example, various price indexes like the CPI or the personal consumption expenditure (PCE) deflator are only noisy measures of underlying inflationary pressures. The second challenge is that when facing changes in economic conditions, agents need to distinguish in real time between temporary factors and low frequency changes in fundamentals. The latter capture structural shifts in the economy, as for example, changes in potential GDP growth, the long-term mean of inflation or the natural rate of interest. Precisely, our model is a generalization of the well-known noisy information model of Sims (2003) and Woodford (2003). We add to this setup the assumption that the imperfectly observed state is the sum of two unobserved components: a transitory one and a permanent one, following a long tradition in macroeconomics that goes back at least to Kydland and Prescott (1982). Decompositions into permanent and transitory components also play an important role in finance, in particular the literature on long-run risk models (e.g., Bansal and Yaron, 2004).

We calibrate our multivariate model to the data using a maximum penalized likelihood approach. We choose the models’ parameters to match, as closely as possible, the empirical properties of realized real output growth, CPI inflation and the federal funds rate jointly with selected sample moments from the BCFF survey forecasts. Specifically, we use the volatility of average (i.e., consensus) forecasts from the BCFF survey at different forecasting horizons and the one-quarter ahead disagreement as our moments. The calibration exercise reveals that the model is able to replicate the shapes of the term structure of disagreement that we observe in the data (see Figure 1). Both the presence of a slow-moving drift in each variable and the dynamic interaction between variables, as captured in a multivariate framework, are key in obtaining the results. The unobserved slow-moving drift component in the model contributes to the time variation in forecast disagreement at all horizons, and produces a positive level of disagreement even in the very long run. The multivariate setup of the model is required to generate the different shapes of the term structures of disagreement that we observe in the data. Most importantly, a univariate version of our model cannot generate upward-sloping disagreement for reasonable parameter values. Note that the multivariate dimension is also needed to generate disagreement for variables that are perfectly observed such as the federal funds rate.

Perhaps surprisingly, our model’s modest departure from the homogeneous perfect information rational expectations setup proves to be sufficient to explain short- and long-

term disagreement. Notably, in our generalized version of the noisy information model no agent is endowed with “better” information than any other agent and they all know the true data-generating process (DGP). This stands in contrast to models where agents observe more informative signals either because they have more precise priors or higher signal-to-noise ratios. It also contrasts with models which feature persistent disagreement about the true DGP, either because agents can never fully learn about the true DGP or have immutable priors. The symmetry of agents is an appealing property of our model as it is then consistent with the well-documented fact that the consensus forecast is difficult to beat, i.e., that no individual forecaster has systematically better forecast performance (e.g., Bauer, Eisenbeis, Waggoner, and Zha, 2003; Watson and Stock, 2004). Moreover, Coibion and Gorodnichenko (2012a,b) show, by studying the conditional response of consensus forecast errors to shocks, that noisy information models with the aforementioned asymmetries are not consistent with the data.

There is a large literature in macroeconomics which studies survey data. The properties of consensus or median survey forecasts have been widely documented. In particular, numerous papers have discussed the bias and the efficiency of consensus forecasts (see, for example, Pesaran and Weale, 2006 for a survey) or have used consensus forecasts in model evaluation and estimation (e.g., Roberts, 1995; Adam and Padula, 2011; Del Negro and Eusepi, 2011). More recently, the focus has shifted towards the cross-section of forecasters and in particular the evolution of short-term disagreement as well as forecast uncertainty at the individual level. Mankiw, Reis, and Wolfers (2003) emphasize that disagreement about short-term inflation forecasts in different surveys of the US economy is time varying and somewhat correlated with changes in macroeconomic variables such as inflation and output growth. Dovern, Fritsche, and Slacalek (2012) rely on the Consensus Economics survey to document that the properties of near-term disagreement about real output growth, inflation, and short-term interest rates differ across variables and across G7 countries. A number of papers discuss the relationship between disagreement about US inflation and measures of inflation uncertainty as implied by density forecasts, see for example Rich and Tracy (2010). Lahiri and Sheng (2008) and Patton and Timmermann (2010) study disagreement for fixed-target forecasts up to two years ahead also using the Consensus Economics survey. All of these papers have in common that they investigate the properties of forecast disagreement only up to horizons of at most two years.

Our paper is also linked to the growing literature that uses survey data to understand the

formation of expectations. Mankiw, Reis, and Wolfers (2003) relate the properties of a sticky inflation model to the observed forecast disagreement for future inflation. Carroll (2003) uses consensus forecasts from households and professional forecasters to validate an epidemiological model of expectations. Branch (2004), Coibion and Gorodnichenko (2012a,b) and Andrade and Le Bihan (2013) use survey data to discriminate among various models of expectations including sticky and noisy information models. All these papers confront the implications of the various models with the properties of short-term survey forecasts only. In addition, the existing literature has mostly relied on univariate models (see, for example, the comprehensive study by Coibion and Gorodnichenko, 2012b) with Andrade and Le Bihan (2013) an exception.

Our paper is organized as follows. Section 2 provides a detailed description of the BCFF data and our new set of facts. In Section 3 we introduce our model, discuss its properties and describe how we calibrate it to the data. Our main results are presented in Section 4. Section 5 concludes.

2 Stylized Facts about Disagreement

2.1 Data

We study a unique collection of individual forecasts of real output growth, CPI inflation, and the federal funds rate from the Blue Chip Financial Forecasts (BCFF) survey. This survey, conducted monthly since 1982, asks participants ranging from broker-dealers to economic consulting firms to provide forecasts of the quarterly average of a variety of economic and financial variables for specific calendar quarters as far as six quarters in the future. Importantly, since 1986, this survey has also been collecting information on professional forecasts from as far as 11 years-ahead.

The survey is typically released on the first day of the month, and is based on participants' responses that have been collected during the last week of the previous month. Interest rate forecasts are reported as the average over the target period at an annual rate. Real output and CPI targets are period-over-period percent changes at an annual rate. Real output forecasts are measured with respect to forecasts for real GNP prior to April 1992 and with respect to real GDP thereafter. Since its inception in November 1982, each monthly survey

compiles individual forecasts for horizons of one quarter ahead to at least five quarters ahead. We collect the one-quarter ahead through four-quarters ahead forecasts as the four-quarters ahead forecast is the longest horizon forecast available in every month. Beginning in 1986, twice a year, participants were also surveyed on their longer-term forecasts for a selected set of financial and macroeconomic variables for upcoming calendar years between two and five years ahead along with an average value for a six-to-ten-years ahead horizon. Because the longer-term forecasts refer to specific calendar years and are collected biannually, the forecast horizons vary somewhat across surveys. For example, the horizon we refer to as two-years ahead (2Y) is either six or eight quarters ahead depending on whether we are using the survey taken later in the year or earlier in the year, respectively. When we calibrate the model we mimic this sampling scheme to ensure we are consistent with the survey data. Between March 1986 and March 1996 long-run forecasts are provided in the March and October surveys. From December 1996 onward, long-run forecasts are provided in the June and December releases.³ Beginning in the December 1997 survey, the longest-horizon 5-year-average forecasts shifted from 6-to-10 years ahead to 7-to-11 years ahead. We combine these time series for our analysis to approximately double the number of observations and label the series as the “6-11 years ahead” (6-11Y) forecast for simplicity.

Unfortunately, individual long-run forecasts are not available. Instead the BCFF survey reports the top-10 average long-run forecast and the bottom-10 average long-run forecast. Consequently, at all horizons we use the difference between the top-10 and bottom-10 average forecast as our measure of disagreement. For the shorter term forecasts up to five quarters ahead for which we observe individual forecasts, this measure of disagreement is almost perfectly correlated with the cross-sectional standard deviation of forecasts and highly correlated with the interquartile range of individual forecasts which have both been used as measures of disagreement in the literature.

Although the survey begins in late 1982, our data sample starts with the March 1986 survey and ends with the July 2013 survey, which guarantees we have no missing observations for consensus forecasts or disagreement at all horizons. All data are quarterly where we choose the January, April, July, and October surveys for the short-horizon forecasts matched with the nearest monthly survey which includes long-run forecasts.⁴ This results in 110

³There is one exception to this rule. Long-run forecasts were provided in the January 2003 survey instead of the December 2002 survey.

⁴Recall that surveys are taken at the end of the month previous to the publication date. We choose these survey months as they are based on the maximum amount of information about the current quarter available

observations for nine reported forecast horizons (Q1, Q2, Q3, Q4, 2Y, 3Y, 4Y, 5Y, 6-11Y).

2.2 Three Novel Facts about Forecaster Disagreement

We use the dataset to establish a novel set of stylized facts about forecasters' disagreement. Figure 2 shows the time series of forecaster disagreement for real output growth (upper panel), CPI inflation (middle panel), and the federal funds rate (lower panel) for two forecast horizons: the very short term (one quarter ahead) and the very long term (6-to-11 years ahead). The time series of long term forecast disagreement was already shown in the right panel of Figure 1, but we contrast it here with the equivalent time series for short-term disagreement.

The charts along with Figure 1 document three novel facts about forecaster disagreement. First, forecasters disagree both about the short term but also the medium- and long-run prospects of the economy. Second, the disagreement among forecasters is strongly time varying, even for long term forecasts. Third, the shape of the term structure of disagreement differs markedly across variables. While disagreement at both short- and long-horizons is time varying for all three variables, the ordering of the level of disagreement across horizons differs for each variable. While the professional forecasters in the Blue Chip survey have disagreed more about output growth in the near term than in the long term over the entire sample from 1986 through 2013, the opposite is true about their forecasts of the federal funds rate. Indeed, while there is typically little disagreement about the federal funds rate in the next quarter, forecasters disagree substantially about the level of short term interest rates in the very long run. Interestingly, for CPI inflation disagreement about the short and long term was at similar levels in the late 1980s and the 1990s, but forecasters started to disagree more about near-term than long-term inflation since around the year 2000.

While we only show the time series of disagreement for two different forecast horizons here for simplicity, the left panel of Figure 1 documents the term structures of average disagreement across all forecast horizons. In summary, our data show striking differences across variables: the term structure of disagreement is downward sloping for real output growth, relatively flat for inflation, and upward sloping for the federal funds rate. At first sight, the results for real output growth and inflation appear to be at odds with the findings of Lahiri and Sheng (2008) and Patton and Timmermann (2010) who have studied forecast disagreement up to

to survey participants.

two years into the future using the Consensus Economics survey. These authors argue that disagreement increases with the forecast horizon for both variables. In order to understand the difference between their findings and the ones reported here, it is important to highlight the differences between the two sources of survey data. In the Consensus Economics survey the forecast target, i.e. the value of a variable in a particular calendar year, is held fixed across twenty four consecutive monthly forecasts. This implies that in this survey the forecast horizon is shrinking while time passes. In other words, the information set available to forecasters decreases with the forecast horizon. This is in contrast to the Blue Chip survey that we study which asks participants for forecasts at constant horizons. Hence, when interpreting the empirical findings of Lahiri and Sheng (2008) and Patton and Timmermann (2010), it is important to keep in mind that in these two studies, by the nature of the survey they are based on, the information available to forecasters is not the same across forecast horizons. In contrast, in this paper we take the more conventional view that the information set available to forecasters is fixed in any given period and that based on this same information set forecasts at various horizons into the future are made.

In addition to Lahiri and Sheng (2008) and Patton and Timmermann (2010), a few other papers have studied certain aspects of the disagreement among forecasters. Mankiw, Reis, and Wolfers (2003) document that the disagreement about US inflation expectations up to 17 months ahead from various surveys of consumers and professional forecasters (excluding the Blue Chip survey) is time varying. They also study the correlation of inflation disagreement with changes in macroeconomic variables such as inflation and GDP growth and find weak evidence of such correlations. Doern, Fritsche, and Slacalek (2012) study the behavior of forecasts for real GDP growth, inflation, and short-term interest rates over the next year for the G7 countries. Their analysis is based on the Consensus Economics survey of professional forecasters which is also employed by Lahiri and Sheng (2008) and Patton and Timmermann (2010) for the US. Since that survey does not provide fixed-horizon forecasts for real GDP and inflation, Doern, Fritsche, and Slacalek (2012) approximate these using the reported fixed-event forecasts. Based on their constructed series, they conclude that short-term disagreement about the three variables differs across variables and across G7 countries. Wright (2011) documents that disagreement of one-year ahead inflation forecasts from the Consensus Economics survey is correlated with nominal term premia in a number of countries. He measures disagreement as the cross-sectional standard deviation of individual inflation forecasts and argues that this variable captures inflation uncertainty. Using data on individual point as well as density forecasts from the US Survey of Professional Forecasters,

Rich and Tracy (2010) show that disagreement about US inflation is not systematically related to measures of inflation uncertainty. Boero, Smith, and Wallis (2008) study the relationship between forecast uncertainty and disagreement up to two years into the future for a UK survey of professional forecasts and find a sustained reduction of inflation uncertainty after the introduction of a formal inflation targeting regime by the Bank of England.

One common thread among the papers cited above is that they all study disagreement at horizons of at most two years into the future. To the best of our knowledge, this paper is the first documenting facts about disagreement in the very long term. As we will argue in the next section, these new facts about very long-term forecasts provide information that is important for differentiating the plausibility of various models of expectation formation.

3 Modeling Disagreement

3.1 A Generalized Noisy Information Model

The true state of the macroeconomy is captured by the random vector $z_t = (g_t, \pi_t, i_t)'$ representing real output growth, g_t , inflation, π_t , and the central-bank policy rate i_t . The data generating process for these state variables is,

$$z_t = (I_3 - \Phi) \mu_t + \Phi z_{t-1} + v_t^z, \quad (3.1)$$

$$\mu_t = \mu_{t-1} + v_t^\mu, \quad (3.2)$$

with initial conditions z_0 and μ_0 . We assume all of the eigenvalues of the matrix Φ are inside the unit circle and v_t^z and v_t^μ are *i.i.d.* Gaussian innovations which are mutually independent with variance matrices Σ^z and Σ^μ , respectively. Consequently, the variable μ_t plays the role of the “long-run” component in the sense that $\lim_{h \rightarrow \infty} \mathbb{E}[z_{t+h} | z_t, \mu_t, z_{t-1}, \mu_{t-1} \dots] = \mu_t$. In the following sections we will compare our model to that of one without shifting endpoints (i.e., equation (3.1) with $\mu_t = \mu \forall t$).

The unobserved data can then be written in the compact form,

$$X_t = F X_{t-1} + \epsilon_t, \quad (3.3)$$

where $X_t = (z'_t, \mu'_t)'$, and ϵ_t are *i.i.d.* Gaussian innovations with variance matrix Σ^ϵ and

$$F = \begin{bmatrix} \Phi & (I_3 - \Phi) \\ 0 & I_3 \end{bmatrix}, \quad \Sigma^\epsilon = \begin{bmatrix} I_3 & (I_3 - \Phi) \\ 0 & I_3 \end{bmatrix} \begin{bmatrix} \Sigma^z & 0 \\ 0 & \Sigma^\mu \end{bmatrix} \begin{bmatrix} I_3 & (I_3 - \Phi) \\ 0 & I_3 \end{bmatrix}'. \quad (3.4)$$

There are N agents in our model. Each agent i observes the data $\{y_{it} : t = 1, \dots, T\}$ where

$$y_{it} = H'X_t + \eta_{it}, \quad (3.5)$$

where $H = [I_3 \mid 0_{3 \times 3}]'$ and η_{it} are *i.i.d.* Gaussian observation noise with diagonal variance matrix Σ^η . In words, each agent receives a noisy signal about z_t and uses the current and past history of y_{it} to construct forecasts of each variable. In particular, we assume each agent has full knowledge of the structural parameters defined in equation (3.4) and produces forecasts for z_{t+h} , $h \in \mathbb{Z}_+$, conditional on $\Omega_{it} = \{y_{it}, y_{i(t-1)}, \dots\}$ based on the Kalman filter,

$$\mathbb{E}[z_{t+h} | \Omega_{it}] = F^h X_{t|it}, \quad (3.6)$$

where

$$X_{t|it} = X_{t|i(t-1)} + P_{t|(t-1)} H (H' P_{t|(t-1)} H + \Sigma^\eta)^{-1} (y_{it} - H' X_{t|i(t-1)}), \quad (3.7)$$

$$P_{t|t} = P_{t|(t-1)} - P_{t|(t-1)} H (H' P_{t|(t-1)} H + \Sigma^\eta)^{-1} H' P_{t|(t-1)}, \quad (3.8)$$

$$P_{t|(t-1)} = F P_{(t-1)|(t-1)} F' + \Sigma^\epsilon, \quad (3.9)$$

with initial conditions $X_{0|i0}$ and $P_{0|0}$.

From these recursive formulas it is clear that at any point in time, disagreement about the current state of the macroeconomy depends only on the current realized observation error η_{it} and all past realizations through the difference between the realized state X_t and the agents' previous period forecast $X_{t|i(t-1)}$. Each agent shares the same model and initial condition, $P_{0|0}$, and receives noisy signals drawn from an identical distribution. As a result, each agent has the same (optimal) Kalman gain and no agent will produce systematically better forecasts than any other. At each point in time, each agent updates their estimate of the true state of the macroeconomy, which requires disentangling the “short-term” component from the “long-term” component.

3.2 Model Properties and Predictions

3.2.1 Discussion of the Model

The model incorporates two important informational constraints that forecasters face. First, agents in the model do not perfectly observe the “true state” of the economy, as represented by the vector X_t in the previous section. Second, agents have to infer to what extent changes in the observed variables are due to transitory shocks, as represented by the innovation v_t^z , or reflect changes to the slow-moving permanent components, as captured by the innovation v_t^μ .

The first informational constraint is easily motivated by observing that economic variables such as real GDP and CPI inflation are released with a delay and, feature sizable and significant future revisions (at least in the case of GDP). Hence, the macroeconomic releases that agents observe about these variables in real time are noisy signals of the true state of the economy. More broadly, the noisy measures y_{it} that agents observe can be interpreted as signals about the state of the economy, obtained by using private information or reflecting different weights given to available information. This information friction induces disperse beliefs in a rational expectations framework and has been widely used in many macroeconomic and finance models, see, among others, Morris and Shin (2002), Sims (2003), Woodford (2003), Lorenzoni (2009), Maćkowiak and Wiederholt (2009), and Nimark (2012).

The second constraint implies that agents optimally use different components of the signals they observe for short-term versus long-term forecasts. In particular, they need to filter from the observed data the highly volatile temporary factors from the permanent components of the variables of interest. This decomposition in permanent and transitory elements has a long and widespread tradition in theoretical and empirical macroeconomic research. For instance, the seminal real-business cycle model in Kydland and Prescott (1982) considers such a decomposition of productivity growth. More recently, Sbordone and Cogley (2008) model inflation as having a permanent and a transitory components. Gürkaynak, Sack, and Swanson (2005) study the consequences of such a specification of inflation for the term structure of interest rates.

An important aspect of our model is that no agent has informational advantages over any other, as they each draw from the same distribution of noisy signals. In addition, in our

model agents do not disagree about the model of the economy. Not only do they agree about the model, but they do not have different priors about the model's parameters. As a result, our model implies that in a long enough sample no agent will systematically forecast better than other agents. We think that these are appealing properties in light of the widely documented result that it is difficult to beat consensus forecasts of both survey participants and econometric models (see e.g. Bauer, Eisenbeis, Waggoner, and Zha (2003) or Watson and Stock (2004)).

The choice of this specific model of information frictions is also motivated by its ability to explain important properties of consensus forecasts and the economic variables these forecasts are based upon. Noisy information models have been shown to be able to account for the sluggish adjustment of the consensus forecasts to macroeconomic shocks (Coibion and Gorodnichenko, 2012b). They also show that noisy information models with heterogeneity in priors about long-run means or different signal-to-noise ratios have implications which are not strongly supported in the data. Moreover, several studies show that a time-varying drift captures well both the dynamics properties of variables such as real GDP growth (Stock and Watson, 1989; Cogley and Sargent, 2005; Laubach and Williams, 2003), the inflation rate (Stock and Watson, 2007; Cogley, Primiceri, and Sargent, 2010), and the federal funds rate (Kozicki and Tinsley, 2001; Gürkaynak, Sack, and Swanson, 2005) as well as the slow movements of their consensus long-term expectations (Edge, Laubach, and Williams, 2007; Kozicki and Tinsley, 2012).

3.2.2 Predictions of the Model about Forecast Disagreement

Here we review the main implications of the model presented in the previous section. In order to simplify the discussion, let us assume that the economy is populated by a continuum of forecasters. Rewriting equation (3.7) from above, we see that the h -step ahead optimal forecast of agent i is given by

$$\begin{aligned} z_{t+h|i,t} = & H' F^h X_{t|i(t-1)} \\ & + F^h P_{t|(t-1)} H \left(H' P_{t|(t-1)} H + \Sigma^\eta \right)^{-1} \left(H' (X_t - X_{t|i(t-1)}) + \eta_{it} \right). \end{aligned} \quad (3.10)$$

Then the steady-steady disagreement, captured by the cross-sectional variance of forecasts in the model can be described by this simple expression:

$$V_h^z = H' F^h [(I - GH') \mathbf{V}_1^X (I - GH')' + G \Sigma^\eta G'] (F^h)' H \quad (3.11)$$

where $G = F^{-1} \bar{K}$ with $\bar{K} = F \bar{P} H (H' \bar{P} H + \Sigma^\eta)^{-1}$ denotes the steady-state Kalman gain and \bar{P} the steady-state mean squared error matrix, and where \mathbf{V}_1^X stands for the (time-invariant) cross-sectional variance of agents' predictions in $t - 1$ for the state vector at date t , $X_{t|it-1}$.

We now review how the model can potentially explain the observed term structures of disagreement. We discuss the implications for the time variation of disagreement at the end of this section. Let us start from the simplest possible model and progressively add features as needed to explain the facts. Consider a simple univariate model without shifting endpoints, so that all the terms in equation (3.11) are scalars and $|F| < 1$. Then, it is immediate to see that: (i) for $h \rightarrow \infty$, disagreement tends to zero and (ii) the term structure of disagreement is monotonically decreasing with the forecasting horizon ($F^{2h} \rightarrow 0$). If we add shifting endpoints, the maximum eigenvalue of F is now equal to one. From equation (3.11) it is easy to see that disagreement in the long-run is positive. However, it can be shown that this model can only generate an upward sloping term structure of disagreement for unreasonably large values of the variance of the innovation to the long-run component. Instead, it appears a more natural restriction to assume that the diagonal elements of Σ^z are much larger than those of Σ^μ in a pointwise comparison, since the long-term component is meant to capture a slow moving trend. Thus, a univariate model would not be able to generate the different observed shapes of the term structure under reasonable assumptions.

Consider instead a multivariate model without shifting endpoints. As apparent already from the discussion above this model cannot generate long-term disagreement. However, specific choices for F and Σ^ϵ can deliver any shape of disagreement in the short-run. Intuitively, as the forecasting horizon h increases, some of the off-diagonal elements of F may increase or decrease, generating different patterns of disagreement for different variables. Finally, augmenting this model with shifting endpoints would then inherit these properties along with generating positive fundamental disagreement. One remarkable feature of a multivariate framework is that it does not require idiosyncratic observation noise for all variables in the system in order to generate disagreement for all the variables. This is an appealing feature because for some economic variables such as interest rates and stock prices,

it is difficult to argue that they are imperfectly observed by economic agents.

Finally, we discuss the implications for the time variation in disagreement. Equation (3.11) shows that in this model, taking the limit as the number of forecasters grows, disagreement is constant over time. This is inconsistent with the evidence from the BCFF survey and points to a possible limitation of the model. However, the stylized facts we introduce are, of course, derived for a small number of forecasters. In this case the model does predict some time variation in disagreement and correspondingly some potential correlation in the disagreement measures across different variables. The next section evaluates to what extent the fixed number of forecasters considered here is sufficient to match the second moments observed in the data.

3.3 Calibration

The generalized noisy information model introduced in the previous section appears to have the ability to replicate the key features of our new set of facts. However, it is important to assess the performance of the model in reproducing these facts when parameter values are “reasonable” in the sense of being consistent with the properties of our data. In order to do so, define $\theta = (\Phi, \Sigma^z, \Sigma^\mu)$ and consider the following criterion function,

$$\mathcal{C}(\theta, \Sigma_\eta, \tilde{\Sigma}_\eta; \alpha) = \mathcal{L}(\theta, \tilde{\Sigma}_\eta; \mathcal{Y}_1, \dots, \mathcal{Y}_T) + \alpha \cdot \mathcal{P}(\theta, \Sigma_\eta; \mathcal{S}_1, \dots, \mathcal{S}_T),$$

where \mathcal{Y}_t are the actual output, inflation and interest rate data and \mathcal{S}_t are the BCFF survey data at time t and \mathcal{L} is the negative of the Gaussian likelihood function,

$$\begin{aligned} \mathcal{L}(\theta, \tilde{\Sigma}_\eta; \mathcal{Y}_1, \dots, \mathcal{Y}_T) &= -(2\pi)^{-3/2} \left| H' \tilde{P}_{t|(t-1)} H + \tilde{\Sigma}^\eta \right|^{-1/2} \times \\ &\exp \left\{ -\frac{1}{2} \left(\mathcal{Y}_t - H' \tilde{X}_{t|(t-1)} \right)' \left(H' \tilde{P}_{t|(t-1)} H + \tilde{\Sigma}^\eta \right)^{-1} \left(\mathcal{Y}_t - H' \tilde{X}_{t|(t-1)} \right) \right\}, \end{aligned}$$

where a tilde denotes a variable pertaining to the econometrician (as opposed to the agents in the model). Specifically, we allow for a difference in information available to the econometrician versus the agents via the variables $\tilde{\Sigma}_\eta$ and Σ_η , respectively.⁵ Our

⁵As a robustness check, we also considered specifications where the econometrician has no observation error or when $\Sigma_\eta = \tilde{\Sigma}_\eta$. In both of these cases the model still captures the observed term structures of disagreement.

interpretation of the model is that neither the econometrician nor the agents ever observe the “true” state of the macroeconomy. Instead, agents can observe other measures of economic activity and price pressures whereas we, as econometricians, only use available data on real output growth, headline CPI inflation and the federal funds rate.⁶ Conversely, we do observe (for most of the sample) the revised versions of these variables which the agents do not observe when making their forecasts.

The second term in the criterion function is a penalization term of the observed moments from the survey forecasts relative to the corresponding model-implied moments,

$$\mathcal{P}(\theta, \Sigma_\eta; \mathcal{S}_1, \dots, \mathcal{S}_T) = (g(\theta, \Sigma_\eta) - g_S(\mathcal{S}_1, \dots, \mathcal{S}_T))' W (g(\theta, \Sigma_\eta) - g_S(\mathcal{S}_1, \dots, \mathcal{S}_T)),$$

where W is a positive semi-definite weighting matrix and $g_S(\mathcal{S}_1, \dots, \mathcal{S}_T)$ is a collection of moments from the data:

- We use 15 sample moments (5 sample moments for each of the three variables)
 - Our disagreement measure for the one-quarter ahead forecast only;
 - The standard deviation of consensus forecast for one- and four-quarters ahead along with two-years ahead and six-to-eleven years ahead.

The corresponding model-implied statistics are constructed by the function $g(\theta, \Sigma_\eta)$ via a simulation approach:

- We simulate the model using $T = 120$ (approximately the length our sample period) and choose $N = 50$ (consistent with number of participants in the survey) across 100 simulations in our optimization procedure;
- We choose a diagonal weighting matrix which places a weight of 1 on the disagreement related moments and a weight of 0.1 on standard deviation related moments.

The moments and weight matrix are selected so as to choose parameter values such that, as closely as possible, the level of the model-implied one-quarter ahead disagreement is consistent with the data. We can then evaluate the performance of the model to match the

⁶We use real GNP and GDP data as provided by the Bureau of Economic Analysis, headline CPI from the Bureau of Labor Statistics and the federal funds rate from the H.15 data provided by the Board of Governors of the Federal Reserve.

term structure of disagreement using the least amount of disagreement data to do so. Very loosely, we are “normalizing” the model so that the shortest-horizon forecast disagreement is approximately correct.

We solve $\min_{\theta, \Sigma_\eta, \tilde{\Sigma}_\eta} \mathcal{C}(\theta, \Sigma_\eta, \tilde{\Sigma}_\eta)$. The calibrated parameters are provided in Tables 1 and 2 for a value of $\alpha = 1$ and $\alpha = 50$, respectively.⁷ We choose to show two disparate values for the penalty parameter to illustrate that model-implied disagreement is relatively insensitive to large shifts in α . As we will discuss in the next section in more detail we view $\alpha = 50$ as our “baseline” calibration as it ensures that the volatility of model-implied consensus forecasts match the data well across horizons. Note first that the observation errors facing the econometrician are smaller than those facing the agents for real output growth and CPI inflation, and the observation error for the federal funds is approximately zero for both calibrations.⁸ These differences for output and inflation line up with the intuition for our model as the econometrician uses revised data without publication lag for the majority of the sample, but never observes the true state of the macroeconomy.

For both values of the penalty parameter, Φ has a maximum eigenvalue near (but below) one. The calibrated Φ matrix for $\alpha = 50$ has a higher maximum eigenvalue than that of $\alpha = 1$. The counterpart of this result is that the volatility of the long-run component, as measured by the maximum eigenvalue of Σ^μ , is noticeably lower whereas there is only a modest difference in the volatility of the short-run component. In accordance with our interpretation of the role of the long-term component, we see that the maximum eigenvalue of Σ^μ is considerably lower than that of Σ^z . Figure 3 shows the filtered variables as processed by the econometrician for real output growth, CPI inflation and the federal funds rate based on the two values of the penalty parameter. Despite the aforementioned differences in calibrated parameters between the two values of α , there are only small differences in the filtered variables z_t . However, as we show in the next section, the implications for the forecasts produced by the agents will be different for these two sets of parameters.

The final input necessary to the model calibration is the choice of initial conditions. Conceptually, the choice hinges on what is the most reasonable assumption about the recent past of the state of the economy. If the economy has not featured any recent structural changes, then it is reasonable to believe that expectations in the beginning of the sample are generated by the invariant (steady-state) distribution of the model. However, if the

⁷Results for other values of α and other specifications are provided in an online supplemental appendix.

⁸All reported standard deviations in Tables 1 and 2 are in annualized terms.

economy has been subject to recent structural changes, then the sample information could reflect a transition to the new steady-state distribution. In the years leading up to our sample period, it appears uncontroversial to say that, at the least, there were structural changes in the conduct of monetary policy. This is captured by the gradual decline in both disagreement and the consensus forecast for CPI inflation and the federal funds rate. However, we have the advantage that we can observe, at least partially, the dispersion in agents’ expectations at the beginning of the sample. In order to calibrate the model we use information from the March 1986 BCFF survey to provide initial conditions for both the $z_{it|t}$ and the $\mu_{it|t}$. For the former, we use the forecasts for the first quarter of 1986 from the March 1986 BCFF survey as a “nowcast”. For the latter, we do not observe individual long-term forecasts, so instead we scale the initial conditions from the nowcast to replicate the 6-to-11 years ahead disagreement measured by the top-10 average minus the bottom-10 average in the same survey.⁹

We want to emphasize that we do not interpret the variation in initial conditions as a reflection in different priors about the structural parameters of the economy, but rather as a result of past observation errors that occurred prior to our sample period. Regardless, in the next section we show that removing the influence of the initial conditions does not alter the main conclusions implied by the model.

4 Main Results

4.1 Baseline Results

We start by discussing the model’s implications for the term structure of disagreement. For each of the three variables, we report the observed term structure and its model-implied counterpart obtained from 5,000 simulations. Confidence intervals presented in the figures are constructed using the appropriate simulation-based quantiles for each forecast horizon. Recall that the calibrated variance for the observation error for the federal funds rate was very close to zero for the econometrician and the agents. Throughout the simulation experiments we set it equal to exactly zero so that the short-term interest rate is perfectly observed. In order to document that our results are insensitive to the relative weight the

⁹The March 1986 survey only includes forecasts for 45 participants. The additional 5 agents in our model are endowed with initial conditions equal to the median of the survey data.

estimation procedure attributes to the target moments of cross-sectional forecast dispersion and standard deviation of consensus forecasts, we present the results for two choices of the penalty parameter, $\alpha = 1$ and $\alpha = 50$. It is important to reiterate that only the disagreement of one-quarter ahead forecasts is used in the penalization function so that our calibration does not constrain the disagreement at other horizons. In fact, the calibrated model replicates both the level and the shape of the term structure of disagreement observed for the three variables, as Figure 4 illustrates: the model-implied term structure for real output growth is downward sloping, for CPI inflation it is approximately flat, and for the federal funds rate it is upward sloping. Moreover, these results hold for both values of the penalization parameter α . In the remainder of this section, we use the parameters implied by $\alpha = 50$ as the baseline and defer further discussion of the robustness with respect to the penalty to Section 4.2.

Which feature of our model drives these results? The left panel of Figure 5 compares the previous term structures of disagreement (for $\alpha = 50$) with the ones obtained from the equivalent noisy information model without shifting endpoints. We calibrate this model using the same method as for the generalized model. We use the same moments of disagreement (one-quarter ahead only for each variable) but only the one-quarter and four-quarter ahead standard deviations of consensus forecasts, as this model cannot generate variability in long-term forecasts. The model without shifting endpoints clearly falls short at explaining disagreement for all but the shortest horizons. As expected, for long horizons the disagreement implied by this model approaches zero for all variables. Of note, the model with shifting endpoints provides a better fit to disagreement at horizons above one year without compromising the fit of short-term disagreement. This improved performance is not entirely obvious as the model with shifting endpoints has six more parameters, but the calibration imposes six additional restrictions which discipline the volatility of the model-implied longer-term consensus forecasts.

These results confirm our analysis in Section 3.2.2. In summary, the introduction of shifting endpoints results in a dramatic improvement of the fit of the term structure of disagreement, especially for horizons above one year.

We next compare the performance of the two models in terms of the variability of model-implied consensus forecasts, which is also used in the calibration. The right panel of Figure 5 presents the standard deviation of consensus forecasts from the BCFF survey along with the model-implied standard deviations. The model with shifting endpoints captures well the term structure of consensus forecast volatility for all three variables. Moreover, the BCFF

term structure of consensus forecast volatility is within the bands implied by the model with shifting endpoints for all three variables with the exception of the three-quarter ahead and four-quarter ahead forecasts of real output growth. Except for CPI inflation, the model without shifting endpoints has a comparable fit at short to intermediate horizons. In sum, the model with shifting endpoints does a much better job at replicating the two first facts about forecast disagreement that we document. We next turn to a discussion of the third fact related to the time variance and cross-correlation of disagreement.

The left column of Figure 6 shows the time variance of disagreement from the BCFF data and its model-implied counterpart. Although our model qualitatively matches the different shapes of the term structure of the time variance of disagreement, it is not capable of explaining the levels observed in the data. This is not surprising as we showed in Section 3.2.2 that the time variance of disagreement goes to zero when the number of forecasters goes to infinity. In the right column of the figure we show the pairwise time series correlations of disagreement for different horizons for both the model and the data. We start by briefly discussing the properties of these correlations in the data. First, note that in the survey there is a substantial degree of correlation among the three time series of disagreement at various horizons. An interesting exception is the correlation between disagreement about CPI inflation and the federal funds rate at the one-quarter ahead horizon which almost exactly equals zero. In contrast, the same correlation between the federal funds rate and output growth disagreements equals 60 percent in the data. Consequently, the small amount of short-term disagreement observed for the federal funds rate appears to be to a large extent driven by disagreement about near term-growth prospects. At long horizons, the correlation between disagreement about the federal funds rate and CPI inflation (real output growth) forecasts is more than 80 (60) percent correlated in the data. Hence, long term disagreement about the federal funds rate is clearly driven by disagreement about the determinants of interest rate policy. Finally, note that the disagreement about real output growth and inflation are positively correlated at all forecast horizons including the long term. This is likely driven by the heterogeneity of views about common underlying determinants of potential output and long-run inflation.¹⁰ Moreover, these results provide an additional motivation for adopting a multivariate framework to study the evolution of survey forecasts. While the model does not exactly replicate these correlations at all horizons, the qualitative

¹⁰Carvalho and Nechio (2013) show that at least some subgroups of survey participants form their expectations about the future path of interest rates, inflation, and unemployment in a way that is consistent with simple monetary policy rules.

features of the bivariate correlation term structures are still well captured.

In the following figures we show some other summary measures and compare them to their model-based counterparts. We start by assessing whether disagreement is correlated with the business cycle, as proxied by real GDP growth. To do so, we generate model-implied target variables observed by the econometrician by calculating $\tilde{\mathcal{Y}}_t = z_t + \tilde{\eta}_t$ where $\tilde{\eta}_t$ is the *i.i.d.* Gaussian observation error of the econometrician (i.e., $\tilde{\eta}_t \sim \mathcal{N}(0, \tilde{\Sigma}_\eta)$). The left column of Figure 7 shows the correlation between real output growth and disagreement for all three variables in the data and in the model. Since disagreement in the model is generated by an observation error that has a constant variance the model-implied correlations are all centered around zero. In contrast, the correlations observed in the data are non-zero for some horizons. In particular, the disagreement about one-quarter ahead inflation forecasts is negatively correlated with real output growth. However, we find that our model is largely consistent with the data as for almost all horizons and variables the sample correlations are comfortably within the 95% confidence bands implied by the model.

We perform a comparable exercise to study whether disagreement about a variable is correlated with the level of that same variable, see the right column of Figure 7. Again, the data correlations are within the 95% bands of the corresponding model-implied correlations with the exception of the federal funds rate which shows non-zero positive correlations at horizons up to one year. The likely explanation is that both the level of the federal funds rate as well as the disagreement were high at the beginning of the sample and then both trended down over the first fifteen years of our sample.

In sum, our model produces disagreement that matches all of the qualitative properties and most of the quantitative properties of the new facts as well as some other features of the data.

4.2 Robustness

In this section, we assess the robustness of our findings relative to the choice of penalty parameter and the initial conditions of the filter. Figure 8 reports the model-implied term structures of disagreement and standard deviations of consensus forecasts along with their data counterparts for different values of the penalty parameter: $\alpha = 1$, $\alpha = 5$, and $\alpha = 50$. We first note that the choice of penalty does not substantially affect the model-implied term

structures of disagreement. Notice, however, that small penalty parameters imply standard deviations of consensus forecasts that do not match the survey data as well as when $\alpha = 50$. This reflects the fact that for a small penalty parameters the estimated volatility of the long-term component is higher than under the baseline.

Figure 9 displays the sensitivity of our results with respect to the initial conditions. More precisely we report the same moments as above, simulating the model for 240 periods (twice the original sample size) and discarding the first 120 observations of the simulated paths. We do so in order to assess the impact of the initial levels of disagreement on our calibration results. The long-run disagreement is slightly lower for all three variables in this exercise, see the left column of Figure 9. This result is consistent with the term structure of disagreement in our sample which generally shifts lower if the early years of our sample are excluded.

The right column of the figure shows the impact of the initial conditions on the model-implied standard deviation of consensus forecasts. Across all three variables, the mean values are nearly indistinguishable highlighting that the initial conditions do not affect these predictions of the model.

To summarize, the results in this section have shown that our results are robust to changes in the penalization parameter as well as to changes in the initial conditions of the underlying filter.

5 Conclusion

This paper documents a novel set of facts about disagreement among professional forecasters: (i) forecasters disagree at all horizons including the very long run; (ii) the shape of the term structure of disagreement differs markedly across variables: the term structure is downward sloping for real output growth, relatively flat for CPI inflation, and upward sloping for the federal funds rate; (iii) disagreement is time varying at all horizons including the very long run. We introduce a model of expectation formation that is able to broadly replicate these results. The model incorporates two important informational constraints that forecasters face: first, agents in the model do not perfectly observe the “true state” of the economy; second, agents have to allocate changes in the observed variables to changes in transitory and long-term factors.

An important aspect of our model is that no agent has informational advantages over any other and agents have rational expectations and full knowledge about the structure of the economy. Our model captures the main features of the term structure of disagreement well; however, it does not generate enough time variation in disagreement compared to the survey data. Several extensions to our model could be introduced to overcome this limitation. One approach might be to relax our strict assumptions and assume that agents do not have full knowledge of the DGP, for example, if they must learn about the parameters. An alternative approach would be to endow the model with endogenously generated time variation in the precision of signals that depends on the state of the economy as in Van Nieuwerburgh and Veldkamp (2006). Finally, having shown that our proposed model of expectation formation matches various important features of survey forecasts, one avenue for future research would be to embed the model in a general equilibrium setup.

References

- ADAM, K., AND M. PADULA (2011): “Inflation Dynamics And Subjective Expectations In The United States,” *Economic Inquiry*, 49(1), 13–25.
- ANDRADE, P., AND H. LE BIHAN (2013): “Inattentive Professional Forecasters,” *Journal of Monetary Economics*, forthcoming.
- ANGELETOS, G.-M., AND J. LA’O (2013): “Sentiments,” *Econometrica*, 81(2), 739–779.
- BANSAL, R., AND A. YARON (2004): “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *Journal of Finance*, 59(4), 1481–1509.
- BAUER, A., R. A. EISENBEIS, D. F. WAGGONER, AND T. ZHA (2003): “Forecast evaluation with Cross-Sectional data: The Blue Chip Surveys,” Federal Reserve Bank of Atlanta *Economic Review*, Q2, 17–31.
- BOERO, G., J. SMITH, AND K. F. WALLIS (2008): “Uncertainty and Disagreement in Economic Prediction: The Bank of England Survey of External Forecasters,” *Economic Journal*, 118(530), 1107–1127.
- BRANCH, W. A. (2004): “The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations,” *Economic Journal*, 114(497), 592–621.
- BURNSIDE, C., M. EICHENBAUM, AND S. REBELO (2013): “Understanding Booms and Busts in Housing Markets,” Working Paper 16734, National Bureau of Economic Research.
- CARROLL, C. D. (2003): “Macroeconomic Expectations of Households and Professional Forecasters,” *Quarterly Journal of Economics*, 118(1), 269–298.
- CARVALHO, C., AND F. NECHIO (2013): “Do People Understand Monetary Policy?,” Textos para discussão, Department of Economics PUC-Rio (Brazil).
- COGLEY, T., G. E. PRIMICERI, AND T. J. SARGENT (2010): “Inflation-Gap Persistence in the US,” *American Economic Journal: Macroeconomics*, 2(1), 43–69.
- COGLEY, T., AND T. J. SARGENT (2005): “Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S,” *Review of Economic Dynamics*, 8(2), 262–302.
- COIBION, O., AND Y. GORODNICHENKO (2012a): “Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts,” Working paper.
- (2012b): “What Can Survey Forecasts Tell Us About Informational Rigidities?,” *Journal of Political Economy*, 120(1), 116–159.
- DEL NEGRO, M., AND S. EUSEPI (2011): “Fitting Observed Inflation Expectations,” *Journal of Economic Dynamics and Control*, 35(12), 2105–2131.
- DOVERN, J., U. FRITSCHKE, AND J. SLACALEK (2012): “Disagreement Among Forecasters in G7 Countries,” *Review of Economics and Statistics*, 94(4), 1081–1096.
- EDGE, R. M., T. LAUBACH, AND J. C. WILLIAMS (2007): “Learning and Shifts in Long-run Productivity Growth,” *Journal of Monetary Economics*, 54(8), 2421–2438.
- GÜRKAYNAK, R., B. SACK, AND E. T. SWANSON (2005): “The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models,” *American Economic Review*, 95(1), 425–436.

- HANSEN, L. P. (2007): “Beliefs, Doubts and Learning: Valuing Macroeconomic Risk,” *American Economic Review*, 97(2), 1–30.
- KOZICKI, S., AND P. A. TINSLEY (2001): “Shifting Endpoints in the Term Structure of Interest Rates,” *Journal of Monetary Economics*, 47(3), 613–652.
- (2012): “Effective Use of Survey Information in Estimating the Evolution of Expected Inflation,” *Journal of Money, Credit and Banking*, 44(1), 145–169.
- KYDLAND, F. E., AND E. C. PRESCOTT (1982): “Time to Build and Aggregate Fluctuations,” *Econometrica*, 50(6), 1345–70.
- LAHIRI, K., AND X. SHENG (2008): “Evolution of Forecast Disagreement in a Bayesian Learning Model,” *Journal of Econometrics*, 144(2), 325–240.
- LAUBACH, T., AND J. C. WILLIAMS (2003): “Measuring the Natural Rate of Interest,” *The Review of Economics and Statistics*, 85(4), 1063–1070.
- LORENZONI, G. (2009): “A Theory of Demand Shocks,” *American Economic Review*, 99(5), 2050–2084.
- MAĆKOWIAK, B., AND M. WIEDERHOLT (2009): “Optimal Sticky Prices under Rational Inattention,” *American Economic Review*, 99(3), 769–803.
- MANKIW, N. G., AND R. REIS (2002): “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, 117(4), 1295–1328.
- (2010): “Imperfect Information and Aggregate Supply,” in *Handbook of Monetary Economics*, ed. by B. M. Friedman, and M. Woodford, vol. 3 of *Handbook of Monetary Economics*, chap. 5, pp. 183–229. Elsevier.
- MANKIW, N. G., R. REIS, AND J. WOLFERS (2003): “Disagreement about Inflation Expectations,” in *NBER Macroeconomics Annual 2003*, vol. 18, pp. 209–270. National Bureau of Economic Research.
- MORRIS, S., AND H. S. SHIN (2002): “Social Value of Public Information,” *American Economic Review*, 92(5), 1521–1534.
- NIMARK, K. P. (2012): “Speculative Dynamics in the Term Structure of Interest Rates,” Working paper.
- PATTON, A. J., AND A. TIMMERMANN (2010): “Why do Forecasters Disagree? Lessons from the Term Structure of Cross-sectional Dispersion,” *Journal of Monetary Economics*, 57(7), 803–820.
- PESARAN, M. H., AND M. WEALE (2006): “Survey Expectations,” in *Handbook of Economic Forecasting*, ed. by C. W. G. Graham Elliott, and A. Timmermann, pp. 715–776. North-Holland.
- RICH, R., AND J. TRACY (2010): “The Relationships Between Expected Inflation, Disagreement and Uncertainty: Evidence from Matched Point and Density Forecasts,” *Review of Economics and Statistics*, 92(1), 200–207.
- ROBERTS, J. M. (1995): “New Keynesian Economics and the Phillips Curve,” *Journal of Money, Credit, and Banking*, 27, 975–984.
- RONDINA, G., AND T. B. WALKER (2012): “Information Equilibria in Dynamic Economies with Dispersed Information,” Working paper.

- SARGENT, T. J. (2008): “Evolution and Intelligent Design,” *American Economic Review*, 98(1), 5–37.
- SBORDONE, A. M., AND T. COGLEY (2008): “Trend Inflation, Indexation and Inflation Persistence in the New Keynesian Phillips Curve,” *American Economic Review*, 98(5), 2101–2126.
- SCHEINKMAN, J. A., AND W. XIONG (2003): “Overconfidence and Speculative Bubbles,” *Journal of Political Economy*, 111(6), 1183–1220.
- SIMS, C. A. (2003): “Implications of Rational Inattention,” *Journal of Monetary Economics*, 50(3), 665–690.
- STOCK, J. H., AND M. W. WATSON (1989): “New Indexes of Coincident and Leading Economic Indicators,” in *NBER Macroeconomics Annual 1989*, vol. 4, pp. 351–409. National Bureau of Economic Research.
- (2007): “Why Has U.S. Inflation Become Harder to Forecast?,” *Journal of Money, Credit and Banking*, 39(s1), 3–33.
- VAN NIEUWERBURGH, S., AND L. VELDKAMP (2006): “Learning Asymmetries in Real Business Cycles,” *Journal of Monetary Economics*, 53(4), 753–772.
- WATSON, M. W., AND J. H. STOCK (2004): “Combination Forecasts of Output Growth in a Seven-Country Data Set,” *Journal of Forecasting*, 23(6), 405–430.
- WOODFORD, M. (2003): “Imperfect Common Knowledge and the Effects of Monetary Policy,” in *Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, ed. by P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, pp. 25–58. Princeton University Press.
- WRIGHT, J. H. (2011): “Term Premia and Inflation Uncertainty: Empirical Evidence from an International Panel Dataset,” *American Economic Review*, 101(4), 1514–34.

Table 1: **Results of Calibration for $\alpha = 1$**

Φ	$ \text{eig}(\Phi) $	$\text{sqrt}(\text{diag}(\tilde{\Sigma}^\eta))$	$\text{sqrt}(\text{diag}(\Sigma^\eta))$
$\begin{bmatrix} 0.455 & -0.613 & -0.009 \\ 0.142 & 1.004 & -0.072 \\ 0.157 & 0.127 & 0.905 \end{bmatrix}$	$\begin{bmatrix} 0.946 \\ 0.730 \\ 0.730 \end{bmatrix}$	$\begin{bmatrix} 2.461 \\ 1.585 \\ 0.000 \end{bmatrix}$	$\begin{bmatrix} 3.812 \\ 2.193 \\ 0.000 \end{bmatrix}$
$(\Sigma^z)^{1/2}$	$\text{eig}(\Sigma^z)$	$(\Sigma^\mu)^{1/2}$	$\text{eig}(\Sigma^\mu)$
$\begin{bmatrix} 1.546 & 0 & 0 \\ 0.270 & 0.698 & 0 \\ 0.393 & 0.505 & 0.598 \end{bmatrix}$	$\begin{bmatrix} 2.729 \\ 0.795 \\ 0.191 \end{bmatrix}$	$\begin{bmatrix} 0.140 & 0 & 0 \\ 0.310 & 0.043 & 0 \\ 0.587 & -0.049 & 0.04 \end{bmatrix}$	$\begin{bmatrix} 0.461 \\ 0.004 \\ 0.000 \end{bmatrix}$

Table 2: **Results of Calibration for $\alpha = 50$**

Φ	$ \text{eig}(\Phi) $	$\text{sqrt}(\text{diag}(\tilde{\Sigma}^\eta))$	$\text{sqrt}(\text{diag}(\Sigma^\eta))$
$\begin{bmatrix} 0.473 & -0.711 & 0.115 \\ 0.107 & 0.925 & -0.030 \\ 0.144 & 0.080 & 0.939 \end{bmatrix}$	$\begin{bmatrix} 0.976 \\ 0.696 \\ 0.696 \end{bmatrix}$	$\begin{bmatrix} 2.378 \\ 1.587 \\ 0.072 \end{bmatrix}$	$\begin{bmatrix} 4.377 \\ 2.969 \\ 0.002 \end{bmatrix}$
$(\Sigma^z)^{1/2}$	$\text{eig}(\Sigma^z)$	$(\Sigma^\mu)^{1/2}$	$\text{eig}(\Sigma^\mu)$
$\begin{bmatrix} 1.697 & 0 & 0 \\ 0.322 & 0.770 & 0 \\ 0.344 & 0.542 & 0.583 \end{bmatrix}$	$\begin{bmatrix} 3.209 \\ 0.926 \\ 0.195 \end{bmatrix}$	$\begin{bmatrix} 0.068 & 0 & 0 \\ 0.161 & 0.115 & 0 \\ 0.146 & -0.040 & 0.011 \end{bmatrix}$	$\begin{bmatrix} 0.056 \\ 0.011 \\ 0.000 \end{bmatrix}$

Figure 2: **Time Series of Disagreement**

This figure shows the time series of forecaster disagreement as measured by the average forecast of the highest ten responses minus that of the lowest ten responses for the shortest and longest forecast horizons from the Blue Chip Financial Forecasts survey. The sample period is from 1986Q1 - 2013Q2.

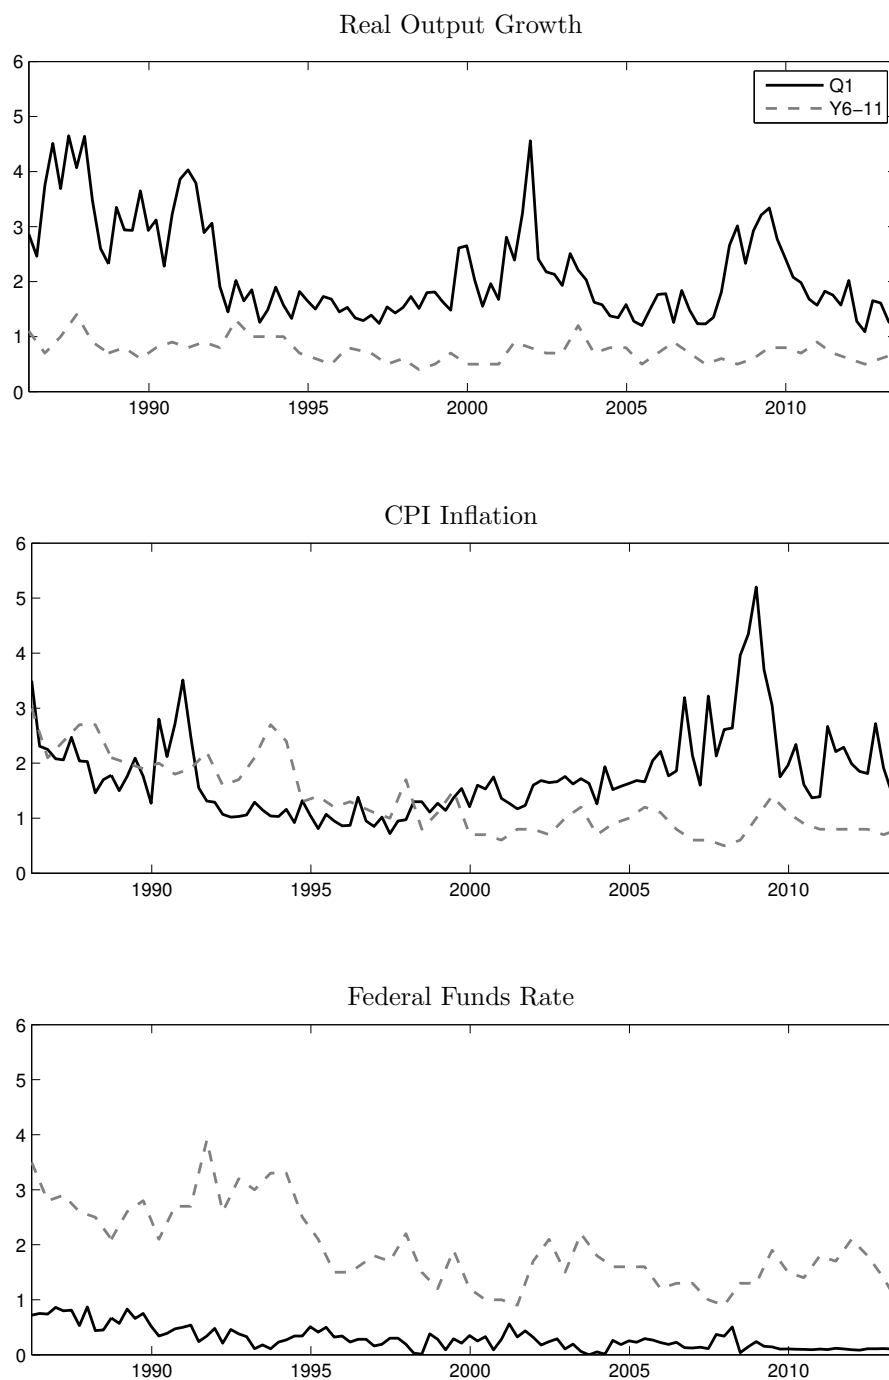


Figure 3: **Filtered Macroeconomic Variables**

This figure shows the filtered z_t (one-sided) as processed by the econometrician for $\alpha = 1$ (light blue) and $\alpha = 50$ (dark blue) along with the observed data (red). The sample period is from 1986Q1 - 2013Q2.

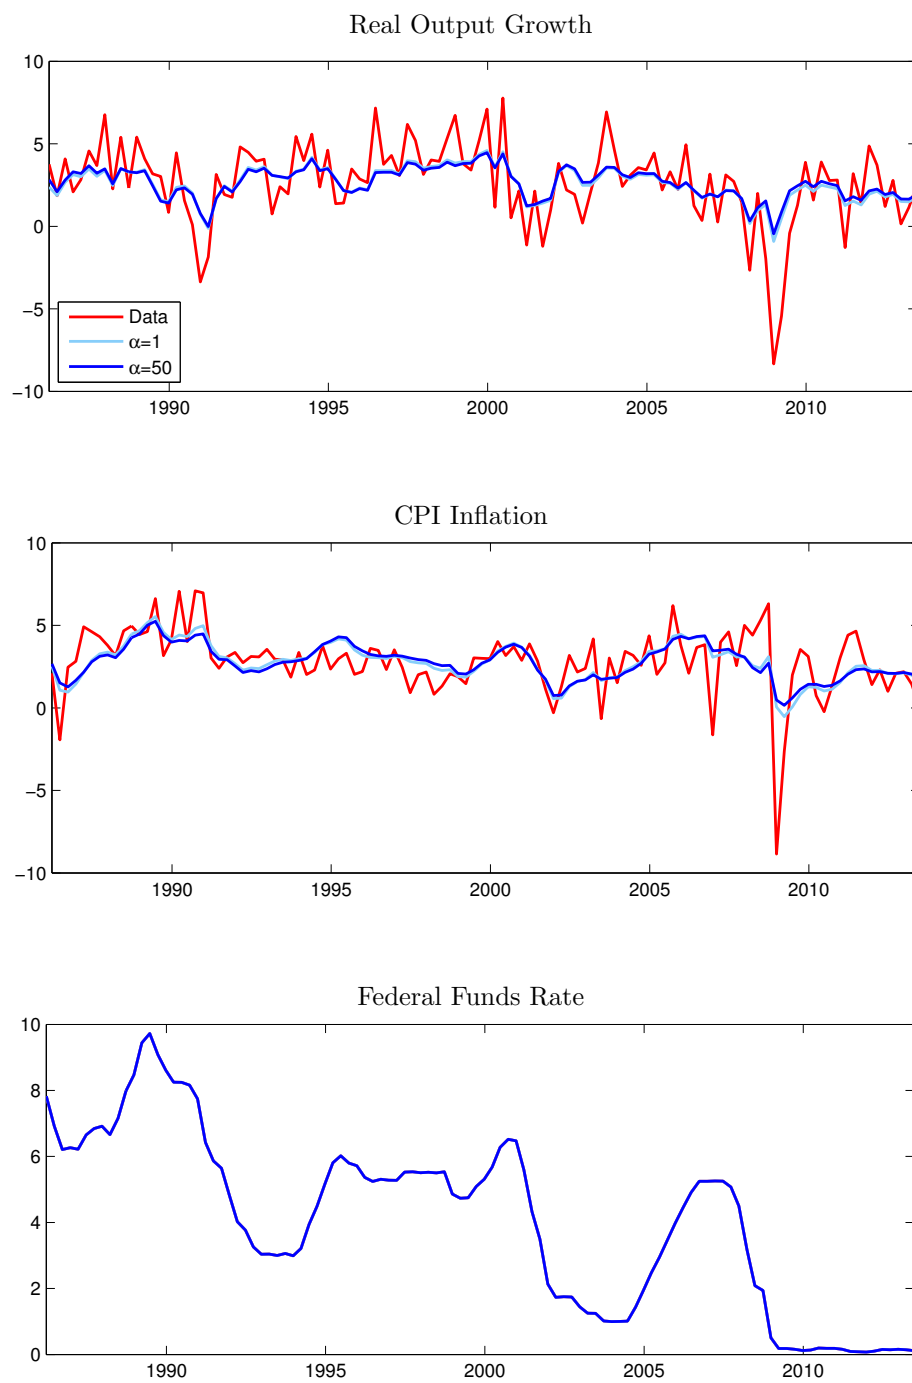


Figure 4: **Term Structure of Disagreement**

This figure displays the model-implied (time) average of disagreement across different horizons for the generalized noisy information model calibrated with $\alpha = 1$ (light blue) and $\alpha = 50$ (dark blue) along with the Blue Chip Financial Forecasts survey (red). Model-implied 95% confidence intervals for the calibrated parameters based on $\alpha = 1$ and $\alpha = 50$ are designated by dotted lines and shaded regions, respectively.

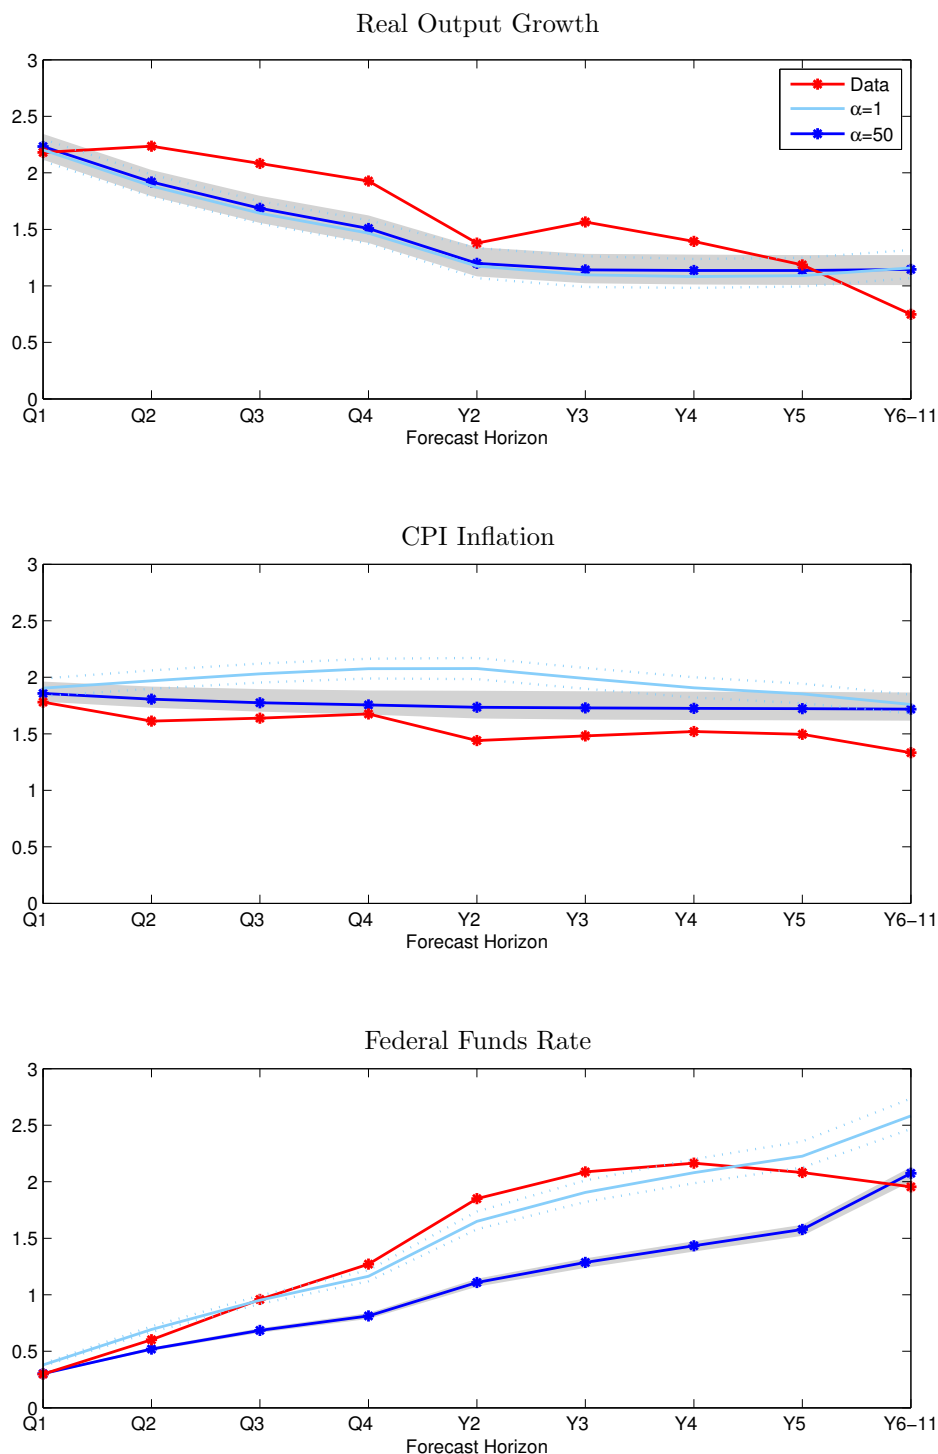


Figure 5: **Disagreement and Standard Deviation of Forecasts**

The first column displays the model-implied disagreement for the generalized noisy information model calibrated with $\alpha = 50$ (blue) and the noisy information model without shifting endpoints calibrated with $\alpha = 50$ (green) along with the Blue Chip Financial Forecasts survey (red). The second column displays the corresponding standard deviation of consensus forecasts. Model-implied 95% confidence intervals for the model with and without shifting endpoints are designated by shaded regions and dotted lines, respectively.

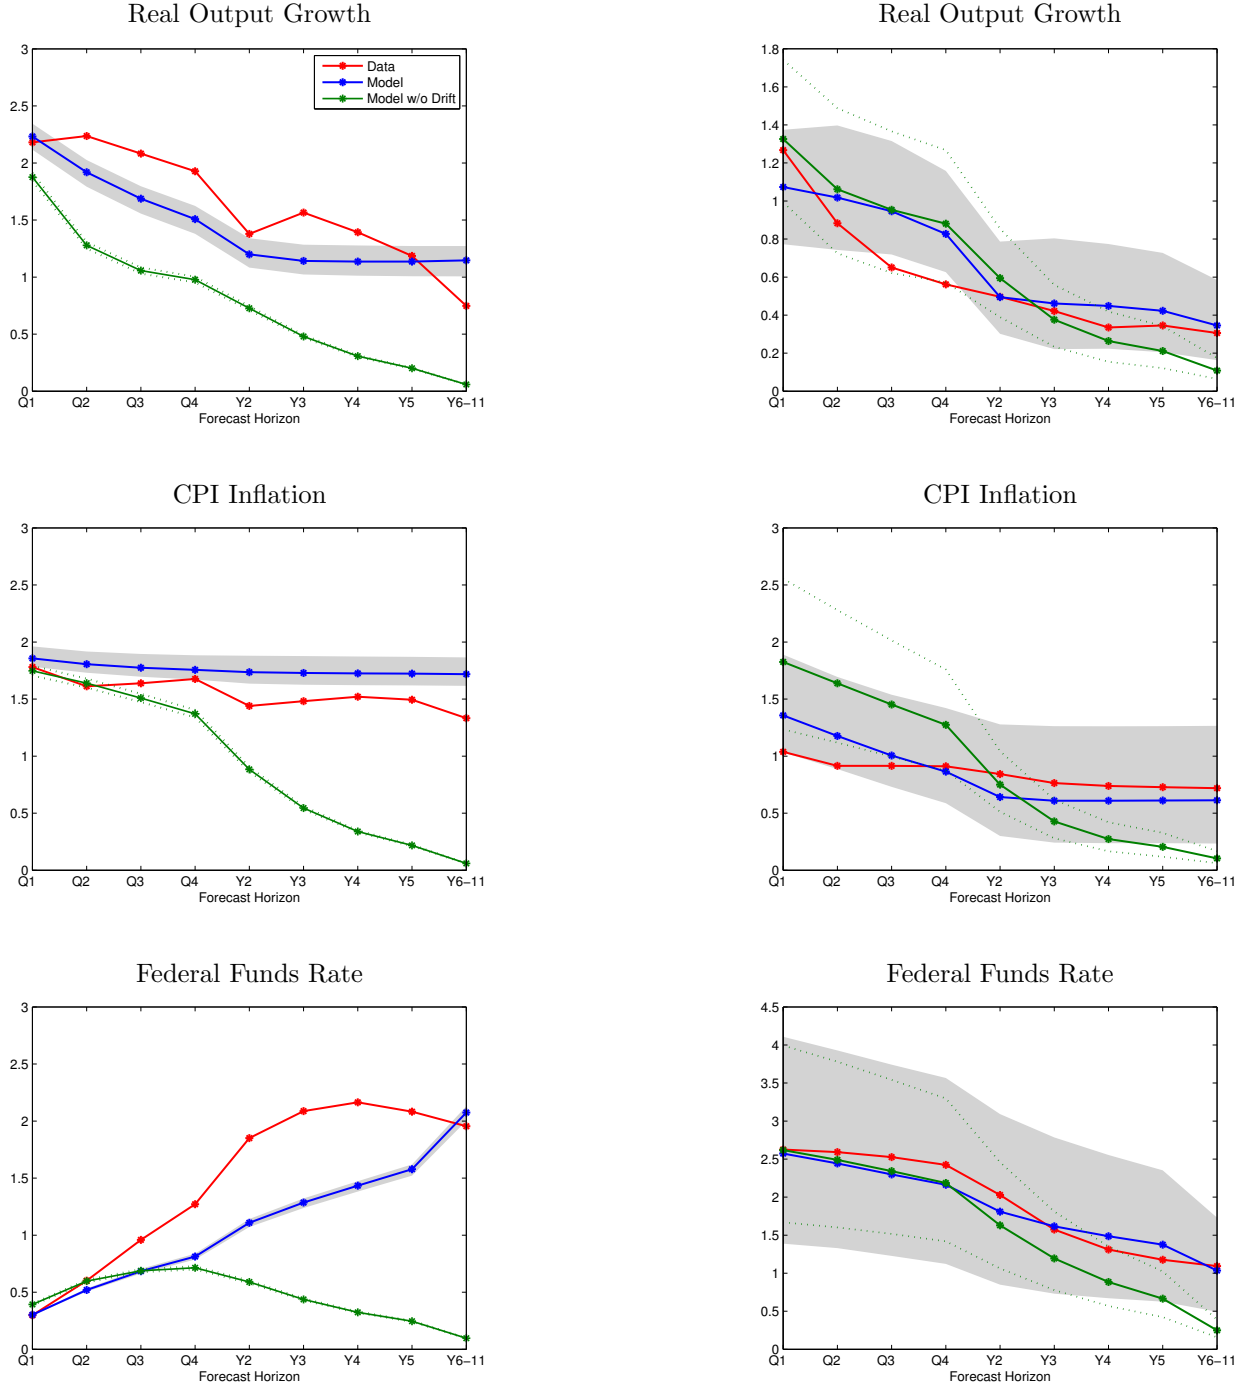


Figure 6: **Second Moments of Disagreement**

The first column displays the model-implied (time) variance of disagreement for the generalized noisy information model calibrated with $\alpha = 50$ (blue) along with the Blue Chip Financial Forecasts survey (red). The second column displays the corresponding correlation of disagreement between variables. Model-implied 95% confidence intervals are designated by shaded regions.

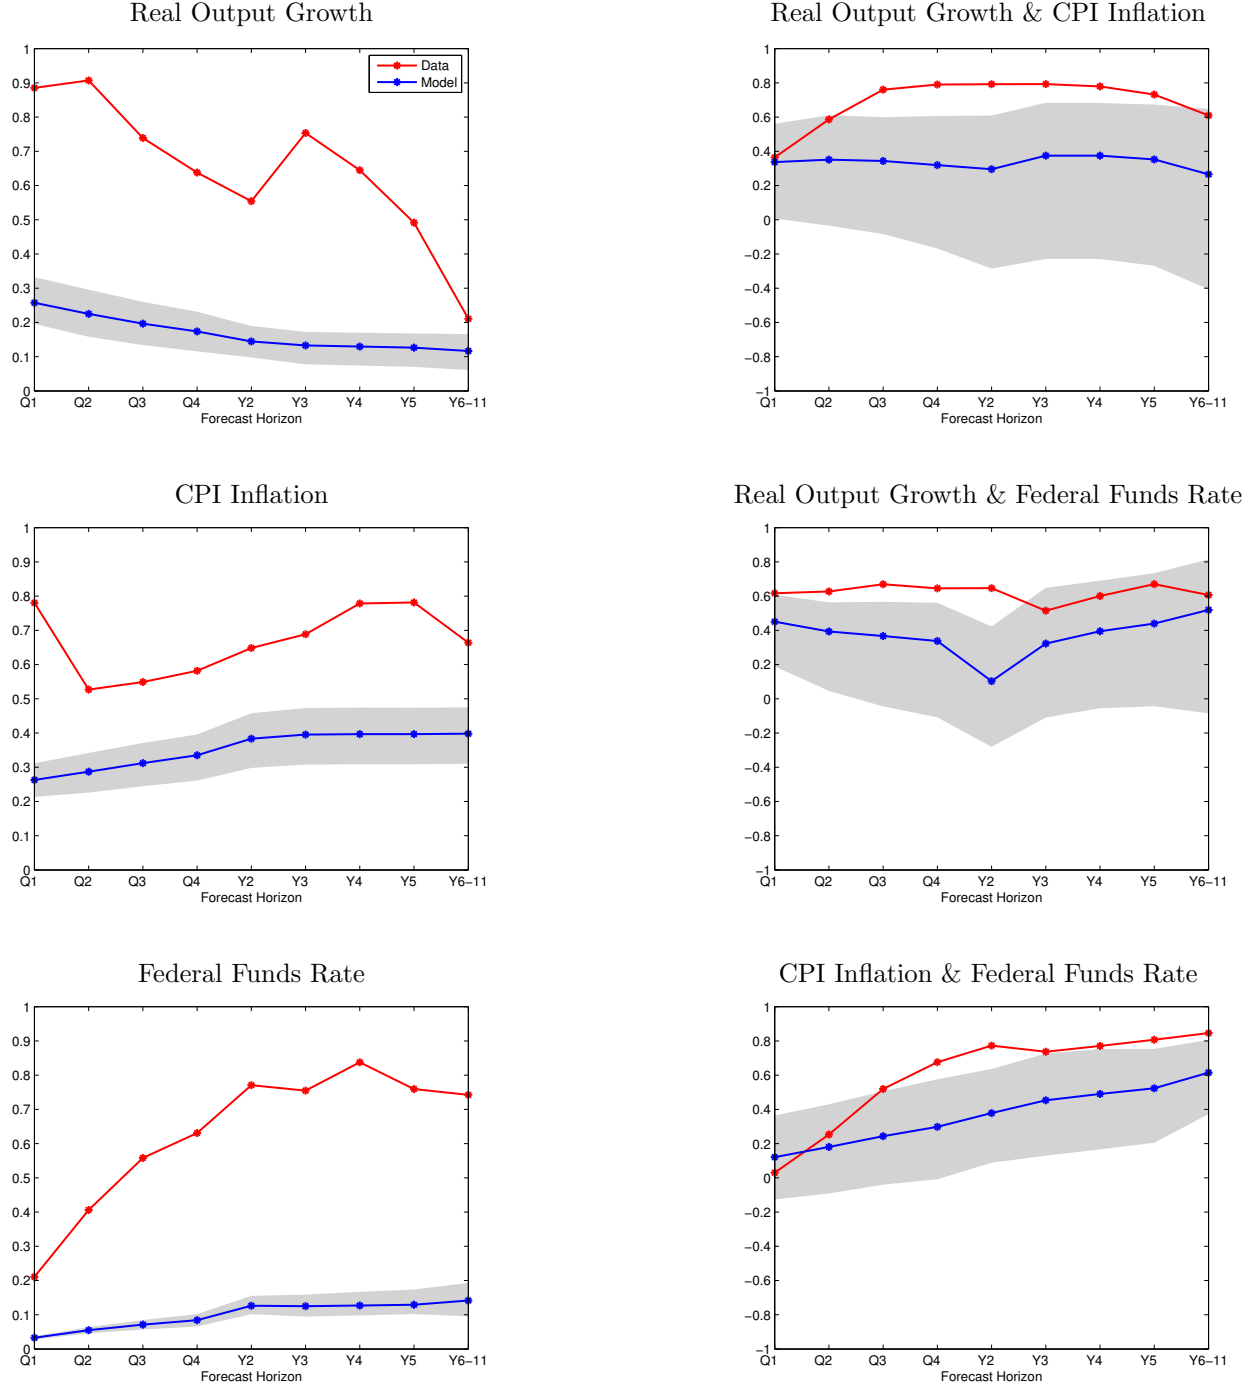


Figure 7: **Other Summary Statistics**

The first [second] column displays the correlation between model-implied real output growth [target variable] observed by the econometrician and model-implied disagreement for the generalized noisy information model calibrated with $\alpha = 50$ (blue) along with the correlation between real output growth [target variable] and disagreement from the Blue Chip Financial Forecasts survey (red). Model-implied 95% confidence intervals are designated by shaded regions.

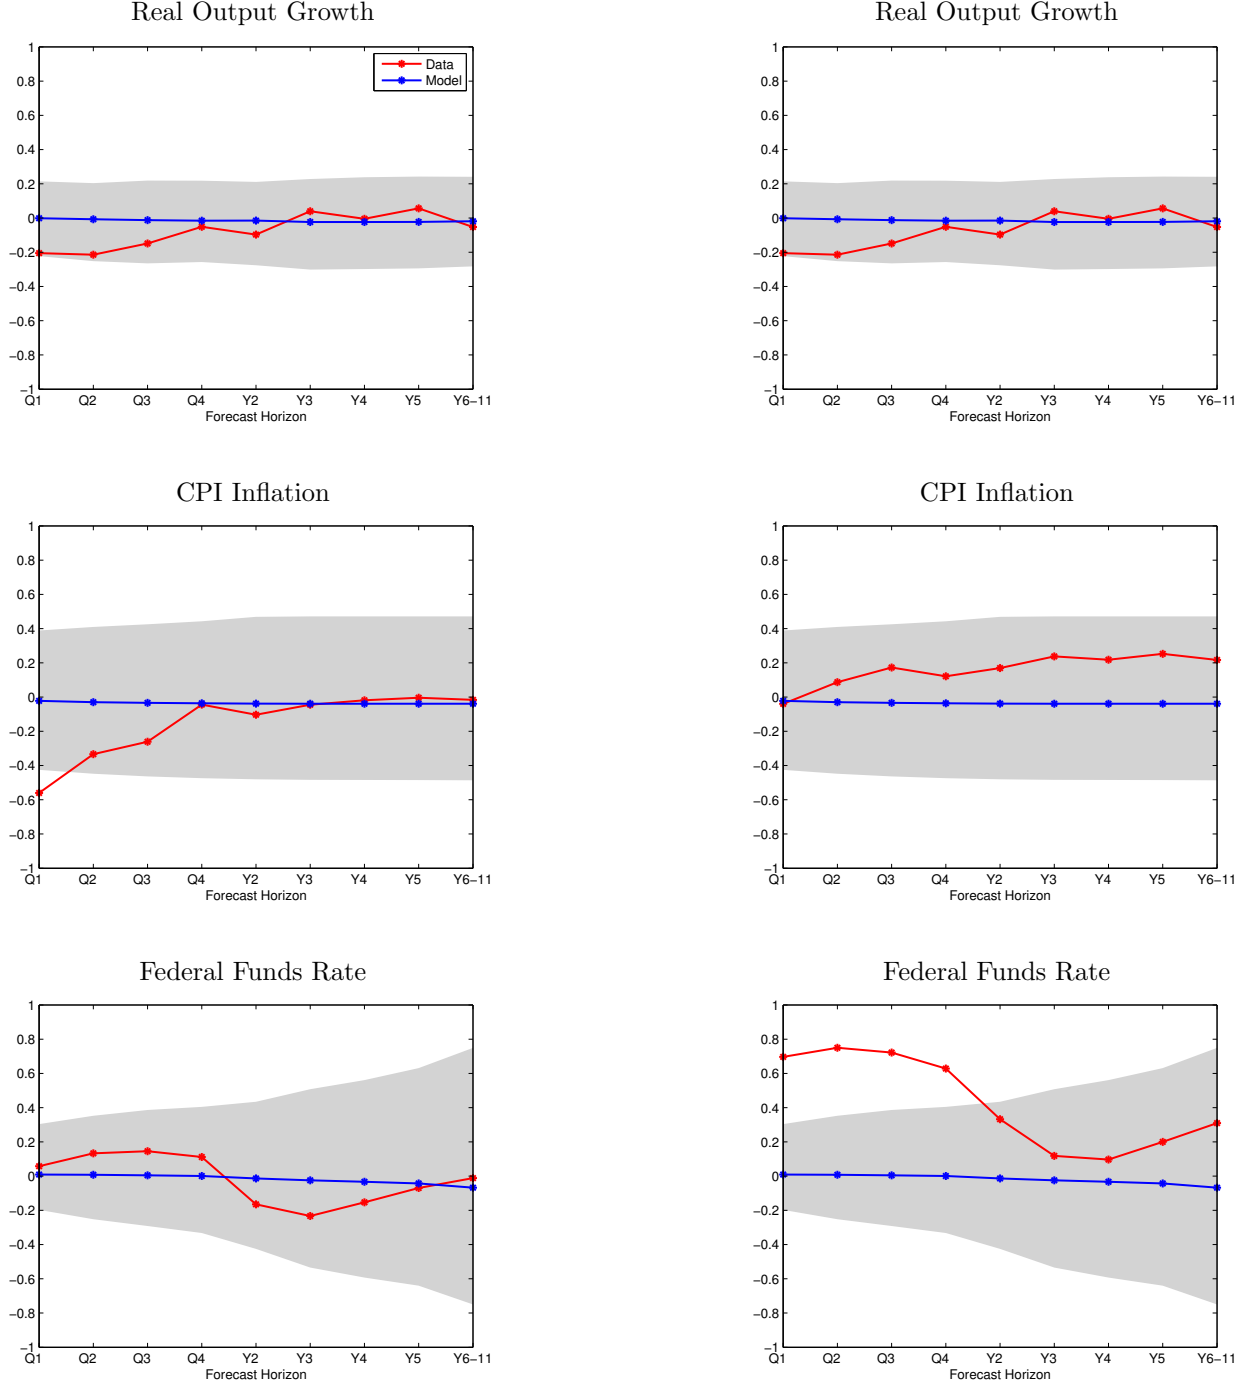


Figure 8: **Results for Different Values of α**

The first column displays the model-implied disagreement for the generalized noisy information model for different values of the penalty parameter α along with the Blue Chip Financial Forecasts survey (red). The second column displays the corresponding standard deviation of consensus forecasts.

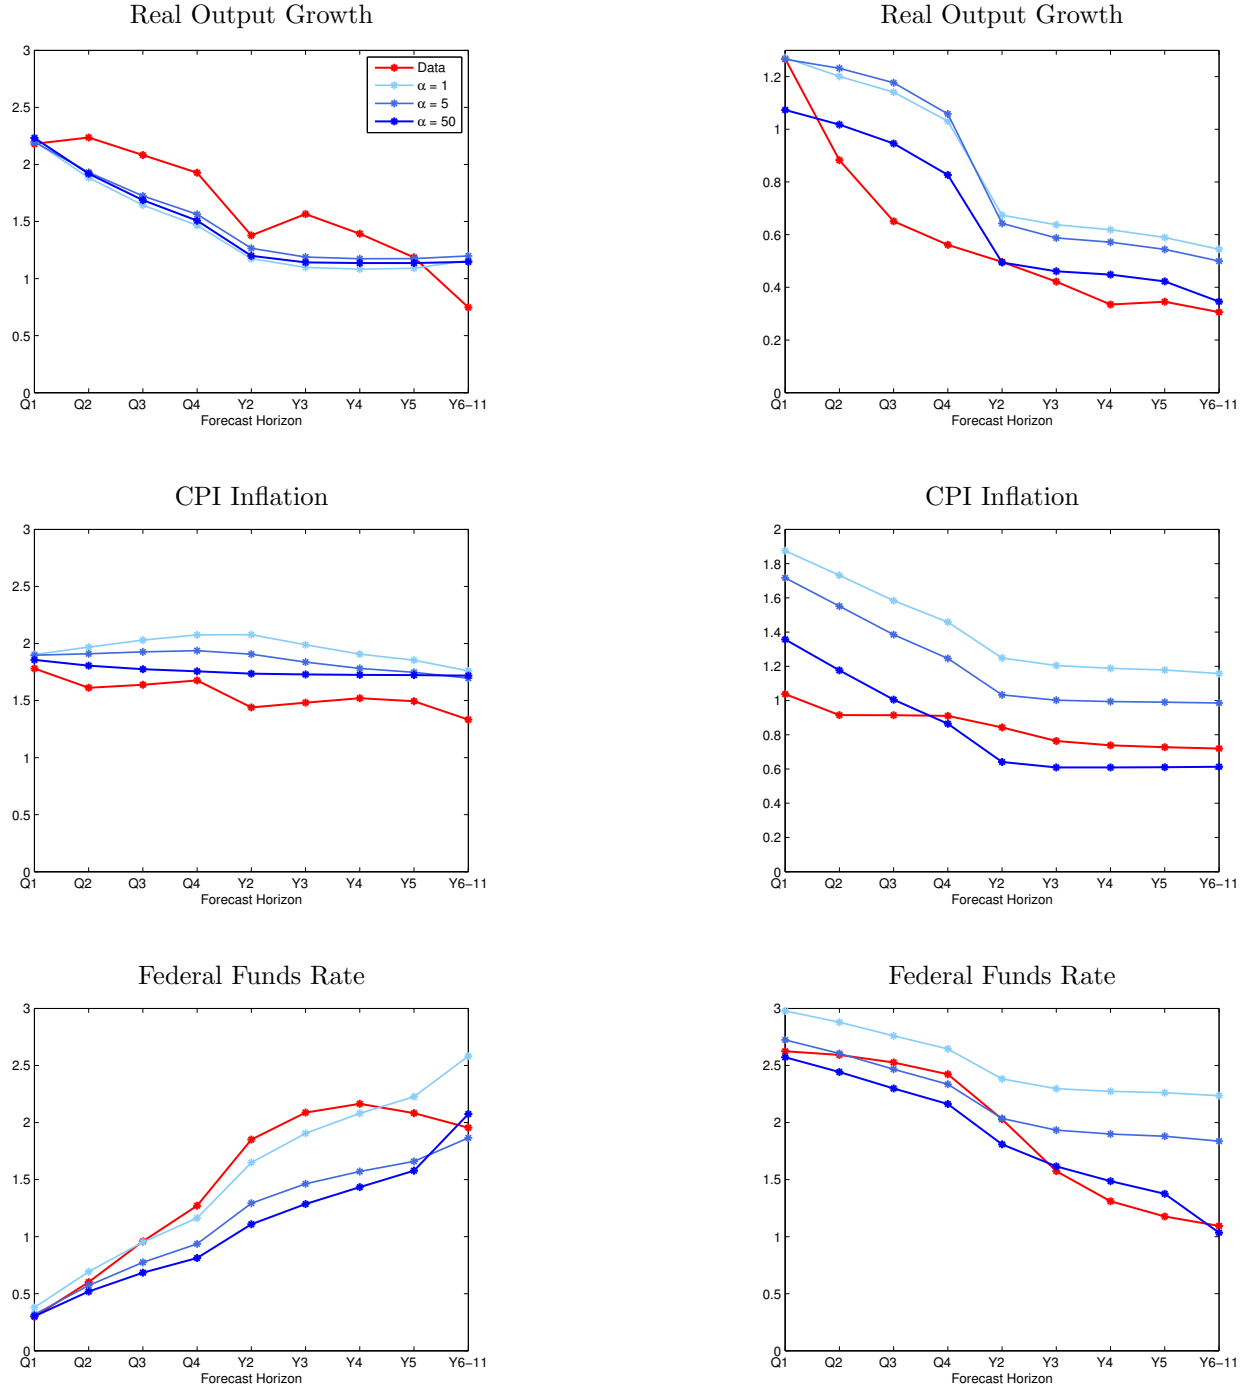


Figure 9: **Comparison to Results with Burn In**

The first column displays the model-implied disagreement for the generalized noisy information model calibrated with $\alpha = 50$ with and without burn in of 120 observations (purple and blue, respectively) along with the Blue Chip Financial Forecasts survey (red). The second column displays the corresponding standard deviation of consensus forecasts. Model-implied 95% confidence intervals for the model with and without burn in are designated by dotted lines and shaded regions, respectively.

