The Factor Structure of Disagreement

Edward Herbst and Fabian Winkler*

September 2019

Abstract

We document how disagreement in macroeconomic expectations comoves across variables. We directly use individual response data in the Survey of Professional Forecasters to estimate a three-dimensional dynamic factor model with Bayesian methods. The extracted factors are interpreted through a semi-structural model where heterogenous expectations arise because of dispersed information. Up until the Great Moderation, the factors describe disagreement about the supply side of the economy, while in recent years and particularly during the Great Recession, disagreement about the demand side of the economy has become more important. Disagreement about the course of monetary policy seems to play a minor role in the data.

Keywords: Disagreement, Heterogeneous Expectations, Dynamic Factor Model, Survey Data.

1 Introduction

People disagree about nearly all aspects of the future. Survey measures of macroeconomic expectations reveal that individuals have different expectations for all forecast variables elicited in these surveys. This heterogeneity in expectations matters because can lead to inertia in price dynamics (Woodford, 2002; Mackowiak and Wiederholt, 2009), non-fundamental driven

*Board of Governors of the Federal Reserve System, 20th St and Constitution Ave NW, Washington DC 20551. Our email addresses are edward.p.herbst and fabian.winkler, both @frb.gov. We thank Travis Berge, Andrew Chen, Neil Ericsson, Nathan Foley-Fischer, Kirsten Hubrich, Elmar Mertens, Anna Orlik, as well as seminar participants at Detusche Bundesbank, George Washington University, Singapore Management University and FAU Nuremberg for helpful comments. Sarah Baker, Carter Bryson and Rahul Kasar provided excellent research assistance. The views herein are those of the authors and do not represent the views of the Board of Governors of the Federal Reserve System or the Federal Reserve System.

business cycle fluctuations (Lorenzoni, 2009; Angeletos and La'O, 2013; Ilut and Schneider, 2014), as well as speculative dynamics and booms and busts in asset prices (Scheinkman and Xiong, 2003; Barillas and Nimark, 2013; Burnside, Eichenbaum, and Rebelo, 2016).

Most empirical studies of forecast disagreement focus on measures of forecast dispersion that compress the distribution of individual forecasts about each variable like inflation or GDP into a single summary statistic, such as the standard deviation or the inter-quartile range (e.g. Capistran and Timmermann, 2009; Dovern, Fritsche, and Slacalek, 2012; Andrade, Crump, Eusepi, and Moench, 2016). Dispersion measures are an important summary of disagreement and has been used e.g. to relate inflation disagreement to the level of inflation (Mankiw, Reis, and Wolfers, 2003). But survey data contain much more information about the structure of disagreement. In particular, they contain information about the comovement of disagreement between different forecast variables. By this, we mean the tendency of forecasts to be correlated across variables in the cross-section of individual forecasters. This multivariate comovement, which isn't captured in univariate measures of dispersion, can tell us a lot about the aspects of the economy that forecasters disagree about. For example, if forecasters disagreed mainly about supply side shocks in the economy, individual forecasts of output and inflation would tend to be negatively correlated in the cross-section. Or, if forecasters disagreed mainly about monetary policy, one would expect positive comovement of inflation, output and interest rates. Such correlation patterns are also informative for models of heterogeneous expectations in which agents' disagreement about the entire economy usually stems from one or two idiosyncratic signals. Despite their potential, thorough examinations of the multivariate distribution of individual-level forecasts have not been put forward in the literature.

In this paper, we explicitly model the multivariate structure of disagreement using a dynamic factor model (DFM) suited to three-dimensional panel data that we estimate with Bayesian techniques.[1] Each forecaster is endowed with a small number of factors that describe the deviation of his forecasts of all variables from the consensus forecast. The factors are assumed to be independent across forecasters, but the factor loadings are identical for all forecasters. This particular factor structure is set up so as to flexibly recover the most important comovement relationships between different forecast variables. The use of a factor structure makes it easy to simultaneously analyze a large number of forecast variables, and can naturally incorporate unbalanced panels and missing data that are ubiquitous in survey data.

[1] Our Bayesian approach is well-suited for the type of three-dimensional panel data with time and cross-sectional correlations present in survey data. We note that frequentist estimation of factor models in three-dimensional panels is an active area of research (Lu and Su, 2018).

We use this model to document new facts about disagreement in the Philadelphia Fed's Survey of Professional Forecasters (SPF). We find that the comovement of disagreement can be described by one factor that captures supply-side disagreement, and a second factor that captures demand-side disagreement. The supply side factor moves medium-term inflation and output expectations in opposite directions. It also moves expectations of long-run GDP growth, productivity and the natural rate of unemployment. The demand side factor moves inflation and output expectations in the same direction, and does not affect long-run forecasts except for inflation and interest rates. Demand-side disagreement is estimated to be much more persistent in time than supply-side disagreement.

Variance decompositions reveal that disagreement about the demand side has become more important over time and has been particularly important during the Great Recession. We also find no evidence of sizable disagreement about the course of monetary policy—which would manifest itself a combination of lower interest rate, higher inflation and higher output forecasts—as these comovement patterns are not picked up by the factors.

While our DFM is a reduced-form model, we offer an interpretation through a semi-structural model of heterogeneous expectations. In this model, forecasters predict a dynamic multivariate data-generating process by filtering noisy signals about structural shocks and measurement error. Forecast disagreement arises because some of the signals contain idiosyncratic noise. We map the forecast structure generated by this model into our reduced form factor model and show that the reduced-form factor loadings are identified with the impulse responses of the shocks that forecasters disagree about. We further illustrate this mapping using the standard New-Keynesian model. In that model, disagreement about the supply side and the demand side economy, as well as disagreement about monetary policy, can be identified through the signs of the factor loadings on inflation, output, and interest rates.

Our empirical findings can serve to discipline structural models of heterogeneous expectations. In many existing models with heterogeneous information, agents receive signals about aggregate and/or sector-specific total factor productivity (e.g. Lorenzoni, 2009; Angeletos and La'O, 2013; Nimark, 2014). According to our results, this type of disagreement only accurately describes the data before the Great Moderation. Melosi (2014) incorporates idiosyncratic signals about both monetary policy and productivity shocks into a model of heterogeneous expectations. However, our estimation finds that disagreement about monetary policy plays a negligible role in the data. We hope that future research will make use of the results presented here to identify models consistent with the structure of disagreement as it is documented here.

Our paper relates to a large literature that uses survey data to inform models of expectations. A large part of the literature either uses consensus forecasts (e.g. Coibion and Gorodnichenko, 2015) or examines dispersion statistics to summarize disagreement. A few papers go beyond dispersion measures: Patton and Timmermann (2010) study the persistence of the disagreement of an individual forecaster from the consensus over time. Rich and Tracy (2017) study how the extent of individual disagreement predicts forecast revisions and forecast accuracy, and relate disagreement to individual uncertainty measured in density forecasts. Bordalo, Gennaioli, Ma, and Shleifer (2018) use individual forecasts to study the predictability of forecast errors by forecast revisions. All of these studies focus on one variable at a time. By contrast, our explicit goal is to study the structure of disagreement across forecast variables.

The paper most closely related to our analysis is Dovern (2015). In addition to studying time-varying univariate forecast dispersion, he also examines the cross-sectional correlations between individual forecasts of output growth, unemployment and inflation. He concludes that these correlations are not particularly strong in the data, in contrast to what is predicted by most theoretical models of forecast disagreement. Consistent with Dovern, we also find that the idiosyncratic components in our factor model explain more than half of the variance of forecaster disagreement. In our semi-structural model, this implies that disagreement about variable-specific measurement error is driving most of total disagreement. But to us, the glass is half-full rather than half-empty: The correlations in the data are still informative for inferring structural sources of disagreement.

Our paper also relates to a small literature that examines "theory-consistency" of forecasts. For example, Carvalho and Nechio (2014) examine to what extent individual forecasts in the Michigan survey of consumers conform to a Taylor rule-type equation, and Dräger, Lamla, and Pfajfar (2016) apply similar tests using the Fisher equation, the Taylor rule and the Phillips curve, and show that central bank communication can improve consistency of expectations with these theoretical relations. In this paper, we do not aim to test for particular restrictions on multivariate forecasts, but rather we let the data speak as freely as possible about the most important aspects that produce disagreement among forecasters.

The remainder of this paper is structured as follows: Section 2 describes the data. Section 3 draws the distinction between the comovement of dispersion and the comovement of disagreement, and takes a first look at the structure of disagreement through simple summary statistics. Section 4 spells out our dynamic factor model. Section 5 offers a semi-structural interpretation through a model of heterogeneous signals. Section 6 discusses our estimation results. Section 7.2 applies our estimation to subsamples of the data as well as to the consensus forecast. Section 8 concludes.

2 Data

The data we use are from the Survey of Professional Forecasters (SPF). The survey is the longest-running quarterly survey of macroeconomic forecasts in the United States and has been conducted since 1968:Q4. We include data through 2017:Q3. Since 1990, the survey is run by the Philadelphia Fed.

In the middle of each quarter, participants are asked to forecast a wide range of variables for the current quarter and each of the following quarters, up to four quarters out. In addition, they are also asked to report a number of longer-run forecasts, as well as recession probabilities.[2] The full set of variables and their abbreviations are documented in the appendix for reference.

Importantly for our purposes, the SPF is a panel data set in which each forecaster can be tracked over time.[3] The number of forecasters who participate averages about 40. However, the panel contains a very large number of missing values. About 79 percent of entries are missing compared to a balanced panel. There are several reasons for the sparsity of the data. First, forecast variables have been introduced into the survey at various points in time. While the 1968 survey only asked about eight variables, 34 variables were included in the 2016 survey. Second, many forecasters drop in and out of the sample. Over the sample, 583 forecasters have participated in the survey, but on average only about 40 have participated in each quarter. Third, even forecasters that do participate in a given quarter do not always respond to all questions in the survey. While the more prominent forecasts such as real GDP are almost always filled in, some less prominent ones like corporate profits are only filled in by about two thirds of respondents, on average.

Figure 1 is an illustration of these patterns. The figure shows when forecasters are present in the sample over time, and how many combined pairs of variables and forecast horizons they enter into the survey. The number of participating forecasters is highly variable. In the first few years of the survey, there were more than 50 participants each quarter, a number that subsequently trended down. A big wave of new participants entered the sample in 1990, when the Philadelphia Fed took over the survey. Since then, the number of forecasters in the sample has been relatively stable. The figure also documents that some forecasters deliver more extensive forecasts than others who only forecast a few select variables. Even The

[2]The survey also includes forecasts at annual horizons, as well as density forecasts later in the sample. We do not consider these additional forecasts here.

[3]There are caveats with the use of SPF individual identifiers, as noted by the Philadelphia Fed. First, It is impossible to reconstruct whether mistakes have been made in assigning the same identifier to different forecasters prior to 1990. Second, when an individual that switches employers but continues to participate in the survey, the decision whether to retain the same identifier for that individual involves judgment.

Figure 1: Missingness patterns in the SPF.

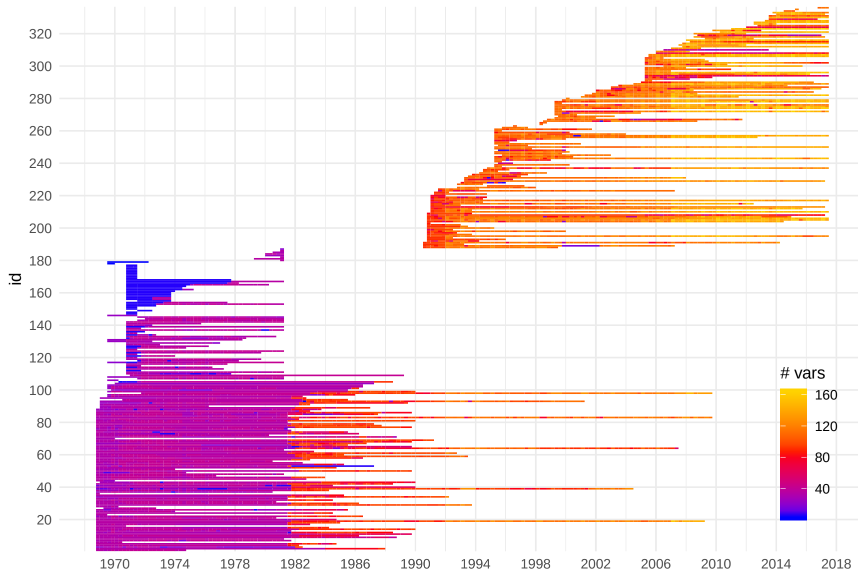

Note: Forecaster ids on vertical axis. For readability, the figure excludes an additional 102 forecasters that participate in the survey for a single period only. Colors represent the number of responses for variables and forecast horizonsof for each forecaster and time period.

number of forecasts made by a single individual can vary over time. Any estimation based on these data would suffer from substantial information loss if it relied on complete cases, and instead needs to be able to handle missing data explicitly.

In this paper, we do not use any data other than SPF forecasts. In particular, we do not consider realizations of the forecast variables, because our interest is only to characterize forecast disagreement. This eliminates the need to take into account data revisions and other measurement problems, since all our data are real-time data by definition.

3 Summary statistics

In this section, we take a first look at the structure of disagreement using simple summary statistics. We highlight the difference between the comovement of dispersion and the comovement of disagreement and show that the two behave quite differently in the data.

The data come in form of a panel where each observation $\hat{y}_{jt+h|it}$ is a forecast of individual i made during period t about variable j, concerning the realization $\tilde{y}_{jt+h}$ at time $t+h$. $h = 0$ corresponds to the current quarter nowcast, $h = 1$ is the one quarter-ahead forecast and so on. The cross-sectional average forecast is given by

$$\bar{y}_{jht} = \frac{1}{\mid N_t \mid} \sum_{i \in N_t} \hat{y}_{jt+h|it} \tag{1}$$

where N_t is the subset of forecasters who respond to the forecast of variable j at horizon h in period t. This average forecast is called the *consensus* forecast. We define *disagreement* as the distribution of individual forecasts around the consensus. An individual's disagreement relative to consensus, $y_{ijt+h|t}$, is simply the difference of the individual's forecast and the consensus:

$$y_{ijht} = \hat{y}_{jt+h|it} - \bar{y}_{jht}. \tag{2}$$

Table 1 provides summary statistics for each of the variables that we consider in this paper. It also lists the transformations we apply to the individual forecasts, which follow Stock and Watson (2002).

Existing studies of disagreement typically compress the distribution of disagreement into the cross-sectional standard deviations for each variable j, time period t and forecast horizon h, which is called the *dispersion* of forecasts:

$$\hat{\sigma}_{jht} = \sqrt{\frac{1}{\mid N_t \mid} \sum_{i \in N_t} y_{ijht}^2}. \tag{3}$$

It has been established (e.g. Andrade, Crump, Eusepi, and Moench, 2016) that dispersion is highly correlated across variables over time. At times when forecasters disagree more about output, they tend to disagree more about everything. This fact is documented in Figure 2.

The figure displays correlation coefficients of dispersion between real GDP and other variables in gray bars. The correlation is taken over the time dimension. The covariance dispersion between two variables j, k used for the computation is:

$$\widehat{\text{Cov}}_t \left(\hat{\sigma}_{jht}, \hat{\sigma}_{kht} \right) = \frac{1}{T} \sum_{t=1}^{T} \left(\hat{\sigma}_{jht} - \frac{1}{T} \sum_{\tau=1}^{T} \hat{\sigma}_{jh\tau} \right) \left(\hat{\sigma}_{kht} - \frac{1}{T} \sum_{\tau=1}^{T} \hat{\sigma}_{kh\tau} \right). \tag{4}$$

From Figure 2, it is immediately clear that dispersion correlates positively with real GDP for almost all series in the survey. The correlation is particularly high for consumption and GDP deflator-based inflation, while it is lower for interest rate spreads and long-term forecasts.

The focus of this paper, however, is not the comovement of dispersion, but the comovement

Table 1: SUMMARY STATISTICS.

Variable y_{ijht}	format	s.d.	autocorr.	# time periods	average # forecasters
GDP components					
RGDP4	5	1.10	0.52	191	36.8
RCONSUM4	5	0.80	0.59	145	32.5
RNRESIN4	5	2.52	0.64	145	31.8
RRESINV4	5	4.61	0.66	145	31.7
RSLGOV4	5	1.18	0.62	145	30.8
RFEDGOV4	5	2.26	0.63	145	30.8
RCBI4	5	0.21	0.47	145	31.5
REXPORT4	5	0.42	0.55	145	31.8
Other real activity					
INDPROD4	5	1.78	0.61	191	35.1
HOUSING4	4	10.36	0.67	191	35.6
CPROF4	5	5.96	0.50	191	28.1
Inflation					
PGDP4	6	1.20	0.12	191	36.7
CPI4	6	0.65	0.47	145	33.3
CORECPI4	6	0.41	0.73	43	38.0
COREPCE4	6	0.37	0.74	43	35.2
Labor market					
UNEMP4	1	0.40	0.66	191	37.4
EMP4	5	0.44	0.58	56	35.9
Interest rates					
TBILL4	2	0.64	0.63	145	32.7
TBOND_TBILL4	1	0.44	0.54	103	35.0
BOND_TBOND4	1	0.31	0.47	103	30.4
BAABOND_BOND4	1	0.24	0.23	31	27.5
Recession prob.					
RECESS4	1	12.82	0.60	192	36.4
Long-term forecasts					
RGDP10	5	0.35	0.85	26	34.1
UBAR	1	0.42	0.93	22	18.4
CPI10	6	0.48	0.79	104	34.0
TBILL10	2	0.83	0.78	26	29.4
PROD10	4	0.47	0.86	26	30.5
STOCK10	4	1.99	0.81	26	25.1

Note: Sample runs from 1968Q4 through 2017Q2. "s.d." is the standard deviation of y_{ijht} and "autocorr." is its first autocorrelation. "# time periods" describes the number of time periods for which the variable could be constructed from available responses from at least one forecaster. "average # forecasters" describes the average number of forecasters per time period for which the variable could be constructed from available responses. Transformations: 1 is plain levels. 2 is the difference of the time t-forecast level for $t+h$ minus the "forecast" level for $t-1$ (realizations at $t-1$ are known at t). 4 is log levels. 5 is the time t-forecast annualized log growth rate between $t+h$ and $t-1$. 6 is the difference of the time t-forecast of the annualized log growth rate between $t-1$ and the "forecast" annualized log growth rate between $t-1$ and $t-2$. Transformations are applied to individual forecasts before differencing out the consensus. Forecast horizons are $h=4$ except for UBAR which does not have a horizon and long-term forecasts for which the horizon is $h=40$.

Figure 2: Correlation of Dispersion And Disagreement With Real GDP.

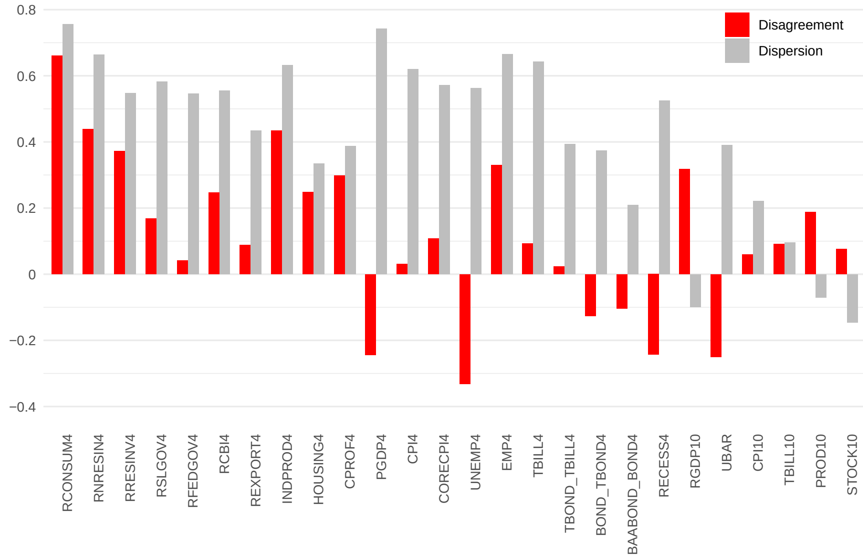

Note: Gray bars represent correlation coefficients of forecast dispersion between real GDP and other SPF variables, over time. Red bars represent correlation coefficients of disagreement between real GDP and other SPF variables, over individuals and time. All forecasts are log level forecasts made for a horizon of 4 quarters ahead, except for RCBI and REXPORT which are transformed to fractions of RGDP. Sample runs from 1968Q4 through 2017Q2.

of disagreement. Here, the question is: When someone forecasts higher output relative to consensus, what does that imply for their relative forecast in another variable? We compute the covariance between relative disagreement in two variables j, k at horizons h, ℓ according to:.

$$\widehat{\text{Cov}}_{it}\left(\hat{y}_{jt+h|it}, \hat{y}_{kt+\ell|it}\right) = \frac{1}{T}\frac{1}{\mid N_t \mid}\sum_{t=1}^{T}\sum_{i\in N_t} y_{ijht}y_{ik\ell t}. \tag{5}$$

Figure 2 displays the corresponding correlation coefficients for real GDP and other variables in red bars. It is immediate that disagreement correlations are very different from dispersion correlations. In fact, the magnitude of correlation is generally smaller for disagreement than for dispersion. This difference is perhaps most striking for 3-month Treasury bill yields, suggesting that the reasons why forecasters disagree about short-term rates and output are quite different, even though forecast dispersion for short-term rates and output move strongly together over time.

Looking in a little more detail at the figure, one can see that disagreement correlation of real GDP with all NIPA components is positive. This fact is somewhat surprising for real net exports since they correlate negatively with output in realized U.S. data (Stock and Watson, 1999). It is also notable that forecasts of government expenditure correlate more strongly with output at the state level than at the federal level. Less surprisingly, the disagreement correlation of unemployment and recession risk with output are negative, so that forecasters who expect higher output also expect lower unemployment and lower recession risk.

We can already see from the figure that the disagreement correlation of inflation measured by the GDP deflator and real GDP is negative: Forecasters who expect higher output tend to expect lower inflation. This fact is inconsistent with a world in which forecasters only disagree about demand shocks which move output and inflation in the same direction, but it could be consistent with disagreement about supply shocks. At the same time, the correlation is much weaker when the consumer price index (CPI) is used as the inflation measure. This difference could be resulting from structural breaks in the comovement over time, since GDP deflator forecasts are collected since 1968, while CPI forecasts only start in 1981. We will examine this issue in further detail below.

4 Factor model of multivariate disagreement

In this section, we construct a time series model for disagreement. We assume that the random variables driving disagreement are identical and independent across individuals. Thus, it is sufficient to describe the probability model for an arbitrary individual i's disagreement. The aggregate model is merely the collection of individual models. Recall from (2) that at t an individual's i disagreement for the h-step forecast of variable j is denoted by y_{ijht}. Let $\mathcal{J}$ denote the set of all variables considered and $\mathcal{H}$ denote the set of all horizons considered. Then we can write individual i's disagreement at time t as a vector of length $n = |\mathcal{J}| \times |\mathcal{H}|$,

$$y_{it} = [y_{ijht}]_{j \in \mathcal{J}, h \in \mathcal{H}}.$$

For each individual i, we assume the elements of this vector disagreement can be decomposed into into systematic component which represent that part of disagreement that is common across variables and/or horizons and an orthogonal, idiosyncratic component. Thus, disagreement can be written as a factor model,

$$y_{it} = \Lambda f_{it} + \xi_{it}. \tag{6}$$

Here Λ is an $n \times \kappa$ matrix of factor loadings, f_{it} is a $\kappa \times 1$ vector of common factors, and ξ_{it} is an $n \times 1$ vector idiosyncratic errors. The κ factors each follow independent autoregressive processes (AR) of order 1. Let ϕ be the $\kappa \times 1$ vector corresponding to the autoregressive coefficients associated with these processes and $\Phi = \text{diag}(\phi)$, where the $\text{diag}(\cdot)$ operator places the vector ϕ on the diagonal of $\kappa \times \kappa$ matrix whose other elements are zero. Then we can write the dynamics of the factors as a vector autoregression (VAR):

$$f_{it} = \Phi f_{it-1} + u_{it}, \quad u_{it} \overset{\text{iid}}{\sim} N(0, I_\kappa - \Phi^2) \tag{7}$$

Each factor is normalized to have an unconditional variance of 1.

Finally, the idiosyncratic components also follow independent AR processes of order 1. Let ρ be the $n \times 1$ vector of autoregressive coefficients of these processes and $P = \operatorname{diag}(\rho)$. Similarly, let σ be the vector of standard deviations of their innovations and $\Sigma = (\operatorname{diag}(\sigma))^2$. We can write the VAR for the idiosyncratic term as:

$$\xi_{it} = P\xi_{it-1} + v_{it}, \quad v_{it} \overset{\text{iid}}{\sim} N(0, \Sigma). \tag{8}$$

Equations (6), (7), and (8) form a dynamic factor model for individual i's disagreement. It's worth mentioning a few comments on the characteristics of such a model. First, for almost all Λ, the factors are already identified up to sign and label by the assumption that they are independent from one another. This identifying assumption makes interpretation of the factors easier. Moreover, as shown in section 5, it is consistent with popular models of heterogeneous information. One drawback to this approach is that it precludes the use of the two-step estimator in (Doz, Giannone, and Reichlin, 2012). Second, we again emphasize that each forecaster i is described by the same econometric model, with differences only occurring in the realization of the factors and idiosyncratic terms. That is, the parameters $\theta = [\Lambda, \phi, \rho, \sigma]$ are identical across forecasters. Let $Y_i = [y_{i1}, \ldots, y_{iT}]'$ be the matrix of the time series of disagreement for individual i. Given a vector of parameters θ, the likelihood function for individual i, $p(Y_i|\theta)$, can be evaluated using the Kalman filter using the state space representation implied by (6), (7), and (8). This approach has the advantage, noted by (Banbura and Modugno, 2014), of efficiently handling missing observations.

A different way of writing the model is to stack all forecasters $i = 1, \ldots, m$ in one large vector to recast the model in a standard two-dimensional structure:

$$\overset{mn\times 1}{\begin{pmatrix} y_{1t} \\ y_{2t} \\ \vdots \\ y_{mt} \end{pmatrix}} = \overset{mn\times \kappa m}{\begin{pmatrix} \Lambda & 0 & \cdots & 0 \\ 0 & \Lambda & & 0 \\ \vdots & \vdots & & \ddots \\ 0 & 0 & \cdots & \Lambda \end{pmatrix}} \overset{\kappa m\times 1}{\begin{pmatrix} f_{1t} \\ f_{2t} \\ \vdots \\ f_{mt} \end{pmatrix}} + \overset{mn\times 1}{\begin{pmatrix} \xi_{1t} \\ \xi_{2t} \\ \vdots \\ \xi_{mt} \end{pmatrix}} \tag{9}$$

$$\begin{pmatrix} f_{1t} \\ f_{2t} \\ \vdots \\ f_{mt} \end{pmatrix} = \begin{pmatrix} \Phi & 0 & \cdots & 0 \\ 0 & \Phi & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & \Phi \end{pmatrix} \begin{pmatrix} f_{1t-1} \\ f_{2t-1} \\ \vdots \\ f_{mt-1} \end{pmatrix} + \begin{pmatrix} u_{1t} \\ u_{2t} \\ \vdots \\ u_{mt} \end{pmatrix}, \; u_t \sim \mathcal{N}\left(0, I_m \otimes I_\kappa\right) \tag{10}$$

$$\begin{pmatrix} \xi_{1t} \\ \xi_{2t} \\ \vdots \\ \xi_{mt} \end{pmatrix} = \begin{pmatrix} P & 0 & \cdots & 0 \\ 0 & P & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & P \end{pmatrix} \begin{pmatrix} \xi_{1t-1} \\ \xi_{2t-1} \\ \vdots \\ \xi_{mt-1} \end{pmatrix} + \begin{pmatrix} v_{1t} \\ v_{2t} \\ \vdots \\ v_{mt} \end{pmatrix}, \; v_t \sim \mathcal{N}\left(0, I_m \otimes \Sigma\right). \tag{11}$$

Expressed in this form, we can interpret our model as a standard dynamic factor model with $\kappa \cdot m$ factors together with strong restrictions on the factor loadings that assign each factor to one particular forecaster.

As mentioned above, aggregation is simple because the factors and idiosyncratic terms are independent across individuals. Let $Y = [Y_1', \ldots, Y_I']'$. Then the likelihood function for the entire set of disagreement is given by:

$$p(Y|\theta) = \prod_{i=1}^{I} p(Y_i|\theta). \tag{12}$$

Given this likelihood function, one could estimate the model via MLE as in (Stock and Watson, 1989). Instead, we follow a Bayesian approach. The central object of Bayesian inference in the posterior distribution, $p(\theta|Y)$, which is just a combination of the likelihood and a prior distribution, $p(\theta)$, specifying initial beliefs about θ, using Bayes rule:

$$p(\theta|Y) = \frac{p(Y|\theta)p(\theta)}{p(Y)}. \tag{13}$$

In the Bayesian approach, the calculus of probability characterizes how the state of knowledge or degree of beliefs about some object (for example the parameters θ) changes in light of the data. On advantage in this application is that, the posterior distribution completely characterizes the uncertainty about an object of interest, without reference to potentially inaccurate asymptotic approximations or tedious bootstrapping.

Priors. Our current prior distribution reflects the desire for the mode of the posterior to reflect relatively little influence of the prior distribution. Thus, for the elements of matrix of factor loadings Λ we use independent uniform priors over wide intervals the contain any *a priori* reasonable values for Λ. We likewise use uniform priors of over the unit interval for the elements of ϕ and the parameter ρ. We likewise parameterize the prior for the variance of the idiosyncratic component as follows. We set $\sigma_{j,h} = \alpha_{j,h}\bar{\sigma}_{j,h}$, where $\bar{\sigma}_{j,h}$ is the sample standard deviation of the disagreement of variable j at horizon h. We set the prior of $\alpha_{j,h}$ to be a uniform over the unit interval.[4]

Estimation. The posterior distribution of the parameters is not available in closed form and thus we must rely on simulation methods to estimate the model. For Bayesian dynamic factor models, this has typically been accomplished using Gibbs sampling following (Geweke and Zhou, 1996) and (Otrok and Whiteman, 1998). We depart from this tradition and estimate the model using Sequential Monte Carlo (SMC) methods, following (Chopin, 2002). Specifically, we use the algorithm presented in (Herbst and Schorfheide, 2014) and extended by Del Negro et al. While a detailed exposition of SMC methods is not possible in this paper, we list a few advantages of such an approach. First, relative to Gibbs Samplers, we do not require conjugate priors (and the tedious derivation of conditional posteriors). Moreover, we avoid need to simulate from the posterior of factors which can be time consuming. Second, SMC is embarrassingly parallelizable, which means we can leverage multiprocessing to obtain estimates in reasonable (clock) time, even for extremely cumbersome likelihoods such as the one here.

5 A semi-structural interpretation

The factor structure described in the previous section is a purely reduced-form description of disagreement. Here, we offer an interpretation through a semi-structural model of heterogeneous information. In this model, disagreement among agents comes from signals about structural shocks. The reduced form structure of disagreement is exactly the one in the previous section, and the loadings of each factor are identified (up to a scalar) with the impulse responses to a structural shock in the economy.

[4] Note that this prior, for instance, is not uniform in standard deviation units. Estimation under alternative priors is in progress; moreover, the likelihood dominates the estimation of these terms.

5.1 Model setup

In the semi-structural model, the economy is thought to evolve according to the following data-generating process:

$$\tilde{y}_t = C\tilde{x}_t + \text{diag}\left(\sigma_{\tilde{\eta}}\right)\tilde{\eta}_t,\ \tilde{\eta}_t \sim \mathcal{N}\left(0, I_n\right) \tag{14}$$

$$\tilde{x}_t = A\tilde{x}_{t-1} + B\tilde{\varepsilon}_t,\ \tilde{\varepsilon}_t \sim \mathcal{N}\left(0, I_K\right). \tag{15}$$

The vector of variables $\tilde{y}_t \in \mathbb{R}^n$ is the set of measurements of the economy. Each variable $\tilde{y}_{jt}$ depends on a number of state variables $\tilde{x}_t$ and a noise term $\tilde{\eta}_{jt}$ which is unrelated across variables. The state variables $\tilde{x}_t$ follow an autoregressive process with uncorrelated structural shocks $\tilde{\varepsilon}_t$. This data-generating process can be thought of as the reduced form of a DSGE model, where $\tilde{x}_t$ are the state variables and y_t is the set of observables. But it equally nests backward-looking models used in the 1970s, or non-structural time series models often used in forecasting.

A set of forecasters $i = 1, \ldots, N$ makes forecasts about the economy using the above model, subject to imperfect and heterogeneous information. At time t, forecasters observe the state of the economy in the previous period x_{t-1}.[5] However, they only observe noisy private signals of the current state $\tilde{x}_t$ and of the measurements $\tilde{y}_t$. Disagreement arises from idiosyncratic differences in the signals received by forecasters.

Forecasters receive signals about the structural shocks $\tilde{\varepsilon}_{kt}$, $k = 1, \ldots, M$. These signals have an iid and a persistent noise component:

$$s_{\varepsilon ikt} = \tilde{\varepsilon}_{kt} + \tilde{u}_{\varepsilon ikt} + \tilde{\omega}_{\varepsilon ikt}$$

$$\tilde{u}_{\varepsilon ikt} = \rho_{\varepsilon k}\tilde{u}_{\varepsilon ikt-1} + \tilde{v}_{\varepsilon ikt}.$$

The disturbances $\tilde{\omega}_{\varepsilon ikt}$ and $\tilde{\nu}_{\varepsilon ikt}$ are normally distributed with mean zero. The signals are also identically distributed across forecasters, though they can be correlated. In fact, to obtain a factor structure we will assume that they are perfectly correlated for all but the first p of the K structural shocks they are associated with. Optimal filtering of $\tilde{\varepsilon}_{kt}$ takes into account the information content of the signals, as well as past values of ε_{kt} revealed through observation of past values of the state up to x_{t-1}. Even though the $\tilde{\varepsilon}_{kt}$ are iid, knowing their true past values is informative for the correlated noise $\tilde{u}_{\varepsilon ikt}$. It can be shown that the

[5] This assumption implies that forecasters agree about the past state of the economy, which is clearly a simplification of reality. The assumption can easily be relaxed at the expense of including lagged factors $f_{t-\ell}$ to the reduced-form observation equation (6).

filtering formula are:

$$\hat{\varepsilon}_{kt|it} = g_{\varepsilon k1}\left(s_{\varepsilon ikt} - \rho_{\varepsilon k}\hat{u}_{\varepsilon ikt-1|it-1+}\right) \tag{16}$$

$$\hat{u}_{\varepsilon ikt|it+} = (1 - g_{\varepsilon k2})\,\rho_{\varepsilon k}\hat{u}_{\varepsilon ikt-1|it-1} + g_{\varepsilon k2}\left(s_{\varepsilon ikt} - \tilde{\varepsilon}_{kt}\right) \tag{17}$$

where the gains $g_{\varepsilon k1}$ and $g_{\varepsilon k1}$ are functions of the signal-to-noise ratios in the signal process. Forecasters disagree only about the first p structural shocks $\tilde{\varepsilon}_{1t}, \ldots .\tilde{\varepsilon}_{pt}$, but agree on the values of the other $K-p$ shocks. Denoting $\varepsilon_{ikt} = \hat{\varepsilon}_{kt|it} - 1/N\sum_{j=1}^{N}\hat{\varepsilon}_{kt|jt}$, $\omega_{\varepsilon ikt} = \tilde{\omega}_{\varepsilon ikt} - 1/N\sum_{j=1}^{N}\tilde{\omega}_{\varepsilon ijt}$ etc., we can describe disagreement about the kth shock by a simple ARMA(1,2) process:

$$\varepsilon_{ikt} = (1 - g_{\varepsilon k2})\,\rho_{\varepsilon k}\varepsilon_{ikt-1} + g_{\varepsilon k1}\left(v_{\varepsilon ikt} + \omega_{\varepsilon ikt} - \rho_{\varepsilon k}\omega_{\varepsilon ikt-1}\right). \tag{18}$$

Similarly, forecasters at time t receive signals about current and future measurement shocks η_{jt+h}, $j = 1, \ldots, J$, $h = 0, 1, 2, \ldots$. These signals take the same form as those of the structural shocks described above:

$$s_{\eta ijht} = \tilde{\eta}_{jt+h} + \tilde{u}_{\eta ijht} + \tilde{\omega}_{\eta ijht}$$
$$\tilde{u}_{\eta ijht} = \rho_{\eta jh}\tilde{u}_{\eta ijht-1} + \tilde{v}_{\eta ijht}.$$

Again, the innovations are assumed to be normally distributed with mean zero, iid in time and identically distributed across forecasters, though we do allow for correlated signals in the cross-section. Compared to the structural shocks, the optimal filtering problem has to take into account the fact that the signals are correlated not only over time but also across time horizons h, because s_{ijht} and $s_{ijh+1t-1}$ both inform about $\tilde{\eta}_{jt+h}$. It can be shown that the filtering formula are given by:

$$\hat{\eta}_{jt+h|it} = g_{\eta jh1}\left(s_{\eta ijht} - \hat{u}_{\eta ijht|it}\right) \tag{19}$$

$$\hat{u}_{\eta ijht|it} = (1 - g_{\eta jh2})\,\rho_{\eta jh}\hat{u}_{\eta ijht-1|it-1} + g_{\eta jh2}\left(s_{\eta ijht} - \hat{\eta}_{jt+h|it-1}\right) \tag{20}$$

where the gains $g_{\eta jh1}$ and $g_{\eta jh2}$ are again functions of the signal-to-noise ratios in the signal processes. Disagreement at time t about the set of current and future measurement shocks $\eta_{jt+h}, h = 0, 1, 2, \ldots$ is now described by a VARMA(1,2) process with the autoregressive matrix being different from zero only on the main diagonal and the first superdiagonal:

$$\eta_{ijht} = (1 - g_{\eta jh2})\, \rho_{\eta jh} \eta_{ijht-1} + g_{\eta jh1} g_{\eta jh2} \eta_{ijh+1t-1} \\ + g_{\eta jh1} (1 - g_{\eta jh2}) (v_{\eta ijht} + \omega_{\eta ijht} - \rho_{\eta jh} \omega_{\eta ijht-1}) \,. \tag{21}$$

5.2 Mapping to reduced-form model

With this information structure, the forecasts of the variables y_t in the economy can be cast as a factor model that has the same form as the reduced-form model (6),(8),(7). The observation equation is:

$$y_{ijt+h|t} = \delta_{jht} + \sum_{k=1}^{p} \lambda_{jhk} \varepsilon_{ikt} + \eta_{ijht}. \tag{22}$$

In the semi-structural model, the common component δ_{jht} is simply identified with the average of all forecasts, which incorporates knowledge of the past state x_{t-1} and the common components of the signals received by the forecasters. The idiosyncratic components e_{ijht} are identified with disagreement about the measurement shocks η_{ijht}. As in the reduced-form model, the idiosyncratic components are uncorrelated across forecast variables j. They follow an approximate AR(1) structure described by (21). Finally, the factors f_{ikt}, $k = 1, \ldots, p$ are identified with disagreement about the first p structural shocks ε_{ikt} that are the source of agents' systematic disagreement across variables. The factors follow an approximate AR(1) structure described by (18). Importantly, the factor loadings λ_{jhk} are identified with the impulse response functions to the p structural shocks that agents disagree about:

$$\lambda_{jhk} = C_{j\cdot} A^h B_{\cdot k} \,. \tag{23}$$

Thus, the semi-structural model interpretation allows us to see the factor loadings as the effects on the economy of the particular shocks that forecasters disagree most about.

5.3 Illustration with the New-Keynesian model

As an illustration, consider the standard three-equation New-Keynesian model describing inflation π_t, output y_t and nominal interest rates i_t by a Phillips curve, an IS curve, and a

Table 2: Signs of Factor Loadings in the New-Keynesian model.

Shock	y	π	i
supply u_t	(+)	(−)	(−)
demand r_t^n	(+)	(+)	(+)
monetary policy e_t	(+)	(+)	(−)

Taylor rule:

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa y_t + u_t \tag{24}$$

$$y_t = \mathbb{E}_t y_{t+1} - \frac{1}{\sigma} \left(i_t - \mathbb{E}_t \pi_{t+1} - r_t^n \right) \tag{25}$$

$$i_t = \phi_\pi \pi_t + \phi_y y_t + e_t. \tag{26}$$

The parameters of the model are the discount factor β, the slope of the Phillips curve κ, the intertemporal elasticity of substitution σ, and the reaction coefficients $\phi_\pi$and ϕ_y in the Taylor rule. The three shocks are a cost-push shock u_t, a shock to the natural real interest rate r_t^n, and a shock to the Taylor rule e_t. We can think of these three shocks as a supply, demand, and monetary policy shock, respectively.[6] The left panel of Table 2 sums up the signs of the factor loadings corresponding to disagreement about each of the three shocks.[7] Because the sign patterns are different for each shock, they can be used to identify the shocks that agents disagree about from the factor loadings in the estimated reduced form of the model. For example, if the loadings of the dominant factor have the same signs on inflation, output and interest rates, then the factor can be identified with disagreement about a demand shock.

6 Estimation results

We apply the dynamic factor model to the cross-section of SPF variables at fixed horizon. We include all four-quarter and ten-year forecasts ($h = 4$ or $h = 40$) ; all forecasts are in log levels. We set the number of factors to $\kappa = 2$. Our full sample runs from 1968:Q4 through 2017:Q2, and has 7,595 observations.

[6] We think of the shock to the natural real rate as a pure demand shock, e.g. a discount factor shock, that does not move the natural level of output. This justifies writing the model in terms of the level of output instead of the output gap.

[7] Productivity shocks, which simultaneously affect the natural real rate and the natural level of output, would behave similarly to the cost-push shock. One way of differentiating between the two could be that for technology shocks, profits increase with output, while for markup shocks they decrease with output.

6.1 Parameter estimates

Tables 3 describes the estimated posterior means, 5th, and 95th percentiles for Λ, grouped by component. The two sets of columns refer to the loadings on the first and second factor, respectively.

Because the sample size is an order of magnitude larger than in aggregate macroeconomic data, the precision of the estimated posterior parameters is high. A unit increase in the first factor is associated with an increase in an individual's four-quarter real GDP forecast, relative to the consensus, of about 0.8 percentage points. Consistent with this positive response, forecasts of key components of spending, consumption and investment, as well as other measures of real activity also increase. On the other hand, forecasts of inflation measures all *fall*. A unit increase in the first factor leads to about 0.3 decline in the forecast for four-quarter-ahead headline CPI inflation. The forecast of 10-year ahead real GDP growth and productivity growth both increase, while the forecast of the natural rate in unemployment falls.

Taken together, the estimated loadings of the first factor describe disagreement about the supply side of the economy. The loadings on real GDP–and other other indicators of real activity–and inflation have opposite signs. Moreover, the first factor also raises forecasts of long-term GDP growth and productivity, while it lowers the prediction of the natural rate of unemployment by more than the four-quarter prediction of the actual unemployment rate. This pattern is consistent with disagreement about permanent productivity shocks. It is not consistent with disagreement about temporary markup shocks, all the more since the factor is associated with both higher corporate profits and lower inflation.

Turning to the posterior for the second factor, individual forecasts of the components of real activity, save federal government spending and net exports, are all positive to related to movements in the second factor. Unlike the coefficients associated with the first factor, however, expectations of four-quarter-ahead inflation increases in response to an increase in the second factor. Thus, the second factor can be said to describe disagreement about the demand side of the economy. Interestingly, the factor also raises forecasts of long-term predictions of inflation, nominal interest rates and (nominal) stock returns.

Neither of the factors describes disagreement about the course of monetary policy. Such disagreement would manifest itself in a factor that lowers nominal interest rate forecasts, but raises output and inflation forecasts (or vice-versa). Even if we increase the number of factors to $\kappa = 3$ or $\kappa = 4$, we do not obtain a factor with such loadings. We therefore conclude that disagreement about the course of monetary plays only a minor role in the

Table 3: Posterior of Λ

Variable	Mean $\Lambda_{\cdot 1}$	[5, 95]	Mean $\Lambda_{\cdot 2}$	[5, 95]
GDP components				
RGDP4	0.79	[0.76, 0.81]	0.41	[0.37, 0.45]
RCONSUM4	0.56	[0.52, 0.59]	0.31	[0.28, 0.34]
RNRESIN4	0.92	[0.82, 1.03]	0.76	[0.68, 0.83]
RRESINV4	1.50	[1.32, 1.67]	1.12	[0.98, 1.26]
RSLGOV4	0.27	[0.22, 0.31]	0.15	[0.11, 0.18]
RFEDGOV4	0.26	[0.17, 0.35]	-0.06	[-0.12, 0.00]
RCBI4	0.03	[0.02, 0.04]	0.04	[0.03, 0.04]
REXPORT4	0.07	[0.05, 0.09]	-0.02	[-0.03, 0.00]
Other real activity				
HOUSING4	0.50	[0.31, 0.70]	1.54	[1.35, 1.73]
INDPROD4	0.20	[0.15, 0.24]	0.61	[0.57, 0.65]
CPROF4	0.41	[0.26, 0.56]	1.37	[1.24, 1.50]
Labor market				
UNEMP4	-0.39	[-0.41, -0.37]	0.12	[0.10, 0.15]
EMP4	-0.27	[-0.30, -0.24]	0.18	[0.16, 0.20]
Inflation	-0.25	[-0.29, -0.22]	0.18	[0.16, 0.20]
PGDP4	-0.23	[-0.26, -0.20]	0.17	[0.15, 0.18]
CPI4				
CORECPI4	-0.03	[-0.04, -0.02]	-0.12	[-0.13, -0.12]
COREPCE4	0.15	[0.12, 0.18]	0.12	[0.10, 0.14]
Interest rates				
TBILL4	-0.03	[-0.05, -0.00]	0.10	[0.08, 0.12]
TBONDTBILL4	0.04	[0.02, 0.06]	-0.01	[-0.03, 0.00]
BONDTBOND4	-0.04	[-0.05, -0.02]	-0.04	[-0.05, -0.03]
BAABONDBOND4	-0.03	[-0.05, 0.00]	-0.01	[-0.03, 0.01]
Recession prob.				
RECESS4	-0.98	[-1.24, -0.72]	-1.88	[-2.13, -1.63]
Long-term forecasts				
UBAR	-0.14	[-0.20, -0.08]	-0.03	[-0.07, 0.00]
STOCK10	0.04	[-0.16, 0.27]	0.12	[0.01, 0.24]
PROD10	0.13	[0.08, 0.17]	0.03	[0.00, 0.05]
RGDP10	0.18	[0.14, 0.21]	0.05	[0.03, 0.07]
TBILL10	-0.07	[-0.16, 0.03]	0.13	[0.08, 0.18]
CPI10	-0.09	[-0.11, -0.07]	0.08	[0.07, 0.10]

The table shows the posterior mean, 5th, and 95th percentile.

Table 4: Posterior of Factor Autoregressive Parameters

Variable	mean	[5, 95]
ϕ_1	0.45	[0.42, 0.49]
ϕ_2	0.80	[0.78, 0.82]

The table shows the posterior mean, 5th, and 95th percentile.

structure of disagreement. That said, our results do not preclude that forecasters might disagree about the *effects* of monetary policy, insofar as the demand-side disagreement in our second factor could stem from disagreement about the strength of aggregate demand that is induced by monetary policy surprises.

The estimation also reveals that disagreement is persistent, and more so for demand-side disagreement. This persistence for the two factors are $\phi_1 = 0.47$ and $\phi_2 = 0.83$, respectively, at the posterior means. The decay of the idiosyncratic components is lower, with a posterior mean of $\rho = 0.59$. The posterior estimates for the idiosyncratic components are relegated to the appendix.

6.2 Variance decompositions

How much disagreement is explained by the factors varies considerably by variable, as Figure 3 shows. The figure plots the sample variance decomposition of disagreement for each variable j and horizon h, which is computed with the formula:

$$1 = \sum_{k=1}^{p} \frac{\sum_{i,t} y_{ijht} \Lambda_{ijk} \hat{f}_{kit}}{\sum_{i,t} y_{ijht}^2} + \frac{\sum_{i,t} y_{ijht} \hat{\xi}_{ijht}}{\sum_{i,t} y_{ijht}^2}. \tag{27}$$

The factors and error terms are the mean smoothed values at the estimated parameter values.

On average across variables, the two factors explain about 20 percent of the variance of disagreement, so that most of the forecast disagreement in the SPF is estimated to be uncorrelated across forecast variables. This finding is consistent with Dovern (2015). Nevertheless, the factors pick up almost all of the variance of real GDP disagreement and a about three fifths of disagreement in core CPI and core PCE inflation. At the other end of the spectrum, bond spreads are almost entirely explained by their idiosyncratic components, indicating that disagreement about spreads does not systematically comove with most other forecast variables.

Because our estimation yields estimates of the time series of the factors for each forecaster,we

Figure 3: Full-sample Variance Decompositions.

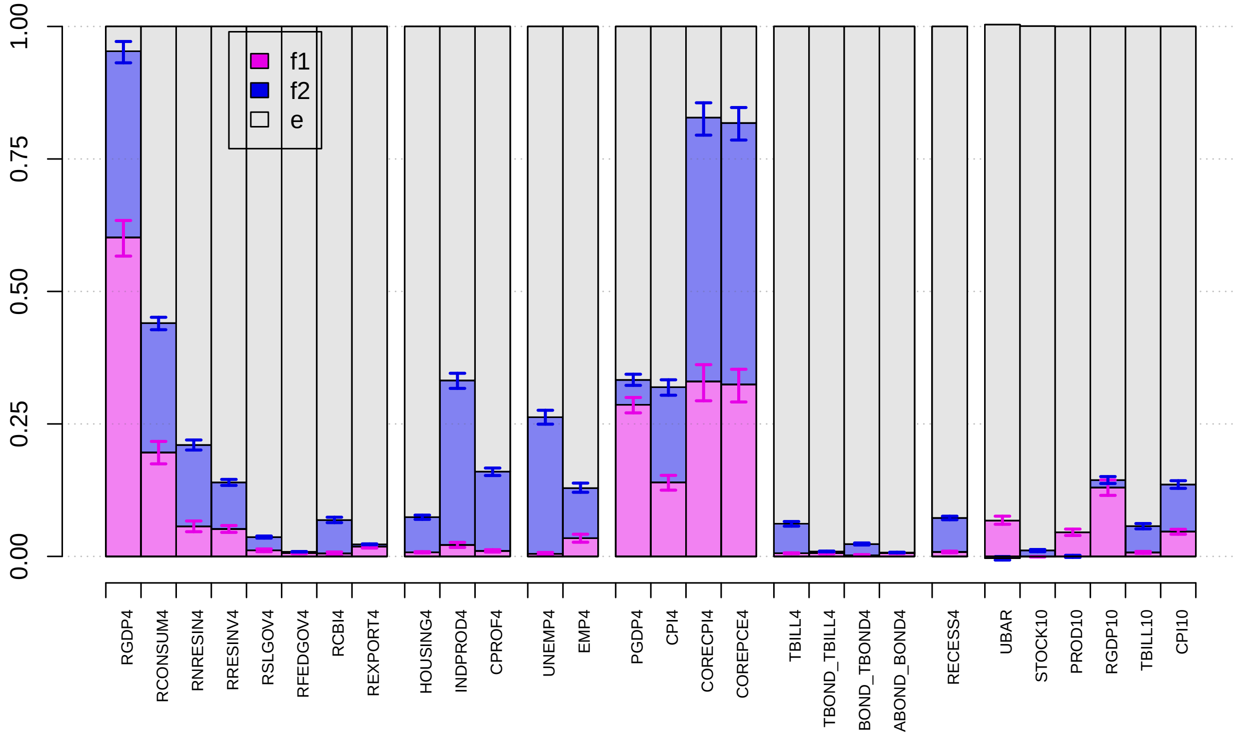

Note: Decomposition of the sample variance of disagreement across forecasters and time into sample covariances with both common components and the idiosyncratic component. Fraction of total variance. Factors and errors terms are smoothed mean values. Bars indicate decompositions at the posterior mean parameter values while whiskers indicate 5th and 95th percentiles across the posterior parameter distribution.

can even decompose disagreement at each period in time, by repeating the decomposition in (27) for each variable j, horizon h and time period t:

$$\widehat{\mathrm{Var}}_t\left(y_{ijht}\right) = \frac{1}{m}\sum_{k=1}^{p}\sum_{i=1}^{m} y_{ijht}\Lambda_{ijk}\hat{f}_{kit} + \frac{1}{m}\sum_{k=1}^{p}\sum_{i=1}^{m} y_{ijht}\beta_{ijk}\hat{\xi}_{ijht},\ t = 1,\ldots,T. \quad (28)$$

In Figure 4, we plot this decomposition for real GDP and (headline) CPI inflation. The appendix contains the full set of decompositions for all variables.

The black lines in the figure are the standard deviations of forecast disagreement, which traces the decline of output disagreement since the late 1960s. The bars show the fraction of disagreement explained by the factors and the idiosyncratic component.[8] For real GDP, the idiosyncratic component is estimated to be very small, in contrast to most other variables. Intuitively, most revisions to GDP are associated with revisions to other forecast variables but not the other way around. The figure also shows that the contribution of the second (demand-side) factor to output disagreement has been rising over time, since the yellow bars representing the covariance of the factor with forecasts in the cross-section remain stable in size, but the total variance of forecasts decreases over time. A similar picture emerges for CPI inflation, although a more sizeable portion of dispersion in that variable is left unexplained by the factors.

6.3 Factor dispersion

While the variance of the factors is normalized to unity in the estimation, fitting the data requires matching time-varying changes in forecast dispersion across variables. In Figure 5, we plot the cross-sectional dispersion of the estimated factors across time $\sum_{i=1}^{m}\hat{f}_{kit}^{2}/m$ against time. The figure reveals that the dispersion in each of the factors follows a distinct pattern. For the first factor describing supply-side disagreement, dispersion is largest in the earlier part of the sample, before the Great Moderation. Dispersion increases during the Great Recession but still remains below the levels experienced in the 1970s. By contrast, the second factor describing demand-side disagreement has distinct spikes during the Volcker disinflation in the early 1980s and during the Great Recession around 2009. This suggests that disagreement about the demand side of the economy was more pronounced during these periods.

[8] In order to provide a better visual comparison to the literature which uses the standard deviation rather than the variance, we show the decomposition in standard deviation space by taking the signed squared root, where the signed squared root of x equals $\mathrm{sgn}\left(x\right)\sqrt{\mid x\mid}$.

Figure 4: Decomposition of Dispersion over time.

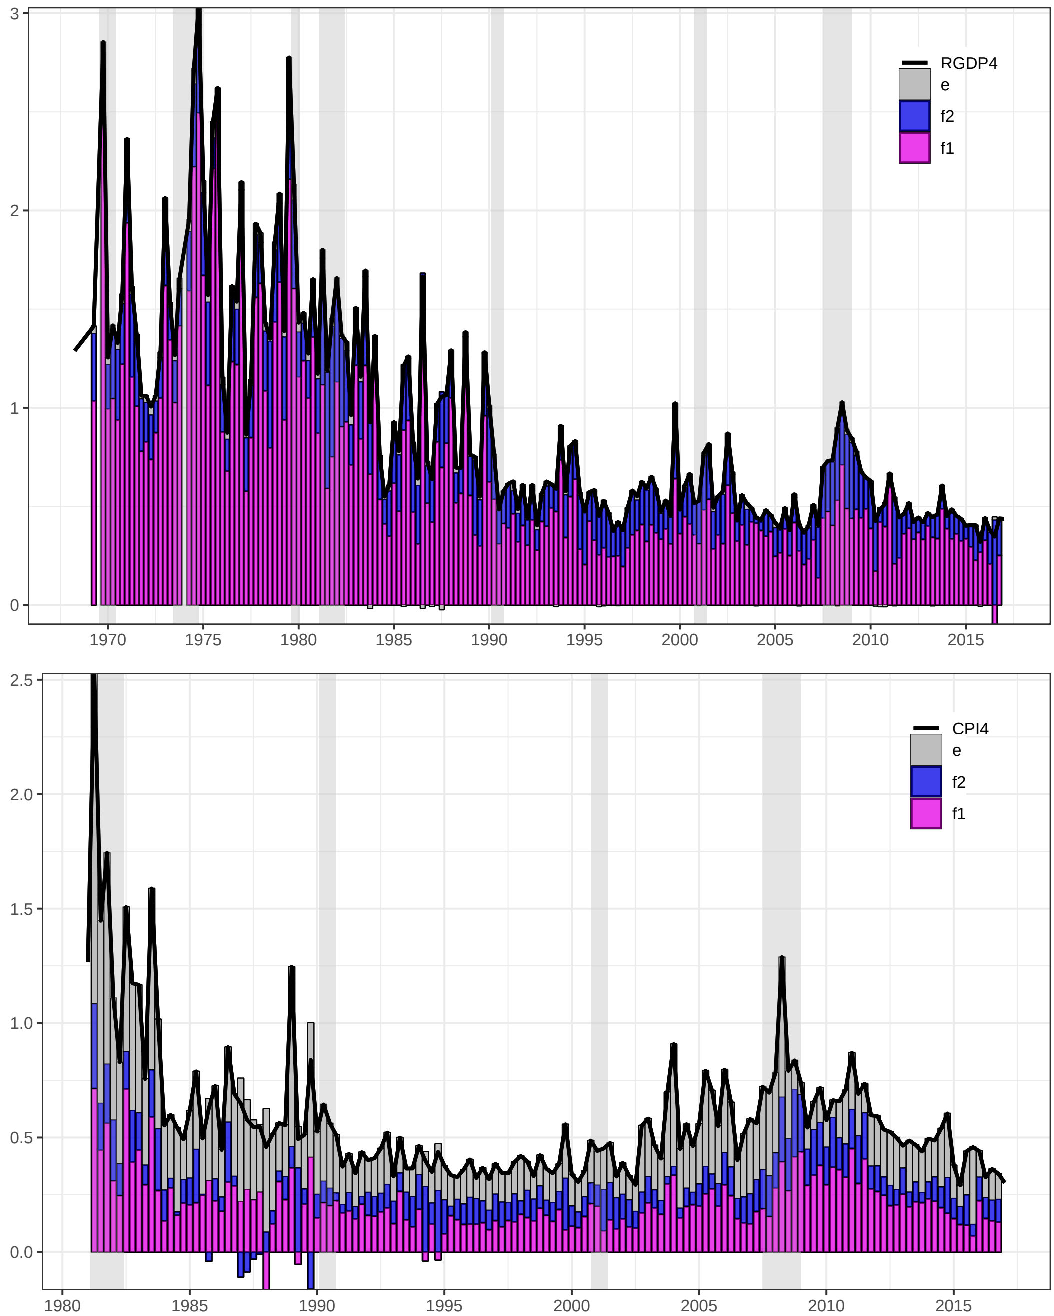

Note: Decomposition of the sample variance of disagreement across forecasters into sample covariances with both common components and the idiosyncratic component. Full sample, RGDP4 and CPI4 variables. Factors and error terms are smoothed mean values at the posterior mean parameters. Variances and covariances are transformed by taking the signed squared root. Black line is the sample standard deviation.

Figure 5: Cross-Sectional Dispersion of Factors over time.

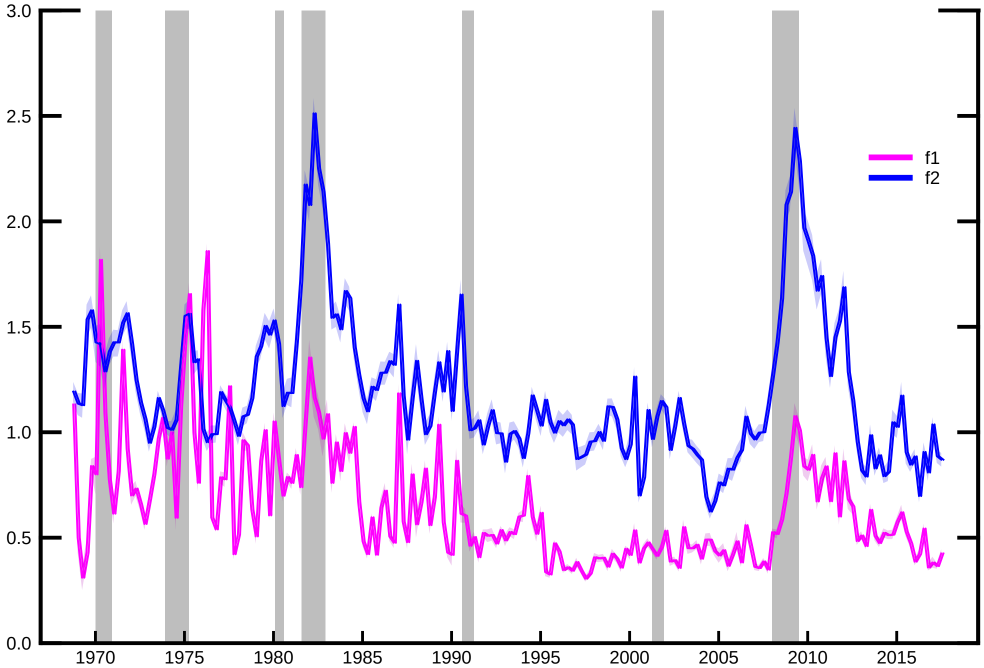

Note: Cross-sectional standard deviation over time of the mean smoothed factors. Solid lines indicate values at the posterior mean parameters while shaded areas indicate values between the 5th and 95th percentiles across the posterior parameter distribution.

7 Subsamples and consensus forecast

In this section, we apply the factor model estimation to subsamples of the data, uncovering substantial time variation in the estimated loadings. We also discuss a factor model to the consensus forecast.

7.1 Subsample analysis

To examine how the structure of disagreement has changed over time, we repeat the estimation on three subsamples that are divided by the (approximate) start of the Great Moderation and the start of the Great Recession. Table 5 summarizes factor loadings for key variables. The full set of loadings can be found in the appendix.

Table 5 reveals that the structure of disagreement has changed substantially over time. Before the Great Moderation, both extracted factors both look like supply-side disagreement, as their loadings induce a negative correlation between real GDP and inflation forecasts. The loadings of the first factor also have the same sign for inflation and corporate profit forecasts, while those of the second factor have opposite signs. Through the lens of the New-Keynesian

Table 5: Summary of estimated factor loadings. Subsamples, cross-section of variables.

	1968Q3–1984Q4		1985Q1–2008Q2		2008Q3–2017Q2	
mean	$\Lambda_{\cdot 1}$	$\Lambda_{\cdot 2}$	$\Lambda_{\cdot 1}$	$\Lambda_{\cdot 2}$	$\Lambda_{\cdot 1}$	$\Lambda_{\cdot 2}$
RGDP4	0.73*	0.77*	0.32 *	0.31*	0.22*	0.27*
CPROF4	-0.55*	2.26*	0.95*	0.77*	0.80*	0.77*
UNEMP4	0.01	-0.18*	-0.03*	-0.10*	-0.01	-0.11*
PGDP4	-0.81*	-0.10*	-0.06*	0.03*	-0.16*	0.16*
TBILL4	0.02	-0.09	-0.23*	0.32*	-0.04*	0.07*

Note: Number of observations: 2,854 (1st subsample), 3,182 (2nd subsample), 1,559 (3rd subsample). * indicate that the interval between the 5th and 95th percentile of the posterior does not include zero.

model described earlier, the two factors could be interpreted as disagreement about cost-push shocks and productivity shocks, respectively. The loadings on interest rates are relatively imprecisely estimated, which is largely due to the fact that interest rate forecasts are only elicited starting in 1981.

After the Great Recession, both extracted factors both look like demand-side disagreement as their loadings induce a positive correlation between real GDP and inflation forecasts. In the second subsample, during the Great Moderation, this correlation is weak, but then strengthens again in the third subsample, after the Great Recession. Interest rate disagreement is strongly associated with the factors during the Great Moderation, consistent with forecasts of a strong expected reaction of monetary policy to inflation and output. After the Great Recession, the factor loadings on interest rate forecasts diminish in size, consistent with short-term interest rates being constrained by the zero lower bound during most of this subsample.

7.2 Factors in the consensus forecast

The data we use to estimate our factor model are deviations of individual forecasts from the consensus (mean) forecast. In doing so, we effectively ignore throw away all forecast movements that affect all forecasters alike. But these movements are of course also informative for the way in which expectations are formed. In this section, we ask what factors drive consensus forecasts, and whether these factors exhibit similarities to those that drive disagreement.

We apply the factor model (6)–(8) on consensus forecasts $\bar{y}_{jht}$ of our sample, treating the observations as if they were coming from a single forecaster ($m = 1$). The priors are the same as before. The consensus forecasts constitute a standard two-dimensional panel in time and forecast variables. Moreover, the consensus forecasts are closely tied to the realizations of the variables. For these two reasons, the estimation of the factor model on consensus forecasts becomes quite similar to the DFM models on realized data e.g. in Stock and Watson (2002).

The posterior estimates of the factor loadings Λ are tabulated in Table 6. The factor loadings are much less precisely estimated due to the lower number of observations (192 instead of 7,965 for the disagreement data). The loadings do not closely resemble those estimated from the disagreement data, and in fact there is no reason to expect such a correspondence a priori. The first factor loads almost exclusively on inflation variables, picking up the strong comovement between these variables that seems unrelated to other forecasts (see also Figure 8). The second factor loads on a broad range of variables. The loadings broadly resemble those of the demand factor in the disagreement data, since loadings of GDP and its main components, inflation measures, and interest rates all have the same sign.

8 Conclusion

We have estimated a dynamic factor model that captures in a parsimonious way the comovement of disagreement, and offered an interpretation of of the extracted factors through a semi-structural model of heterogeneous expectations. In the model, forecasters predict a data-generating process that evolves independently of their predictions. Forecast disagreement arises because either forecasters receive idiosyncratic noisy signals about the state of the economy, or because they use slightly different forecasting models. Using the New-Keynesian model, we have categorized disagreement about the supply side and the demand side economy, as well as disagreement about monetary policy, by the signs of the factor loadings on inflation, output, and interest rates.

The estimation revealed that the comovement of disagreement is best described by one factor that captures supply-side disagreement, and one that captures demand-side disagreement. The supply side factor moves expectations of long-run GDP growth, productivity and the natural rate of unemployment while the demand side factor does not. Variance decompositions reveal that disagreement about the demand side has become more important over time and has been particularly important during the Great Recession. Disagreement about the stance of monetary policy, which would manifest itself a combination of lower interest rate, higher inflation and higher output forecasts, plays only a limited role throughout our sample. Our empirical findings can serve to discipline structural models of heterogeneous expectations, and we hope that future research will make use of the results presented here to identify models consistent with the structure of disagreement as it is documented here.

Our estimation is based on forecasts at the four-quarter horizon and certain long-run forecasts

Table 6: Posterior of Λ, Consensus Forecast Model

Variable	Mean $\Lambda_{\cdot 1}$	[5, 95]	Mean $\Lambda_{\cdot 2}$	[5, 95]
GDP components				
RGDP4	-0.03	[-0.10, 0.06]	0.56	[0.51, 0.62]
RCONSUM4	0.01	[-0.06, 0.08]	0.41	[0.35, 0.47]
RNRESIN4	-0.15	[-0.38, 0.07]	1.41	[1.17, 1.66]
RRESINV4	-0.16	[-0.43, 0.12]	1.11	[0.72, 1.52]
RSLGOV4	0.02	[-0.02, 0.07]	0.08	[0.01, 0.16]
RFEDGOV4	0.03	[-0.13, 0.18]	0.05	[-0.20, 0.30]
RCBI4	0.01	[-0.00, 0.03]	0.06	[0.04, 0.08]
REXPORT4	-0.02	[-0.07, 0.03]	0.01	[-0.08, 0.09]
Other real activity				
HOUSING4	-0.50	[-1.08, 0.08]	2.28	[1.65, 2.92]
INDPROD4	-0.02	[-0.18, 0.13]	1.01	[0.90, 1.12]
CPROF4	0.06	[-0.40, 0.53]	2.54	[2.19, 2.91]
Labor market				
UNEMP4	0.03	[-0.01, 0.07]	-0.10	[-0.15, -0.06]
EMP4	-0.04	[-0.11, 0.04]	0.42	[0.33, 0.51]
Inflation				
PGDP4	0.95	[0.66, 1.22]	0.13	[-0.02, 0.27]
CPI4	1.96	[1.74, 2.20]	0.07	[-0.20, 0.34]
CORECPI4	0.23	[0.14, 0.33]	0.19	[0.05, 0.33]
COREPCE4	0.20	[0.12, 0.27]	0.17	[0.06, 0.29]
Interest rates				
TBILL4	0.03	[-0.04, 0.11]	0.28	[0.18, 0.38]
TBONDTBILL4	-0.02	[-0.07, 0.03]	-0.03	[-0.12, 0.06]
BONDTBOND4	0.00	[-0.03, 0.03]	-0.10	[-0.14, -0.06]
BAABONDBOND4	0.04	[0.00, 0.08]	-0.06	[-0.14, 0.02]
Recession prob.				
RECESS4	-0.03	[-0.51, 0.44]	-0.92	[-1.37, -0.47]
Long-term forecasts				
UBAR	0.12	[-0.03, 0.28]	-0.08	[-0.24, 0.08]
STOCK10	0.02	[-0.25, 0.29]	0.17	[-0.15, 0.49]
PROD10	0.07	[-0.03, 0.17]	0.13	[0.03, 0.24]
RGDP10	0.08	[-0.00, 0.16]	0.16	[0.08, 0.25]
TBILL10	0.17	[-0.15, 0.50]	0.03	[-0.34, 0.42]
CPI10	1.99	[1.76, 2.24]	-0.10	[-0.39, 0.20]

The table shows the posterior mean, 5th, and 95th percentile for the 2 factor model estimated for the consensus sample. The simulation uses 3,000 draws.

in the SPF, but we can flexibly extend it to cover all forecast horizons. Doing so would allow us to uncover not only comovement of disagreement across variables but also across forecast horizons, thereby identifying more of the dynamics of expectations. It would also be interesting to apply the methodology to other forecasting panel data such as the Blue Chip or Consensus Forecast datasets which, although having shorter overall samples, feature monthly surveys and different sets of forecast variables.

References

Andrade, P., R. K. Crump, S. Eusepi, and E. Moench (2016): "Fundamental disagreement," *Journal of Monetary Economics*, 83, 106 – 128.

Angeletos, G.-M., and J. La'O (2013): "Sentiments," *Econometrica*, 81(2), 739–779.

Banbura, M., and M. Modugno (2014): "Maximum Likelihood Estimation of Factor Models on Datsets With Arbitrary Pattern of Missing Data," *Journal of Applied Econometrics*, 29(1), 133–160.

Barillas, F., and K. Nimark (2013): "Speculation, Risk Premia and Expectations in the Yield Curve," Working Papers 659, Barcelona Graduate School of Economics.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2018): "Over-reaction in Macroeconomic Expectations," NBER Working Papers 24932, National Bureau of Economic Research, Inc.

Burnside, C., M. Eichenbaum, and S. Rebelo (2016): "Understanding Booms and Busts in Housing Markets," *Journal of Political Economy*, 124(4), 1088–1147.

Capistran, C., and A. Timmermann (2009): "Disagreement and Biases in Inflation Expectations," *Journal of Money, Credit and Banking*, 41(2-3), 365–396.

Carvalho, C., and F. Nechio (2014): "Do people understand monetary policy?," *Journal of Monetary Economics*, 66, 108 – 123.

Chopin, N. (2002): "A Sequential Particle Filter for Static Models," *Biometrika*, 89(3), 539–551.

Coibion, O., and Y. Gorodnichenko (2015): "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, 105(8), 2644–78.

Dovern, J. (2015): "A multivariate analysis of forecast disagreement: Confronting models of disagreement with survey data," *European Economic Review*, 80(C), 16–35.

Dovern, J., U. Fritsche, and J. Slacalek (2012): "Disagreement Among Forecasters in G7 Countries," *The Review of Economics and Statistics*, 94(4), 1081–1096.

Doz, C., D. Giannone, and L. Reichlin (2012): "A Quasi-Maximum Likelihood Approach for Large, Approximate Dynamic Factor Models," *The Review of Economics and Statistics*, 94(4), 1014–1024.

Dräger, L., M. J. Lamla, and D. Pfajfar (2016): "Are survey expectations theory-consistent? The role of central bank communication and news," *European Economic Review*, 85, 84 – 111.

Geweke, J., and G. Zhou (1996): "Measuring the pricing error of the arbitrage pricing theory," *The review of financial studies*, 9(2), 557–587.

Herbst, E., and F. Schorfheide (2014): "SEQUENTIAL MONTE CARLO SAMPLING FOR DSGE MODELS," *Journal of Applied Econometrics*, 29(7), 1073–1098.

Ilut, C. L., and M. Schneider (2014): "Ambiguous Business Cycles," *American Economic Review*, 104(8), 2368–2399.

Lorenzoni, G. (2009): "A Theory of Demand Shocks," *American Economic Review*, 99(5), 2050–84.

Lu, X., and L. Su (2018): "Estimation and Inference for Three-Dimensional Factor Models," Working paper.

Mackowiak, B., and M. Wiederholt (2009): "Optimal Sticky Prices under Rational Inattention," *American Economic Review*, 99(3), 769–803.

Mankiw, N. G., R. Reis, and J. Wolfers (2003): "Disagreement about inflation expectations," *NBER macroeconomics annual*, 18, 209–248.

Melosi, L. (2014): "Estimating Models with Dispersed Information," *American Economic Journal: Macroeconomics*, 6(1), 1–31.

Nimark, K. P. (2014): "Man-Bites-Dog Business Cycles," *The American Economic Review*, 104(8), 2320–2367.

Otrok, C., and C. H. Whiteman (1998): "Bayesian Leading Indicators: Measuring and Predicting Economic Conditions in Iowa," *International Economic Review*, 39(4), 997–1014.

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.

Rich, R. W., and J. Tracy (2017): "The behavior of uncertainty and disagreement and their roles in economic prediction: a panel analysis," Staff Reports 808, Federal Reserve Bank of New York.

Scheinkman, J. A., and W. Xiong (2003): "Overconfidence and Speculative Bubbles," *Journal of Political Economy*, 111(6), 1183–1219.

Stock, J. H., and M. W. Watson (1989): "New indexes of coincident and leading economic indicators," *NBER macroeconomics annual*, 4, 351–394.

Stock, J. H., and M. W. Watson (1999): "Business cycle fluctuations in us macroeconomic time series," vol. 1 of *Handbook of Macroeconomics*, chap. 1, pp. 3 – 64. Elsevier.

Stock, J. H., and M. W. Watson (2002): "Macroeconomic Forecasting Using Diffusion Indexes," *Journal of Business & Economic Statistics*, 20(2), 147–162.

Woodford, M. (2002): "Imperfect Common Knowledge and the Effects of Monetary Policy," in *Knowledge, Information and Expectations in Modern Macroeconomics*, ed. by P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, no. 8673. Princeton University Press.

A Additional tables and figures

Table 7: List of SPF Variables.

Symbol	Variable	Forecast horizon	Freq.	Included since
RGDP	Real GDP (s.a.)	0–4 q., 10 y.*	Q	1968:Q4
NGDP	Nominal GDP (s.a.)	0–4 q.	Q	1968:Q4
PGDP	GDP price index (s.a.)	0–4 q.	Q	1968:Q4
UNEMP	Unemployment rate (s.a.)	0–4 q.	Q	1968:Q4
INDPROD	Industrial production (s.a.)	0–4 q.	Q	1968:Q4
CPROF	Corporate profits after tax (s.a.)	0–4 q.	Q	1968:Q4
HOUSING	Housing starts (s.a.)	0–4 q.	Q	1968:Q4
RECESS	Probability of recession	0–4 q.	Q	1968:Q4
RCONSUM	Real personal consumption expenditures (s.a.)	0–4 q.	Q	1981:Q3
RNRESIN	Real nonresidential fixed investment (s.a)	0–4 q.	Q	1981:Q3
RRESINV	Real residential fixed investment (s.a.)	0–4 q.	Q	1981:Q3
RFEDGOV	Real federal govt. cons. and gross inv. (s.a.)	0–4 q.	Q	1981:Q3
RSLGOV	Real state and local govt. cons. and gross inv. (s.a.)	0–4 q.	Q	1981:Q3
RCBI	Real change in private inventories (s.a.)	0–4 q.	Q	1981:Q3
REXPORT	Real net exports (s.a.)	0–4 q.	Q	1981:Q3
CPI	Headline CPI inflation (s.a.)	0–4 q., 5&10 y.*	Q	1981:Q3
TBILL	Average yield on 3-month Treasury bills	0–4 q., 10 y.*	Q	1981:Q3
BOND	Average yield on Moody's Aaa corporate bonds	0–4 q.	Q	1981:Q3
TBOND	Average yield on 10-year Treasury bonds	0–4 q., 10 y.*	Q	1992:Q1
STOCK	Average return on S&P500	10 y.	A	1992:Q1
PROD	Average productivity growth	10 y.	A	1992:Q1
UBAR	Natural rate of unemployment	n/a	A	1996:Q3
EMP	Non-farm payroll employment (s.a.)	0–4 q.	Q	2003:Q4
CORECPI	Core CPI inflation (s.a.)	0–4 q.	Q	2007:Q1
PCE	Headline PCE inflation (s.a.)	0–4 q., 5&10 y.	Q	2007:Q1
COREPCE	Core PCE inflation (s.a.)	0–4 q.	Q	2007:Q1
BAABOND	Average yield on Moody's Baa corporate bonds	0–4 q.	Q	2010:Q1

Note: List is sorted by date of first inclusion in the survey. The list excludes probability density forecasts that were started in 2007 and forecasts at annual horizons for which quarterly horizons are also available.
* For CPI, 10-year forecasts start in 1991:Q4 and 5-year forecasts in 2005:Q3. For RGDP, TBILL and TBOND, 10-year forecasts start in 1992:Q1 and are conducted annually.

Table 8: Posterior of ρ and σ^2

Variable	Mean ρ	[5, 95]	Mean σ^2	[5, 95]
GDP components				
RGDP4	0.43	[0.26, 0.61]	0.23	[0.21, 0.26]
RCONSUM4	0.56	[0.52, 0.60]	0.51	[0.50, 0.52]
RNRESIN4	0.63	[0.61, 0.64]	1.75	[1.72, 1.78]
RRESINV4	0.67	[0.66, 0.68]	3.25	[3.19, 3.30]
RSLGOV4	0.64	[0.62, 0.66]	0.90	[0.89, 0.92]
RFEDGOV4	0.67	[0.66, 0.69]	1.71	[1.68, 1.74]
RCBI4	0.48	[0.39, 0.57]	0.18	[0.17, 0.18]
REXPORT4	0.59	[0.54, 0.64]	0.34	[0.33, 0.35]
Other real activity				
HOUSING4	0.66	[0.65, 0.67]	7.46	[7.35, 7.57]
INDPROD4	0.55	[0.53, 0.57]	1.23	[1.21, 1.25]
CPROF4	0.47	[0.46, 0.49]	4.88	[4.80, 4.96]
Labor market				
UNEMP4	0.63	[0.58, 0.68]	0.27	[0.26, 0.27]
EMP4	0.55	[0.48, 0.62]	0.33	[0.32, 0.34]
Inflation				
PGDP4	0.58	[0.55, 0.62]	0.56	[0.54, 0.57]
CPI4	0.55	[0.49, 0.60]	0.48	[0.47, 0.49]
CORECPI4	0.54	[0.36, 0.70]	0.14	[0.13, 0.15]
COREPCE4	0.52	[0.34, 0.68]	0.13	[0.13, 0.14]
Interest rates				
TBILL4	0.70	[0.66, 0.73]	0.48	[0.48, 0.49]
TBONDTBILL4	0.54	[0.49, 0.59]	0.36	[0.36, 0.37]
BONDTBOND4	0.47	[0.40, 0.54]	0.28	[0.27, 0.28]
BAABONDBOND4	0.36	[0.23, 0.48]	0.23	[0.22, 0.24]
Recession prob.				
RECESS4	0.61	[0.60, 0.62]	9.83	[9.68, 9.97]
Long-term forecasts				
UBAR	0.78	[0.66, 0.88]	0.23	[0.19, 0.26]
STOCK10	0.81	[0.79, 0.84]	1.10	[1.03, 1.17]
PROD10	0.79	[0.72, 0.85]	0.27	[0.24, 0.29]
RGDP10	0.71	[0.60, 0.81]	0.21	[0.19, 0.24]
TBILL10	0.77	[0.72, 0.81]	0.51	[0.47, 0.55]
CPI10	0.80	[0.76, 0.85]	0.27	[0.26, 0.27]

The table shows the posterior mean, 5th, and 95th percentile for the 2 factor model estimated over the entire sample. The simulation uses 3,000 draws.

Figure 6: Decompositions of Forecast Dispersion over time (I).

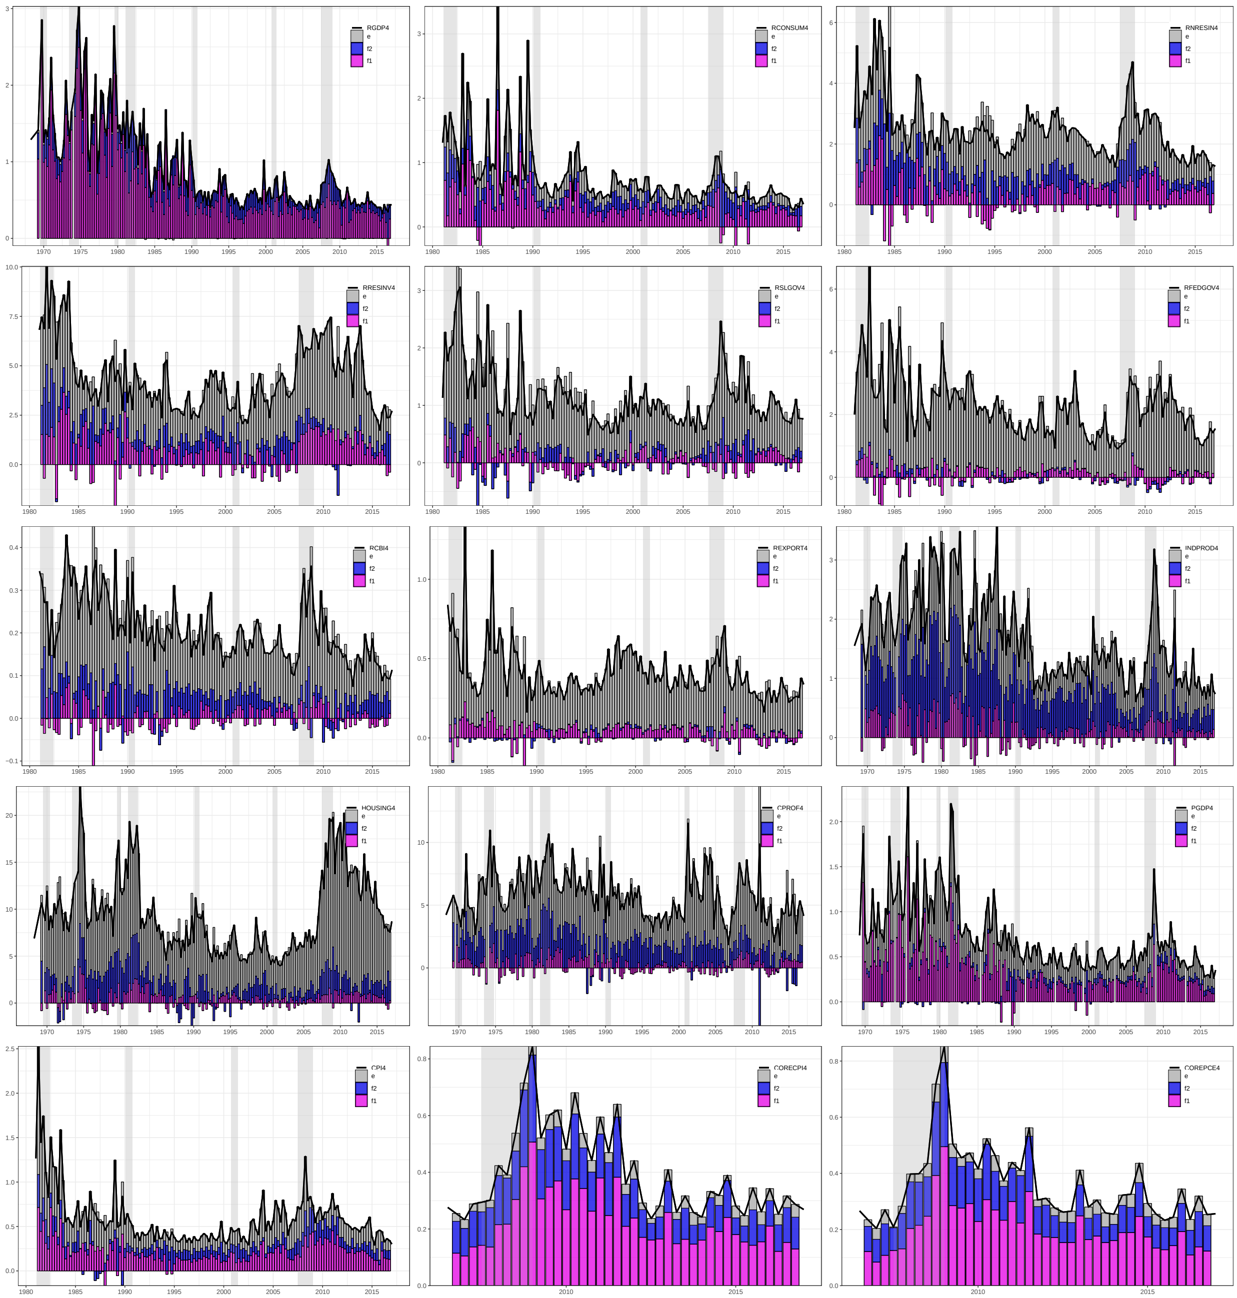

Note: Decompositions of the sample variance of disagreement across forecasters into sample covariances with both common components and the idiosyncratic component. Full sample. For each variable, factors and error terms are mean smoothed values at estimated parameter values. Variances and covariances are transformed by taking the signed squared root. Black line is the sample standard deviation.

Figure 7: Decompositions of Forecast Dispersion over time (II).

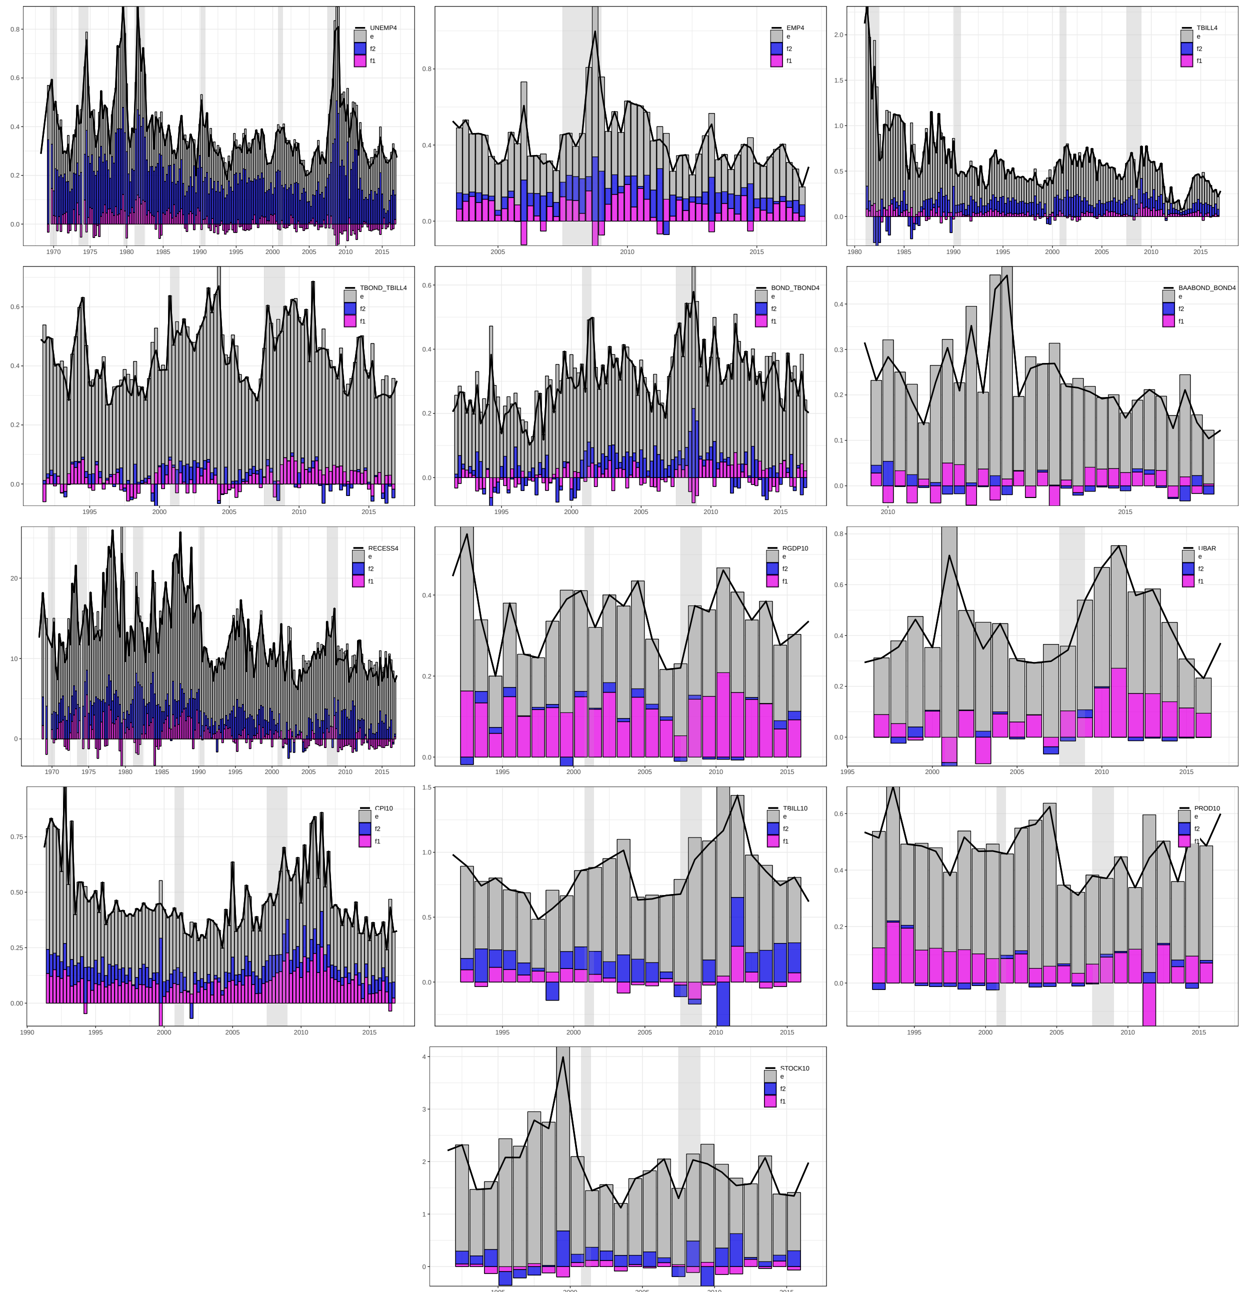

Note: Decompositions of the sample variance of disagreement across forecasters into sample covariances with both common components and the idiosyncratic component. Full sample. For each variable, factors and error terms are mean smoothed values at estimated parameter values. Variances and covariances are transformed by taking the signed squared root. Black line is the sample standard deviation.

Table 9: Posterior of Λ for Subsamples

	1968Q3–1984Q4		1985Q1–2008Q2		2008Q3–2017Q2	
mean	$\Lambda_{\cdot 1}$	$\Lambda_{\cdot 2}$	$\Lambda_{\cdot 1}$	$\Lambda_{\cdot 2}$	$\Lambda_{\cdot 1}$	$\Lambda_{\cdot 2}$
GDP components						
RGDP4	0.73*	0.77*	0.32*	0.31*	0.22*	0.27*
RCONSUM4	0.33*	0.64*	0.30*	0.24*	0.15*	0.20*
RNRESIN4	0.91*	1.32*	0.35*	0.57*	0.53*	0.64*
RRESINV4	-0.40	2.83*	0.73*	0.51*	0.99*	0.86*
RSLGOV4	0.41*	0.34*	0.11*	0.11*	0.05	0.15*
RFEDGOV4	1.39*	-0.17	0.02	0.13*	0.06	0.11*
RCBI4	0.06*	0.08*	0.01*	0.03*	0.01	0.02*
REXPORT4	0.09	0.02	0.02*	-0.00	0.02*	-0.01
Other real activity						
HOUSING4	-0.21	2.90*	0.84*	0.46*	1.36*	1.91*
INDPROD4	-0.16*	1.08*	0.23*	0.36*	0.16*	0.33*
CPROF4	-0.55*	2.26*	0.95*	0.77*	0.80*	0.77*
Labor market						
UNEMP4	0.01	-0.18*	-0.03*	-0.10*	-0.01	-0.11*
EMP4			0.07*	0.11*	0.04*	0.09*
Inflation						
PGDP4	-0.81*	-0.10*	-0.06*	0.03*	-0.16*	0.16*
CPI4	-0.32*	-0.02	-0.04*	0.06*	-0.20*	0.19*
CORECPI4			-0.01	0.07*	-0.18*	0.17*
COREPCE4			-0.01	0.06*	-0.16*	0.15*
Interest rates						
TBILL4	0.02	-0.09	-0.23*	0.32*	-0.04*	0.07*
TBONDTBILL4			0.18*	-0.16*	0.04*	0.03*
BONDTBOND4			-0.00	-0.03*	-0.02	-0.04*
BAABONDBOND4					-0.01	-0.01
Recession prob.						
RECESS4	-0.66*	-3.37*	-1.56*	-1.20*	-0.66*	-1.21*
Long-term forecasts						
UBAR			-0.05	-0.01	-0.11*	-0.04
STOCK10			0.00	0.07	-0.03	0.15*
PROD10			0.06*	0.02	0.04	0.04*
RGDP10			0.08*	0.05*	0.09*	0.04*
TBILL10			-0.10*	0.11*	-0.06	0.17*
CPI10			0.00	0.03*	-0.09*	0.08*

The table shows the posterior mean for the 2 factor model estimated oversubsamples. * indicate that the interval between the 5th and 95th percentile of the posterior does not include zero.

Table 10: Posterior of ρ and σ^2 for Consensus Forecasts.

Variable	Mean ρ	[5, 95]	Mean σ^2	[5, 95]
GDP components				
RGDP4	0.53	[0.31, 0.76]	0.22	[0.20, 0.25]
RCONSUM4	0.50	[0.28, 0.71]	0.26	[0.23, 0.29]
RNRESIN4	0.88	[0.81, 0.94]	0.95	[0.86, 1.06]
RRESINV4	0.93	[0.90, 0.96]	1.86	[1.69, 2.04]
RSLGOV4	0.82	[0.70, 0.94]	0.34	[0.30, 0.39]
RFEDGOV4	0.83	[0.76, 0.89]	1.17	[1.06, 1.29]
RCBI4	0.53	[0.29, 0.77]	0.10	[0.08, 0.12]
REXPORT4	0.95	[0.90, 0.99]	0.41	[0.37, 0.46]
Other real activity				
HOUSING4	0.98	[0.98, 0.99]	4.70	[4.32, 5.11]
INDPROD4	0.68	[0.53, 0.82]	0.55	[0.49, 0.60]
CPROF4	0.82	[0.78, 0.86]	2.15	[1.96, 2.35]
Labor market				
UNEMP4	0.92	[0.84, 0.98]	0.35	[0.32, 0.39]
EMP4	0.57	[0.36, 0.78]	0.25	[0.21, 0.30]
Inflation				
PGDP4	0.16	[0.05, 0.26]	1.49	[1.33, 1.66]
CPI4	0.50	[0.29, 0.71]	0.41	[0.35, 0.47]
CORECPI4	0.44	[0.24, 0.64]	0.47	[0.39, 0.55]
COREPCE4	0.47	[0.26, 0.67]	0.37	[0.31, 0.43]
Interest rates				
TBILL4	0.66	[0.51, 0.81]	0.47	[0.42, 0.52]
TBONDTBILL4	0.82	[0.70, 0.95]	0.35	[0.30, 0.40]
BONDTBOND4	0.59	[0.38, 0.80]	0.16	[0.13, 0.19]
BAABONDBOND4	0.51	[0.27, 0.75]	0.09	[0.08, 0.11]
Recession prob.				
RECESS4	0.76	[0.73, 0.78]	3.43	[3.15, 3.72]
Long-term forecasts				
UBAR	0.53	[0.31, 0.74]	0.31	[0.24, 0.38]
STOCK10	0.79	[0.64, 0.92]	0.60	[0.44, 0.77]
PROD10	0.52	[0.30, 0.74]	0.24	[0.19, 0.29]
RGDP10	0.51	[0.29, 0.74]	0.20	[0.16, 0.24]
TBILL10	0.76	[0.61, 0.89]	0.72	[0.55, 0.93]
CPI10	0.51	[0.31, 0.72]	0.44	[0.38, 0.50]

The table shows the posterior mean, 5th, and 95th percentile for the 2 factor model estimated over the entire sample.

Table 11: Posterior of Autoregressive Parameters for Consensus Forecasts.

Variable	mean	[5, 95]
ϕ_1	0.22	[0.08, 0.37]
ϕ_2	0.85	[0.78, 0.91]

The table shows the posterior mean, 5th, and 95th percentile for the 2 factor model estimated for the consensus sample.

Figure 8: Variance Decomposition for Consensus Forecasts.

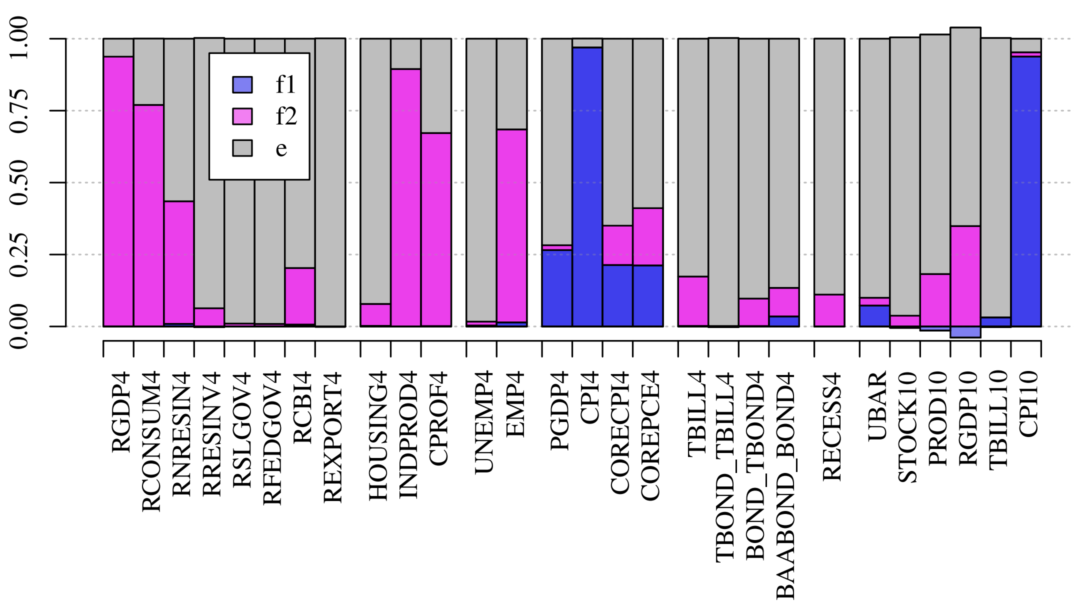

Note: Decomposition of the sample variance of consensus forecasts into covariances with both common components and the idiosyncratic component. Fraction of total variance. Factors and errors terms are mean smoothed values at the estimated parameter values.