

**The Resolution of Macroeconomic Uncertainty:
Evidence from Survey Forecasts***

Andrew J. Patton, University of Oxford

Allan Timmermann, University of California, San Diego

November 20, 2008

Abstract

We develop an unobserved components approach to study surveys of forecasts containing multiple forecast horizons. Under the assumption that forecasters optimally update their beliefs about past, current and future state variables as new information arrives, we use our model to extract information on the degree of predictability of the state variable and the importance of measurement errors on that variable. Empirical estimates of the model are obtained using survey forecasts of annual GDP growth and inflation in the US with forecast horizons ranging from 1 to 24 months, and the model is found to closely match the joint realization of forecast errors at different horizons. Our empirical results confirm several findings in the existing literature that were obtained using very different data sets: professional forecasters face severe measurement error problems for GDP growth in real time, while this is not an problem for inflation; inflation exhibits greater persistence, and thus is predictable at longer horizons, than GDP growth; the persistent component of both of these variables is well-approximated by a low-order autoregressive specification.

Keywords: Fixed-event forecasts, multiple forecast horizons, Kalman filtering, survey data.

*We thank seminar participants at the Board of Governors of the Federal Reserve, Cambridge, City University London, Duke, European Central Bank, London School of Economics, NBER Summer Institute, ESAM08 meetings in Wellington, Oxford, Universite Libre Bruxelles (ECARES), Tilburg and Stanford (SITE workshop) for helpful comments.

Patton: Department of Economics, University of Oxford, Manor Road, Oxford OX1 3UQ, UK.
Email: andrew.patton@economics.ox.ac.uk.

Timmermann: UCSD, 9500 Gilman Drive, La Jolla, CA 92093-0553, USA. Email: atimmerm@ucsd.edu.

1 Introduction

Economic agents' expectations about the state of the economy play an important role in economic analysis. In recent remarks, Chairman of the Federal Reserve Ben Bernanke quotes academic work as showing that “the process of [the public’s] learning can affect the dynamics and even the potential stability of the economy.” He goes on to state that “Undoubtedly, the state of inflation expectations greatly influences actual inflation and thus the central bank’s ability to achieve price stability.” (Bernanke (2007)).

How rapidly uncertainty about macroeconomic variables is resolved through time is an important part of this learning process, partly due to the irreversibility and lags in many economic decisions (Kydland and Prescott (1982) and Dixit and Pindyck (1994)) and partly because of its welfare implications: Ramey and Ramey (1995) link output growth to the degree of uncertainty surrounding it, arguing that firms scale back planned output during periods with high levels of uncertainty.

Consequently, much can be learned by studying how forecasters update their beliefs over time about the same “event”. As a simple illustration, consider Figure 1, which shows consensus forecasts from surveys of US GDP growth for 2002 as this evolved each month from January 2001 (corresponding to a 24-month horizon) to December 2002 (a one-month horizon). A comparison of the initial and final forecasts—at 3.5% and 2.5%, respectively—shows a fairly sizeable reduction in the predicted growth, but fails to reveal the full picture of the dramatic revisions that occurred in the interim. At the beginning of September 2001, the growth forecast for 2002 was 2.7%. Following the events of 9/11, the October 2001 forecast fell to 1.2%, i.e. by a full 1.5%—the largest single-month forecast revision observed in more than a decade. It declined even further to 0.8% in November 2001 before stabilizing. Expectations of 2002 GDP growth then increased by 1.7% from January through April of 2002, from which point the subsequent forecasts were within 0.5% of the actual growth figure, which came in just below 2.5%.

Analysis of “fixed event” survey data such as that presented in Figure 1 is complicated by several factors. First, since forecasts are recorded at both long and short horizons, there is considerable overlap in the forecasts and forecast errors. Second, measurement errors in the underlying variables affect agents’ forecasts introduces a signal extraction problem in agents’ learning process and causes further dependence in forecast errors measured at different horizons. For these reasons, only limited results are available using this type of data, see Nordhaus (1987), Swidler and Ketcher (1990), Davies and Lahiri (1995), Clements (1997), and Isiklar, Lahiri and Loungani (2006).

This paper develops a new approach for extracting information on how rapidly agents learn about the state of the economy and characterizing their views about temporary and persistent components in the predicted variable. Specifically, we develop a framework for studying panels of forecasts containing numerous different forecast horizons (“large H ”) recorded for relatively few time periods (“small T ”). The first contribution of this paper is to analytically reveal the rich information available by studying how forecasts of a variable measured at a low frequency (e.g., annual GDP growth) are updated at a higher frequency (monthly, in our case). We do so by modeling agents’ learning problem—accounting for how they simultaneously backcast, nowcast and forecast past, current and future variables. We then seek to exploit this information using method-of-moments-based estimation techniques to match the properties of forecasts observed across different horizons with the moments implied by our model for agents’ updating process. To conduct inference, we propose a method for simulating standard errors of the moments that are consistent with the underlying model. To our knowledge, this approach for modelling learning and conducting inference has not previously been considered in the literature.

The “large H ” nature of our data enables us to answer a number of interesting questions that are intractable with forecasts of just one or two different horizons, such as the importance of measurement errors, the rate at which uncertainty about macroeconomic variables (as measured by root mean-squared forecast errors, RMSE)

is resolved as the forecast horizon is reduced, and forecasters' beliefs about the current state of the economy (as measured by their "nowcasts" of GDP growth and inflation). *It is of course no surprise that uncertainty about macroeconomic variables declines as the date the variable is revealed draws nearer; the novel aspect of this paper is to propose a model that theoretically explains both the **level** and the **shape** of this uncertainty as a function of the forecast horizon.*

The second contribution of this paper is empirical: we use consensus forecasts of US inflation and GDP growth over 1991-2004, and find many interesting results. Consistent with our model, we find that the rate of uncertainty resolution is faster at short and medium horizons than at long horizons, due in part to the presence of a persistent component in the predicted series, in part to forecasters' access to noisy data on current-period realizations. Measurement error appears to be important in forecasts of GDP growth but not for inflation, a finding that is consistent with other studies of measurement error in macroeconomic variables, but using different data sets, see Croushore and Stark (2001) for example. Comparing the filtered estimates of the persistent component in output growth to the observed outcomes, our analysis reveals that the forecasters were surprised by the strong GDP growth during most of the 1990s. The forecasters were also surprised by the low and declining inflation during the 1990s, with the estimated persistent component generally coming in above the realized rate of inflation.

The plan of the paper is as follows. Section 2 presents our model for how forecasters update their predictions as the forecast horizon shrinks. Section 3 presents empirical results using data from Consensus Economics over the period 1990-2004 and Section 4 concludes. Proofs are contained in an appendix.

2 Multi-horizon Forecast Errors

We start by developing a model for how forecasters update their beliefs about macroeconomic variables such as output growth and inflation. Our analysis makes use of

the rich information available in high frequency revisions of forecasts of a variable observed at a lower frequency, e.g., monthly revisions to forecasts of annual inflation. Since we shall be concerned with flow variables that agents gradually learn about as new information arrives prior to and during the period of their measurement, the fact that part of the outcome may be known prior to the end of the measurement period (the “event date”) introduces additional complications, and means that the timing of the forecasts has to be carefully considered.

Our analysis assumes that agents have a squared loss function over the forecast error, $e_{t,t-h} \equiv z_t - \hat{z}_{t,t-h}$, where z_t is the predicted variable, $\hat{z}_{t,t-h}$ is the forecast computed at time $t - h$, t is the event date and h is the forecast horizon. Other loss functions have been discussed by, e.g., Patton and Timmermann (2007). One advantage of assuming squared loss is that it is easier to justify focusing on aggregate or consensus forecasts, as we shall be doing here, computed as an average of the individual forecasts. Under this loss function, the optimal h -period forecast is simply the conditional expectation of z_t given information at time $t - h$, \mathcal{F}_{t-h} :

$$\hat{z}_{t,t-h}^* = E[z_t | \mathcal{F}_{t-h}]. \quad (1)$$

We test the assumption of forecast rationality under squared loss empirically in Section 3.1 and find that it cannot be rejected.

To study agents’ learning process we keep the event date, t , fixed and vary the forecast horizon, h . As illustrated in Figure 1, this allows us to track how agents update their beliefs through time.

2.1 A Benchmark Model

Since the predicted variable in our application is measured less frequently than the forecasts are revised, it is convenient to describe the target variable as a rolling sum of a higher-frequency variable. To this end, let y_t denote the single-period variable (e.g., monthly log-first differences of GDP or a log-price index tracking inflation), while the rolling sum of the 12 most recent single-period observations of y is denoted

z_t :

$$z_t = \sum_{j=0}^{11} y_{t-j}. \quad (2)$$

Our model is based on a decomposition of y_t into a persistent (and thus predictable) first-order autoregressive component, x_t , and a temporary component, u_t :

$$y_t = x_t + u_t \quad (3)$$

$$x_t = \phi x_{t-1} + \varepsilon_t, \quad -1 < \phi < 1$$

$$u_t \sim iid(0, \sigma_u^2), \varepsilon_t \sim iid(0, \sigma_\varepsilon^2), E[u_t \varepsilon_s] = 0 \forall t, s.$$

Here ϕ measures the persistence of x_t , while u_t and ε_t are innovations assumed to be both serially uncorrelated and mutually uncorrelated. Setting y_t to be a combination of an AR(1) process and an unpredictable process implies that y_t follows an ARMA(1,1), see Granger and Newbold (1986) for example. Without loss of generality, we assume that the unconditional mean of x_t , and thus y_t and z_t , is zero.

Our use of a variable tracking monthly changes in GDP (y_t) is simply a modelling device: US GDP figures are currently only available quarterly. Economic forecasters, however, can almost certainly be assumed to employ higher frequency data when constructing their monthly forecasts of GDP. Giannone, et al. (2008), for example, propose methods to incorporate into macroeconomic forecasts news about the economy between formal announcement dates. When we take our model to data we focus on those aspects of the model that have empirical counterparts.

The assumption that the predicted variable contains a first-order autoregressive component, while clearly an approximation, is able to capture the presence of a persistent component in most macroeconomic data. For example, much of the dynamics in the common factors extracted from large cross-sections of macroeconomic variables by Stock and Watson (2002a) is captured by low-order autoregressive terms. It is straight-forward to allow more lags or other observed variables to enter in the forecasting model, although the latter approach is complicated by the existence of literally hundreds of economic state variables that could be adopted in such mod-

els, (Stock and Watson (2002b, 2006)), “real time” revisions to such data (Diebold and Rudebusch (1991)) and uncertainty about which models agents actually use (Garratt et al. (2003)).

We first present results under simple, but unrealistic, assumptions about the forecasters’ information set in order to reveal some basic properties of the problem. We introduce more realistic assumptions in the next section. Under the assumption that both x_t and y_t are observed at time t , the simplicity of our benchmark model allows an analytic characterization of how the mean squared forecast error (MSE) evolves as a function of the forecast horizon (h):

Proposition 1 *Suppose that y_t can be decomposed into a persistent component (x_t) and a temporary component (u_t) satisfying equation (3) and forecasters minimize the squared loss given the information set $\mathcal{F}_t = \sigma([x_{t-j}, y_{t-j}], j = 0, 1, 2, \dots)$. Then the mean squared forecast error as a function of the forecast horizon is given by:*

$$E[e_{t,t-h}^2] = \begin{cases} 12\sigma_u^2 + \frac{1}{(1-\phi)^2} \left(12 - 2\frac{\phi(1-\phi^{12})}{1-\phi} + \frac{\phi^2(1-\phi^{24})}{1-\phi^2} \right) \sigma_\varepsilon^2 \sigma_\varepsilon^2 \\ \quad + \frac{\phi^2(1-\phi^{12})^2(1-\phi^{2h-24})}{(1-\phi)^3(1+\phi)}, & \text{for } h \geq 12 \\ h\sigma_u^2 + \frac{1}{(1-\phi)^2} \left(h - 2\frac{\phi(1-\phi^h)}{1-\phi} + \frac{\phi^2(1-\phi^{2h})}{1-\phi^2} \right) \sigma_\varepsilon^2, & \text{for } h < 12 \end{cases}$$

Proposition 1 is proved in the Appendix and is simple to interpret: The first term in the expression for the mean squared error (MSE) captures the unpredictable component, u_t . The second term captures uncertainty about shocks to the remaining values of the persistent component, x_t , over the measurement period. The additional term in the expression for $h \geq 12$ comes from having to predict x_{t-11} , the initial value of the persistent component at the beginning of the measurement period.

As $h \rightarrow \infty$, the optimal forecast converges to the unconditional mean of z_t (normalized to zero in our model). This forecast generates the upper bound for the

MSE of an optimal forecast, which is the unconditional variance of z_t :

$$\lim_{h \rightarrow \infty} E [e_{t,t-h}^2] = 12\sigma_u^2 + \frac{\sigma_\varepsilon^2}{(1-\phi)^2} \left(12 - 2\frac{\phi(1-\phi^{12})}{1-\phi} + \frac{\phi^2(1-\phi^{24}) + \phi^2(1-\phi^{12})^2}{1-\phi^2} \right). \quad (4)$$

To illustrate Proposition 1, Figure 2 plots the root mean squared error (RMSE) for $h = 1, 2, \dots, 24$ using parameters similar to those we obtain in the empirical analysis for U.S. GDP growth. Holding the unconditional variance of annual GDP growth, σ_z^2 , and the ratio of the transitory component variance to the persistent component variance, σ_u^2/σ_x^2 , fixed we show the impact of varying the persistence parameter, ϕ . The figure shows the large impact that this parameter has on the shape of the RMSE function. The variance of the forecast error grows linearly, for $h < 12$, as a function of the length of the forecast horizon if y has no persistent component ($\phi = 0$). Conversely, the presence of a persistent component gives rise to a more gradual decline in the forecast error variance as the horizon is reduced. In effect uncertainty is resolved more gradually, the higher the value of ϕ . Notice also how the change in RMSE gets smaller at the longest horizons, irrespective of the value of ϕ .

2.2 Measurement Errors

Proposition 1 is helpful in establishing intuition for the drivers of how macroeconomic uncertainty gets resolved through time. However, it also has some significant shortcomings. Most obviously, it assumes that forecasters observe both the predicted variable, y , and its persistent component, x , without error, and so uncertainty vanishes completely as $h \rightarrow 0$. Macroeconomic variables are, however, to varying degrees, subject to measurement errors as reflected in data revisions and changes in benchmark weights. Such errors are less important for survey-based inflation measures such as the consumer price index (CPI). Revisions are, however, very common for measures of output, such as GDP (see Croushore and Stark (2001), Mahadeva and Muscatelli (2005) and Croushore (2006) for example).

Measurement errors make the forecasters' problem more difficult and introduces a signal extraction problem: the greater the measurement error, the noisier are past observations of y and hence the less precise the forecasters' readings of the state of the economy. They also mean that forecasters cannot simply "plug in" observed values of past y 's during the measurement period ($h < 12$): these quantities must also be estimated.

To account for these effects, we cast our original model in state-space form with a state equation

$$\begin{aligned} \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} y_t \\ x_t \end{bmatrix} &= \begin{bmatrix} 0 & 0 \\ 0 & \phi \end{bmatrix} \begin{bmatrix} y_{t-1} \\ x_{t-1} \end{bmatrix} + \begin{bmatrix} u_t \\ \varepsilon_t \end{bmatrix} \\ \begin{bmatrix} u_t \\ \varepsilon_t \end{bmatrix} &\sim iid \left(0, \begin{bmatrix} \sigma_u^2 & 0 \\ 0 & \sigma_\varepsilon^2 \end{bmatrix} \right). \end{aligned} \tag{5}$$

That is, we leave the data generating process as it was in the benchmark model. Next we assume that agents only observe y_t with error, and that x_t is unobserved. This setup is far more realistic for economic data which are often subject to measurement error and whose persistent components are not directly observable. The measurement equation for this system then becomes:

$$\tilde{y}_t = y_t + \eta_t, \quad \eta_t \sim iid(0, \sigma_\eta^2). \tag{6}$$

Faust et al. (2005) find that revisions to U.S. GDP figures are essentially unpredictable, motivating the simple iid noise structure used above. Since we are focusing on the aggregate, or consensus, forecast, we shall not be concerned with heterogeneity across individual forecasters' information sets. See *Lahiri and Sheng (2008) and Patton and Timmermann (2008)* for an analysis that explicitly focuses on modeling cross-sectional dispersion in forecasts.

Despite its simplicity, this model does not yield a formula for the term structure of RMSE-values that is readily interpretable. The key difficulty that arises is best illustrated by considering "current-year" forecasts ($1 \leq h < 12$). When producing

a current-year forecast at time $t - h$, economic agents must use past and current information to “backcast” realizations $y_{t-11}, \dots, y_{t-h-1}$; they must also produce a “nowcast” for the current month y_{t-h} , and, finally, must predict future realizations, y_{t-h+1}, \dots, y_t . When the persistent component, x_t , is not observable, the resulting forecast errors will generally be serially correlated even after conditioning on all information that is available to the agents. For example, a large positive realization of the measurement error, η_{t-h} , will not only lead to overly optimistic projections for current and future values of y , but will increase the entire sequence of backcast values. Handling this problem is difficult and requires expressing past, current and future forecast errors in terms of the primitive shocks, u_t, ε_t and η_t , which are serially uncorrelated. We next show how to accomplish this.

We assume that our forecasters know the form and parameters of the data generating process, presented in equations (5) and (6), and we further assume that they use the Kalman filter to optimally predict (forecast, nowcast and backcast) the values of y_t needed for the forecast of the annual variable, z_t . Thus the learning problem faced by the forecasters in our model relates to the latent state of the economy (measured by x_t and y_t), but not to the parameters of the model. This simplification is necessitated by our short time series of data.

Given our focus on z_t in equation (2), it is convenient to extend the state vector presented in equation (5) to also include 11 lags of y_t , and define:

$$\begin{aligned} \boldsymbol{\xi}_t &\equiv [x_t, y_t, y_{t-1}, \dots, y_{t-11}]' \\ \mathbf{F} &\equiv \begin{bmatrix} \phi & \mathbf{0}'_{12} \\ \phi & \mathbf{0}'_{12} \\ \mathbf{0}_{11} & I_{11} & \mathbf{0}_{11} \end{bmatrix} \end{aligned} \tag{7}$$

where now

$$z_t = \boldsymbol{\omega}' \boldsymbol{\xi}_t,$$

and $\boldsymbol{\omega} = [0, \boldsymbol{\iota}'_{12}]'$. Here $\mathbf{0}_k$ is a $k \times 1$ vector of zeros, $\boldsymbol{\iota}_p$ is a $p \times 1$ vector of ones, and

I_k is a $k \times k$ identity matrix. The state equation is:

$$\boldsymbol{\xi}_t = \mathbf{F}\boldsymbol{\xi}_{t-1} + \mathbf{v}_t, \quad (8)$$

where $\boldsymbol{\nu}_t \equiv [\varepsilon_t, \varepsilon_t + u_t, \mathbf{0}'_{11}]'$, and the measurement equation is

$$\tilde{\mathbf{y}}_t = \mathbf{H}'\boldsymbol{\xi}_t + \mathbf{w}_t. \quad (9)$$

In our application the measurement variable is a scalar, $\tilde{y}_t = y_t + \eta_t$, and so $\mathbf{H} = [0, 1, \mathbf{0}'_{11}]'$, but we will present our theoretical framework for the general case that $\tilde{\mathbf{y}}_t$ is a vector. The innovations to the state and measurement equations are:

$$\begin{aligned} \boldsymbol{\nu}_t &\sim iid N(0, \mathbf{Q}) \\ \mathbf{Q} &= \begin{bmatrix} \sigma_\varepsilon^2 & \sigma_\varepsilon^2 & \mathbf{0}'_{11} \\ \sigma_\varepsilon^2 & \sigma_u^2 + \sigma_\varepsilon^2 & \mathbf{0}'_{11} \\ \mathbf{0}_{11} & \mathbf{0}_{11} & \mathbf{0}_{11 \times 11} \end{bmatrix} \\ \mathbf{w}_t &\sim iid N(0, \mathbf{R}), \end{aligned} \quad (10)$$

where $\mathbf{0}_{k \times p}$ is a $k \times p$ matrix of zeros. In our application $\mathbf{R} = \sigma_\eta^2$. Further, we assume

$$E[\mathbf{v}_t \mathbf{w}'_s] = 0 \quad \forall s, t. \quad (11)$$

We also assume that the forecaster has been using the Kalman filter long enough that all updating matrices are at their steady-state values. This is done simply to remove any “start of sample” effects that may or may not be present in the data.

Let:

$$\begin{aligned} \tilde{\mathcal{F}}_t &= \sigma(\tilde{\mathbf{y}}_t, \tilde{\mathbf{y}}_{t-1}, \dots, \tilde{\mathbf{y}}_1) \\ \hat{\boldsymbol{\xi}}_{t|t-1} &\equiv E[\boldsymbol{\xi}_t | \tilde{\mathcal{F}}_{t-1}] \equiv E_{t-1}[\boldsymbol{\xi}_t] \\ \hat{\mathbf{y}}_{t|t-1} &\equiv E[\tilde{\mathbf{y}}_t | \tilde{\mathcal{F}}_{t-1}] \equiv E_{t-1}[\tilde{\mathbf{y}}_t]. \end{aligned}$$

Following Hamilton (1994), define the following matrices

$$\begin{aligned}
\mathbf{P}_{t+1|t} &\equiv E \left[\left(\boldsymbol{\xi}_{t+1} - \hat{\boldsymbol{\xi}}_{t+1|t} \right) \left(\boldsymbol{\xi}_{t+1} - \hat{\boldsymbol{\xi}}_{t+1|t} \right)' \right] \\
&= (\mathbf{F} - \mathbf{K}_t \mathbf{H}') \mathbf{P}_{t|t-1} (\mathbf{F}' - \mathbf{H} \mathbf{K}_t') + \mathbf{K}_t \mathbf{R} \mathbf{K}_t' + \mathbf{Q} \rightarrow \mathbf{P}_1^* \\
\mathbf{K}_t &\equiv \mathbf{F} \mathbf{P}_{t|t-1} \mathbf{H} (\mathbf{H}' \mathbf{P}_{t|t-1} \mathbf{H} + \mathbf{R})^{-1} \rightarrow \mathbf{K}^* \\
\mathbf{P}_{t|t} &\equiv E \left[\left(\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t} \right) \left(\boldsymbol{\xi}_t - \hat{\boldsymbol{\xi}}_{t|t} \right)' \right] \\
&= \mathbf{P}_{t|t-1} - \mathbf{P}_{t|t-1} \mathbf{H} (\mathbf{H}' \mathbf{P}_{t|t-1} \mathbf{H} + \mathbf{R})^{-1} \mathbf{H}' \mathbf{P}_{t|t-1} \\
&\rightarrow \mathbf{P}_1^* - \mathbf{P}_1^* \mathbf{H} (\mathbf{H}' \mathbf{P}_1^* \mathbf{H} + \mathbf{R})^{-1} \mathbf{H}' \mathbf{P}_1^* \equiv \mathbf{P}_0^*.
\end{aligned} \tag{12}$$

The convergence of $\mathbf{P}_{t|t-1}$, $\mathbf{P}_{t|t}$ and \mathbf{K}_t to their steady-state values relies on $|\phi| < 1$, and we impose this in the estimation. To initialize these matrices we use their unconditional equivalents, $\mathbf{P}_{1|0} \equiv E \left[(\boldsymbol{\xi}_t - E[\boldsymbol{\xi}_t]) (\boldsymbol{\xi}_t - E[\boldsymbol{\xi}_t])' \right]$ and $\hat{\boldsymbol{\xi}}_{1|0} = E[\boldsymbol{\xi}_t]$. Estimates of the state variables are updated via

$$\hat{\boldsymbol{\xi}}_{t|t} = \hat{\boldsymbol{\xi}}_{t|t-1} + \mathbf{P}_{t|t-1} \mathbf{H} (\mathbf{H}' \mathbf{P}_{t|t-1} \mathbf{H} + \mathbf{R})^{-1} \left(\tilde{\mathbf{y}}_t - \mathbf{H}' \hat{\boldsymbol{\xi}}_{t|t-1} \right), \tag{13}$$

while the multi-step prediction error uses

$$\begin{aligned}
\hat{\boldsymbol{\xi}}_{t+s|t} &= \mathbf{F}^s \hat{\boldsymbol{\xi}}_{t|t} \\
\mathbf{P}_{t+s|t} &\equiv E \left[\left(\boldsymbol{\xi}_{t+s} - \hat{\boldsymbol{\xi}}_{t+s|t} \right) \left(\boldsymbol{\xi}_{t+s} - \hat{\boldsymbol{\xi}}_{t+s|t} \right)' \right] \\
&= \mathbf{F}^s \mathbf{P}_{t|t} (\mathbf{F}')^s + \sum_{j=0}^{s-1} \mathbf{F}^j \mathbf{Q} (\mathbf{F}')^j \rightarrow \mathbf{P}_s^*, \text{ for } s \geq 1.
\end{aligned} \tag{14}$$

The full set of MSE-values across different horizons can now be extracted from \mathbf{P}_s^* :

$$\begin{aligned}
\hat{z}_{t,t-h}^* &\equiv E \left[z_t | \tilde{\mathcal{F}}_{t-h} \right] = \boldsymbol{\omega}' \hat{\boldsymbol{\xi}}_{t|t-h} \\
\text{yielding } E \left[(z_t - \hat{z}_{t,t-h}^*)^2 \right] &= \boldsymbol{\omega}' \mathbf{P}_h^* \boldsymbol{\omega}, \text{ for } h \geq 0.
\end{aligned} \tag{15}$$

Note that for $h < 12$ the optimal forecast $\hat{z}_{t,t-h}^*$ will involve a combination of forecasts, $E \left[y_{t-h+j} | \tilde{\mathcal{F}}_{t-h} \right]$ for $j > 0$, nowcasts, $E \left[y_{t-h} | \tilde{\mathcal{F}}_{t-h} \right]$, and backcasts, $E \left[y_{t-h-k} | \tilde{\mathcal{F}}_{t-h} \right]$ for $k > 0$. Our use of an extended state equation means that these terms are all captured in the above expressions without having to handle them separately.

Figure 3 uses these equations to illustrate the impact of measurement error on the RMSE-values at different horizons. The degree of measurement error is described as $\sigma_\eta^2 = k^2\sigma_u^2$, so k measures the size of the measurement error in terms of the innovation variance for y . In the absence of measurement errors the RMSE will converge to zero as $h \rightarrow 0$, whereas in the presence of measurement error the RMSE will converge to some positive quantity. As the horizon, h , shrinks towards zero, the relative importance of measurement errors grows. Moreover, the RMSE function gets flatter as the size of the measurement error increases. Note, however, that measurement error plays no part for long-horizon forecasts, since its impact on overall uncertainty is small relative to other sources of uncertainty, and so Figure 3 resembles Figure 2 for long horizon forecasts. This also shows that the persistence (ϕ) and measurement error (σ_η^2) parameters are separately identified by jointly considering long and short horizon forecast errors, and illustrates the rich information contained in survey forecasts covering multiple forecast horizons.

The analytical results in this section show that a simple model of the forecasting environment faced by macroeconomic forecasters in practice can accommodate a rich set of empirical phenomena: with just four free parameters a variety of RMSE patterns is obtained. Further, by studying such a model in detail we gain some quantitative insight into the key drivers of macroeconomic forecast errors. We next move on to matching the parameters of our model to data.

2.3 Estimation

Our strategy for estimation is to choose the parameters that enable the model to match the observed forecast errors as closely as possible. To this end, we estimate the parameters using the Generalized Methods of Moments (GMM), see Hansen (1982), based on the moment conditions obtained by matching the sample MSE, $T^{-1} \sum_{t=1}^T e_{t,t-h}^2$ at various forecast horizons to the population mean squared errors,

$MSE_h(\boldsymbol{\theta})$, implied by our model. Our parameter estimates are obtained from:

$$\hat{\boldsymbol{\theta}}_T \equiv \arg \min_{\boldsymbol{\theta}} \mathbf{g}_T(\boldsymbol{\theta})' \mathbf{W}_T \mathbf{g}_T(\boldsymbol{\theta}) \quad (16)$$

$$\mathbf{g}_T(\boldsymbol{\theta}) \equiv \frac{1}{T} \sum_{t=1}^T \begin{bmatrix} e_{t,t-1}^2 - MSE_1(\boldsymbol{\theta}) \\ e_{t,t-2}^2 - MSE_2(\boldsymbol{\theta}) \\ \vdots \\ e_{t,t-H}^2 - MSE_H(\boldsymbol{\theta}) \end{bmatrix} \quad (17)$$

where $\boldsymbol{\theta} \equiv [\sigma_u^2, \sigma_\varepsilon^2, \sigma_\eta^2, \phi]'$ and $MSE_h(\boldsymbol{\theta})$ is obtained using Proposition 1 or the updating equations leading to (15).

It is also possible to estimate our model using maximum likelihood, though doing so requires making stronger assumptions on the distribution of the various innovation terms in our model, making it more susceptible to model mis-specification. We opt to use GMM so that the parameters we obtain are interpretable as those that yield a predicted term structure of RMSE that most closely matches what we observe in the data.

In situations with large H there are several over-identifying restrictions, and so the choice of weighting matrix, \mathbf{W}_T , in the GMM estimation is important. We use the identity matrix as the weighting matrix so that all horizons get equal weight in the estimation procedure; this is not fully efficient, but is justified by our focus on modeling the entire term structure of forecast errors. Nevertheless, we still require the covariance matrix of the sample moments to compute standard errors and a test of the over-identifying conditions. As we shall see, our sample is only $T = 14$ years long while we have $H = 24$ forecast horizons, and so it is not feasible to estimate this matrix directly from the data since this would require controlling for the correlation between the sample moments induced by overlaps across the 24 horizons. Fortunately, given the simple structure of our model, for a given parameter value we can compute a model-implied covariance matrix of the sample moments. Under the assumption that the model is correctly specified, this matrix captures the correlation between sample moments induced by overlaps and serial persistence.

To obtain \mathbf{P}_1^* , \mathbf{P}_0^* , and \mathbf{K}^* we simulate 100 non-overlapping years of data and update $\mathbf{P}_{t|t-1}$, $\mathbf{P}_{t|t}$ and \mathbf{K}_t following Hamilton (1994). We use these matrices at the end of the 100th year as estimates of \mathbf{P}_1^* , \mathbf{P}_0^* , and \mathbf{K}^* . To obtain the covariance matrix of the moments, used to compute standard errors and the test of over-identifying restrictions, we use the model-implied covariance matrix of the moments, based on the parameter estimate from the first-stage GMM parameter estimate. This matrix is not available in closed-form and so we simulate 1,000 non-overlapping years of data to estimate it, imposing that the innovations to these processes ($\boldsymbol{\nu}_t$ and \mathbf{w}_t) are Normally distributed, and using the expressions given above to obtain the Kalman filter forecasts.

We use only six forecast horizons ($h = 1, 3, 6, 12, 18, 24$) in the estimation, rather than the full set of 24, in response to studies of the finite-sample properties of GMM estimates (see Tauchen, 1986, for example) which find that using many more moment conditions than required for identification leads to poor approximations from the asymptotic theory, particularly when the moments are highly correlated, as in our application. We have also estimated the models presented in this paper using the full set of 24 moment conditions and the results were qualitatively similar.

3 Empirical Results

After a brief introduction to the data, we present estimation results both for the simple model that ignores measurement errors and for the extended model that accounts for such errors. Finally, we use our estimates to shed light on the persistent and transitory components of GDP growth and inflation implied by the forecasts.

3.1 Data

Our data is taken from the Consensus Economics Inc. forecasts which comprise polls of more than 600 private sector forecasters and are widely considered by organizations such as the IMF and the U.S. Federal Reserve. Each month participants

are asked about their views of a range of variables for the major economies and the consensus (average) forecast is recorded. Our analysis focuses on US real GDP growth and Consumer Price Index (CPI) inflation for the current and subsequent year. This gives us 24 monthly next-year and current-year forecasts over the period 1991-2004 or a total of $24 \times 14 = 336$ monthly observations. Naturally our observations are not independent draws but are subject to a set of tight restrictions across horizons, as revealed by the analysis in the previous section. We use data from the IMF [*ALLAN - CAN YOU CHECK WHETHER THIS IS EQUIVALENT TO BLS DATA?*] to measure the realized value of the target variable (GDP growth or inflation), and we follow Romer and Romer (2000) and use second release of this data. Our results are very similar when the first release is used instead as recommended by Corradi, Fernandez and Swanson (2007).

Our analysis takes the target variable, z_t , as the December-on-December change in the log-level of US real GDP or the Consumer Price Index, which can of course be written as the sum of the month-on-month changes in the log-levels of these series, denoted y_t , as in equation (2). The Consensus Economics survey formally defines the target variable slightly differently to this but the impact of this difference on our results is negligible.

As a prelude to our analysis of the RMSE function, we initially undertook a range of statistical tests that check for biases and serial correlation in the forecast errors. We tested for biases in the forecasts by testing whether the forecast errors were mean zero and by estimating “Mincer-Zarnowitz” (1969) regressions

$$y_t = \beta_0^h + \beta_1^h \hat{y}_{t,t-h} + \epsilon_{t,t-h} \quad (18)$$

and testing that $\beta_0^h = 0, \beta_1^h = 1$ for $h = 1, \dots, 24$, which is requirement for unbiased forecasts. The results are presented in Table 1. For GDP growth, there was no evidence of a bias and none of the Mincer-Zarnowitz F -tests rejected the null. For inflation, the test for bias did reject for some horizons, whereas the Mincer-Zarnowitz F -test did not reject the optimality of these forecasts. Hence we shall proceed to

estimate the parameters of our model under the assumption that forecasters use information efficiently.

3.2 Parameter Estimates and Tests

The simple benchmark model contains just three free parameters, namely the variance of the innovations in the temporary (σ_u^2) and persistent (σ_ε^2) components, and the persistence parameter, ϕ , for the predictable component. The expressions for the MSE as a function of h , stated in Proposition 1 for the benchmark model and in equation (15) for the model that allows for measurement error, enable us to use GMM to estimate the unknown parameters given a panel of forecast errors measured at various horizons. These parameters are not separately identifiable if forecasts for a single horizon are all that is available so access to multi-horizon forecasts is crucial to our analysis. Since the variance of the h -period forecast error grows linearly in σ_u^2 while σ_ε^2 and ϕ generally affect the MSE in a non-linear fashion, these parameters can be identified from a sequence of MSE-values corresponding to different forecast horizons, h , provided at least three different horizons are available.

Figure 4 plots the RMSE-values for output growth and inflation at the 24 different horizons. In the case of output growth the RMSE shrinks from about 1.8% at the 24-month horizon to 1% at the 12-month horizon and 0.5% at the 1-month horizon. For inflation it ranges from 0.8% at the two-year horizon to 0.4% at the 12-month horizon and less than 0.1% at the 1-month horizon. Forecast precision improves systematically as the forecast horizon is reduced, as expected. Moreover, consistent with Proposition 1, the rate at which the RMSE declines is smaller in the next-year forecasts ($h \geq 12$) than in current-year forecasts ($h < 12$).

The fitted values from the model without measurement error, also shown in Figure 4, clearly illustrate the limitation of this specification. This model assumes that forecasters get a very precise reading of the outcome towards the end of the current year and hence forces the fitted estimate of the RMSE to decline sharply at

short forecast horizons. This property is clearly at odds with the GDP growth data and means that the benchmark model without measurement error does not succeed in simultaneously capturing the behavior of the RMSE at both the short and long horizons. For inflation forecasts the assumption of zero measurement error appears consistent with the data. Table 2 presents the parameter estimates and provides a formal test of this model. Unsurprisingly, in view of Figure 4, the model is strongly rejected for GDP growth but not for inflation.

The model extended to allow for measurement errors introduces an extra parameter, σ_η^2 , which reflects the magnitude of measurement errors. Even though both σ_η^2 and σ_u^2 are well-identified in theory, in practice they are difficult to estimate separately. We therefore set σ_η to be proportional to σ_u : $\sigma_\eta = k \cdot \sigma_u$ and estimate the model for $k = \{0.01, 0.25, 0.5, 1, 2, 3, 4, 5, 10\}$. The goodness-of-fit of the model (as measured by Hansen’s (1982) J -test of over-identifying restrictions) is generally robust for $1 \leq k \leq 4$. We set $k = 2$ in the estimation.

Table 3 presents parameter estimates for the model with measurement errors fitted to the consensus forecasts. This model passes the specification tests for both GDP growth and inflation and thus there is little statistical evidence against our simple specification even for output growth, once measurement errors are considered. Table 3 reveals that the predictable component of inflation is slightly more persistent than that in output growth.

Figure 4 shows that the specification with measurement errors does a much better job at matching the decay pattern in observed RMSE for US output growth as the forecast horizon, h , shrinks to zero. In the case of US inflation, the models with and without measurement error are identical, as the best estimate of the variance of the measurement error is zero. This is consistent with Croushore and Stark (2001) who report that revisions in reported GDP figures tend to be larger than those in reported inflation figures, *and with Giannone et al. (2008) who note that “nowcasting” GDP in real time is a difficult statistical task, whereas it is less so for inflation as reliable estimates of inflation are available at a monthly frequency.*

Our study of the term structure of RMSE for this survey of professional forecasters thus confirms several empirical findings in the existing literature that were obtained using very different data sets: (1) Professional forecasters face severe measurement error problems for GDP growth in real time, while this is not an issue for inflation; (2) inflation exhibits greater persistence, and thus is predictable at longer horizons, than GDP growth; and (3) the persistent component of these variables is well-approximated by a low-order autoregressive specification.

3.3 Components of GDP Growth and Inflation

Our model for the multi-horizon patterns in forecast errors is based on a decomposition of the target variable, GDP growth or inflation, into a persistent component, x_t , and an unpredictable component, u_t . The GMM estimation procedure used above does not require the estimation of the sample paths for x_t and u_t , but with the estimated parameter vector and the panel of forecasts we are able to infer the forecasters' estimated values of these variables. Without a model for the persistent and transitory components of these series it would be impossible to extract estimates of the forecasters' beliefs about the persistent components of GDP growth and inflation. We can use the long horizon forecasts ($h \geq 12$) to infer the forecasters' estimate of the persistent component, and the short horizon ($h < 12$) forecasts to infer the forecasters' estimate of the unpredictable component. Intuitively, one can think of our estimates of these two components as an alternative representation of the two forecasts the forecasters make at each point in time (the "next year", $h \geq 12$, and the "current year", $h < 12$, forecasts). We can obtain both of these components without needing to make any further identifying assumptions, and without needing to employ any data other than the collection of forecasts. The most economically interesting piece is the persistent component and below we focus on that.

To show how we derive estimates of the persistent components, it is convenient to consider a smaller state variable, $\ddot{\xi}_t \equiv [y_t, x_t]'$. For concreteness, consider estimates

based on the first row in our panel, so the annual target variable is z_{25} , and the first forecast is $\hat{z}_{25,1}$. Let

$$\boldsymbol{\chi}_{25} \equiv \sum_{j=0}^{11} \ddot{\boldsymbol{\xi}}_{25-j} = \begin{bmatrix} \sum_{j=0}^{11} y_{25-j} \\ \sum_{j=0}^{11} x_{25-j} \end{bmatrix}$$

then

$$\hat{\boldsymbol{\chi}}_{25,1} = \sum_{j=0}^{11} \hat{\ddot{\boldsymbol{\xi}}}_{25-j,1} = \left(\sum_{j=0}^{11} \mathbf{F}^{24-j} \right) \hat{\ddot{\boldsymbol{\xi}}}_{1,1} \equiv \mathbf{F}^{13} \mathbf{F}^{(11)} \hat{\ddot{\boldsymbol{\xi}}}_{1,1},$$

where $\mathbf{F}^{(k)} \equiv \sum_{j=0}^k \mathbf{F}^j$. Thus

$$\begin{aligned} \hat{z}_{25,1} &= \mathbf{e}'_1 \hat{\boldsymbol{\chi}}_{25,1} \\ &= \mathbf{e}'_1 \mathbf{F}^{13} \mathbf{F}^{(11)} \hat{\ddot{\boldsymbol{\xi}}}_{1,1} \\ &= \frac{\phi^{13} (1 - \phi^{12})}{1 - \phi} E \left[x_1 | \tilde{\mathcal{F}}_1 \right], \\ \text{so } E \left[x_1 | \tilde{\mathcal{F}}_1 \right] &= \frac{1 - \phi}{\phi^{13} (1 - \phi^{12})} \cdot \hat{z}_{25,1}, \end{aligned}$$

where $\mathbf{e}_1 \equiv [1, 0]'$. Since the 24-month forecast is proportional to the “nowcast” of the predictable component, with the proportionality constant being a simple function of the parameter of the data generating process, we can back out the forecaster’s “nowcast” of the predictable component from the forecast. This same step holds for all “long horizon” forecasts:

$$\begin{aligned} \hat{\boldsymbol{\chi}}_{25,25-h} &= \mathbf{F}^{h-11} \mathbf{F}^{(11)} \hat{\ddot{\boldsymbol{\xi}}}_{25-h,25-h}, \text{ for } h \geq 12 \\ \text{and so } E \left[x_{25-h} | \tilde{\mathcal{F}}_{25-h} \right] &= \frac{1 - \phi}{\phi^{h-11} (1 - \phi^{12})} \cdot \hat{z}_{25,25-h}, \text{ for } h \geq 12. \end{aligned}$$

Thus using the long-horizon forecasts we can extract the filtered estimate of the predictable component of the target variable. This is, of course, available monthly, which is more frequently than data is available on GDP growth, although some inflation series are available monthly.

Our model above assumed, without loss of generality for the study of RMSE, that all variables have zero mean. This of course is not true in reality, and does

have implications for our estimates of x_t . The expressions derived above can be re-interpreted as expressions for $E \left[x_t - \mu | \tilde{\mathcal{F}}_t \right]$ and $E \left[y_t - \mu | \tilde{\mathcal{F}}_t \right]$ when $\mu \neq 0$. Modifying the subsequent calculations the forecasts become:

$$\chi_{25} \equiv \sum_{j=0}^{11} \xi_{1,1,25-j} = \left[\begin{array}{c} \sum_{j=0}^{11} (y_{25-j} - \mu) \\ \sum_{j=0}^{11} (x_{25-j} - \mu) \end{array} \right] = \left[\begin{array}{c} \sum_{j=0}^{11} y_{25-j} \\ \sum_{j=0}^{11} x_{25-j} \end{array} \right] - 12\mu.$$

Thus we can simply de-mean the observed forecasts (using, for example, one-twelfth the sample mean of the z_t series), compute $E \left[x_t | \tilde{\mathcal{F}}_t \right]$ and $E \left[y_t | \tilde{\mathcal{F}}_t \right]$ as above, and then add back the means to the estimates. This corresponds to estimating the parameter μ by GMM, using simply the sample mean of the z_t series.

In Figure 5 we present for each month in our sample the estimated persistent components of GDP growth and inflation, as implied by the observed consensus forecasts and the parameters of our model. For reference we plot both the “filtered” estimates, which are estimates of $E \left[x_t | \tilde{\mathcal{F}}_t \right]$, and the “smoothed” estimates, which are estimates of $E \left[x_t | \tilde{\mathcal{F}}_T \right]$. The estimates for GDP growth reveal that the forecasters in the survey estimated the level of GDP growth in the early 1990s quite well, but were surprised by the strong GDP growth in the mid to late 1990s: the estimated persistent component of GDP growth hovered around 1.5% annualized, whereas the actual GDP growth in that period was closer to 4%. Since the 2001 recession the persistent component has consistently been above the realized values of GDP growth.

Similarly, the forecasters in the survey were surprised by the declining inflation of the 1990s, with our estimated persistent component generally coming out above the realized values of inflation. In the latter part of the sample the estimated persistent component is more in line with realized inflation, consistent with the view that forecasters took some time to adjust their views on long-run inflation in the US.

4 Conclusion

This paper studied how macroeconomic uncertainty is resolved over time. To this end we considered survey forecasts of macroeconomic variables which hold the “event” date constant, while reducing the length of the forecast horizon. We proposed a simple, parsimonious unobserved components model and developed tools for estimation and inference based on simulation methods that account for agents’ learning problem. Our methods can be used to estimate the size of measurement errors in the underlying variables and the degree of persistence in the data generating process. They can also be used to extract information on forecasters’ beliefs about the underlying state of the economy and thus to characterize the types of forecast errors that agents made over a given historical sample.

5 Appendix

Proof of Proposition 1. Since $z_t = \sum_{j=0}^{11} y_{t-j}$ and $y_t = x_t + u_t$, where x_t is the persistent component, forecasting z_t given information h months prior to the end of the measurement period, $\mathcal{F}_{t-h} = \{x_{t-h}, y_{t-h}, x_{t-h-1}, y_{t-h-1}, \dots\}$, requires accounting for the persistence in x . Note that

$$\begin{aligned} x_{t-h+1} &= \phi x_{t-h} + \varepsilon_{t-h+1} \\ x_{t-h+2} &= \phi^2 x_{t-h} + \phi \varepsilon_{t-h+1} + \varepsilon_{t-h+2} \\ &\vdots \\ x_t &= \phi^h x_{t-h} + \sum_{j=0}^{h-1} \phi^j \varepsilon_{t-j}. \end{aligned}$$

Adding up these terms we find that, for $h \geq 12$,

$$\begin{aligned} z_t &= \sum_{j=0}^{11} x_{t-j} + \sum_{j=0}^{11} u_{t-j} \\ &= \frac{\phi(1 - \phi^{12})}{1 - \phi} x_{t-12} + \frac{1}{1 - \phi} \sum_{j=0}^{11} (1 - \phi^{12-j}) \varepsilon_{t-12+1+j} + \sum_{j=0}^{11} u_{t-j}. \end{aligned}$$

Thus the optimal forecast for $h \geq 12$ is

$$\hat{z}_{t,t-h}^* \equiv E[z_t | \mathcal{F}_{t-h}] = \sum_{j=0}^{11} E[y_{t-j} | \mathcal{F}_{t-h}] = \sum_{j=0}^{11} E[x_{t-j} | \mathcal{F}_{t-h}] = \sum_{j=0}^{11} \phi^{h-j} x_{t-h},$$

so $\hat{z}_{t,t-h}^* = \frac{\phi^{h-11} (1 - \phi^{12})}{1 - \phi} x_{t-h}$, for $h \geq 12$.

For the current-year forecasts ($h < 12$) the optimal forecast of z_t makes use of those realizations of y that have already been observed. Thus the optimal forecast is:

$$\hat{z}_{t,t-h}^* = \sum_{j=0}^{11} E[y_{t-j} | \mathcal{F}_{t-h}] = \sum_{j=h}^{11} y_{t-j} + \sum_{j=0}^{h-1} E[x_{t-j} | \mathcal{F}_{t-h}] = \sum_{j=h}^{11} y_{t-j} + \sum_{j=0}^{h-1} \phi^{h-j} x_{t-h},$$

so $\hat{z}_{t,t-h}^* = \sum_{j=h}^{11} y_{t-j} + \frac{\phi(1 - \phi^h)}{1 - \phi} x_{t-h}$, for $h < 12$.

Using these expressions for the optimal forecasts we can derive the forecast error, $e_{t,t-h} \equiv z_t - \hat{z}_{t,t-h}^*$, as a function of the forecast horizon. For $h \geq 12$,

$$\begin{aligned} e_{t,t-h} &= \sum_{j=0}^{11} u_{t-j} + \sum_{j=0}^{11} x_{t-j} - \frac{\phi^{h-11} (1 - \phi^{12})}{1 - \phi} x_{t-h} \\ &= \sum_{j=0}^{11} u_{t-j} + \sum_{j=0}^{11} \frac{1 - \phi^{j+1}}{1 - \phi} \varepsilon_{t-j} + \sum_{j=12}^{h-1} \frac{\phi^{j-11} (1 - \phi^{12})}{1 - \phi} \varepsilon_{t-j}. \end{aligned}$$

In computing the variance of $e_{t,t-h}$ we exploit the fact that u and ε are independent of each other at all lags. For $h \geq 12$,

$$\begin{aligned} E[e_{t,t-h}^2] &= \sum_{j=0}^{11} E[u_{t-j}^2] + \sum_{j=0}^{11} \frac{(1 - \phi^{j+1})^2}{(1 - \phi)^2} E[\varepsilon_{t-j}^2] + \sum_{j=12}^{h-1} \frac{\phi^{2j-22} (1 - \phi^{12})^2}{(1 - \phi)^2} E[\varepsilon_{t-j}^2] \\ &= 12\sigma_u^2 + \frac{\sigma_\varepsilon^2}{(1 - \phi)^2} \sum_{j=0}^{11} (1 - \phi^{j+1})^2 + \frac{(1 - \phi^{12})^2}{(1 - \phi)^2} \sigma_\varepsilon^2 \sum_{j=12}^{h-1} \phi^{2j-22} \\ &= 12\sigma_u^2 + \frac{\sigma_\varepsilon^2}{(1 - \phi)^2} \left(12 - 2 \frac{\phi(1 - \phi^{12})}{1 - \phi} + \frac{\phi^2(1 - \phi^{24})}{(1 - \phi^2)} \right) \\ &\quad + \frac{\phi^2(1 - \phi^{12})^2(1 - \phi^{2h-24})}{(1 - \phi)^3(1 + \phi)} \sigma_\varepsilon^2, \end{aligned}$$

as presented in the proposition. For $h < 12$ we have:

$$\begin{aligned}
e_{t,t-h} &= \sum_{j=0}^{11} y_{t-j} - \sum_{j=h}^{11} y_{t-j} - \frac{\phi(1-\phi^h)}{1-\phi} x_{t-h} \\
&= \sum_{j=0}^{h-1} u_{t-j} + \sum_{j=0}^{h-1} \frac{1-\phi^{j+1}}{1-\phi} \varepsilon_{t-j} \\
\text{so } E[e_{t,t-h}^2] &= \sum_{j=0}^{h-1} E[u_{t-j}^2] + \sum_{j=0}^{h-1} \frac{(1-\phi^{j+1})^2}{(1-\phi)^2} E[\varepsilon_{t-j}^2] \\
&= h\sigma_u^2 + \frac{\sigma_\varepsilon^2}{(1-\phi)^2} \left(h - 2\frac{\phi(1-\phi^h)}{1-\phi} + \frac{\phi^2(1-\phi^{2h})}{1-\phi^2} \right).
\end{aligned}$$

■

References

- [1] Bernanke, B.S., 2007, Inflation Expectations and Inflation Forecasting. Remarks by Chairman Ben S. Bernanke, Federal Reserve Board.
- [2] Clements, M.P., 1997, Evaluating the Rationality of Fixed-Event Forecasts. *Journal of Forecasting* 16, 225-239.
- [3] Corradi, V., A. Fernandez and N.R. Swanson, 2007, Information in the Revision Process of Real-time Data. Mimeo, Rutgers.
- [4] Croushore, D., 2006, Forecasting with Real-Time Macroeconomic Data. Pages 961-982 in G. Elliott, C. Granger and A. Timmermann (eds.) *Handbook of Economic Forecasting*, North Holland: Amsterdam.
- [5] Croushore, D. and T. Stark, 2001, A Real-Time Data Set for Macroeconomists. *Journal of Econometrics* 105, 111-130.
- [6] Davies, A. and K. Lahiri, 1995, A New Framework for Analyzing Survey Forecasts using Three-Dimensional Panel Data, *Journal of Econometrics*, 68, 205-227.
- [7] Diebold, F.X. and G.D. Rudebusch, 1991, Forecasting Output with the Composite Leading Index: A Real-Time Analysis, *Journal of the American Statistical Association*, 86, 603-610.
- [8] Dixit, A.X. and R.S. Pindyck, 1994, *Investment under Uncertainty*. Princeton University Press: Princeton N.J.

- [9] Faust, J., J.H. Rogers and J.H. Wright, 2005, News and Noise in G-7 GDP Announcements, *Journal of Money, Credit, and Banking*, 37, 403-417.
- [10] Garratt, A., K. Lee, M. H. Pesaran, and Y. Shin, 2003, Forecast Uncertainties in Macroeconometric Modelling: An Application to the UK Economy. *Journal of the American Statistical Association*, 98, 464, 829-838.
- [11] Giannone, D., L. Reichlin and D. Small, 2008, Nowcasting GDP: The Real Time Informational Content of Macroeconomic Data Releases, *Journal of Monetary Economics*, 55, 665-676.
- [12] Granger, C.W.J. and P. Newbold, 1986, *Forecasting Economic Time Series*, Second Edition, Academic Press.
- [13] Hamilton, J.D., 1994, *Time Series Analysis*, Princeton University Press, Princeton, New Jersey.
- [14] Hansen, L.P., 1982, Large Sample Properties of Generalized Method of Moments Estimators, *Econometrica*, 50, 1029-1054.
- [15] Isiklar, G., K. Lahiri and P. Loungani, 2006, How Quickly do Forecasters Incorporate News? Evidence from Cross-Country Surveys. *Journal of Applied Econometrics* 21, 703-725.
- [16] Kydland, F.E. and E.C. Prescott, 1982, Time to Build and Aggregate Fluctuations. *Econometrica* 50, 1345-1370.
- [17] Lahiri, K. and X. Sheng, 2008, Evolution of Forecast Disagreement in a Bayesian Learning Model, *Journal of Econometrics*, 144, 325-340.
- [18] Mahadeva, L. and A. Muscatelli, 2005, National Accounts Revisions and Output Gap Estimates in a Model of Monetary Policy with Data Uncertainty, External MPC Unit Discussion Paper 14, Bank of England.
- [19] Mincer, J., and V. Zarnowitz, 1969, The Evaluation of Economic Forecasts, in J. Mincer (ed.) *Economic Forecasts and Expectations*, National Bureau of Economic Research, New York.
- [20] Newey, W.K., and K.D. West, 1987, A Simple, Positive Semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica*, 55, 703-708.
- [21] Nordhaus, W.D., 1987, Forecasting Efficiency: Concepts and Applications. *Review of Economics and Statistics* 69, 667-674.
- [22] Patton, A.J., and A. Timmermann, 2007, Testing Forecast Optimality under Unknown Loss. *Journal of the American Statistical Association* 102, 1172-1184.

- [23] Patton, A.J., and A. Timmermann, 2008, Why do Professional Forecasters Disagree? Lessons from the Term Structure of Cross-sectional Dispersion. Unpublished manuscript, Oxford University and UCSD.
- [24] Ramey, G. and V. Ramey, 1995, Cross-Country Evidence on the Links between Volatility and Growth. *American Economic Review* 85, 1138–51.
- [25] Romer, C.D. and D.H. Romer, 2000, Federal Reserve Information and the Behavior of Interest Rates, *American Economic Review*, 90, 429-457.
- [26] Stock, J.H. and M.W. Watson, 2002a, Macroeconomic Forecasting using Diffusion Indexes. *Journal of Business and Economic Statistics* 20, 147-162.
- [27] Stock, J.H. and M.W. Watson, 2002b, Forecasting Using Principal Components from a Large Number of Predictors. *Journal of the American Statistical Association* 97, 1167-1179.
- [28] Stock, J.H. and M.W. Watson, 2006, Forecasting with Many Predictors. Pages 515-554 in G. Elliott, C.W.J. Granger and A. Timmermann (eds.), *Handbook of Economic Forecasting*, Elsevier.
- [29] Swidler, S. and D. Ketcher, 1990, Economic Forecasts, Rationality, and the Processing of New Information over Time. *Journal of Money, Credit, and Banking* 22, 65-76.
- [30] Tauchen, G., 1986, Statistical Properties of Generalized Method of Moments Estimators of Structural Parameters obtained from Financial Market Data, *Journal of Business and Economic Statistics*, 4, 397-416.

Table 1: Testing rationality of consensus forecasts of US GDP growth and Inflation

Horizon	Bias		MZ p -values	
	GDP growth	Inflation	GDP growth	Inflation
1	0.11	0.01	0.99	1.00
2	0.07	0.02	0.99	1.00
3	0.03	0.04*	0.99	1.00
4	0.02	0.06*	0.99	1.00
5	0.01	0.08*	0.99	1.00
6	0.05	0.08*	0.99	1.00
7	0.07	0.11*	0.98	1.00
8	0.07	0.08	0.97	1.00
9	-0.06	0.01	0.98	0.99
10	-0.16	-0.03	0.98	0.97
11	-0.26	-0.05	0.97	0.96
12	-0.34	0.03	0.95	0.96
13	-0.33	0.11	0.95	0.96
14	-0.35	0.16	0.94	0.96
15	-0.29	0.22	0.95	0.97
16	-0.14	0.28	0.95	0.97
17	-0.10	0.26	0.99	0.98
18	-0.06	0.29	0.99	0.98
19	-0.07	0.33*	0.99	0.98
20	-0.07	0.36*	0.99	0.98
21	-0.08	0.33*	0.99	0.97
22	-0.07	0.32	0.99	0.97
23	-0.07	0.35*	0.99	0.97
24	-0.09	0.37*	0.98	0.98

Notes: * indicates that the bias is significant at the 5% level based on Newey-West (1987) autocorrelation and heteroskedasticity robust standard errors. The first two columns report the average bias in the forecast, for each variable and each horizon, which should be zero for a rational forecast. The final two columns give the p -values from a joint test that $\beta_0^h = 0 \cap \beta_1^h = 1$ in the Mincer-Zarnowitz regression of the realized value of the target variable on the forecast: $y_t = \beta_0^h + \beta_1^h \hat{y}_{t,t-h} + \epsilon_{t,t-h}$, for each horizon h .

**Table 2: GMM parameter estimates of the consensus forecast model:
No measurement error**

	σ_u	σ_ε	ϕ	J p-val
GDP growth	0.229 (0.405)	0.050 (0.031)	0.930 (0.033)	0.000
Inflation	0.000 (—)	0.023 (0.007)	0.953 (0.047)	0.935

Notes: The table reports the GMM estimates of the parameters of the Kalman filter model fitted to the consensus forecasts with standard errors in parentheses, estimated imposing that there is no measurement error (i.e., that $\sigma_\eta = 0$). p -values from the test of over-identifying restrictions are given in the column titled “ J p-val”.

**Table 3: GMM parameter estimates of the consensus forecast model:
Allowing for measurement error**

	σ_u	σ_ε	ϕ	σ_η	J p-val
GDP growth	0.063 (0.012)	0.054 (0.013)	0.936 (0.034)	0.126 (—)	0.403
Inflation	0.000 (—)	0.023 (0.007)	0.953 (0.047)	0.000 (—)	0.935

Notes: The table reports the GMM estimates of the parameters of the Kalman filter model fitted to the consensus forecasts with standard errors in parentheses. p -values from the test of over-identifying restrictions are given in the column titled “ J p-val”. The parameter σ_η was fixed at $2\sigma_u$ and is reported here for reference only.




Figure 1: *Evolution in consensus forecasts for US GDP growth in 2002, for horizons ranging from 24 months (January 2001) to 1 month (December 2002).*

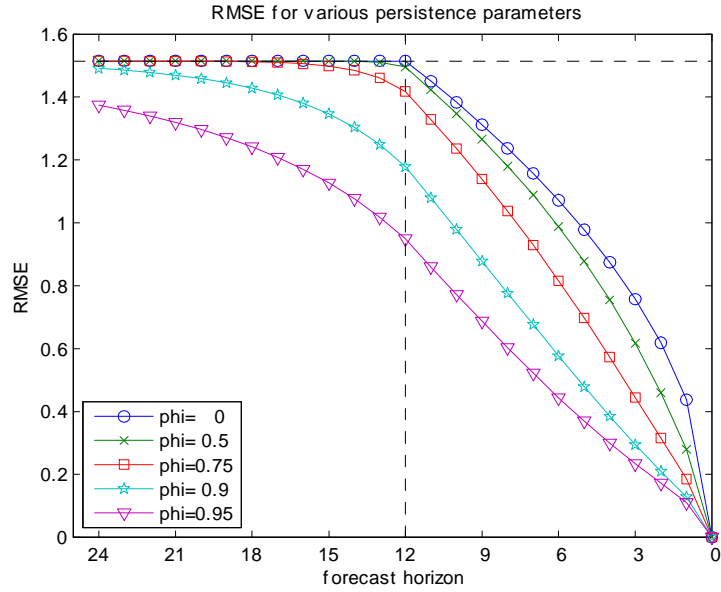


Figure 2: *Root-mean squared forecast errors as a function of the forecast horizon (h) for various degrees of persistence (ϕ) in the predictable component*




Figure 3: *Root-mean squared forecast errors as a function of the forecast horizon (h) for various degrees of measurement error in the predicted variable. In this example, the degree of measurement error is described as $\sigma_\eta = k\sigma_u$, where σ_u is the standard deviation of the unpredictable component of y_t .*

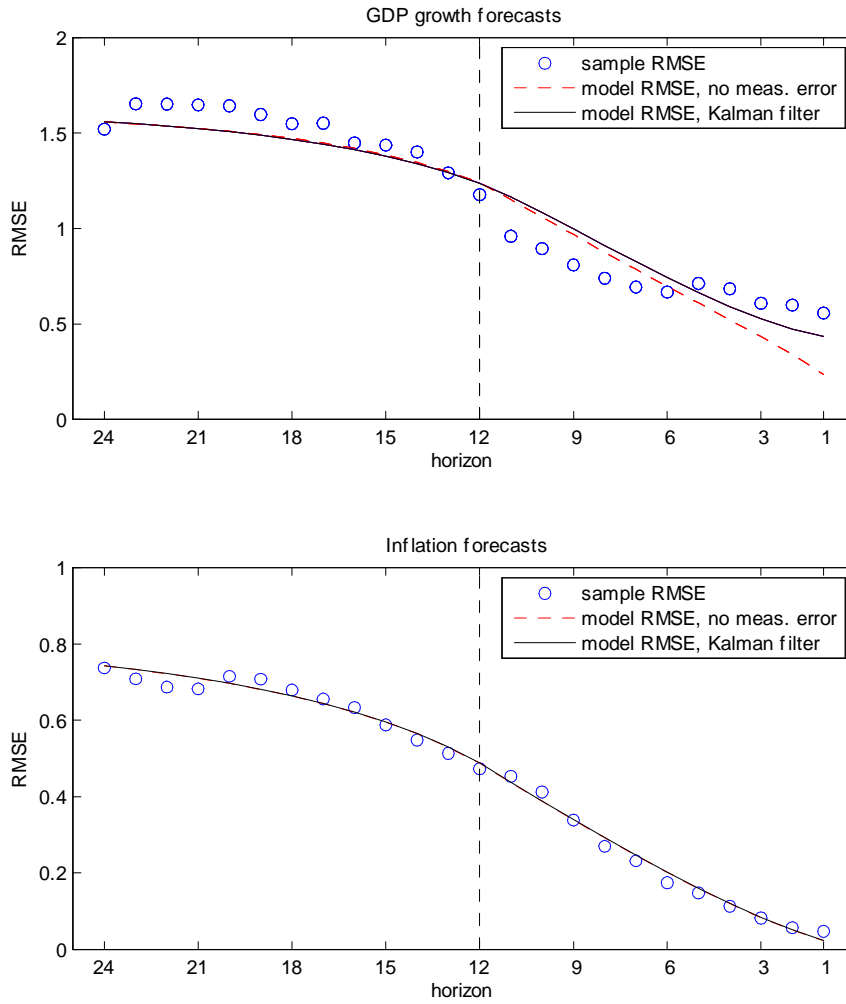


Figure 4: *Root mean squared forecast errors for US GDP growth and Inflation as a function of the forecast horizon.*




Figure 5: *Estimates of the persistent component ($xhat$) of GDP growth and inflation for each month in the sample period, as implied by the observed forecasts and the estimated model for the multi-horizon forecast errors.*