AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Queen's Economics Department Working Paper No. 1316

Measuring the Slowly Evolving Trend in US Inflation with Professional Forecasts

James M. Nason
North Carolina State University

Gregor W. Smith
Queen's University

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

10-2013

Measuring the Slowly Evolving Trend in US Inflation with Professional Forecasts

James M. Nason and Gregor W. Smith[†]

October 2013
first version September 2012

Abstract

Much research studies US inflation history with a trend-cycle model with unobserved components. A key feature of this model is that the trend may be viewed as the Fed's evolving inflation target or long-horizon expected inflation. We provide a new way to measure the slowly evolving trend and the cycle (or inflation gap), based on forecasts from the *Survey of Professional Forecasters*. These forecasts may be treated either as rational expectations or as adjusting to those with sticky information. We find considerable evidence of inflation-gap persistence and some evidence of implicit sticky information. But statistical tests show we cannot reconcile these two widely used perspectives on US inflation and professional forecasts, the unobserved-components model and the sticky-information model.

JEL classification: E31, E37.

Keywords: US inflation, professional forecasts, sticky information, Beveridge-Nelson

[†]Nason: Department of Economics, North Carolina State University; jmnason@ncsu.edu. Smith: Department of Economics, Queen's University; smithgw@econ.queensu.ca. Earlier versions of this paper were entitled 'Reverse Kalman filtering US inflation with sticky professional forecasts.' We thank the Social Sciences and Humanities Research Council of Canada and the Bank of Canada research fellowship programme for support of this research. For helpful comments we thank Drew Creal, Elmar Mertens, Jonathan Wright, seminar participants at the European University Institute (EUI), and conference participants at the Canadian Economics Association meetings, the Federal Reserve Bank of Philadelphia's 2012 Conference on Real-Time Data Analysis, Methods, and Applications, the EFWW Workshop at the NBER Summer Institute, the 2013 NBER-NSF Time Series Conference, and the Canadian Econometrics Study Group. Nason thanks the Department of Economics of the EUI for its support and hospitality while working on this project. The views expressed herein are not necessarily those of the Bank of Canada and are the authors' alone.

1. Introduction

For the past thirty years the unobserved-components (UC) model has been an informative lens through which economists have viewed US inflation dynamics. That statistical model decomposes inflation into permanent and transitory components. The permanent component or trend usually (and in this paper) is identified with the Beveridge-Nelson (1981) decomposition, meaning that it is a random walk. This decomposition has been widely adopted in forecasting inflation. For example, Stock and Watson (2007) use it to isolate changes in the variances of the components and hence in the overall persistence and forecastability of inflation over time. Faust and Wright (2013), in their review of inflation forecasting, list the many studies that feature a slowly evolving trend. But this decomposition also sheds light on inflation history. For example, Cogley, Primiceri, and Sargent (2010) examine time variation in the persistence of the “inflation gap,” defined as the transitory component from this decomposition. A key feature of this model is that the trend component serves as a measure of long-horizon inflation expectations, an indicator of the Fed’s credibility as well as a constraint on the effect of policy.

We use professional forecasts to measure the two components of inflation. A simple example illustrates the idea. First, suppose that inflation is the sum of a random-walk trend, τ_t , and an inflation gap, ϵ_t , that is white noise. Thus τ_t also is the expectation of next period’s inflation. Second, suppose that professional forecasters report their rational expectations. Thus, their one-step-ahead forecasts directly provide an estimate of τ_t , and ϵ_t can be found by subtracting that trend estimate from actual inflation. We then show how to extend this idea to allow for a persistent inflation gap, to incorporate information from forecasts at multiple horizons, and to integrate sticky forecasts into these settings.

As this simple example shows, this approach requires a view on the connection between unobservable, h -step-ahead, rational-expectations (RE) forecasts of inflation, denoted $E_t\pi_{t+h}$, and the mean, reported inflation forecasts of professional forecasters, denoted $F_t\pi_{t+h}$, in our case from the *Survey of Professional Forecasters (SPF)*. We consider two possibilities. First, one way to extract information from the *SPF* is to assume that their mean forecast coincides with a prediction from the UC model with some information set. We first estimate and test under that assumption. But, second, considerable recent research on panels of professional forecasts suggests that they are not full-information, rational expectations but rather exhibit bias. How can professional forecasts be useful if they are biased? Precisely because the pattern of forecast errors is systematic, these surveys provide information on true expectations. One way to describe the evidence is that forecasts are sticky and can be modeled using the sticky-information (SI) frame-

work of Mankiw and Reis (2002). This second description of forecasts also is of interest because it has been widely used to close macroeconomic models, for example in studies of the New Keynesian Phillips curve. Coibion and Gorodnichenko (2011) use the SI model to link reported forecasts to the actual conditional expectations of professional forecasters. We employ this model to estimate the parameter describing stickiness along with those of the UC model, and also apply estimates from Coibion and Gorodnichenko. Under either assumption about professional forecasts, then, we develop procedures to connect mean, reported forecasts to the UC model.

We also study whether we can reconcile the two statistical models jointly with the time-series properties of actual inflation and the mean h -step-ahead prediction of inflation from the *SPF*. This procedure comes with several consistency tests: joint tests of the link between reported *SPF* forecasts and unobserved expectations and of the econometrician's statistical model of inflation. For example, we can test whether the implied stochastic trend in inflation follows a martingale, whether persistence in the implied inflation gap matches that estimated indirectly through the properties of forecasts, and whether forecasts are unbiased (or the extent to which they are sticky). Consistency implies that the parametric models of inflation and of inflation forecasts can be reconciled. If we find such consistency, then we have an easy and informative way to filter US inflation, by outsourcing much of the work to the participants in the *Survey of Professional Forecasters*. The estimated components depend only on reported forecasts and so automatically are available in real time.

If we do not find consistency, then either (a) forecasters are not using the UC model (with *any* information set; this is not a test that the information sets of economists and forecasters coincide), or (b) we do not have the correct model of forecast reporting and so cannot yet reliably use it to extract information from the *SPF*. We cannot know which of these conclusions holds because the approach jointly relies on the UC model and the assumptions about forecasts.

There are four main findings. First, detrending after assuming that mean forecasts coincide with rational expectations provides a direct way to track the historical mixture of the shocks to the trend and inflation gap, with time-varying volatilities. Second, we also find considerable evidence of inflation-gap persistence. Third, though, estimation from forecast data leads to a trend-cycle decomposition with a trend shock (and, when we measure inflation with the GDP deflator, also an inflation-gap shock) that is predictable or persistent, a finding that is inconsistent with the underlying assumptions of the UC model. Fourth, when we allow for stickiness in forecasts, and either estimate it using reported

forecasts or use Coibion and Gorodnichenko’s estimate, the joint model continues to fail the consistency tests. The combined model cannot reproduce unpredictable innovations in the two components of the UC model along with the predictable pattern in forecast errors. So far, then, we cannot reconcile these two widely adopted perspectives on US inflation—the unobserved-components model and the sticky-information model—and fit the time-series properties of both historical inflation and the term structure of *SPF* forecasts.

2. The Trend-Cycle Model

The first element in our study is a variation on the Beveridge-Nelson-type decomposition of inflation. For simplicity we refer to this as the unobserved-components (UC) model or the SW (for Stock and Watson) UC model. Suppose that inflation, π_t , evolves as a sum of two components: a stochastic trend τ_t and a stationary component ϵ_t . In this environment the stochastic trend component follows a driftless random walk, with innovation η_t . Thus:

$$\begin{aligned}\pi_t &= \tau_t + \epsilon_t \\ \tau_t &= \tau_{t-1} + \eta_t.\end{aligned}\tag{1}$$

The stationary component ϵ_t and the trend-innovation η_t are martingale difference series. But they may be correlated and may have time-varying volatilities.

This decomposition has been fruitful in studies of several aspects of inflation dynamics. For example, Ireland (2007) estimates the Federal Reserve’s implicit, time-varying inflation target with a Beveridge-Nelson trend. Cogley and Sbordone (2008) use a similar, stochastic trend around which to estimate a New Keynesian Phillips curve. Stock and Watson (2007) interpret the changing persistence and forecastability of US inflation with the UC model with changes in shock variances. Cogley, Primiceri, and Sargent (2010) use the model to identify changes in the persistence of the inflation gap, ϵ_t .

Estimation and forecasting with the UC model require one to use the Kalman filter to extract the unobserved components. The filter is applied beginning with orthogonality assumptions (for example, a zero covariance between η_t and ϵ_t) and a set of covariates in observation equations. Examples of studies that apply the Kalman filter to this model include Cogley and Sargent (2005), Nason (2006), Stock and Watson (2007), Cogley, Primiceri, and Sargent (2010), Mertens (2011), and Shephard (2013). To take one example, Mertens applies the Kalman filter to a wide-range of macroeconomic data with the assumption that actual inflation, inflation surveys, and nominal interest rates share a common stochastic trend. He allows for a correlation between the trend and gap shocks as well as stochastic volatility in the trend-shock, η_t , that itself follows a random walk.

The filter allows the joint estimation of parameters (through the prediction-error decomposition of the likelihood function) and extraction of the components. In familiar notation, we denote by $\tau_{t|t}$ the estimate of τ_t with information at time t (*i.e.*, the filtered value) and similarly for $\epsilon_{t|t}$. The h -step-ahead forecast of inflation then is

$$E_t\pi_{t+h} = \tau_{t|t}, \quad (2)$$

for $h \geq 1$. This formula yields the Beveridge-Nelson (1981) result that

$$E_t\pi_{t+\infty} = \tau_{t|t}, \quad (3)$$

so that the trend estimate also is the estimate of expected inflation at the infinite horizon.

We reverse the last steps of this sequence. We begin with the reported forecast, denoted $F_t\pi_{t+h}$, which (in the simplest example described in the introduction) coincides with the expectation $E_t\pi_{t+h}$. Thus $\tau_{t|t} = F_t\pi_{t+h}$. It is obvious that we cannot then continue and uncover a unique, underlying information set and a set of orthogonality assumptions. But we also do not require a zero covariance between the shocks or restrictions on their variances to measure the two components and later estimate inflation-gap persistence. Morley, Nelson, and Zivot (2003) show that restricting this covariance can have a large impact on trend-cycle decomposition. Any pattern of time-varying volatility in the innovations in τ_t also is possible, so that the importance of the non-stationary component can vary over the sample. Our method uses only reported, professional forecasts and actual inflation. It is possible to study inflation forecasting and trend-cycle decomposition without any covariates because their assessment and selection implicitly are outsourced to the forecasters.

Our study can be thought of as a sequel to that of Kozicki and Tinsley (2012), who estimate the parameters of an unobserved-components model of CPI inflation using actual inflation and the long span of observations from the Livingston survey, allowing for higher-order dynamics to fit seasonally unadjusted data, and under the assumption that forecasts are conditional expectations. They provide a detailed discussion of the interpretation and need for a shifting endpoint, $\tau_{t|t}$, for inflation forecasts. Henzel (2013) similarly combines SPF forecasts with the UC model to estimate inflation expectations. He also contrasts the speed of adjustment (or Kalman gain) in SPF forecasts with that estimated for the UC model alone. These two studies combine information sources using the Kalman filter (in sub-sections 4.3 and 4.4 below we provide a similar exercise for the *SPF*, with findings complementary to theirs). They also provide a test of the combined statistical model's

ability to fit observed survey forecasts. We use the *SPF* with multiple, short horizons and quarterly observations and also consider the alternative assumption that forecasts are sticky. This extension potentially reconciles the method with the bias in mean *SPF* forecasts yet still allows us to use those forecasts to estimate the UC model and measure inflation expectations.

This project also is related to several other recent studies that jointly analyze survey-based inflation expectations and time-series models of actual inflation. Clark and Davig (2011) include one-year-ahead and ten-year-ahead inflation expectations (from the *SPF*) in a VAR with time-varying parameters and stochastic volatility. They document the decline in the volatility of long-term inflation expectations and find that this is due largely to shocks to expectations themselves. Del Negro and Eusepi (2011) examine whether the observed properties of professional forecasts are consistent with a New Keynesian DSGE model. They find the closest match when there is time variation in the Fed’s implicit target for inflation. But their test of over-identifying restrictions shows there is not a complete reconciliation between the forecast data and the expectations predicted in the economic model. Jain (2011) applies a state-space model with a persistent but stationary unobserved component to the forecasts of individual forecasters in the *SPF*. She uses the properties of forecast revisions to deduce the persistence implied in these forecasters’ views of the underlying state variables and finds that this persistence has declined over time for many forecasters.

3. Inflation Forecast Data

The forecast data come from the *Survey of Professional Forecasters*, organized by the Federal Reserve Bank of Philadelphia. The survey was conducted by the ASA/NBER prior to the summer of 1990. We use the mean forecast for the annualized rate of CPI inflation, measured quarterly from 1981:3 to 2012:3, yielding 124 observations. The survey reports forecasts from zero (the nowcast) to four quarters ahead.

The upper panel of figure 1 shows actual US inflation, given by the annualized quarter-to-quarter growth rate in the CPI for all urban consumers and all items, series `cpiaucsl` from FRED at the Federal Reserve Bank of St. Louis. The lower panel of figure 1 shows the mean *SPF* forecasts at a common date of origin (rather than for a common target) with the five different horizons, on the same scale as for realized inflation. As the horizon rises, the volatility of the forecast decreases strikingly.

The *SPF* also contains data on long-term inflation forecasts, specifically over the next year and the next ten years. The one-year forecast is the average of the median forecasts

for $h = 1$ to $h = 4$. The ten-year forecast is the annual average inflation rate predicted for this period. Thus it is not a long-horizon forecast that can directly measure $\tau_{t|t}$. Moreover, this survey information has been collected only since 1991.

Forecasts for inflation are available for a longer time span, beginning in 1968:4, if we study the inflation rate in the GDP deflator rather than the CPI. These forecasts are for seasonally adjusted levels of the deflator, defined as (a) the GNP deflator prior to 1992, (b) the GDP deflator from 1992 to 1995, and (c) the chain-weighted price index for GDP from 1996 to the present. Then implicit mean forecasts for the annualized growth rate in the deflator are from `mean_PGDP_Growth.xls`. The corresponding realized inflation rate (unlike the CPI) is subject to revisions. We use the most recent observation, `realized5` from `Data_SPF_Error_Statistics_PGDP_3.AIC.xls` in the *SPF*.

The upper panel of figure 2 shows the realized inflation rate measured with the GDP deflator, quarterly, at annual rates. The lower panel of figure 2 shows the mean forecasts at the same dates, from the nowcast ($h = 0$) to four quarters ahead ($h = 4$). Again, these are on the same vertical scale as that for realized inflation. The span of years, of course, now includes the high-inflation years of the 1970s. Again, the volatility of the forecasts decreases at longer horizons.

4. Mean Forecasts as Rational Expectations

The second element in our study is a description of forecast data, and we begin with the simplest assumption: The cross-forecaster mean coincides with the rational expectation of future inflation. Unbiasedness of professional forecasts constitutes indirect evidence in favor of this coincidence. Keane and Runkle (1990) provide early evidence of the unbiasedness of price forecasts using disaggregated data from the Livingston Survey. Ang, Bekaert, and Wei (2007) describe an inflation-forecasting tournament in which the median professional forecast is the best predictor of annual inflation. Gil-Alana, Moreno, and Pérez de Gracia (2012) find similarly favorable results for survey-based expectations of quarterly inflation and, specifically, the mean CPI inflation forecasts from the *SPF*. Croushore (2010) demonstrates the general lack of bias in the *SPF* forecasts using real-time measures of target variables. Overall, as Faust and Wright (2013, p. 5) note, “Subjective forecasts of inflation seem to outperform model-based forecasts in certain dimensions, often by a wide margin.” Winning tournaments based on mean-squared error, of course, does not imply unbiasedness, but it at least rules out some systematic biases, for otherwise a time-series model would incorporate those and improve upon the professional forecasts. To quote Faust and Wright again (p. 21), “A useful way of assessing models [thus] is by their ability

to match survey measures of inflation expectations.” Section 6 adopts a more general description of survey-based expectations that admits some bias in forecasts, but meanwhile we first show how to apply our method using the preliminary assumption that the mean forecast coincides with the unobserved expectation of inflation.

In this section, we first use the method outlined in section 2 to decompose CPI inflation using forecasts for a single horizon or averaged over horizons. We also extend the method to allow for persistence in the inflation gap, in the form of first-order autocorrelation. Next, we use the Kalman filter to extract the stochastic trend τ from a set of forecasts at multiple horizons. Finally, we also examine evidence for the inflation rate measured with the GDP deflator.

4.1 The Basics

Suppose that the unobserved expectation of inflation one quarter ahead coincides with the mean professional forecast, denoted $F_t\pi_{t+1}$. From the trend-cycle model, then, the trend is simply

$$\tau_{t|t} = F_t\pi_{t+1}, \quad (4)$$

so the inflation gap is

$$\epsilon_{t|t} = \pi_t - F_t\pi_{t+1}. \quad (5)$$

Our procedure comes with three consistency tests. First, the extracted, stochastic trend should follow a random walk, so its difference should be unpredictable by its own past values: $\Delta\tau_{t|t}$ should be white noise. Second, the extracted inflation gap, $\epsilon_{t|t}$, should also be white noise. Third, the mean forecast should be unbiased, so $\pi_{t+h} - F_t\pi_{t+h}$ should be unpredictable for all horizons h . (Notice that this forecast error differs from the estimated inflation gap, which is $\pi_t - F_t\pi_{t+h}$.)

Table 1 gives the sample variances s^2 of each innovation. The variances of $\epsilon_{t|t}$ and $\eta_{t|t}$ are comparable to the estimates Stock and Watson (2007) report using their UC model. The correlation between $\epsilon_{it|t}$ and $\eta_{it|t}$, denoted $r(\epsilon, \eta)$, is 0.35.

Table 1 next gives $Q(j)$, the Ljung-Box Q -statistic with j lags, and its p -value, for each innovation. These show that there is little evidence of autocorrelation in the inflation gap, but some evidence of autocorrelation in the innovation to the trend. We did not impose the martingale-difference-series property on these series in estimation, so these statistics provide tests of the consistency of the UC model with the *SPF* data (assuming that $F_t\pi_{t+h} = E_t\pi_{t+h}$), something that does not automatically hold. Overall, these two

widely used ways of studying inflation forecasts do seem to be approximately consistent. The exception is some evidence of persistence in the inflation-trend innovations, $\eta_{t|t}$.

To document possible changes over time in these moments, the lower panel of table 1 also shows sample variances s^2 for both the inflation gap and the difference in the trend, but now for three sub-samples of approximately a decade each. The break dates are one quarter after NBER-dated troughs and roughly line up with the break dates implied by rolling estimates of the Stock-Watson UC model by Nason (2006). Table 1 also reports the sample correlation $r(\epsilon, \eta)$ for each time period.

The volatility of each component declines from the 1980s to the 1990s, then increases after 2002. Grassi and Proietti (2010) and Creal (2012) estimate the SW-UC model with stochastic volatility. They find that the volatility of CPI inflation has increased recently, with the increased volatility attributed by the estimates to the transitory rather than the permanent component of the SW UC model. Table 1 leads to a similar conclusion. It also shows that the correlation between the two components follows a similar pattern over time—falling then rising—but remains positive.

Of course, the UC model also implies that $F_t \pi_{t+h} = \tau_{t|t}$, for all $h > 1$. The UC model implies a singularity that is not present in the forecasts. In this preliminary exercise we appeal informally to measurement error, perhaps due to differences in the composition of the *SPF* panel reports across horizons, to suggest an alternate estimator that averages forecasts across horizons:

$$\tau_{t|t} = \frac{1}{4} \sum_{h=1}^4 F_t \pi_{t+h}, \quad (6)$$

with $\epsilon_{t|t}$ again given by subtraction (5). (Another possibility would be to inverse-weight the multi-step forecasts by their variances.)

Table 2 provides statistics on the inflation components extracted by this second method. The top line of table 2 shows that the main difference between the two trend-estimates (in equations (4) and (6)) is the lesser volatility of $\eta_{t|t}$ in the second case. By averaging over horizons, the second method produces a smoother trend. The variance of $\epsilon_{t|t}$ rises and the variance of $\eta_{t|t}$ falls when we include information at longer horizons, which suggests more uncertainty about transitory inflation shocks than permanent ones at a moment in time at the longer forecast horizons.

The Q -statistics in table 2 again show more evidence of persistence in $\eta_{t|t}$ than in $\epsilon_{t|t}$. And in the lower panel, the volatility of both components again falls, then rises across the three sub-samples.

Figure 3 shows the CPI inflation rate along with $\tau_{t|t}$ from one-step-ahead forecasts (in the upper panel). The trend formed from averaging across horizons is very similar and is not shown. Both estimated trends are smoother than the actual inflation rate and centered on it. The trend estimated only from the one-horizon forecasts is more volatile than the one based on averaging forecasts over horizons, but it is striking that they are very similar for substantial periods of time. Also, the positive correlation between innovations in tables 1 and 2 naturally fits with the observation in figure 3 that the trend is smoother than the cycle.

The lower panel of figure 3 shows the trend and cycle innovations for the first estimation, based on a single horizon: $\eta_{t|t} = \Delta\tau_{t|t}$ (the long-dashed, blue line) and $\epsilon_{t|t}$ (the short-dashed, red line). Little persistence is evident in either series. (The graphs of the corresponding innovations implied by the second detrending method (6) are quite similar.)

The forecast-based decomposition also gives us insights into the history of squared innovations in the trend and inflation gap: roughly speaking, quarterly “realized volatility” or the things we would average over periods of time to estimate time-varying volatilities. It is clear from figure 3 that there is variation over time in volatility: The inflation-gap (the red, short-dashed line) is relatively more volatile prior to 1990 and again after 2005. (Again, results for the other trend method are quite similar and so are not shown.)

4.2. Persistence in the Inflation Gap

The UC model (combined with the assumption that $F_t\pi_{t+h}$ coincides with $E_t\pi_{t+h}$) makes the prediction (2) that forecasts are the same at all horizons, a feature that does not hold in the *SPF*. To avoid that implication, we next follow Cogley, Primiceri, and Sargent (2010) and allow the stationary component of inflation, also known as the inflation gap, to itself be persistent. We work with an AR(1) version:

$$\epsilon_t = \rho\epsilon_{t-1} + v_t, \quad (7)$$

where v_t is a martingale-difference series. But we do not allow time variation in the persistence parameter ρ , unlike Cogley, Primiceri, and Sargent. The infinite-horizon inflation forecast remains $\tau_{t|t}$, but in general

$$E_t\pi_{t+h} = \tau_{t|t} + \rho^h\epsilon_{t|t}. \quad (8)$$

Faust and Wright (2013) note that subjective forecasts often prove superior to econometric forecasts of inflation because they do not simply extrapolate the current value but allow

for a gradual return to some medium-term pattern. Similarly, Kozicki and Tinsley (2012) note that a successful inflation-forecasting model needs to have a role for current inflation at short horizons but not at long ones. The forecasts (8) allow for these patterns.

This sub-section briefly shows how one can estimate ρ , the persistence in the inflation gap, using forecast data. Continue to suppose that $F_t\pi_{t+h} = E_t\pi_{t+h}$. Then we can annihilate the stochastic trend, $\tau_{t|t}$, using the difference across horizons:

$$F_t\pi_{t+h+1} - F_t\pi_{t+h} = F_t\Delta\pi_{t+h+1} = \rho^h(\rho - 1)\epsilon_{t|t}. \quad (9)$$

Multiplying the difference equations (9) by $(1 - \rho L)$ then gives the quasi-differences over time:

$$F_t\Delta\pi_{t+h+1} = \rho F_{t-1}\Delta\pi_{t+h} + \rho^h(\rho - 1)v_{t|t}, \quad (10)$$

which have innovation errors. The forecasts on the right-hand side are dated $t-1$ or earlier, so it is natural to assume that the inflation-gap shock, v_t , is uncorrelated with them. Thus the persistence in the inflation gap, ρ , can be estimated by ordinary least squares in the estimating equations (10). Inflation-gap persistence coincides with the persistence over time in the forecast of the change in inflation. The mean forecast data provide a simple way to estimate the persistence in the inflation gap. Armed with $\hat{\rho}$ we can invert (9) to filter $\epsilon_{t|t}$, then find $\tau_{t|t}$ from the original trend-cycle model.

Notice that the estimation again uses only forecast data; it does not use actual inflation, π_t . Identifying ρ from inflation dynamics would require a restriction on the covariance between the two innovations, v_t and η_t . But the forecast-only estimating equations (10) do not require such a restriction. That is because forecasts at multiple horizons all involve the current trend estimate $\tau_{t|t}$, which thus can be removed by differencing between them.

We interpret the estimated, stochastic trend as the filtered value $\tau_{t|t}$, rather than the smoothed one $\tau_{t|T}$, because it is derived from forecasts observed at time t . Jain (2011) also examines the correlation of revisions to measure perceived inflation persistence but with forecasts from individual forecasters in the *SPF*. Krane (2011) pursues a similar goal using GDP forecast revisions from the Blue Chip survey to identify forecasters' implicit views of shocks to GDP.

We estimated the system (10) without restrictions on the error dispersion matrix (even though those contain information on ρ), using forecasts at individual horizons h and then from pooled estimation over all horizons. The main finding is that the values—pooled or individual—are insignificantly different from zero. Thus, allowing for persistence does not change the findings from sub-section 4.1.

However, both the basic setup and this extension that allows for inflation-gap persistence feature a singularity in that forecasts (8) are perfectly correlated across horizons. Next, we allow for some errors of measurement or reporting and show how to estimate the UC model from the term structure of *SPF* forecasts but without this unrealistic feature.

4.3 The Kalman Filter

So far we have not formally allowed for the fact that the forecasting model has a singularity across horizons, while the reported forecasts do not. In this sub-section, we explicitly allow for differential information on the common, stochastic trend from forecasts at different horizons, and thus allow for an imperfect correlation between them.

Suppose that mean forecasts contain a white-noise error of reporting or composition, labeled ξ_{ht} for horizon h and time t . We also assume that this error is not correlated across horizons. Reported forecasts now are

$$F_t \pi_{t+h} = \tau_{t|t} + \rho^h \epsilon_{t|t} + \xi_{ht}. \quad (11a)$$

To this point we have relied entirely on the *SPF* forecasters to distinguish between the trend and cycle components of inflation; their filtering has been our filtering. Now that there is measurement error, though, we need to do our own filtering of the forecasts. To keep the notation simple, we change the subscripts in the equations for reported forecasts (11a)—for example from $\tau_{t|t}$ to τ_t —to reflect the fact that additional filtering is necessary to estimate the two components of inflation. Thus the mean forecasts can be written

$$F_t \pi_{t+h} = \tau_t + \rho^h \epsilon_t + \xi_{ht}. \quad (11b)$$

To give stationarity, one could take differences across horizons to annihilate the stochastic trend:

$$F_t \Delta \pi_{t+h} = (\rho^h - \rho^{h-1}) \epsilon_t + \xi_{ht} - \xi_{h-1,t}, \quad (12)$$

as in sub-section 4.2. But it makes more sense to subtract actual inflation and so use

$$F_t \pi_{t+h} - \pi_t = (\rho^h - 1) \epsilon_t + \xi_{ht}, \quad (13)$$

because that removes one source of error. Instead of taking quasi-differences to eliminate ϵ_t , as we outlined in sub-section 4.2, we now retain it in each equation to allow filtering and estimation. Equations (13), for $h = 1, 2, 3, 4$, are the observation equations for a state-space model. The transition equation is simply the AR(1) model (7) for ϵ_t , with $\rho \in (-1, 1)$.

We now also need to assume that the innovations to the gap ϵ_t , namely v_t , are homoskedastic, Gaussian shocks: $v_t \sim N(0, \sigma_v^2)$. Place similar assumptions on the white-noise measurement errors: $\xi_{ht} \sim N(0, \sigma_{\xi_h}^2)$, for $h = 1, 2, 3, 4$. Finally, we also assume that the measurement errors are uncorrelated with the shocks to the inflation gap.

In sub-section 4.1 we used a single horizon and averaged over horizons to illustrate our method. But once we include more than one horizon, it makes sense to extract the components using the Kalman filter to avoid the implausible singularity implied by the UC model without measurement error. The filter in this section also now uses actual inflation to enhance efficiency in extracting the components. It still does not require orthogonality between the components (or their innovations), but it does require homogeneity (constant variances). The Kalman filter is used to construct the prediction-error form of the likelihood function (see the appendix for details).

The top panel of table 3 gives maximum-likelihood estimates (MLEs) for the parameters of this UC model of CPI inflation and the *SPF* term structure of its forecasts, based on the 1981Q3 to 2012Q3 sample. The standard errors in parentheses are computed using the Hessian evaluated at the MLEs.

Two key findings emerge from implementing this stochastic environment in table 3. First, we now find a small, but highly significant estimate $\hat{\rho}$, the measured persistence in the inflation gap. As Cogley, Primiceri, and Sargent (2010) note, persistence in the inflation gap describes the speed with which inflation returns to its slowly evolving target and so is a measure of policy. Edge and Gürkaynak (2010) observe that it would be odd, then, if deviations from the target were highly persistent. Second, the shock to the inflation gap is much more variable than any of the measurement errors, which is a reassuring feature.

Recall that we can measure the two components of the UC model without restricting the correlation between them. The correlation of $\eta_{t|t}$ and $\epsilon_{t|t}$ is 0.16 (with a p -value of 0.08) and the correlation of $\eta_{t|t}$ and $v_{t|t}$ is 0.11 (with a p -value of 0.22). The bottom panel of table 3 replaces tables 1 and 2 and reports on the volatility and persistence of the two innovations, $v_{t|t}$ and $\eta_{t|t}$. A key finding is the stronger evidence that $\tau_{t|t}$ is not a martingale, given the sample persistence in $\eta_{t|t}$.

The upper panel of figure 4 shows the filtered trend and cycle along with actual CPI inflation. We also constructed the smoothed estimates of these unobserved components, but they are virtually identical to the filtered estimates and so are not shown. The lower panel of figure 4 shows the estimated innovations. Here, a second key diagnosis is that there is some evidence of time-varying volatility in σ_v^2 , a moment assumed constant in

the state-space model. We do not do sub-sample estimation because there are only 40 time-series observations per sub-sample, yet six parameters. But time variation in ρ or in the variances could be investigated with extensions to the state-space model.

We also backed out the implied measurement errors, $\hat{\xi}_{h,t}$. For horizons $h = 1$ and 2 , the measurement errors have small AR(1) coefficients (less than 0.3). For $h = 3$, the AR(1) coefficient is 0.75 and for $h = 4$ it is 1. As h rises, the forecasts $F_t\pi_{t+h}$ become smoother so that the persistence in π_t shows up in the associated measurement error. Thus, this evidence, too, calls into question the specification of the joint UC/rational expectations model.

4.4 GDP Deflator Evidence

Before reformulating the model in section 5, we next examine its implications for a longer time span, which is available if we study the inflation rate in the US GDP deflator rather than the CPI. However, some forecast observations are missing for the late 1960s and early 1970s, and the results also are sensitive to the starting date. Table 4 presents a particular example of estimation with starting date 1974Q4. Again, the innovation to the inflation cycle is much more volatile than the measurement error at any horizon, which provides some support for the usefulness of the model. The upper panel of figure 5 shows the estimated trend and cycle, along with realized inflation in the GDP deflator since 1975. The lower panel shows the innovations $\eta_{t|t}$ and $v_{t|t}$. Table 4 also gives statistical properties of the innovations to the trend and cycle components of inflation.

In contrast to our results with CPI inflation, now there is evidence of higher-order dynamics in $v_{t|t} \equiv \epsilon_{t|t} - \hat{\rho}\epsilon_{t-1|t-1}$, as shown by the large Q -statistics. Obviously, modeling the dynamics of the inflation gap ϵ_t with an AR(1) model with constant coefficients and homoskedastic shocks is not sufficient to capture its dynamics, given the random-walk model of the trend and the information in the *SPF* term structure of inflation forecasts.

5. Sticky Forecasts

Notwithstanding our earlier citations to research showing that professional forecasts are unbiased, a number of statistical studies have found that forecast errors contain predictable components. Next, a specific pattern of predictability, using mean forecast revisions, leads to an alternative, parametric model of observed, mean forecasts.

We work with the sticky-information (SI) model, as introduced by Mankiw and Reis (2002), or the closely related predetermined pricing model of Devereux and Yetman (2003). The SI model has been applied to professional forecasters by Mankiw, Reis, and Wolfers

(2004) and Coibion and Gorodnichenko (2011, 2012). There also is very interesting, recent research that characterizes the stickiness in household forecasts. For example, Carroll (2003) uses the sticky-information model to describe how households in the Michigan survey adjust their forecasts toward the *SPF*. But we focus on extracting information from professional forecasts, in part because they have the key feature of being available over multiple horizons.

There also is ongoing research on alternatives to the sticky-information model for professional forecasters that seeks to fit additional features of forecasts. For example, Capistrán and Timmermann (2009) show the effects of asymmetric loss functions. Patton and Timmermann (2010) study the term structure of the dispersion in professional forecasts and conclude that heterogeneity in models or priors is necessary to explain it. Andrade and Le Bihan (2010) study individual forecasts from the European *SPF* and find that the SI model is not sufficient to capture all their patterns. Nonetheless, we work with the SI model because of its tractability in this application and also because it has been used to close and estimate macroeconomic models, not just to describe forecasts. For example, Reis (2006), Khan and Zhu (2006), Kiley (2007), and Coibion (2010) test versions of the SI model applied to price-setting and hence to aggregate inflation.

Following Coibion and Gorodnichenko (2011), suppose that forecasters update their information with probability $1 - \lambda$, so that λ measures the degree of stickiness in information. Recall that $F_t\pi_{t+h}$ is the cross-forecaster mean forecast at time t for inflation h steps ahead. Coibion and Gorodnichenko (2011) show that this average forecast is a weighted average of the rational expectation and the previous period's mean, reported forecast:

$$F_t\pi_{t+h} = (1 - \lambda)E_t\pi_{t+h} + \lambda F_{t-1}\pi_{t+h}. \quad (14)$$

Define the non-sticky-information forecast error:

$$\vartheta_{t+h} = \pi_{t+h} - E_t\pi_{t+h}. \quad (15)$$

Subtracting each side of this pattern in reported, mean forecasts (14) from realized inflation gives

$$\begin{aligned} \pi_{t+h} - F_t\pi_{t+h} &= \lambda(E_t\pi_{t+h} - F_{t-1}\pi_{t+h}) + (\pi_{t+h} - E_t\pi_{t+h}) \\ &= \frac{\lambda}{1 - \lambda}(F_t\pi_{t+h} - F_{t-1}\pi_{t+h}) + \vartheta_{t+h}. \end{aligned} \quad (16)$$

Because ϑ_{t+h} has the properties of an econometric error, this link (16) can be used to estimate λ , by regressing the observed, mean forecast error on the mean forecast revision.

Coibion and Gorodnichenko note that this pattern of predictability applies to mean forecasts, not individual ones. So finding stickiness in mean forecasts is consistent with empirical evidence on the unbiasedness of individual, professional forecasts. A further key feature of their regression (16) is that it allows them to test full-information rational expectations (the null hypothesis that $\lambda = 0$) but also to parametrize stickiness with the estimate of λ . They also find that additional regressors, in the form of past, realized values of macroeconomic variables, are not significant in explaining forecast errors once mean revisions are included.

Coibion and Gorodnichenko show that a non-zero value of λ can be consistent with information that is noisy rather than sticky. Under sticky information, but not noisy information, though, the parameter λ is constant across horizons h , a feature we adopt in estimation. Coibion and Gorodnichenko find evidence that this parameter is indeed stable across horizons in professional forecasts of most macroeconomic variables, including inflation. But they also note that λ appears to vary over time, depending on macroeconomic conditions.

Using *SPF* data and pooling across all horizons they find (in their figure 1) $\hat{\lambda} = 0.42$ for the GDP deflator inflation rate and $\hat{\lambda} = 0.39$ for the CPI inflation rate. These point estimates imply that forecasters update every five months on average. Overall, then, their findings serve as a third consistency test of the approach in section 4 and one that statistically rejects the hypothesis that mean forecasts are conditional expectations. But they also give us a well-tested and parametrized alternative to apply.

5.1 The Basics

We shall use their estimator but also consider an alternate estimator that uses only forecast data. We again begin with CPI inflation and do not yet allow for persistence in the inflation gap. Combining the forecast implication of the trend-cycle model (2) with the description of forecast updating (14) gives

$$F_t \pi_{t+h} = (1 - \lambda) \tau_{t|t} + \lambda F_{t-1} \pi_{t+h}. \quad (17)$$

Next, take differences over time to give estimating equations:

$$F_t \pi_{t+h} - F_{t-1} \pi_{t+h-1} = \lambda (F_{t-1} \pi_{t+h} - F_{t-2} \pi_{t+h-1}) + (1 - \lambda) \eta_{t|t}, \quad (18)$$

which can be used to estimate λ . Also notice that the reported forecasts are no longer predicted to be equal at all horizons, although the shocks are still perfectly correlated in the system (18).

We estimated λ using the Coibion-Gorodnichenko projection (16) and, alternately, using our forecast-only projection (18). And we estimated horizon-by-horizon and pooled across horizons. The estimates varied considerably depending on the estimating equation and the horizon, with values ranging from near 0 to 0.4. Standard errors also varied, but a number of the estimates were imprecise. We conclude that there is considerable uncertainty about the value, depending on the horizon and information used in estimation, so we present detrending results for several illustrative values for λ .

Given a value for $\hat{\lambda}$ we can invert the forecast updating rule (17) for the UC model to give the estimated trend based on forecasts at horizon h . And again we can average over horizons to form an estimate, or use a weighted average based on fit. Here we use the simple average:

$$\tau_{t|t} = \frac{1}{4} \sum_{h=1}^4 \frac{F_t \pi_{t+h} - \lambda F_{t-1} \pi_{t+h}}{(1 - \lambda)}, \quad (19)$$

and two trial values for λ , 0.2 and 0.4. Table 5 gives statistics for the corresponding shocks to the inflation trend and inflation gap, in CPI data since 1981. The first set of columns pertains to $\lambda = 0.2$ and the last set to $\lambda = 0.4$.

It is clear that allowing for stickiness does not help with the borderline predictability of the trend-innovation $\eta_{t|t}$ in tables 1–3 and in fact worsens this syndrome dramatically as λ rises from 0. (Though the results are not shown, we also find that $\epsilon_{t|t}$ no longer has mean 0 as λ rises.) The trend that is jointly implied by the UC and sticky-forecast models and the *SPF* data does not have unpredictable changes.

The upper panel of figure 6 graphs actual CPI inflation and the two implied trends corresponding to the two candidate values of λ . The lower panel then shows the innovations to the trend and cycle, for the case of $\lambda = 0.4$. It is obvious from figure 6 that this trend does not run through the middle of the realized inflation series. Indeed, the estimated trend slopes up over time for large enough values of λ , whereas the slope of the path of actual CPI inflation is negative overall for these three decades. At higher frequency, figure 6 also shows that the estimated trends are very different during the recent recession, for example, and so the inferences for expected inflation also are very sensitive to the value of λ . The forecast “quasi-updates” $F_t \pi_{t+h} - \hat{\lambda} F_{t-1} \pi_{t+h}$ that yield trend estimates of $\tau_{t|t}$ from equation (19) do not appear to have a random-walk component. The stickiness that fits the partial predictability of forecast errors (at least for some horizons from estimating equations (14) or (19) or from Coibion and Gorodnichenko’s findings) leads to properties that are clearly inconsistent with the assumptions of the UC model. Although we have not incorporated the requirements from table 5 in our estimators (*i.e.*, the estimators do

not require innovation properties in the implied $\epsilon_{t|t}$ and $\eta_{t|t}$), it seems fairly clear so far that no reasonable value of λ can allow the joint statistical model to pass the consistency tests.

5.2 Identifying Stickiness and Persistence

Section 4.2 showed the features of forecast persistence that identify the persistence parameter, ρ , in the inflation gap. One might wonder whether the parameter λ —when estimated strictly from forecasts—is perhaps measuring persistence in the inflation gap rather than stickiness in forecasts. We next show the characteristics of forecasts that separately identify these two parameters. The analysis of this section also shows why we cannot automatically fit the properties of the term structure of forecasts with the UC model.

Combining the forecast implication of the trend-cycle model with a persistent inflation gap (8) with the description of forecast updating (17) gives

$$\frac{F_t \pi_{t+h} - \lambda F_{t-1} \pi_{t+h}}{1 - \lambda} = \tau_{t|t} + \rho^h \epsilon_{t|t}, \quad (20)$$

or

$$F_t \pi_{t+h} = \lambda F_{t-1} \pi_{t+h} + (1 - \lambda) \tau_{t|t} + (1 - \lambda) \rho^h \epsilon_{t|t}. \quad (21)$$

We again use the fact that forecasts at time t for any horizon involve the random-walk component τ_t and so we difference out that unobserved variable over horizons. Leading the horizon gives

$$F_t \pi_{t+h+1} = \lambda F_{t-1} \pi_{t+h+1} + (1 - \lambda) \tau_{t|t} + (1 - \lambda) \rho^{h+1} \epsilon_{t|t}, \quad (22)$$

so that the difference across horizons is

$$F_t \Delta \pi_{t+h+1} = \lambda F_{t-1} \Delta \pi_{t+h+1} + (1 - \lambda) (\rho^{h+1} - \rho^h) \epsilon_{t|t}. \quad (23)$$

Suppose that the inflation gap, ϵ_t , follows the AR(1) process (7) so that $\epsilon_t(1 - \rho L) = v_t$. Multiplying the difference equations (23) by $(1 - \rho L)$ gives

$$F_t \Delta \pi_{t+h+1} = \lambda F_{t-1} \Delta \pi_{t+h+1} + \rho F_{t-1} \Delta \pi_{t+h} - \rho \lambda F_{t-2} \Delta \pi_{t+h} + (1 - \lambda) (\rho^{h+1} - \rho^h) v_{t|t}. \quad (24)$$

The forecasts on the right-hand side are dated $t - 1$ or earlier, so it is natural to assume that the inflation-gap shock, v_t , is uncorrelated with them. Thus the persistence in the inflation gap, ρ , and the stickiness in inflation forecasts, λ , can be jointly estimated by

ordinary least squares in the estimating equation (24) for a given horizon h . The stickiness and persistence parameters are separately identified, from distinct sources of dynamics in forecasts. Persistence, ρ , is estimated from the role for lagged, constant-horizon forecasts, while stickiness, λ , is identified from lagged, constant-target (*i.e.*, longer horizon) forecasts. Identification also should be aided by the “common factor” restriction, for there are three right-hand-side variables but only two parameters.

Notice that the system (24) consists of four equations, for $h = 1, 2, 3, 4$, and so has 12 coefficients yet involves only two parameters in the conditional means: ρ and λ . This degree of over-identification from the use of multiple horizons shows that one cannot necessarily easily reverse-engineer fitting the term structure of *SPF* forecasts by adding higher-order dynamics in the inflation gap. For example, modeling ϵ_t as an AR(4) process would add three parameters in its law of motion, but the analogue of the system (24) would still involve more coefficients than underlying parameters. Furthermore, higher-order dynamics in the inflation gap naturally also would imply higher-order dynamics in realized inflation. But Stock and Watson (2007) do not reject the null hypothesis that the IMA(1,1) model (implied by the UC model) holds against the alternative hypothesis of higher-order dynamics.

Overall, then, the findings are reminiscent of those derived from modeling the term structure of interest rates, where researchers find, for example, that a one-factor model of the short rate cannot fit both the persistence of that return and the average slope of the yield curve. Here we find that, so far, we cannot fit all of (a) inflation dynamics, (b) the properties of forecast errors, and (c) the term structure of professional inflation forecasts. However, our basic model with sticky forecasts (19) and the system (24) still have a singularity across horizons, so sub-section 5.3 will estimate the parameters using the Kalman filter—with measurement error—amended to allow sticky information.

5.3 Kalman Filtering with Sticky Forecasts

Next, we outline how to extract the inflation gap $\epsilon_{t|t}$ and hence also the trend $\tau_{t|t}$ uniquely from a set of possibly sticky forecasts. As in sub-section 4.3, anytime we see $F_t\pi_{t+h}$ we shall instead see $F_t\pi_{t+h} - \xi_{ht}$ and assume that these errors are uncorrelated across time and horizons. Allowing for measurement error at each horizon in Coibion and Gorodnichenko’s forecast-updating equation gives

$$F_t\pi_{t+h} = (1 - \lambda)E_t\pi_{t+h} + \lambda F_{t-1}\pi_{t+h} + \xi_{ht} - \lambda\xi_{h+1t-1}. \quad (25)$$

Thus, sticky information with measurement error adds a composite but serially uncorrelated disturbance to the average forecast. Using the rational-expectations forecasts (8)

from the UC model, $E_t\pi_{t+h} = \tau_{t|t} + \rho^h\epsilon_{t|t}$, thus gives

$$F_t\pi_{t+h} = (1 - \lambda)[\tau_{t|t} + \rho^h\epsilon_{t|t}] + \lambda F_{t-1}\pi_{t+h} + \xi_{ht} - \lambda\xi_{h+1t-1}. \quad (26a)$$

As in sub-section 4.3, we change the subscripts (replacing $\tau_{t|t}$ and $\epsilon_{t|t}$ by τ_t and ϵ_t , respectively) as a reminder that we (unlike the professional forecasters) have not yet filtered the components:

$$F_t\pi_{t+h} = (1 - \lambda)[\tau_t + \rho^h\epsilon_t] + \lambda F_{t-1}\pi_{t+h} + \xi_{ht} - \lambda\xi_{h+1t-1}. \quad (26b)$$

Again we want to annihilate τ_t but, unlike in the non-sticky case (13), subtracting actual inflation π_t no longer removes the trend because it has a weight of only $(1 - \lambda)$ in (26). Taking differences across adjacent horizons gives an MA(1) error in the observation equation. For example:

$$F_t\Delta\pi_{t+2} = (1 - \lambda)(\rho^2 - \rho)\epsilon_t + \lambda F_{t-1}\Delta\pi_{t+2} + \xi_{2t} - \lambda\xi_{3t-1} + \xi_{1t} - \lambda\xi_{2t-1}. \quad (27)$$

The trick to avoid this, yet still use all available information, is to take differences across *two* horizons. So start with

$$F_t(\pi_{t+2} - \pi_t) = (1 - \lambda)(\rho^2 - 1)\epsilon_t + \lambda[F_{t-1}(\pi_{t+2} - \pi_t)] + \xi_{2t} - \lambda\xi_{3t-1} - \xi_{0t} + \lambda\xi_{1t-1}, \quad (28)$$

where we subtract the nowcast. Notice that equation (28) has no dependence in the composite error term. Then there is a second observation equation:

$$F_t(\pi_{t+3} - \pi_{t+1}) = (1 - \lambda)(\rho^3 - \rho)\epsilon_t + \lambda[F_{t-1}(\pi_{t+3} - \pi_{t+1})] + \xi_{3t} - \lambda\xi_{4t-1} - \xi_{1t} + \lambda\xi_{2t-1}. \quad (29)$$

We do not have an observation equation for $F_t(\pi_{t+4} - \pi_{t+2})$ because we do not have data on $F_{t-1}\pi_{t+4}$, a five-step-ahead forecast; we lose a lead because of the stickiness. Thus the observation equations are

$$\begin{aligned} F_t(\pi_{t+h} - \pi_{t+h-2}) &= (1 - \lambda)(\rho^h - \rho^{h-2})\epsilon_t + \lambda[F_{t-1}(\pi_{t+h} - \pi_{t+h-2})] \\ &\quad + \xi_{ht} - \lambda\xi_{h+1t-1} - \xi_{h-2t} + \lambda\xi_{h-1t-1}, \end{aligned} \quad (30)$$

for $h = 2, 3$.

Finally, these observation equations appear to have an errors-in-variables problem because the lagged, longer-horizon forecasts now appear on the right-hand side. A solution is to adopt the method we used previously in sub-section 5.1: consider several candidate

estimates λ_0 from the Coibion-Gorodnichenko regressions. Then we rearrange to take all the forecasts to the left-hand side:

$$F_t(\pi_{t+h} - \pi_{t+h-2}) + \lambda_0 [F_{t-1}(\pi_{t+h} - \pi_{t+h-2})] = (1 - \lambda_0)(\rho^h - \rho^{h-2})\epsilon_t + \xi_{ht} - \lambda\xi_{h+1t-1} - \xi_{h-2t} + \lambda\xi_{h-1t-1}, \quad (31)$$

for $h = 2, 3$ and a candidate value λ_0 . Having only two observation equations, we can identify only two measurement error variances, so we relabel the combinations seen in (31):

$$F_t(\pi_{t+h} - \pi_{t+h-2}) + \lambda_0 [F_{t-1}(\pi_{t+h} - \pi_{t+h-2})] = (1 - \lambda_0)(\rho^h - \rho^{h-2})\epsilon_t + \omega_{ht}, \quad (32)$$

and estimate the variances of the two composite errors ω_{ht} (ignoring possible covariance) for $h = 2, 3$. Denote the standard error of the composite error ω_{ht} as σ_{ω_h} . The state-space version of the UC model, with sticky information and measurement error, consists of the AR(1) model of the inflation gap:

$$\epsilon_t = \rho\epsilon_{t-1} + v_t, \quad (33)$$

as the state equation, along with the two observation equations (32).

We estimate the state-space system of equations (32) and (33) using maximum likelihood. However, only the parameters ρ , σ_v , σ_{ω_2} , and σ_{ω_3} are estimated. The sticky-information coefficient λ is calibrated on the grid $\lambda_0 = [0.10, 0.95]$ spaced by increments of 0.05. The Kalman filter is applied to the state-space system of equations (32) and (33) to generate its log likelihood conditional on a value for λ_0 . Across the grid of λ_0 , the MLEs of ρ , σ_v , σ_{ω_2} , and σ_{ω_3} associated with the maximal log likelihood of the state-space system are reported below. Label the sticky-information coefficient that achieves the maximal log likelihood $\bar{\lambda}_0$. Robust standard errors are presented with the MLEs because the sticky-information parameter is calibrated and standard errors of composite errors are estimated. (Again, the appendix contains the details.)

Table 6 contains the results, using CPI forecasts since 1981. The best fit involves a calibrated value $\bar{\lambda}_0 = 0.10$, which is at the lower limit of the range of stickiness parameters we considered. There is significant persistence in the inflation gap, as $\hat{\rho} = 0.45$ with a standard error of 0.15. Also, $\epsilon_{t|t}$ exhibits little volatility relative to the composite shocks ω_2 , $\hat{\sigma}_{\omega_2}/\hat{\sigma}_v = 0.82$, but not ω_3 , $\hat{\sigma}_{\omega_3}/\hat{\sigma}_v = 0.22$. However, as in our basic example in table 5, there is considerable persistence in $\eta_{t|t}$. The Q -statistics in the lower panel of table 6 easily reject the null hypothesis that this series is white noise.

The upper panel of figure 7 shows the filtered trend and cycle. The trend is now much more volatile (and the cycle smoother) than in our previous applications. The lower panel of figure 7 then shows the innovations where, correspondingly, the innovation to the trend, $\eta_{t|t}$, is now more volatile than the innovation to the inflation gap, $v_{t|t}$. Figure 7 is perhaps suggestive of instability over time in variances. But if we sub-sample in three time periods, we have only 42 observations for the early sample, too few to allow reliable estimation. Sub-section 5.4 studies the Kalman filter applied to the GDP deflator where forecasts are available for a longer time span.

5.4 GDP Deflator Evidence

Table 7 contains estimation results for the filtered, sticky-information model using the GDP deflator since 1974. It presents estimates first for the entire sample from 1974 to 2012, then for two sub-samples with a break in 1985.

These results indicate that excluding the 1970s and early 1980s leads to the conclusion that there is little information stickiness in the *SPF* forecasts for inflation in the CPI and GDP deflator. The calibration of $\bar{\lambda}_0 = 0.10$ produces the maximal log likelihoods of the state-space model (32) and (33) on samples starting in the early to mid-1980s. Only when the sample for forecasts of the GDP deflator inflation rate begins in 1974Q4 is a calibration to sticky information, $\bar{\lambda}_0 = 0.95$, consistent with the data and state-space model.

Inflation-gap persistence is not sensitive in the same way to splitting the data between the Great Inflation and the Great Moderation. The persistence is estimated to be no smaller than 0.74, using PGDP inflation forecasts across the pre- and post-1985Q2 samples in table 7. In contrast, the persistence estimated from the CPI inflation data in table 6 was 0.45.

The lower panel of table 7 reports the *Q*-test statistics that measure the remaining persistence in the innovations $\eta_{t|t}$ and $v_{t|t}$. As with the CPI inflation rate, these have low *p*-values and so readily reject the null hypothesis of white noise, at conventional levels of significance. The post-1985 innovation moments (not shown in table 7) also strongly reject that null hypothesis.

The upper panel of figure 8 shows the realized inflation rate in the GDP deflator with the estimated trend and inflation gap. As in figure 7, allowing for sticky information leads to a more volatile inflation trend. The lower panel of figure 8 again shows that the shock to the trend is more volatile than the shock to the inflation gap, as was the case with the CPI data in figure 7. There also is evidence of changing volatility over time. We find that $\hat{\sigma}_v$ is much lower after 1985 (so are the two measurement-error variances). Specifically, $\hat{\sigma}_{\omega_2}/\hat{\sigma}_v$

$= 0.15$ and $\hat{\sigma}_{\omega_3}/\hat{\sigma}_v = 0.13$ for the entire sample of 1974Q4 to 2012Q2. For the 1974Q4 to 1985Q1 sub-sample, these ratios are 0.14 and 0.14. But for the post-1985 sub-sample, 1985Q2 to 2012Q2, the ratios are 0.49 and 0.35. Thus, the PGDP inflation forecast data suggest a structural break in the volatility of the shock innovations, rather than in inflation gap persistence, of the state-space model. Remember too that moving across the 1985Q1 split date on the PGDP inflation samples results in $\bar{\lambda}_0$ falling from 0.95 to 0.10.

6. Conclusions

In this paper, we show how to measure the unobserved components in US inflation using the *SPF*, without necessarily assuming that mean, reported forecasts coincide with conditional expectations (and without assuming anything about shock correlations). Both the UC model of inflation and recent descriptions of sticky forecasting restrict unobservable inflation forecasts $E_t\pi_{t+h}$. We combine these statistical models to provide a fast, inexpensive way to filter US inflation into trend and cycle components, with the trend component interpretable as long-term inflation expectations. It is interesting to see the parameter estimates for inflation-gap persistence (ρ) and for information stickiness (λ) implied by estimation with *SPF* forecast data, as well as the implied, historical shock volatilities.

The approach features consistency tests: We study whether we can reconcile the two statistical models jointly with the time-series properties of actual inflation, the term structure of professional forecasts, and mean forecast errors. We find that we cannot. The behavior of mean *SPF* forecasts over multiple horizons apparently cannot be viewed as consistent with the UC model, where trend inflation follows a martingale. The forecast stickiness that seems to be implied by forecast-error properties does not yield a trend-cycle decomposition with unpredictable innovations to the two components.

The unobserved-components model is widely used in forecasting and in reconstructing the history of US inflation. The rational expectations and sticky-information models are widely used in closing macroeconomic models. We hope that showing that they cannot easily be reconciled prompts further research on these important statistical models.

References

- Andrade, Philippe and Hervé Le Bihan (2010) Inattentive professional forecasters. Working paper 307, Banque de France.
- Ang, Andrew, Geert Bekaert, and Min Wei (2007) Do macro variables, asset markets, or surveys forecast inflation better? *Journal of Monetary Economics* 54, 1163-1212.
- Beveridge, Stephen and Charles R. Nelson (1981) A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the “business cycle.” *Journal of Monetary Economics* 7, 151–174.
- Canova, Fabio (2007) *Methods for Applied Macroeconomic Research*. Princeton University Press.
- Carroll, Christopher D. (2003) Macroeconomic expectations of households and professional forecasters. *Quarterly Journal of Economics* 118, 269–298.
- Capistrán, Carlos and Allan Timmermann (2009) Disagreement and biases in inflation expectations. *Journal of Money, Credit and Banking* 41, 365–396.
- Clark, Todd E. and Troy Davig (2011) Decomposing the declining volatility of long-term inflation expectations. *Journal of Economic Dynamics and Control* 35, 981-999.
- Cogley, Timothy and Thomas J. Sargent (2005) Drifts and volatilities: Monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics* 8, 262–302.
- Cogley, Timothy and Argia Sbordone (2008) Trend inflation, indexation, and inflation persistence in the New Keynesian Phillips curve. *American Economic Review* 98, 2101–2126.
- Cogley, Timothy, Giorgio E. Primiceri, and Thomas J. Sargent (2010) Inflation-gap persistence in the U.S. *American Economic Journal: Macroeconomics* 2, 43–69.
- Coibion, Olivier (2010) Testing the sticky information Phillips curve. *Review of Economics and Statistics* 92, 87–101.
- Coibion, Olivier and Yuriy Gorodnichenko (2011) Information rigidity and the expectations formation process: A simple framework and new facts. Mimeo.
- Coibion, Olivier and Yuriy Gorodnichenko (2012) What can survey forecasts tell us about informational rigidities? *Journal of Political Economy* 120, 116–159.
- Creal, Drew (2012) A survey of sequential Monte Carlo methods for economics and finance. *Econometric Reviews* 31, 245-296.
- Croushore, Dean (2010) An evaluation of inflation forecasts from surveys using real-time data. *B.E. Journal of Macroeconomics, Contributions* 10.
- Del Negro, Marco and Stefano Eusepi (2011) Fitting observed inflation expectations. *Journal of Economic Dynamics and Control* 35, 2105–2131.

- Devereux, Michael B. and James Yetman (2003) Predetermined prices and the persistent effects of money on output. *Journal of Money, Credit and Banking* 35, 729–741.
- Edge, Rochelle and Refet S. Gürkaynak (2010) How useful are estimated DSGE model forecasts for central bankers? *Brookings Papers on Economic Activity* 2, 209–244.
- Faust, Jon and Jonathan H. Wright (2013) Forecasting inflation. In Graham Elliot and Allan Timmermann (eds.), *Handbook of Economic Forecasting*, vol. 2, 2–56. New York, NY: Elsevier.
- Gil-Alana, Luis, Antonio Moreno, and Fernando Pérez de Gracia (2012) Exploring survey-based inflation forecasts. *Journal of Forecasting* 6, 524–539.
- Grassi, Stefano and Tommaso Proietti (2010) Has the volatility of U.S. inflation changed and how? *Journal of Time Series Econometrics* 2:1; article 6.
- Henzel, Steffen R. (2013) Fitting survey expectations and uncertainty about trend inflation. *Journal of Macroeconomics* 35, 172–185.
- Ireland, Peter N. (2007) Changes in the Federal Reserve’s inflation target: Causes and consequences. *Journal of Money, Credit and Banking* 39, 1851–1882.
- Jain, Monica (2011) Perceived inflation persistence. Mimeo, Department of Economics, Queen’s University.
- Keane, Michael P. and David E. Runkle (1990) Testing the rationality of price forecasts: New evidence from panel data. *American Economic Review* 80, 714–735.
- Khan, Hashmat and Zhenhua Zhu (2006) Estimates of the sticky-information Phillips curve for the United States. *Journal of Money, Credit and Banking* 38, 195–207.
- Kiley, Michael T. (2007) A quantitative comparison of sticky price and sticky information models of price setting. *Journal of Money, Credit and Banking* 39, 101–125.
- Kozicki, Sharon and P.A. Tinsley (2012) Effective use of survey information in estimating the evolution of expected inflation. *Journal of Money, Credit and Banking* 44, 145–169.
- Krane, Spencer D. (2011) Professional forecasters’ views of permanent and transitory shocks to GDP. *American Economic Journal: Macroeconomics* 3, 184–211.
- Mankiw, N. Gregory and Ricardo Reis (2002) Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve. *Quarterly Journal of Economics* 117, 1295–1328.
- Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers (2004) Disagreement about inflation expectations. *NBER Macroeconomics Annual* 2003 209–248.
- Mertens, Elmar (2011) Measuring the level and uncertainty of trend inflation. Finance and Economics Discussion Series 2011-42, Federal Reserve Board.
- Morley, James C., Charles R. Nelson, and Eric Zivot (2003) Why are the Beveridge-Nelson and unobserved-components decompositions of GDP so different? *Review of Economics and Statistics* 85, 235–243.

- Nason, James M. (2006) Instability in US inflation: 1967–2005. Federal Reserve Bank of Atlanta Quarterly Review 91 (Q2) 39–59.
- Patton, Andrew J. and Allan Timmermann (2010) Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion. *Journal of Monetary Economics* 57, 803–820.
- Reis, Ricardo (2006) Inattentive producers. *Review of Economic Studies* 73, 793–821.
- Shephard, Neil (2013) Martingale unobserved component models. Economics Series Working Papers 644, University of Oxford, Department of Economics.
- Stock, James H. and Mark W. Watson (2007) Why has US inflation become harder to forecast? *Journal of Money, Credit and Banking* 39 (S1), 3–33.

Appendix: The Kalman Filter and MLE of SPF Inflation Term Structure Models

The paper reports maximum likelihood estimates of the rational expectations term structure model of SPF inflation in tables 3 and 4. These estimates are based on the Kalman filter implementation of the prediction error decomposition of the likelihood function of the unobserved-components (UC) model. The Kalman filter is generated by a system of state and observation equations. Canova (2007) provides a description and details of the method. The observation equations are built on equation (13) of the paper:

$$F_t \pi_{t+h} - \pi_t = (\rho^h - 1) \epsilon_t + \xi_{ht}.$$

Equation (13) is the result of subtracting π_t from equation (11b):

$$F_t \pi_{t+h} = \tau_t + \rho^h \epsilon_t + \xi_{ht},$$

with $\xi_{ht} \sim N(0, \sigma_{\xi_h}^2)$. The *SPF* provides an inflation term structure for $h = 1, 2, 3, 4$, so the system of observation equations consists of

$$\begin{bmatrix} F_t \pi_{t+1} - \pi_t \\ F_t \pi_{t+2} - \pi_t \\ F_t \pi_{t+3} - \pi_t \\ F_t \pi_{t+4} - \pi_t \end{bmatrix} = \begin{bmatrix} (\rho - 1) \\ (\rho^2 - 1) \\ (\rho^3 - 1) \\ (\rho^4 - 1) \end{bmatrix} \epsilon_t + \begin{bmatrix} \xi_{1t} \\ \xi_{2t} \\ \xi_{3t} \\ \xi_{4t} \end{bmatrix}, \quad (A.1)$$

with $E(\xi_{jt} \xi_{it}) = 0$ for $j \neq i$. Equation (7) of the paper is the AR(1) transitory component:

$$\epsilon_t = \rho \epsilon_{t-1} + v_t,$$

with $\rho \in (-1, 1)$ and $v_t \sim N(0, \sigma_v^2)$, which serves as the state equation of the rational expectations term structure model of SPF inflation. Given this state equation and the system of observation equations (A.1), Kalman filter predictions of $\epsilon_{t|t}$ are generated along with the prediction errors of the observation vector:

$$\begin{bmatrix} F_t \pi_{t+1} - \pi_t \\ F_t \pi_{t+2} - \pi_t \\ F_t \pi_{t+3} - \pi_t \\ F_t \pi_{t+4} - \pi_t \end{bmatrix} - \begin{bmatrix} (\rho - 1) \\ (\rho^2 - 1) \\ (\rho^3 - 1) \\ (\rho^4 - 1) \end{bmatrix} \epsilon_{t|t}. \quad (A.2)$$

The likelihood is built on $\epsilon_{t|t}$, its mean squared errors, the prediction errors, and the $\sigma_{\xi_h}^2$, $h = 1, 2, 3, 4$. The log likelihood is maximized using the optimizer `csmnwel` of Chris Sims.

There also is a state-space representation of the sticky-information (SI) version of the UC model with measurement error. First, the paper develops equation (30):

$$F_t(\pi_{t+h} - \pi_{t+h-2}) = (1 - \lambda)(\rho^h - \rho^{h-2})\epsilon_t + \lambda[F_{t-1}(\pi_{t+h} - \pi_{t+h-2})] \\ + \xi_{ht} - \lambda\xi_{h+1t-1} - \xi_{h-2t} + \lambda\xi_{h-1t-1},$$

as the observation equations of this state-space model. The state equation of the sticky-information model with measurement error is again the AR(1) characterization of the inflation gap, $\epsilon_t = \rho\epsilon_{t-1} + v_t$. Given the horizons of inflation forecasts available in the *SPF*, the observation equations are the bivariate system for $h = 2$ and 3:

$$\begin{bmatrix} F_t(\pi_{t+2} - \pi_t) \\ F_t(\pi_{t+3} - \pi_{t+1}) \end{bmatrix} - \lambda_0 \begin{bmatrix} F_{t-1}(\pi_{t+2} - \pi_t) \\ F_{t-1}(\pi_{t+3} - \pi_{t+1}) \end{bmatrix} = (1 - \lambda_0)(\rho^2 - 1) \begin{bmatrix} 1 \\ \rho \end{bmatrix} \epsilon_t + \begin{bmatrix} \omega_{2t} \\ \omega_{3t} \end{bmatrix}, \quad (A.3)$$

where λ_0 indicates that the sticky-information parameter is calibrated rather than estimated, and:

$$\omega_{ht} \equiv \xi_{ht} - \lambda\xi_{h+1t-1} - \xi_{h-2t} + \lambda\xi_{h-1t-1},$$

for $h = 2, 3$. Denote the standard deviation of the composite error ω_{ht} as σ_{ω_h} , where $E(\omega_{2t}\omega_{3t}) = 0$.

The state-space system of the AR(1) model of ϵ_t and equations (A.3) is estimated using maximum likelihood. However, only the parameters ρ , σ_v , σ_{ω_2} , and σ_{ω_3} are estimated. The SI coefficient λ is calibrated on the grid $\lambda_0 = [0.10, 0.95]$ spaced by increments of 0.05. The Kalman filter is applied to the system to generate its log likelihood conditional on a value for λ_0 . Across the grid of λ_0 , the MLEs of ρ , σ_v , σ_{ω_2} , and σ_{ω_3} associated with the maximal log likelihood of the state-space system are calculated, again using Chris Sims' optimizer `csminwel`. Robust standard errors are presented with the MLEs because the SI coefficient is calibrated and standard deviations of composite measurement errors are estimated. This task is assigned to the square roots of the diagonal of the sandwich covariance estimator $[\mathcal{H}(\mathcal{G}\mathcal{G}')^{-1}\mathcal{H}]^{-1}$, where $\mathcal{H}$ is the Hessian of the log likelihood function and $\mathcal{G}$ is its Jacobian, and both are evaluated at the MLEs conditional on $\bar{\lambda}_0$, which is the SI coefficient that achieves the maximal log likelihood. Tables 6 and 7 contain the results.

**Table 1: Mean One-Step-Ahead Forecasts as Rational Expectations
Trend and Cycle Innovation Moments**

$$\tau_{t|t} = F_t \pi_{t+1}$$

Time Span	Statistic	$\epsilon_{t t}$	$\eta_{t t}$
1981Q3–2012Q2	s^2	3.02	0.24
	$r(\epsilon, \eta)$	0.35	
	$Q(4)$ (p)	5.43 (0.25)	8.76 (0.06)
	$Q(8)$ (p)	6.46 (0.60)	13.9 (0.09)
1981Q3–1991Q1	s^2	2.71	0.46
	$r(\epsilon, \eta)$	0.35	
1991Q2–2002Q2	s^2	0.72	0.04
	$r(\epsilon, \eta)$	0.15	
2002Q3–2012Q2	s^2	5.59	0.27
	$r(\epsilon, \eta)$	0.49	

Notes: ϵ and η are the estimated innovations to cycle and trend, respectively, and s^2 and r are their sample variance and correlation. $Q(j)$ is the Ljung-Box statistic with j lags. $F_t \pi_{t+1}$ is the mean *SPF* forecast for the annualized CPI inflation rate, one quarter ahead.

**Table 2: Mean Multi-Step Forecasts as Rational Expectations
Trend and Cycle Innovation Moments**

$$\tau_{t|t} = \frac{1}{4} \sum_{h=1}^4 F_t \pi_{t+h}$$

Time Span	Statistic	$\epsilon_{t t}$	$\eta_{t t}$
1981Q3–2012Q2	s^2	3.47	0.10
	$r(\epsilon, \eta)$		0.35
	$Q(4)$ (p)	6.32 (0.18)	8.24 (0.08)
	$Q(8)$ (p)	7.10 (0.52)	15.6 (0.05)
1981Q3–1991Q1	s^2	3.27	0.24
	$r(\epsilon, \eta)$		0.35
1991Q2–2002Q2	s^2	0.79	0.03
	$r(\epsilon, \eta)$		0.13
2002Q3–2012Q2	s^2	6.34	0.05
	$r(\epsilon, \eta)$		0.49

Notes: ϵ and η are the estimated innovations to cycle and trend, respectively, and s^2 and r are their sample variance and correlation. $Q(j)$ is the Ljung-Box statistic with j lags. $F_t \pi_{t+1}$ is the mean *SPF* forecast for the annualized CPI inflation rate, averaged over one to four quarters ahead.

**Table 3: Kalman Filtering with a Persistent Inflation Gap
CPI Evidence: 1981Q3 to 2012Q2**

$$\epsilon_t = \rho\epsilon_{t-1} + v_t$$

$$F_t\pi_{t+h} - \pi_t = (\rho^h - 1)\epsilon_t + \xi_{ht} \quad h = 1, 2, 3, 4$$

Maximum Likelihood Estimates

$\hat{\rho}$	$\hat{\sigma}_v$	$\hat{\sigma}_{\xi_1}$	$\hat{\sigma}_{\xi_2}$	$\hat{\sigma}_{\xi_3}$	$\hat{\sigma}_{\xi_4}$
0.13 (0.01)	1.90 (0.12)	0.29 (0.02)	0.18 (0.01)	0.04 (0.03)	0.15 (0.01)

Notes: Estimation uses actual CPI inflation and mean *SPF* forecasts at horizon h . Standard errors are in brackets.

Trend and Cycle Innovation Moments

Statistic	$v_{t t}$	$\eta_{t t}$
s^2	3.51	0.10
$r(\epsilon, \eta)$		0.11
$Q(4)$ (p)	3.60 (0.46)	10.56 (0.03)
$Q(8)$ (p)	6.19 (0.62)	18.7 (0.02)

Notes: v and η are the estimated innovations to cycle and trend, respectively, and s^2 and r are their sample variance and correlation. $Q(j)$ is the Ljung-Box statistic with j lags.

**Table 4: Kalman Filtering with a Persistent Inflation Gap
GDP Deflator Evidence: 1974Q4 to 2012Q2**

$$\begin{aligned}\epsilon_t &= \rho\epsilon_{t-1} + v_t \\ F_t\pi_{t+h} - \pi_t &= (\rho^h - 1)\epsilon_t + \xi_{ht} \quad h = 1, 2, 3, 4\end{aligned}$$

Maximum Likelihood Estimates

$\hat{\rho}$	$\hat{\sigma}_v$	$\hat{\sigma}_{\xi_1}$	$\hat{\sigma}_{\xi_2}$	$\hat{\sigma}_{\xi_3}$	$\hat{\sigma}_{\xi_4}$
0.26 (0.03)	1.14 (0.07)	0.33 (0.02)	0.19 (0.02)	0.26 (0.02)	0.36 (0.03)

Notes: Estimation uses actual GDP deflator inflation and mean *SPF* forecasts at horizon h . Standard errors are in brackets.

Trend and Cycle Innovation Moments

	$v_{t t}$	$\eta_{t t}$
s^2	1.06	0.13
$r(\epsilon, \eta)$		-0.17
$Q(4)$ (p)	24.8 (0.00)	9.37 (0.05)
$Q(8)$ (p)	49.1 (0.00)	15.2 (0.06)

Notes: v and η are the estimated innovations to cycle and trend, respectively, and s^2 and r are their sample variance and correlation. $Q(j)$ is the Ljung-Box statistic with j lags.

**Table 5: Trend and Cycle Moments with Sticky Forecasts
CPI 1981:3 to 2012:2**

$$\tau_{t|t} = \frac{1}{4} \sum_{h=1}^4 \frac{F_t \pi_{t+h} - \lambda F_{t-1} \pi_{t+h}}{(1 - \lambda)}$$

	$\lambda = 0.2$		$\lambda = 0.4$	
	$\epsilon_{t t}$	$\eta_{t t}$	$\epsilon_{t t}$	$\eta_{t t}$
s^2	2.57	0.27	2.79	0.32
$r(\epsilon, \eta)$	0.51		0.37	
$Q(4)$ (p)	6.63 (0.17)	18.9 (0.00)	5.78 (0.21)	26.9 (0.00)
$Q(8)$ (p)	7.95 (0.43)	25.8 (0.00)	9.34 (0.31)	36.4 (0.00)

Notes: ϵ and η are the estimated innovations to cycle and trend, respectively, and s^2 and r are their sample variance and correlation. $Q(j)$ is the Ljung-Box statistic with j lags. $F_t \pi_{t+h}$ is the mean *SPF* forecast for the annualized CPI inflation rate, h quarters ahead.

**Table 6: Kalman Filtering with Sticky Information
CPI Evidence: 1981Q4 to 2012Q3**

$$\epsilon_t = \rho\epsilon_{t-1} + v_t$$

$$F_t(\pi_{t+h} - \pi_{t+h-2}) + \lambda_0[F_{t-1}(\pi_{t+h} - \pi_{t+h-2})] = (1 - \lambda_0)(\rho^h - \rho^{h-2})\epsilon_t + \omega_{ht} \quad h = 2, 3$$

Maximum Likelihood Estimates

$\ln L$	$\bar{\lambda}_0$	$\hat{\rho}$	$\hat{\sigma}_v$	$\hat{\sigma}_{\omega_2}$	$\hat{\sigma}_{\omega_3}$
-42.9	0.10	0.45 (0.15)	0.83 (0.16)	0.68 (0.08)	0.19 (0.03)

Notes: Estimation uses mean *SPF* forecasts for CPI inflation at horizon h . Standard errors are in brackets.

Trend and Cycle Innovation Moments

	$v_{t t}$	$\eta_{t t}$
s^2	0.57	3.72
$r(\epsilon, \eta)$		0.26
$Q(4)$ (p)	3.58 (0.47)	24.53 (0.00)
$Q(8)$ (p)	11.88 (0.16)	29.79 (0.00)

Notes: v and η are the estimated innovations to cycle and trend, respectively, and s^2 and r are their sample variance and correlation. $Q(j)$ is the Ljung-Box statistic with j lags.

**Table 7: Kalman Filtering with Sticky Information
GDP Deflator Evidence: 1974Q4 to 2012Q2**

$$\epsilon_t = \rho\epsilon_{t-1} + v_t$$

$$F_t(\pi_{t+h} - \pi_{t+h-2}) + \lambda_0[F_{t-1}(\pi_{t+h} - \pi_{t+h-2})] = (1 - \lambda_0)(\rho^h - \rho^{h-2})\epsilon_t + \omega_{ht} \quad h = 2, 3$$

Maximum Likelihood Estimates

Time Period	$\ln L$	$\bar{\lambda}_0$	$\hat{\rho}$	$\hat{\sigma}_v$	$\hat{\sigma}_{\omega_2}$	$\hat{\sigma}_{\omega_3}$
1974Q4–2012Q2	-82.17	0.95	0.88 (0.07)	3.30 (1.41)	0.48 (0.02)	0.43 (0.02)
1974Q4–1985Q1	-17.6	0.95	0.83 (0.08)	5.13 (2.62)	0.73 (0.05)	0.72 (0.03)
1985Q2–2012Q2	-163.8	0.10	0.74 (0.02)	0.53 (0.10)	0.26 (0.01)	0.18 (0.01)

Notes: Estimation uses mean *SPF* forecasts for GDP deflator inflation at horizon h . Standard errors are in brackets.

**Trend and Cycle Innovation Moments
1974Q4 to 2012Q2**

	$v_{t t}$	$\eta_{t t}$
s^2	0.59	1.30
$r(\epsilon, \eta)$		-0.41
$Q(4)$ (p)	10.17 (0.04)	36.95 (0.00)
$Q(8)$ (p)	18.62 (0.02)	45.64 (0.00)

Notes: v and η are the estimated innovations to cycle and trend, respectively, and s^2 and r are their sample variance and correlation. $Q(j)$ is the Ljung-Box statistic with j lags.

Figure 1: US CPI Inflation and Mean Forecasts

Note: Inflation is the annualized quarterly growth rate in the CPI, `cpiaucsl`. Forecast series are the means of the zero- to four-step-ahead forecasts of US quarterly CPI inflation, series `CPI2` to `CPI6` from `Mean_CPI_Level.xls` from the *Survey of Professional Forecasters*. The colors change from dark red to dark blue as the horizon increases.

Figure 2: US GDP Deflator Inflation and Mean Forecasts

Notes: Forecasts are from `mean_PGDP_Growth.xls` in the SPF. Realized inflation is from `Data_SPF_Error_Statistics_PGDP_3_AIC.xls`. Inflation is measured in the GNP deflator prior to 1992, the GDP deflator from 1992 to 1995, and the chain-weighted GDP deflator since.

Figure 3: CPI Inflation, Trend, and Innovations

Note: In the upper panel the solid black line is the realized CPI inflation rate while the trend is the mean, one-quarter-ahead forecast. In the lower panel the solid line is the innovation to the trend, η , while the dashed line is the inflation gap, ε .

Figure 4: CPI Inflation and Kalman-Filtered Components

Notes: The upper panel shows CPI inflation (the solid line), the Kalman-filtered trend (the long-dashed line), and the inflation gap (the short-dashed line). The lower panel shows the innovation to the trend (η) and to the gap (v).

Figure 5: GDP Deflator Inflation and Kalman-Filtered Components

Note: The upper panel shows GDP deflator inflation (the solid line), the Kalman-filtered trend (the long-dashed line), and the inflation gap (the short-dashed line). The lower panel shows the innovations to the trend (η) and to the cycle (v).

Figure 6: CPI Inflation and Sticky-Forecast Trends

Note: The upper panel shows CPI inflation (the solid line) and trends formed from the basic, sticky-forecast model with two values of λ . The lower panel shows the innovations to the trend (η) and cycle (ε) with $\lambda=0.4$.

Figure 7: Kalman-Filtered CPI Inflation with Sticky Forecasts

Note: The upper panel shows CPI inflation (the solid line), its trend (the long-dashed line) and its cycle (the short-dashed line). The lower panel shows the innovations to the trend (η) and cycle (v).

Figure 8: Kalman-Filtered GDP Deflator Inflation with Sticky Forecasts

Note: In the upper panel the solid line is actual inflation in the GDP deflator, the long-dashed line is the trend, and the short-dashed line is the cycle. In the lower panel the long-dashed line is the innovation to the trend (η) and the short-dashed line is the innovation to the cycle (v).