Sticky Professional Forecasts and the Unobserved Components Model of US Inflation

James M. Nason and Gregor W. Smith[†]

October 2016

Abstract

Much recent research studies US inflation history with a trend-cycle model with unobserved components (UC) in which the trend coincides with long-horizon inflation expectations. We show that this trend can be isolated using forecasts from the *Survey of Professional Forecasters* even though these forecasts are (a) at short horizons and (b) biased, as described by the sticky information (SI) model. We develop and apply tests of the joint UC-SI model, applicable at any horizon, allowing for stochastic volatility, and without restricting the trend-cycle covariance. The tests find considerable evidence of implicit sticky information. But they show we cannot reconcile these two widely used perspectives on US inflation and professional forecasts.

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, Canada; smithgw@econ.queensu.ca.
We thank the Social Sciences and Humanities Research Council of Canada, the Bank of Canada research fellowship programme, and the Jenkins Family Economics Fund of North Carolina State University for support of this research. For helpful comments we thank Drew Creal, Elmar Mertens, Morten Nielsen, Jonathan Wright, and numerous seminar and conference participants. 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 often (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. Faust and Wright (2013), in their review of inflation forecasting, list the many studies that feature a slowly evolving trend. This decomposition also sheds light on inflation history. Stock and Watson (2007) estimate a UC model to isolate changes in the variances of the trend and transitory components and hence in the overall persistence and forecastability of inflation since the early 1950s. Given its use in forecasting, the UC model is a natural way to measure inflation expectations at long horizons. Those expectations—and how well-anchored they are—are of great interest to monetary policymakers, because expected inflation serves as an indicator of the Fed's credibility as well as a constraint on the effect of policy.

The UC model also restricts forecasts at shorter horizons. In principle, then, one can use use short-horizon forecasts from the *Survey of Professional Forecasters* (*SPF*) to help measure the two components of inflation. However, 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 evidence about rational expectations. One way to describe the evidence is that forecasts are sticky and can be modeled using the sticky-information (SI) framework of Mankiw and Reis (2002). This 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 (2015) use the SI model to link reported forecasts to the actual conditional expectations of professional forecasters. We employ this model to characterize mean forecasts.

It seems promising to combine these insights from the study of professional forecasts with the successful track record of the UC model and so measure long-horizon expectations of inflation. But there is a catch. The UC model of inflation and the SI model of reported, *SPF* forecasts restrict unobservable, rational expectations of in inflation. Combining these models yields restrictions on the properties of inflation and reported forecasts. We provide tests of these restrictions, which are tests of the joint hypothesis that the UC and SI models hold. If we find that they do, we have a simple and informative way to filter US inflation, by outsourcing much of the work to the participants in the *Survey of Professional Forecasters*. Our tests are applicable at any horizon, do not restrict the covariance between the trend and cycle components of inflation, and apply under stochastic volatility of those components.

The tests clearly reject the joint model. Extracting the trend using forecast data leads to a trend-cycle decomposition with a trend shock that is predictable or persistent, a finding that is inconsistent with the underlying assumptions of the UC model. 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 historical inflation and *SPF* forecasts. Thus our paper outlines test criteria and, given their results, a design challenge in this area of research.

2. The Trend-Cycle Model

The first element in our study is an unobserved-components (UC) model. Suppose that inflation, π_t, evolves as a sum of two components: a stochastic trend τ_t and a stationary component or 'inflation gap' 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. They may be correlated and may have time-varying volatilities.

The h-step-ahead, rational-expectations forecast of inflation is

$$E_t \pi_{t+h} = \tau_t, \tag{2}$$

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

$$E_t \pi_{t+\infty} = \tau_t, \tag{3}$$

so that the trend estimate also is the estimate of expected inflation at the infinite horizon. Watson (1986) and Morley, Nelson, and Zivot (2003) explain how the zero conditional mean property in the shocks is necessary for the trend property (2).

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 using the BN trend within a dynamic stochastic general equilibrium (DSGE) model. Cogley and Sbordone (2008) use a similar, stochastic trend around which to estimate a New Keynesian Phillips curve while Ascari and Sbordone (2014) survey this approach and outline its implications for monetary policy.

The UC model combined with rational expectations 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 generalize the basic model (which assumes that gap inflation has no persistence) by allowing the stationary component of inflation, ϵ_t, to follow an AR(1) process:

$$\epsilon_t = \rho \epsilon_{t-1} + \upsilon_t, \tag{4}$$

where υ_t is a martingale-difference series and the persistence parameter $\rho \in (-1, 1)$ is constant. The infinite-horizon inflation forecast remains τ_t, but in general

$$E_t \pi_{t+h} = \tau_t + \rho^h \epsilon_t, \tag{5}$$

so that there now is a term structure of rational-expectations inflation forecasts. 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 (5) allow for these patterns. Clark and Doh (2014) examine the forecast performance of various models of trend inflation, including the UC model with or without a persistent inflation gap, finding that is is difficult to distinguish between them.

Estimation and forecasting with the UC model usually 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 those of Stock and Watson (2007) and Mertens (2015). Mertens uses *SPF* forecasts as one information source in the filter. Grassi and Proietti (2010) and Creal (2012) estimate the UC model with stochastic volatility.

In principle, though, one can measure the components using only 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. Kozicki and Tinsley (2012) estimate the parameters of a UC 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 shifts in trend inflation, τ_t, for reliable 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 also provide a test of the combined statistical model's ability to fit observed survey forecasts. Chan, Clark, and Koop (2015) use long-run forecasts from the *SPF* to help isolate trend inflation in the UC model. They allow for a flexible, time-varying role for those expectations. They conclude that the forecasts help isolate the trend but do not simply coincide with it.

Our study can be thought of as a sequel to that of Kozicki and Tinsley (2012). We use the *SPF* with a short horizon and consider the alternative, empirically supported assumption that forecasts are sticky. This extension might potentially reconcile the approach with the bias in mean *SPF* forecasts yet still allow one to use those forecasts to estimate the UC model and measure long-horizon inflation expectations. In practice, though, we find this reconciliation is not consistent with the joint properties of inflation and *SPF* forecasts.

3. Sticky Information

It is natural to also draw on professional forecasts to measure inflation expectations. 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*. 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. Nevertheless, 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."

The second element in our study is a description of forecast data. We work with the SI model because of its tractability in this application and also because it has been used both to describe professional forecasts—by Mankiw, Reis, and Wolfers (2004) and Coibion and Gorodnichenko (2012, 2015)—and to close and estimate macroeconomic models. 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 (2015), suppose that forecasters update their information with probability $1-\lambda$, so that λ measures the degree of stickiness in information. Let $F_t\pi_{t+h}$ denote the cross-forecaster mean forecast at time t for inflation h steps ahead. Coibion and Gorodnichenko show that this average forecast is a weighted average of the rational expectation and the previous period's mean, reported forecast:

$$F_t\pi_{t+h} = (1-\lambda)E_t\pi_{t+h} + \lambda F_{t-1}\pi_{t+h}. \tag{6}$$

The cross-forecaster mean coincides with the rational expectation of future inflation in the special case when $\lambda = 0$.

Define the non-sticky-information forecast error:

$$\vartheta_{t+h} = \pi_{t+h} - E_t\pi_{t+h}. \tag{7}$$

Subtracting each side of the pattern in reported, mean forecasts (6) from realized inflation gives

$$\begin{aligned}\pi_{t+h} - F_t\pi_{t+h} &= \lambda(E_t\pi_{t+h} - F_{t-1}\pi_{t+h}) + (\pi_{t+h} - E_t\pi_{t+h}) \\ &= \frac{\lambda}{1-\lambda}(F_t\pi_{t+h} - F_{t-1}\pi_{t+h}) + \vartheta_{t+h}.\end{aligned} \tag{8}$$

Because ϑ_{t+h} has the properties of an econometric error, this link (8) 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 (8) 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. Finally, they 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. In our estimation we do not restrict λ across horizons.

4. Tests of the UC-SI Model

Combining the forecast implication of the UC model with a persistent inflation gap (5) with the description of forecast updating (6) gives

$$\frac{F_t\pi_{t+h} - \lambda F_{t-1}\pi_{t+h}}{1-\lambda} = \tau_t + \rho^h\epsilon_t, \tag{9}$$

or

$$F_t\pi_{t+h} = \lambda F_{t-1}\pi_{t+h} + (1-\lambda)\tau_t + (1-\lambda)\rho^h\epsilon_t. \tag{10}$$

We 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. Thus:

$$F_t\Delta\pi_{t+h} = \lambda F_{t-1}\Delta\pi_{t+h} + (1-\lambda)(\rho^h - \rho^{h-1})\epsilon_t. \tag{11}$$

The inflation gap, ϵ_t, follows the AR(1) process (4) so that $\epsilon_t(1-\rho L) = \upsilon_t$. Multiplying the difference equations (11) by $(1-\rho L)$ gives links between the mean forecasts at different dates and horizons that are implied by the joint UC-SI model:

$$F_t\Delta\pi_{t+h} = \lambda F_{t-1}\Delta\pi_{t+h} + \rho F_{t-1}\Delta\pi_{t+h-1} - \rho\lambda F_{t-2}\Delta\pi_{t+h-1} + (1-\lambda)(\rho^h - \rho^{h-1})\upsilon_t. \quad (12)$$

The forecasts on the right-hand side are dated $t-1$ or earlier, so it is natural to assume that the inflation-gap shock, υ_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 (12) 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. However, we do not study restrictions on ρ and λ implied by the conditional variance in the forecast-only regression (12) because we do not assume that υ_t is homoskedastic or that its volatility follows a specific, parametric process.

The estimating equations consist of equations (8) and (12), with intercepts included, and estimated as a system for a given horizon, h:

$$\begin{aligned} F_t\Delta\pi_{t+h} &= \alpha_0 + \lambda F_{t-1}\Delta\pi_{t+h} + \rho F_{t-1}\Delta\pi_{t+h-1} - \rho\lambda F_{t-2}\Delta\pi_{t+h-1} + e_{1t} \\ \pi_{t+h} - F_t\pi_{t+h} &= \beta_0 + \frac{\lambda}{1-\lambda}(F_t\pi_{t+h} - F_{t-1}\pi_{t+h}) + e_{2t+h}, \end{aligned} \quad (13)$$

where α_0 and β_0 are constant terms and e_{1t} and e_{2t+h} are error terms. Under the UC model inflation and its forecasts have unit roots and the system (13) assumes that these variables share the same stochastic trend. We estimate the system (13) by nonlinear least squares, with the cross-equation restriction on $\hat{\lambda}$. There are four slopes in this system yet only two underlying parameters, λ and ρ. The two restrictions are (i) the 'common-factor' restriction across the coefficients derived in equation (12) and (ii) the cross-equation restriction on λ. We test the restrictions implied by the joint UC-SI model with a likelihood ratio test.

We next extract the unobserved components given estimates $\{\hat{\lambda}, \hat{\rho}\}$ at horizon h. Inverting equation (12) gives $\hat{\upsilon}_t$, while inverting equation (11) gives $\hat{\epsilon}_t$. Then the estimated trend is $\hat{\tau}_t = \pi_t - \hat{\epsilon}_t$. We then test two additional predictions. First, the estimated trend should yield a first-difference, $\hat{\eta}_t = \hat{\tau}_t - \hat{\tau}_{t-1}$, that is not autocorrelated. Second, the estimated innovation to the inflation gap, $\hat{\upsilon}_t$, also should not be autocorrelated.

To test these implications we use the Ljung-Box portmanteau test, modified to allow for time-varying volatility in the extracted components. West and Cho (1995) provide an example of the use of this test. Harris and Kew (2014) give further modifications, a history, and citations to studies on this statistic. For a time series $\{y_t\}$, let $\hat{\gamma}_j$, $j = 1, \ldots, k$ denote the sample autocorrelation at lag j. Define a heteroskedasticity-consistent variance estimator:

$$\hat{\nu}_j^2 = \frac{T^{-1} \sum_{t=1}^{T} y_{t-j}^2 y_t^2}{(T^{-1} \sum_{t=1}^{T} y_t^2)^2}. \tag{14}$$

The test statistic at lag length k then is:

$$Q_k^*(y_t) = T(T+2) \sum_{j=1}^{k} \frac{1}{T-j} \frac{\hat{\gamma}_j^2}{\hat{\nu}_j^2}, \tag{15}$$

which is $\sim \chi_k^2$ under the null hypothesis.

5. Data and Graphical Example

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 1981Q3 to 2016Q2, yielding 140 observations. The survey reports forecasts from zero (the nowcast) to four quarters ahead. The corresponding series for actual US inflation is 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. (This series corresponds to the most recent realization in the *SPF* error statistics files.)

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. However, this survey information has been collected only since 1991.

The upper panel of figure 1 shows actual inflation, π_t, and the mean, three-quarter-ahead forecast at the same date, $F_t\pi_{t+3}$. In the special case with no information stickiness ($\lambda = 0$) and no gap persistence ($\rho = 0$) this forecast is an estimate of the trend, $\hat{\tau}_t$, as equation (2) showed. Notice that the forecast data yield a smooth, stochastic trend (shown in grey) through realized inflation. Its difference from inflation is an estimate of the gap, $\hat{\epsilon}_t$. Much of the variation in inflation is attributed to the inflation gap.

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`. For the realized inflation rate 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 GDP deflator inflation quarterly, at annual rates. It also shows the mean, three-quarter-ahead forecast $F_t\pi_{t+3}$ (in grey), and the difference between the two series. The span of years, of course, now includes the high-inflation years of the 1970s.

These two panels illustrate the appeal of using mean forecasts as trend measures. In each case the forecast looks plausible as a real-time estimate of the slowly evolving trend in actual inflation. But these figures assume that there is no information stickiness, an assumption at odds with the properties of mean forecast errors. To informally show the effect on a measured trend of such stickiness, we next measure τ_t with equation (9), still with $\rho = 0$, using $h = 3$ (so that the formula involves three-quarter-ahead and four-quarter-ahead forecasts) and $\lambda = 0.5$. The lower panel of figure 1 shows the implied trend and inflation gap for the CPI inflation rate. The lower panel of figure 2 does the same for GDP deflator inflation. Notice that the measured trends now are much more volatile and do not look as plausible as slowly evolving means. The estimated inflation gap also has some added, high-frequency variation, reflecting a lack of correlation between $\hat{\tau}_t$ and π_t. This visual evidence suggests that there may be an inconsistency between the UC and SI models. We next study this issue while estimating both λ and ρ, using each possible horizon, and with formal statistical tests.

6. Estimates and Tests

Table 1 contains the results of estimating the system (13) and calculating the three test statistics proposed in section 4. The table includes results for each horizon and measure of inflation. There is no row in the table for $h = 4$, the longest horizon in the survey, because the estimating equations (13) use $F_{t-1}\Delta\pi_{t+h}$, which would require 5-quarter-ahead forecasts in that case.

The estimates of the stickiness parameter, $\hat{\lambda}$, lie between 0.35 and 0.55 (depending on the horizon and inflation measure) and each is statistically significant at the 1% level. The evidence on inflation-gap persistence is more nuanced. For CPI inflation both the value $\hat{\rho}$ and its precision are sensitive to the horizon. In contrast, for GDP deflator inflation there is a small, positive persistence coefficient in the inflation gap at each horizon, though that is not estimated very precisely. Clearly the most precise evidence is of a positive value for λ, which is why we focused on the effects of that parameter on the measured trend in our graphical examples in figures 1 and 2.

Table 1 also contains the likelihood ratio test statistics designed to jointly test the over-identifying restrictions. The p-values show that these restrictions implied by the joint UC-SI model are rejected at conventional significance levels, except that for the *SPF* predictions for CPI inflation at $h = 3$ the p-value is 0.06.

The next columns of table 1 show the test statistics for the two components, $Q^*_k(\hat{\eta}_t)$ and $Q^*_k(\hat{v}_t)$, with their asymptotic p-values, at lag length $k = 3$ quarters. The innovations to the inflation gap, $\hat{v}_t$, indeed appear to be innovations, with no evidence of significant persistence. But for the change in the trend, $\hat{\eta}_t$, and for any horizon or measure of inflation, the test statistics reject the hypothesis of no persistence, at the 5% level of significance. Thus, we can reject the combined UC-SI model with these additional consistency tests. The extracted trend does not have the random-walk property.

Reverse-engineering the trend-cycle model to fit the *SPF* forecasts and so pass these tests, for example by adding higher-order dynamics in the inflation gap, is not necessarily an easy task. 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 (12) 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.

The statistical tests reject the joint UC-SI model at any horizon, so they also would reject the model if the horizons were combined, for there would be even more over-identification in that case. We do not estimate the model for all horizons jointly because it has a singularity across horizons, tied to υ_t appearing in each equation. One approach would be to add measurement errors and extract the unobserved components using all horizons with the Kalman filter. Mertens and Nason (2015) and Nason and Smith (2016) explore that possibility. Their approach requires assumptions about the properties of the measurement errors, a restriction on the covariance between shocks, and a parameteric model of the stochastic volatilities of the shocks. We do not require a zero covariance between the shocks. Morley, Nelson, and Zivot (2003) show that restricting this covariance can have a large impact on trend-cycle decomposition. The tests also allow for time-varying volatility in $\{\eta_t, \upsilon_t\}$, but we do not need a parametric model of the variances.

7. Conclusion

Both the UC model of inflation and sticky forecasting restrict unobservable inflation forecasts $E_t\pi_{t+h}$. It thus seems promising to use professional forecasts (alone, or with other information) to help filter inflation, isolate its trend component, and thus measure long-horizon inflation expectations, a measure of monetary policy's credibility. This approach does not require us to assume that mean forecasts coincide with conditional expectations.

This paper derives and applies tests of the consistency of the UC and SI models. We study whether we can reconcile them jointly with the properties of professional forecasts and mean forecast errors. We find that we cannot. The over-identifying restrictions generally are rejected and the implied forecast stickiness does not yield a trend-cycle decomposition with unpredictable innovations to the trend component.

The UC model is widely used in forecasting and in reconstructing the history of US inflation. The SI treatment of expectations is widely used in closing macroeconomic models. We hope that showing that they cannot easily be reconciled prompts further research on these important statistical models.

One promising avenue may be to draw on research on alternatives to the SI 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 (2013) study individual forecasts from the European *SPF* and find that the SI model is not sufficient to capture all their patterns.

References

Andrade, Philippe and Hervé Le Bihan (2013) Inattentive professional forecasters. *Journal of Monetary Economics* 60, 967–982.

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.

Ascari, Guido and Argia Sbordone (2014) The macroeconomics of trend inflation. *Journal of Economic Literature* 52, 679–739.

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.

Capistrán, Carlos and Allan Timmermann (2009) Disagreement and biases in inflation expectations. *Journal of Money, Credit and Banking* 41, 365–396.

Chan, Joshua C.C., Todd Clark, and Gary Koop (2015) A new model of inflation, trend inflation, and long-run inflation expectations. mimeo

Clark, Todd E. and Taeyoung Doh (2014) Evaluating alternative models of trend inflation. *International Journal of Forecasting* 30, 426–448.

Cogley, Timothy and Argia Sbordone (2008) Trend inflation, indexation, and inflation persistence in the New Keynesian Phillips curve. *American Economic Review* 98, 2101–2126.

Coibion, Olivier (2010) Testing the sticky information Phillips curve. *Review of Economics and Statistics* 92, 87–101.

Coibion, Olivier and Yuriy Gorodnichenko (2012) What can survey forecasts tell us about informational rigidities? *Journal of Political Economy* 120, 116–159.

Coibion, Olivier and Yuriy Gorodnichenko (2015) Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review* 105, 2644–2678.

Creal, Drew (2012) A survey of sequential Monte Carlo methods for economics and finance. *Econometric Reviews* 31, 245–296.

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.

Harris, David and Hsein Kew (2014) Portmanteau autocorrelation tests under Q-dependence and heteroskedasticity. *Journal of Time Series Analysis* 35, 203–217.

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.

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.

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 (2015) Measuring the level and uncertainty of trend inflation. forthcoming in the *Review of Economics and Statistics.*

Mertens, Elmar and James M. Nason (2015) Inflation and professional forecast dynamics: An evaluation of stickiness, persistence, and volatility. mimeo,

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. and Gregor W. Smith (2016) Measuring the slowly evolving trend in US inflation with professional forecasts. mimeo, Department of Economics, North Carolina State University

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.

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.

Watson, Mark W. (1986) Univariate detrending methods with stochastic trends. *Journal of Monetary Economics* 18, 49–75.

West, Kenneth D. and Dongchul Cho (1995) The predictive ability of several models of exchange rate volatility. *Journal of Econometrics* 69, 367–391.

Table 1: Estimates and Tests

$$F_t\Delta\pi_{t+h} = \alpha_0 + \lambda F_{t-1}\Delta\pi_{t+h} + \rho F_{t-1}\Delta\pi_{t+h-1} - \rho\lambda F_{t-2}\Delta\pi_{t+h-1} + e_{1t}$$

$$\pi_{t+h} - F_t\pi_{t+h} = \beta_0 + \frac{\lambda}{1-\lambda}(F_t\pi_{t+h} - F_{t-1}\pi_{t+h}) + e_{2t+h}$$

	CPI				
Horizon	$\hat{\lambda}$ (*se*)	$\hat{\rho}$ (*se*)	$\chi^2(2)$ (*p*)	$Q_3^*(\hat{\eta}_t)$ (*p*)	$Q_3^*(\hat{\upsilon}_t)$ (*p*)
$h=3$	0.55***	-0.20	5.7	14.5	1.3
	(0.08)	(0.16)	(0.06)	(0.00)	(0.73)
$h=2$	0.48***	-0.14*	22.1	11.2	2.4
	(0.09)	(0.07)	(0.00)	(0.01)	(0.49)
$h=1$	0.35***	0.30***	33.6	10.1	3.1
	(0.09)	(0.09)	(0.00)	(0.02)	(0.38)

	GDP Deflator				
Horizon	$\hat{\lambda}$ (*se*)	$\hat{\rho}$ (*se*)	$\chi^2(2)$ (*p*)	$Q_3^*(\hat{\eta}_t)$ (*p*)	$Q_3^*(\hat{\upsilon}_t)$ (*p*)
$h=3$	0.35***	0.14*	12.9	10.5	0.4
	(0.05)	(0.08)	(0.00)	(0.01)	(0.95)
$h=2$	0.51***	0.04	8.6	14.5	3.6
	(0.10)	(0.09)	(0.01)	(0.00)	(0.31)
$h=1$	0.40***	0.13**	10.3	23.4	4.0
	(0.11)	(0.06)	(0.00)	(0.00)	(0.26)

Notes: Estimation uses 138 observations on CPI inflation and its mean SPF forecasts from 1981Q3 to 2016Q2 and 189 observations on GDP deflator inflation and its mean SPF forecasts from 1968Q4 to 2016Q2. Brackets contain Newey-West standard errors with 6 lags. The symbols *, **, and *** denote significance at the 10%, 5%, and 1% levels respectively. χ^2 is the likelihood ratio test statistic of the cross-equation and common-factor restrictions in the system (13). Q_3^* is the modified Ljung-Box portmanteau test statistic, distributed as χ_3^2 where 3 is the lag length in quarters.

Figure 1: CPI Inflation and Components

λ = 0

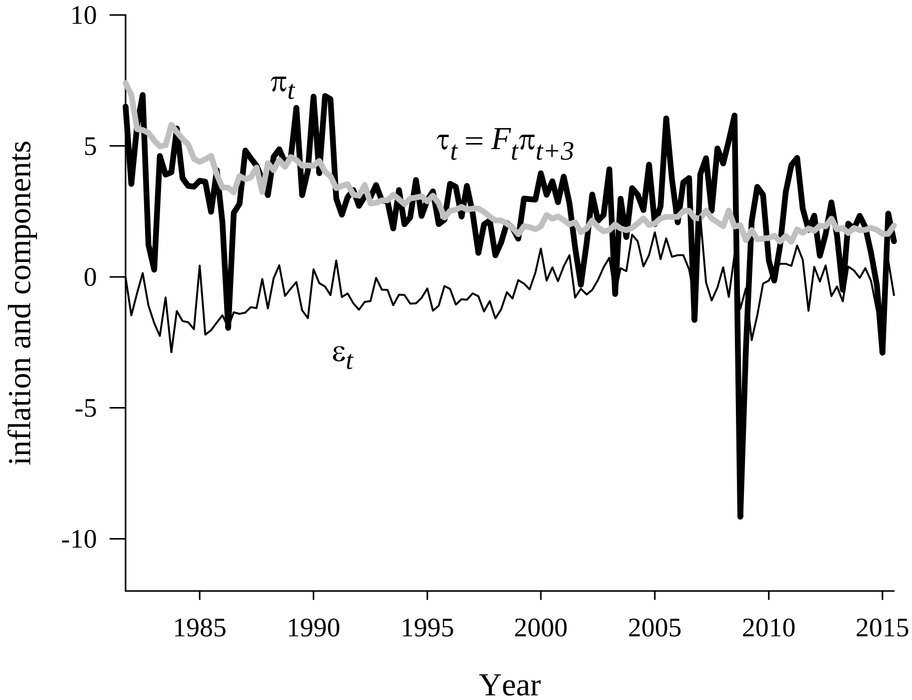

λ=0.5

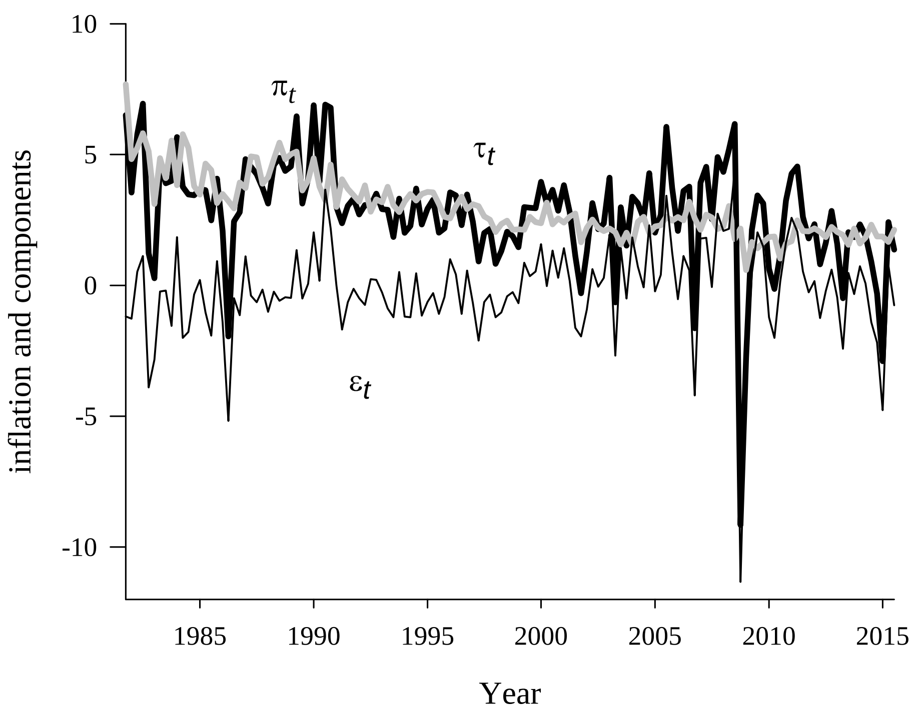

Note: In the upper panel π_t is the annualized rate of CPI inflation, measured quarterly; $F_t\pi_{t+3}$ is the mean, three-quarter-ahead forecast from the *SPF*, and ε_t is their difference. In the lower panel τ_t is the trend measured from equation (9) at $h=3$, $\rho=0$ and $\lambda=0.5$ and ε_t is the inflation gap.

Figure 2: GDP Deflator Inflation and Components

λ = 0

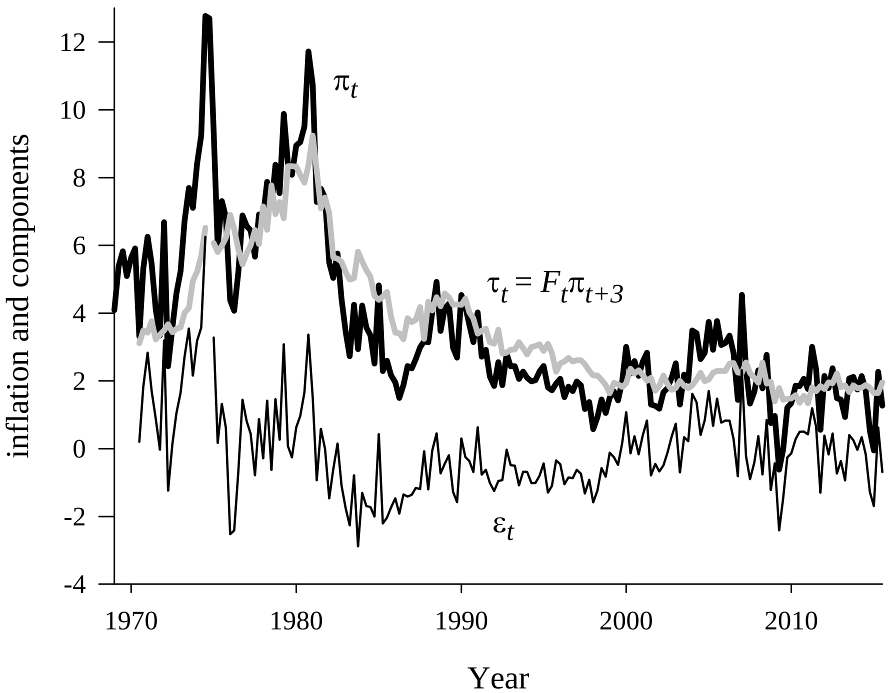

λ=0.5

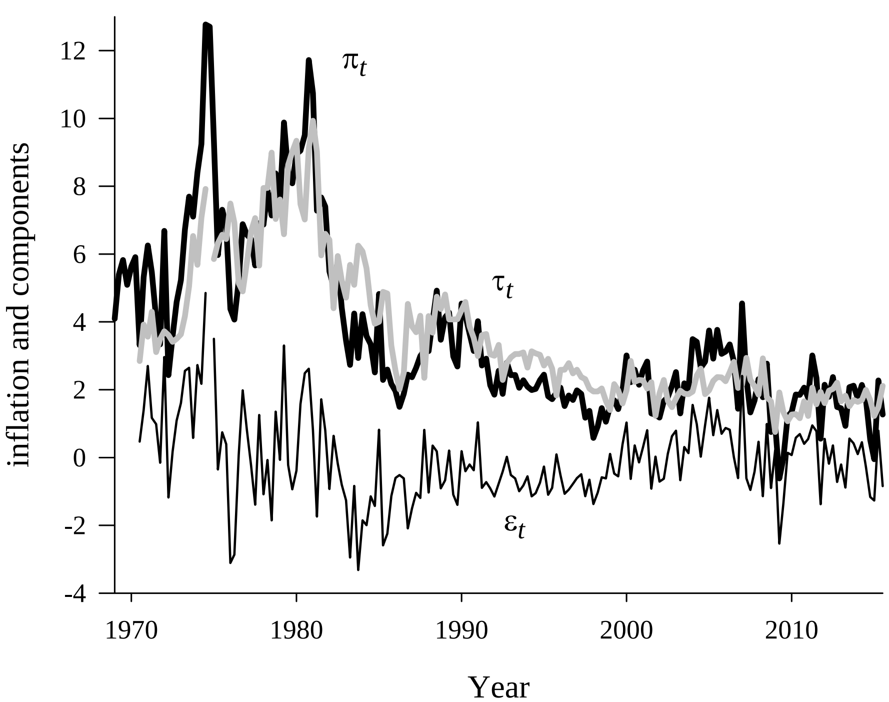

Note: In the upper panel π_t is the annualized, quarterly growth rate in the GDP deflator, $F_t\pi_{t+3}$ is the mean, three-quarter-ahead forecast from the *SPF*, and ε_t is their difference. In the lower panel τ_t is the trend measured with mean forecasts using equation (9) at $h=3$ with $\rho=0$ and $\lambda=0.5$ and ε_t is the inflation gap.