Reverse Kalman Filtering US Inflation with Sticky Professional Forecasts

[preliminary]

James M. Nason and Gregor W. Smith[†]

September 2012

Abstract

We provide a new way to filter US inflation into trend and cycle components, based on extracting long-run forecasts from the *Survey of Professional Forecasters*. We operate the Kalman filter in reverse, beginning with observed forecasts, then estimating parameters, then extracting the stochastic trend in inflation. The trend-cycle model is consistent with numerous studies of US inflation history, and is of interest in part because the trend may be viewed as the Fed's evolving inflation target. The sluggish reporting attributed to forecasters is consistent with evidence on mean forecast errors. We find little evidence of inflation-gap persistence, but considerable evidence of implicit, sticky information. But statistical tests show we cannot reconcile these two perspectives on US inflation forecasts, the unobserved-components model and the sticky information model.

JEL classification: E31, E37.

Keywords: Beveridge-Nelson, US inflation, trend, cycle

[†]Nason: Research Department, Federal Reserve Bank of Philadelphia; jim.nason@.phil.frb.org. Smith: Department of Economics, Queen's University; smithgw@econ.queensu.ca.
We thank the Social Sciences and Humanities Research Council of Canada and the Bank of Canada research fellowship programme for support of this research and seminar participants at the Canadian Economics Association meetings for comments. The views in this paper represent those of the authors alone and are not those of the Bank of Canada, the Federal Reserve Bank of Philadelphia, the Federal Reserve System, or any of its staff.

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 then sheds light on inflation history. 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. 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, a key indicator of the Fed's credibility as well as a constraint on the effect of policy.

Extracting the two unobserved components — the trend and the cycle in inflation — requires statistical assumptions on correlations and information: a filter. In turn, a filter allows estimation of parameters and then yields forecasts of inflation one step or many steps ahead. In this study we explore reversing this process. We begin with professional forecasts and show how they can be used to estimate the parameters of the UC model (with random-walk trend) and hence extract the components. Thus we start with the forecasts and end with a filter.

This approach requires a view on the connection between unobservable, h-step-ahead, inflation forecasts, denoted $E_t\pi_{t+h}$, and mean, reported, inflation forecasts, denoted $F_t\pi_{t+h}$, from the *Survey of Professional Forecasters.* We consider two possibilities. First, one way to extract information from the *SPF* is to assume that the 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 clustering or herding. One way to describe the evidence is that there is too much consensus given how inaccurate the forecasts are. Capistrán and Timmermann

(2009), Andrade and Le Bihan (2010), and Coibion and Gorodnichenko (2011) provide evidence of this pattern. How can professional forecasts be useful if they are biased? Precisely because the pattern of forecast clustering is systematic, it provides information on true expectations. Coibion and Gorodnichenko suggest a parametric model of stickiness in reported forecasts that allows us to link them to actual conditional expectations. We estimate the parameter describing stickiness along with those of the UC model. Under either assumption, then, one can link 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 not (or the extent to which they are sticky). Consistency implies a sort of fixed point, in that the parametric models of inflation and of inflation forecasts can be reconciled. If we find such consistency [and we have yet to work out the sampling theory for this by the way] then we have an easy and informative way to filter US inflation, by out-sourcing much of the work to the participants in the *Survey of Professional Forecasters*. The estimated components (innovations to the trend and cycle in inflation) depend only on observed 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 of a particular coincidence between the information sets of economists and forecasters), 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 hold because the approach jointly relies on the UC model and the assumptions about forecasts.

There are three main, preliminary findings. First, detrending after assuming the mean forecast coincides with the rational expectation leads to a trend-cycle decomposition with trend and cycle shocks that are indeed unpredictable, consistent with the underlying assumptions of the UC model. Hence this method provides a direct way to track the historical mixture of these shocks, with time-varying volatilities.

Second, though, when we allow for forecast stickiness, and either estimate it using reported forecasts or use Coibion and Gordonichenko's estimate, the joint model no longer passes the consistency tests. The combined model cannot reproduce unpredictability in the two components of the UC model along with the predictable pattern in forecast errors. So far, then, we cannot reconcile these two perspectives on US inflation.

Third, we show how to separately identify forecast stickiness and persistence in the inflation gap. There is very little evidence of such persistence implicit in the *SPF*.

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 requires 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), and Mertens (2011). To take one example, Mertens (2011) 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-head forecast of inflation then is:

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

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

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

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

We reverse the last steps of this sequence. Beginning with forecasts $E_t \pi_{t+h}$ (or a sticky function of them described below) we use the martingale property of τ_t to estimate $\tau_{t|t}$ and $\epsilon_{t|t}$. It is obvious that we cannot then continue and uncover a unique, underlying information set and set of orthogonality assumptions. But we also do not require a zero covariance between the shocks or restrictions on their variances to estimate the two components and inflation-gap persistence. For example, any pattern of time-varying volatility in σ^2_η is possible, so that the importance of the non-stationary component can vary over the sample. Moreover, inflation will be more volatile than its trend provided the covariance between the trend and the inflation gap is not too negative. 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 out-sourced to the forecasters.

This project is related to two recent studies that also jointly analyze survey-based inflation expectations and time-series models of actual inflation. Clark and Davig (2011) include 1-year-ahead and 10-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 (2010) 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.

3. Inflation Forecast Data

The forecast data come from the *Survey of Professional Forecasters* conducted by the Federal Reserve Bank of Philadelphia. 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 0 (the nowcast) to 4 quarters ahead.

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. Figure 2 shows the mean *SPF* forecasts at a common date of origin (rather than for a common target) with the five different horizons. As the horizon rises the volatility of the forecast decreases strikingly.

4. Mean Forecasts as Rational Expectations: The Basics

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) provided early evidence of the unbiasedness of price forecasts using disaggregated data from the Livingston Survey. Ang, Bekeart, 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 (2011) find similarly favorable results for survey-based expectations of quarterly inflation and specifically the mean CPI inflation forecasts from the *SPF*. As Faust and Wright (2011, p 2) 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. 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.

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

$$\tau_{1t|t} = F_t\pi_{t+1}, \tag{4}$$

so

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

Of course, the UC model also implies that $F_t\pi_{t+h} = \tau_t$, for all $h > 1$. The UC model implies a singlarity 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:

$$\tau_{2t|t} = \frac{1}{4}\sum_{h=1}^{4} F_t\pi_{t+h}, \tag{6}$$

with $\epsilon_{2t|t}$ again given by subtraction (5). (Another possibility would be to inverse weight the multi-step forecasts by their variances.)

Figure 3 shows the CPI inflation rate along with $\tau_{1t|t}$ (in blue) and $\tau_{2t|t}$ (in red). Both estimated trends are smoother than the actual inflation rate and centered on it. (In the next version we hope to compare it to trends estimated by Cogley, Primiceri, and Sargent (2010) and Mertens (2011).) The trend estimated only from the one-horizon forecast 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.

Reverse filtering 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 $\pi_t - F_t\pi_{t+h}$.)

Figure 4 shows the trend and cycle innovations for the first estimation, $\eta_{1t|t} = \Delta\tau_{1t|t}$ (the solid, black line) and $\epsilon_{1t|t}$ (the blue, dashed line). (The graph of $\eta_{2t|t}$ and $\epsilon_{2t|2}$ is quite similar.) Little persistence is evident in either series.

Table 1 gives the sample variances s^2 of each innovation, as well as $Q(j)$, the Ljung-Box Q-statistic with j lags and its p-value. For each detrending method the correlation between $\epsilon_{it|t}$ and $\eta_{it|t}$, denoted $r(\epsilon, \eta)$, is 0.35. (Recall that is is not restricted by the reverse Kalman filter.) That finding naturally fits with the observation in figure 3 that the trend is smoother than the cycle. In the top line of table 1, the main difference between the two trend-estimates is the greater volatility of $\eta_{1t|t}$ than $\eta_{2t|2}$. By averaging over horizons, the second method produces a smoother trend, as figure 3 shows.

The variances of $\epsilon_{t|t}$ and $\eta_{t|t}$ are comparable to the estimates Stock and Watson (2007) report using their UC model. But notice that the variance of $\eta_{t|t}$ rises and the variance of $\eta_{t|t}$ falls as we include information at longer horizons, which suggests that there is more uncertainty about transitory inflation shocks than permanent ones at a moment in time at the longer forecast horizons.

Table 1: Trend and Cycle Moments
1981:3–2012:2

	$\epsilon_{1t\|t}$	$\eta_{1t\|t}$	$\epsilon_{2t\|t}$	$\eta_{2t\|t}$
s^2	3.02	0.24	3.47	0.10
$r(\epsilon, \eta)$	0.35		0.35	
$Q(4)$	5.43	8.76	6.32	8.24
(p)	(0.25)	(0.06)	(0.18)	(0.08)
$Q(8)$	6.46	13.9	7.10	15.6
(p)	(0.60)	(0.09)	(0.52)	(0.05)

Next, the Q-statistics show that there is little evidence of autocorrelation in the inflation gap (as seen in figure 4) 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 which 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, $\eta_{it|t}$.

The reverse filtering also allows us to comment on the pattern in squared innovations, roughly speaking quarterly 'realized volatility' or the things we would average over periods of time to estimate time-varying volatilities. Squared ϵ can be much larger than squared η so, for ease of viewing, figure 5 graphs the square root of the squared values *i.e.* the absolute value of the two series, again for the first detrending method. (Again results for the other trend method are quite similar and so are not shown.) There is variation over time: the inflation-gap (the blue, dashed line) is relatively more volatile prior to 1990 and again after 2005.

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

Table 2: Sub-Sample Trend and Cycle Variances

Time Span→ Statistic ↓	1981:III–1991:I	1991:II–2002:II	2002:III–2012:II
$s^2(\epsilon_{1t\|t})$	2.71	0.72	5.59
$s^2(\eta_{1t\|t})$	0.46	0.04	0.27
$r(\epsilon_{1t\|t}, \eta_{1t\|t})$	0.35	0.15	0.49
$s^2(\epsilon_{2t\|t})$	3.27	0.79	6.34
$s^2(\eta_{2t\|t})$	0.24	0.03	0.05
$r(\epsilon_{2t\|t}, \eta_{2t\|t})$	0.35	0.13	0.49

The volatility of each component (and under each detrending method) 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 2 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.

The results so far assume that mean forecasts coincide with conditional expectations. One can examine this assumption indirectly by looking at the persistence in h-step-ahead forecast errors. The idea is that the unobserved forecast error relative to the true, conditional expectations, $\pi_{t+h} - E_t\pi_{t+h}$, will not be persistent, so that persistence in the forecast error relative to the mean SPF prediction casts doubt on the equivalence of $E_t\pi_{t+h}$ and

$F_t\pi_{t+h}$. Table 3 provides the Q-statistics for multi-step SPF forecast errors, along with their p-values. The horizons run from the current quarter to 4 quarters ahead.

Table 3: Forecast-Error Persistence
1981:3–2012:2

$h \rightarrow$	0	1	2	3	4
$Q(4)$	4.62	7.59	10.5	13.1	12.8
(p)	(0.33)	(0.11)	(0.03)	(0.001)	(0.01)
$Q(8)$	6.01	9.77	12.7	15.0	13.9
(p)	(0.64)	(0.28)	(0.12)	(0.06)	(0.08)

There is evidence of forecast-error persistence, at conventional levels of significance, for horizons $h = 2, 3, 4$. Table 3 thus offers some evidence against treating the mean SPF forecast as the conditional expectation. In addition, recall the counter-factual prediction of this basic UC model (combined with the assumption that $F_t\pi_{t+h}$ coincides with $E_t\pi_{t+h}$) that forecasts are the same at all horizons. The next section introduces sticky information, which addresses both of these shortcomings of the basic model.

5. Sticky Forecasts

Notwithstanding our earlier citations to research that shows professional forecasts are unbiased, a number of statistical studies (now including table 3 above) have found that forecast errors contain predictable components. Next, a specific pattern of predictability, using forecast revisions, leads to an alternative, parametric model of observed, mean forecasts.

We work with the sticky-information model, as introduced by Mankiw and Reis (2002) and Reis (2006) and applied to professional forecasters by 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:

$$F_t \pi_{t+h} = (1 - \lambda) E_t \pi_{t+h} + \lambda F_{t-1} \pi_{t+h} \tag{7}$$

Define the non-sticky information forecast error:

$$\nu_{t+h} = \pi_{t+h} - E_t \pi_{t+h}. \tag{8}$$

Subtracting each side of this pattern in reported, mean forecasts (7) from realized inflation gives:

$$\begin{aligned} \pi_{t+h} - F_t \pi_{t+h} &= \lambda (E_t \pi_{t+h} - F_{t-1} \pi_{t+h}) + (\pi_{t+h} - E_t \pi_{t+h}) \\ &= \frac{\lambda}{1-\lambda} (F_t \pi_{t+h} - F_{t-1} \pi_{t+h}) + \nu_{t+h} \end{aligned} \tag{9}$$

Because ν_{t+h} has the properties of an econometric error, this link (9) can be used to estimate λ, by regressing the observed forecast error on the forecast revision. Coibion and Gorodnichenko do this using *SPF* inflation forecasts and find $\hat{\lambda} = 0.55$, which implies that forecasters update every 6-7 months on average. They also find that additional regressors, in the form of past, realized values of macroeconomic variables, are not significant in explaining forecast errors once these revisions are included.

We shall use their estimator but also consider an alternate estimator that uses only forecast data. Combining the forecast implication of the trend-cycle model (2) with the description of forecast updating (7) gives:

$$F_t \pi_{t+h} = (1 - \lambda) \tau_{t|t} + \lambda F_{t-1} \pi_{t+h}. \tag{10}$$

Next, take differences over time to give estimating equations:

$$F_t \pi_{t+h} - F_{t-1} \pi_{t+h-1} = \lambda (F_{t-1} \pi_{t+h} - F_{t-2} \pi_{t+h-1}) + (1 - \lambda) \eta_t, \tag{11}$$

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 (11).

Table 4 shows the results of estimating λ horizon-by-horizon and pooled across horizons, with HAC standard errors. The first row uses the Coibion-Gorodnichenko projection (9), while the second row uses our forecast-only equation (11). (We also could combine the two equations.) 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 a several illustrative values for λ.

Table 4: Stickiness Estimates

Equation $\downarrow$	$h \rightarrow$:	0	1	2	3	pooled
(9)	$\hat{\lambda}$	0.37	0.31	-0.12	0.35	0.00
	(se)	(0.07)	(0.20)	(0.78)	(0.26)	(0.11)
(11)	$\hat{\lambda}$:	0.11	0.16	0.13	-0.03	0.64
	(se):	(0.23)	(0.13)	(0.10)	(0.10)	(0.07)

Given $\hat{\lambda}$ we can invert (10) to give the estimated trend based on forecasts at horizon h:

$$\tau_{t|t} = \frac{F_t \pi_{t+h} - \lambda F_{t-1} \pi_{t+h}}{(1 - \lambda)}. \tag{12}$$

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 and two trial values for λ, with $\tau_{3t|t}$ denoting the average for $\lambda = 0.2$ and $\tau_{4t|t}$ the average for $\lambda = 0.4$. (The first subscripts distinguish these from τ_1 and τ_2, the two trends studied in section 4.) Table 5 gives statistics for the corresponding shocks to the inflation trend and inflation gap. The first two columns pertain to $\lambda = 0.2$ and the last two to $\lambda = 0.4$.

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

	$\epsilon_{3t\|t}$	$\eta_{3t\|t}$	$\epsilon_{4t\|t}$	$\eta_{4t\|t}$
s^2	2.57	0.27	2.79	0.32
$r(\epsilon, \eta)$	0.51		0.37	
$Q(4)$	6.63	18.9	5.78	26.9
(p)	(0.17)	(0.00)	(0.21)	(0.00)
$Q(8)$	7.95	25.8	9.34	36.4
(p)	(0.43)	(0.00)	(0.31)	(0.00)

It is clear that allowing for stickiness does not help with the borderline predictability of the trend-innovation $\eta_{t|t}$ in table 1 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.

Figure 6 graphs actual CPI inflation and the two implied trends, $\tau_{3t|t}$ and $\tau_{4t|t}$ corresponding to the two candidate values of λ. It is obvious from the figure 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 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 (12) do not appear to have a random walk component. The stickiness that fits the partial predictability of forecast errors (at least for some horizons in table 4) leads to properties that are inconsistent with the assumptions of the UC model. Although we have not incorporated the requirements from table 5 in our estimator, it seems fairly clear so far that no reasonable value of λ can allow the joint statistical model to pass the consistency tests.

6. Persistence 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 that these two features are separately identified and that stickiness is more significant than inflation-gap persistence in the SPF data.

Following Cogley, Primiceri, and Sargent (2010) we allow the stationary component of inflation, also known as the inflation gap, to itself be persistent. For this edition, we work with an AR(1) version:

$$\epsilon_t = \rho \epsilon_{t-1} + \upsilon_t, \tag{13}$$

where υ_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 τ_t, but in general

$$E_t \pi_{t+h} = \tau_t + \rho^h \epsilon_t. \tag{14}$$

Faust and Wright (2011) note that subjective forecasts often have proved 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; the forecasts (14) allow for that pattern.

First suppose that $F_t \pi_{t+h} = E_t \pi_{t+h}$, as in section 4. Then we can annihilate the stochastic trend, τ_t using the difference across horizons:

$$F_t \pi_{t+h+1} - F_t \pi_{t+h} = F_t \Delta \pi_{t+1} = \rho^h (\rho - 1) \epsilon_t. \tag{15}$$

Multiplying the difference equations (15) by $(1 - \rho L)$ then gives the quasi-differences over time:

$$F_t \Delta \pi_{t+h+1} = \rho F_{t-1} \Delta \pi_{t+h} + \rho^h (\rho - 1) \upsilon_t, \tag{16}$$

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, υ_t, is uncorrelated with them. Thus the persistence in the inflation gap, ρ, can be estimated by ordinary least squares in the estimating equations (16). 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 (15) 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. 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 looks at the correlation of revisions, but with forecasts from individual forecasters in the *SPF*, to measure perceived inflation persistence. Krane (2011) uses GDP forecast revisions from the Blue Chip survey to identify forecasters' implicit views of shocks to GDP.

We first estimate the system (16) without restrictions on the error dispersion matrix (even though those contain information on ρ). We also overlook the singularity in the system for now *i.e.* the fact that a common v_t appears in the equation for each horizon h. Table 6 shows the estimates of the persistence in the inflation gap, $\hat{\rho}$, found using forecasts at individual horizons h and then from pooled estimation over all horizons. Brackets contain robust standard errors. 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 section 4.

Table 6: Persistence Estimates

h:	0	1	2	3	pooled
$\hat{\rho}$:	0.206	-0.046	-0.189	0.004	-0.031
(se):	(0.164)	(0.133)	(0.140)	(0.097)	(0.061)

Notes: Persistence is estimated from equations (16) in *SPF* mean forecasts from 1981–2012.

Next, suppose that forecasts are sticky, as in section 5. Combining the forecast implication of the trend-cycle model (14) with the description of forecast updating (10) gives:

$$\frac{F_t\pi_{t+h} - \lambda F_{t-1}\pi_{t+h}}{1-\lambda} = \tau_t + \rho^h \epsilon_t, \tag{17}$$

or

$$F_t\pi_{t+h} = \lambda F_{t-1}\pi_{t+h} + (1-\lambda)\tau_t + (1-\lambda)\rho^h\epsilon_t. \tag{18}$$

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

$$F_t\pi_{t+h+1} = \lambda F_{t-1}\pi_{t+h+1} + (1-\lambda)\tau_t + (1-\lambda)\rho^{h+1}\epsilon_t, \tag{19}$$

so that the difference across horizons is:

$$F_t\Delta\pi_{t+h+1} = \lambda F_{t-1}\Delta\pi_{t+h+1} + (1-\lambda)(\rho^{h+1} - \rho^h)\epsilon_t. \tag{20}$$

Suppose that the inflation gap, ϵ_t, follows the AR(1) process (13) so that $\epsilon_t(1-\rho L) = \upsilon_t$. Multiplying the difference equations (20) by $(1-\rho L)$ gives:

$$F_t\Delta\pi_{t+h+1} = \lambda F_{t-1}\Delta\pi_{t+h+1} + \rho F_{t-1}\Delta\pi_{t+h} - \rho\lambda F_{t-2}\Delta\pi_{t+h} + (1-\lambda)(\rho^{h+1} - \rho^h)\upsilon_t. \tag{21}$$

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 equations (21).

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.

Using the timing in the estimating equations, the *SPF* thus allows us to study $h = 0$, 1, and 2, because each equation involves the change of inflation and several horizons. Table 7 presents estimates of the stickiness parameter λ and the inflation-gap persistence parameter ρ. The results are consistent with those already reported separately: some evidence of stickiness but no evidence of inflation-gap persistence, at least of this first-order Markov form.

Table 7: Persistence and Stickiness Estimates

h:	0	1	2	pooled
$\hat{\rho}$:	-0.01	-0.05	-0.12	-0.02
(se):	(0.17)	(0.12)	(0.14)	(0.08)
$\hat{\lambda}$:	1.49	0.71	0.69	0.77
(se):	(0.37)	(0.24)	(0.13)	(0.09)

Notes: Parameters are estimated from equations (21).

7. Conclusions

The unobserved-components model of inflation is not identified without auxiliary assumption, as Morley, Nelson, and Zivot (2003) show. In this paper, we have illustrated how to identify the model using the SPF, without necessarily assuming that mean, reported forecasts coincide with conditional expectations (and without assuming anything about shock correlations or volatilities). Both the widely-used, UC model of inflation and recent descriptions of sticky forecasting restrict unobservable inflation forecasts $E_t\pi_{t+h}$. We show how combining these statistical models provides 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 only, as well as the implied, historical shock volatilities.

The approach features over-identification: we study whether we can reconcile the two statistical models jointly with the time-series properties of actual inflation and mean forecasts: $\{\pi_t, F_{t-h}\pi_t\}$. We find that we cannot. So far we do not find that the behavior of mean SPF forecasts over multiple horizons can 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 findings are reminiscent of those from modelling of 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. Our next step is to try to reverse engineer these properties or at least to better document the challenges involved in fitting them jointly.

References

Andrade, Philippe and Le Bihan, Hervé (2010) Inattentive professional forecasters. Working paper 307, Banque de France.

Ang, A., Bekaert, G., Wei, M. (2007) Do macro variables, asset markets or surveys forecast inflation better? *Journal of Monetary Economics* 54, 1163-1212

Beveridge, S. 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.

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, T. and T.J. Sargent (2005) Drifts and volatilities: Monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics* 8, 262–302.

Cogley, T, and A. Sbordone (2008) Trend inflation, indextation, and inflation persistence in the new Keynesian Phillips curve. *American Economic Review* 98, 2101–2126.

Cogley, T., G.E. Primiceri, and T.J. Sargent (2010) Inflation-gap persistence in the U.S. *American Economic Journal: Macroeconomics* 2, 43–69.

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.

Del Negro, Marco and Stefano Eusepi (2010) Fitting observed inflation expectations. Federal Reserve Bank of New York Staff Report no. 476.

Faust, Jon and Jonathan H. Wright (2011) Forecasting inflation. mimeo, Department of Economics, Johns Hopkins University

Gil-Alana, Luis, Antonio Moreno, and Fernando Pérez de Gracia (2011) Exploring survey-based inflation forecasts. *Journal of Forecasting* doi: 10.1002/for.1235

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.

Ireland, P.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.

Krane, Spencer D. (2011) Professional forecasters' views of permanent and transitory shocks to GDP. *American Economic Journal: Macroeconomics* 3, 184–211.

Mankiw, N.G. and R. Reis (2002) Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve. *Quarterly Journal of Economics* 117, 1295–1328.

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 Beverdige-Nelson and 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.

Reis, R. (2006) Inattentive producers. *Review of Economic Studies* 73, 793–821.

Stock, J.H. and M.W. Watson (2007) Why has US inflation become harder to forecast? *Journal of Money, Credit and Banking* 39 (S1), 3–33.

Figure 1: US CPI Inflation

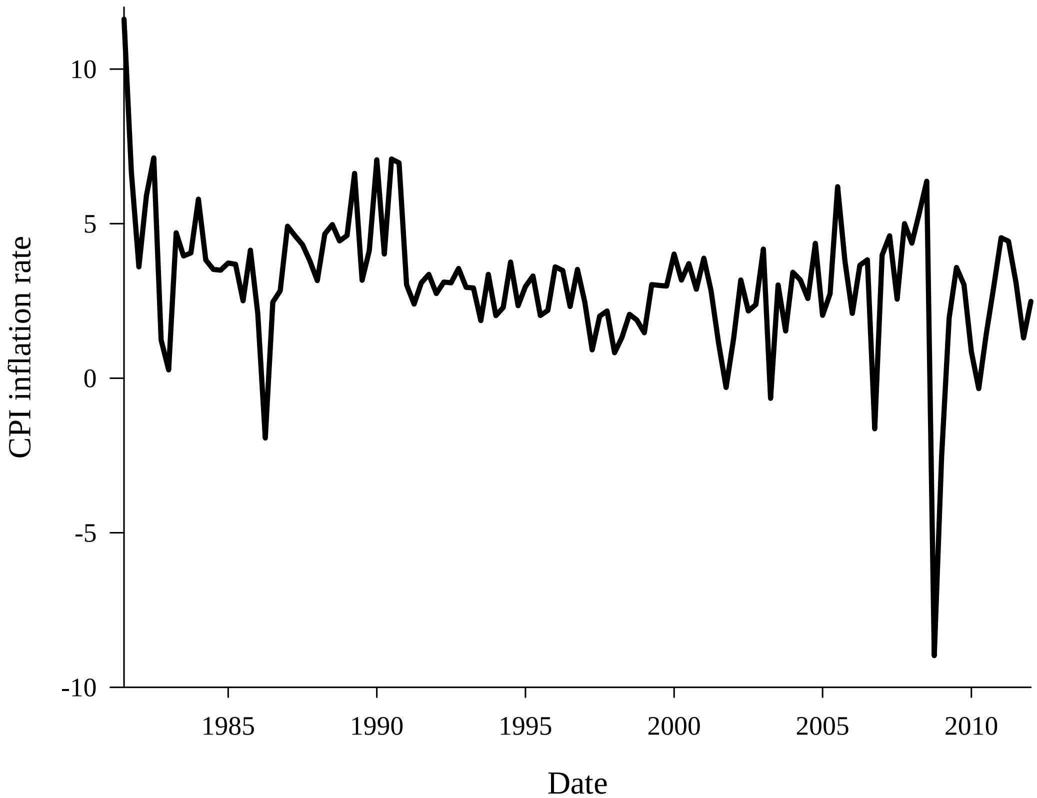

Note: Inflation is the annualized quarterly growth in the CPI for all urban consumers and all items, series `cpiaucsl` from `FRED.`

Figure 2: Mean CPI Inflation Forecasts

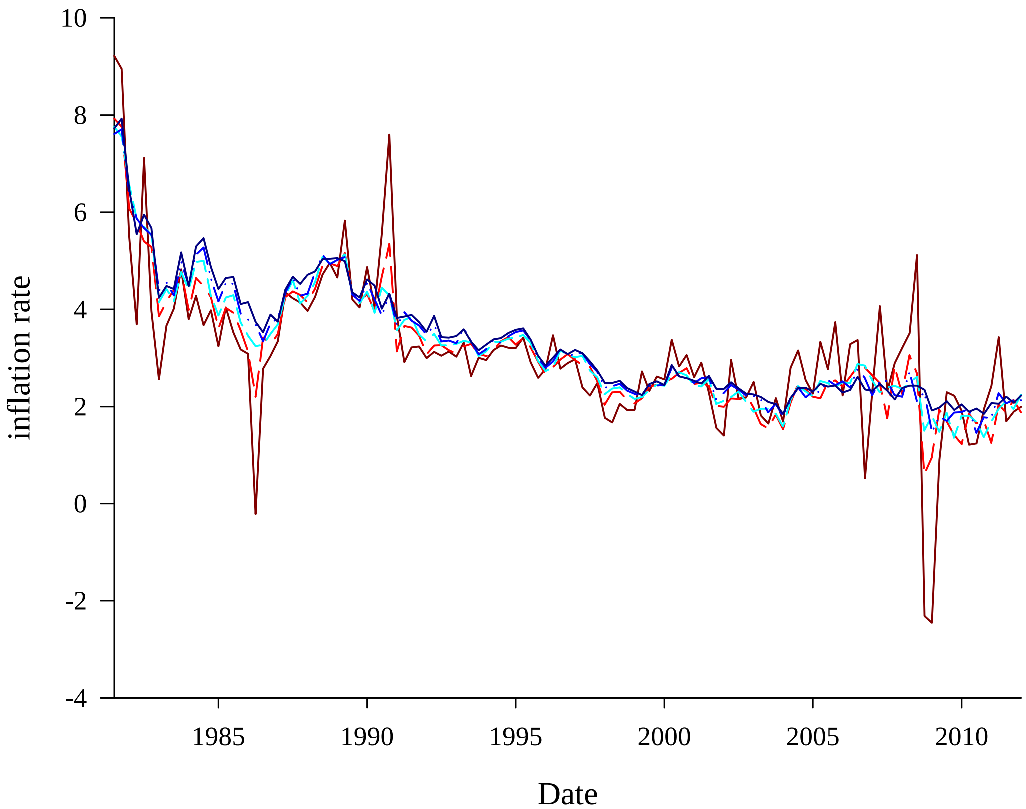

Note: Forecast series are the means of the 0-4 step-ahead forecasts of US quarterly CPI inflation, series CPI2 to CPI6 from `Mean_CPI_Levl.xls` from the *Survey of Professional Forecasters*. The colours change from dark red to dark blue as the horizon increases.

Figure 3: CPI Inflation and Trends

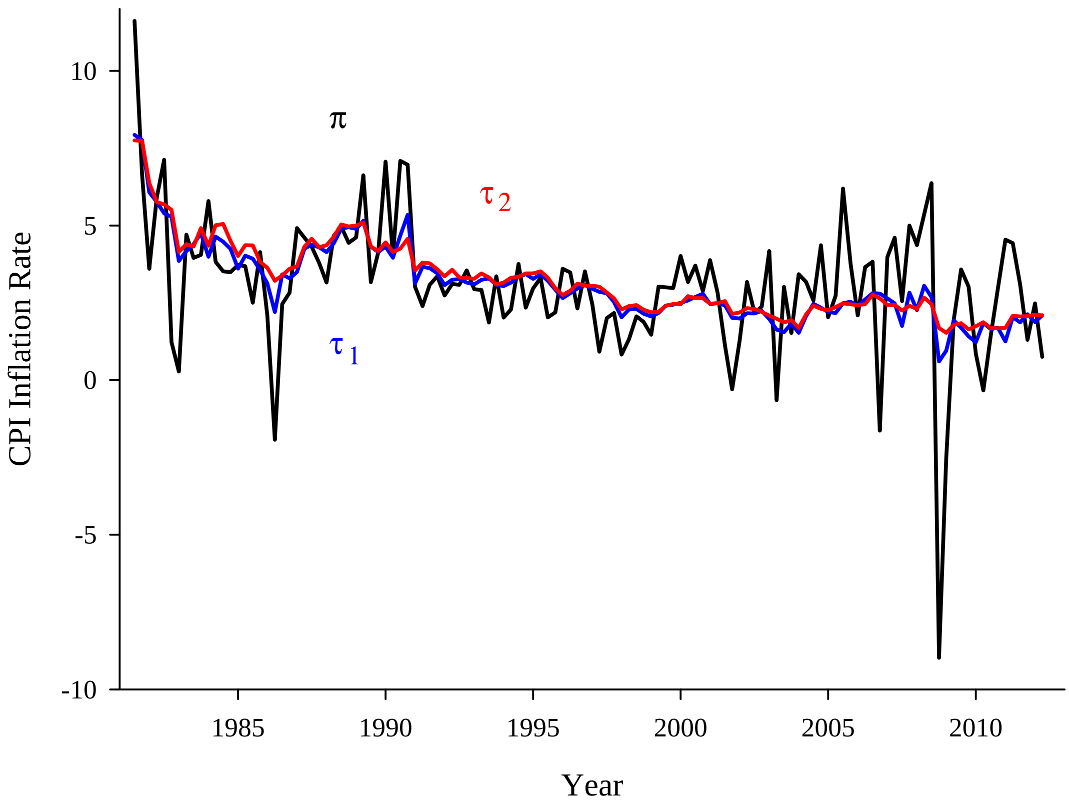

Figure 4: Trend (η) and Cycle (ε) Innovations

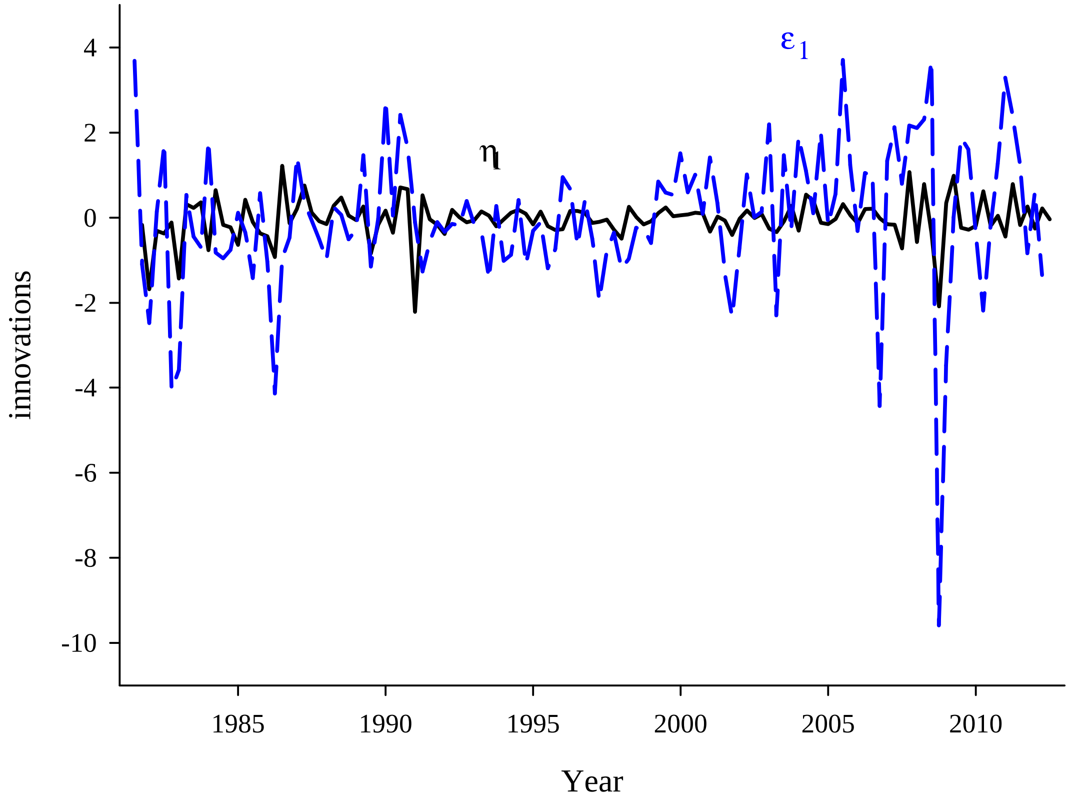

Figure 5: Absolute Values of Trend (η) and Cycle (ε) Innovations

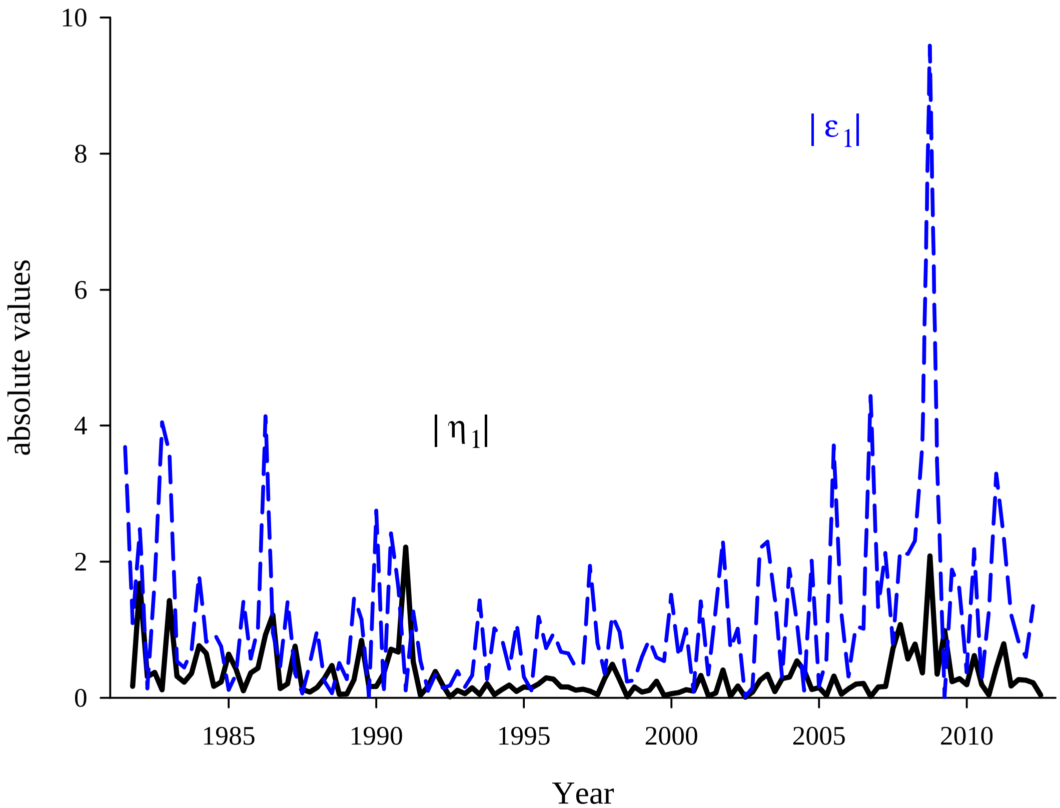

Figure 6: CPI Inflation and Trends