FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Survey Measures of Expected Inflation and the Inflation Process

Bharat Trehan
Federal Reserve Bank of San Francisco

February 2010

Working Paper 2009-10
http://www.frbsf.org/publications/economics/papers/2009/wp09-10bk.pdf

Survey Measures of Expected Inflation and the Inflation Process

Revised: February 2010

Bharat Trehan*
Federal Reserve Bank of San Francisco

Abstract

This paper uses data from surveys of expected inflation to learn how expectations processes have changed following recent changes in the behavior of inflation. Households do not appear to have recognized the change in the process, and are placing substantially more weight than appears warranted on recent inflation data when forming expectations about inflation over the next year. At first glance, professional forecasters do appear to have changed how they predict inflation. But a closer look at the data reveals that professionals are relying on core rather than headline inflation, and are placing too much weight on recent core inflation data. These errors show up in a noticeable (absolute and relative) deterioration in the forecast accuracy of both households and professionals.

* I would like to thank Oscar Jorda and an anonymous referee for helpful comments. Wayne Huang and Puneet Chehal provided high quality research assistance. Any opinions expressed in this paper are those of the author and not those of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

In the first half of 2008, some U.S. surveys showed noticeable increases in expected inflation, leading to concerns about a possible increase in the inflation rate and about the credibility of the Federal Reserve. Such concerns can be justified on the basis of a number of recent studies. For instance, Ang, Bekaert and Wei (2007) show that survey measures of expected inflation provide better forecasts of inflation than any other alternative that they consider, including about a dozen variants each of Phillips curve and term structure models, as well as simple regime switching models. Mehra and Herrington (2008) use a VAR specified by Leduc, Sill and Stark (2007) to examine measures of survey expectations following the change in the monetary policy regime that took place around the end of the 1970s. They find that the expectations process changed in a way that is consistent with the change in the inflation process that took place at about the same time, suggesting that survey participants are able to detect changes in the inflation process relatively quickly.[1] And Bernanke, Laubach, Mishkin and Posen (2001) discuss how the behavior of survey forecasts relative to the monetary authority's inflation target provides information about credibility.[2]

Though the rationality of survey forecasts has been debated (see Croushore, 1998, for a discussion and a defense), they are generally well regarded, especially the forecasts made by the professionals. For instance, Carrol (2003) argues that forecasts from the Society of Professional Forecasters pass all the important tests for rationality and goes on to model households' forecasts as adjusting gradually to the forecasts of professionals. Ang, et. al., (ABW) are positive about both household and professional forecasts: "That the median Livingston and SPF forecasts do well is perhaps not surprising…However, even participants in the Michigan surveys who are consumers, not professionals, produce accurate out-of-sample forecasts, which are only slightly worse than those of the professionals." They go on to speculate that the superior performance of the professionals may result from their ability to recognize structural change more quickly than mechanical model forecasts can.

This paper argues that neither households nor professional forecasters are quite as sophisticated as these arguments make them out to be. The evidence suggests that there have

[1] There is a debate about the nature of the change in the inflation process that took place at this time. This issue is taken up below.

[2] This in no way exhausts the list of uses to which inflation survey data have been put. For instance, Mankiw, Reis and Wolfers (2003) and Orphanides and Williams (2005) use these data to inform aspects of model specification. For an extensive discussion of how various kinds of survey data are used for modeling expectations and testing hypotheses about expectations formation see Pesaran and Weale (2006).

been some changes in the inflation process in recent years, but neither households nor professionals have responded appropriately so far.

More specifically, the evidence suggests that the inflation process has become noticeably less persistent since the beginning of the decade. As with the change in the inflation process around the end of the 1970s, this change could be modeled as a change in the autoregressive coefficients of the inflation process or as a change in the variance of the shocks to the process. In either case, as argued below, the change in the inflation process should show up as a change in the relationship between survey expectations data and realized inflation. But the survey data suggest that there has been little, if any, change at all in the way that households react to inflation data. In particular, it appears that households are placing too large a weight on recent inflation data when forming expectations. Consistent with this finding, household forecasts of inflation are now about the worst of all the alternatives considered below. This contrasts sharply with Ang, et al (2007), who find that--over an earlier sample--household forecasts are among the best.

Professional forecasters, on the other hand, do appear to have changed how they react to recent inflation data, but this change is not fully consistent with the observed change in the inflation process. Specifically, professionals seem to be paying less attention to headline inflation data but are still relying heavily on core inflation data. This is consistent with the position advocated by Blinder and Reis (2005) that it is better to predict headline inflation using lagged core --- rather than headline --- inflation. It turns out, however, that SPF forecasts of headline inflation have deteriorated in recent years as well, and are now worse than forecasts based on lagged headline inflation alone. Surprisingly, SPF forecasts of headline inflation are rather good forecasts of core inflation, which suggests that the professionals may now be implicitly forecasting the core CPI, instead of the headline CPI (which is what they are asked to forecast). Such a switch would be consistent with the argument put forward by Blinder and Reis (henceforth BR) that then-Chairman Greenspan's advocacy of the core inflation concept has shifted U.S. public discourse about inflation from headline to core inflation.

If professional forecasters have indeed begun to pay more attention to core CPI because of Chairman Greenspan's advocacy, then this switch provides unusual evidence on the Federal Reserve's (Fed's) credibility.[3] This finding can also be seen as augmenting the findings of

[3] The SPF forecasts provide more conventional evidence of Fed credibility as well: The 10-year ahead inflation forecast has been quite stable for more than 10 years now.

Orphanides and Williams, whose analysis suggests that the professional forecasters are backward looking (as their forecasts can be approximated with a Kalman filtered version of past inflation data). Our results suggest that the professionals are sensitive to other aspects of the environment as well, though this may not always lead to improved forecast accuracy.

1. The data

The Survey of Consumers was initiated in 1946 and is currently conducted monthly by the University of Michigan's Survey Research Center. Each month, a randomly selected sample of approximately 500 American households are asked (in telephone interviews) about expected changes to key macroeconomic variables such as inflation, interest rates, and unemployment. The sample is designed to be "rotating," in that for any one survey, approximately 60% of respondents are new and the remaining 40% of respondents are interviewed for a second time.

Since 1977, respondents have been asked the following question about inflation:

> "By about what percent do you expect prices to go (up/down) on the average, during the next 12 months?"

Continuous monthly data on the answers to this question are available since January 1978. Quarterly data are available prior to that, but not for every quarter; these data are not used here.

The Survey of Professional Forecasters was first conducted in the fourth quarter of 1968 by the American Statistical Association and the National Bureau of Economic Research. It has been conducted by the Philadelphia Fed since the second quarter of 1990. Sample size has varied noticeably over time; as of this writing, their website identifies more than 50 respondents and there are some anonymous respondents as well. Beginning in 1981Q3, participants were asked to forecast quarterly and annual CPI.

Since the Survey of Professional Forecasters (SPF) is only conducted once a quarter, the analysis of the data from the consumer survey is carried out at a quarterly frequency as well. For the consumer survey, I use data from the third month of each quarter. The implications of this choice for various forecast comparisons are discussed below.

Figure 1 plots data on expected inflation over the next year from both the Michigan and the SPF surveys. In recent years, the SPF forecasts have been perceptibly below the expectations from the Michigan survey. Also noticeable is the increased volatility of the Michigan expectations data towards the end of the sample.

2. The household survey of expected inflation

Do survey respondents use the information in recent inflation data in ways that are consistent with the inflation process? To answer this question, two projections are compared: the projection of expected inflation on recent inflation data and the projection of actual inflation over the same horizon on recent inflation data. Changes over time in the latter provide information about how the inflation process has changed. The next step is to determine whether the projections involving inflation expectations show similar changes.[4]

The starting point is a regression of CPI inflation over the next year on the inflation rate for the current quarter and 7 lags. The estimated equation is

$$\pi_{t,t+4} = \alpha_0 + \sum_{i=1}^{8} \alpha_i \pi_{t-i,t-i+1} + \varepsilon_{1t} \tag{1}$$

where $\pi_{t,t+4}$ measures inflation from t to $t+4$, i.e., over the next year, which is the same horizon as in the Michigan survey. Quarterly data are used here to allow for easy comparison with the results for the SPF forecasts below. Very similar results were obtained when the lag length was varied by four, when monthly data were used instead of quarterly, and when a four quarter average of inflation was used as the explanatory variable.

The first column of Figure 2 shows that it has become harder to predict future inflation over time. Each point on the middle line plotted in the upper left hand panel of Figure 2 is the value of $\sum_{i=1}^{8} \alpha_i$ (that is, the sum of the coefficients on inflation) when the regression described above is estimated over a 15-year window that ends in the quarter against which the point is plotted. Also shown are two-standard-error bands, based on HAC standard errors. The lower left hand panel shows how the fit of this equation changes over time. Taken together, the two panels in the first column reveal that contemporaneous and lagged inflation data contain less and less information about future inflation as time goes on. For sample periods whose endpoint lies within the last two or three years, the point estimate (that is, $\sum_{i=1}^{8} \alpha_i$) is noticeably below zero, and even the upper bound of the confidence interval is not too far above zero. Furthermore, the the adjusted-R^2 is negative or close to zero since the beginning of 2005. Though uneven, the

[4] Under the maintained assumption that inflation is an autoregressive process, Pesando (1975) argues that if expectations are rational, the corresponding coefficients in the two regressions should be equal. Mullineaux (1980) projects expected inflation on lagged inflation and money growth and examines how the coefficients evolve over time.

decline appears to have taken place in two steps; first, over the first half of the 1990s and the second over the last five years or so of the sample. The first decline appears related to the early 1980s dropping out of the sample; for instance, if a 10-year rolling window is used in place of a 15-year window, the first drop in the sum of the coefficients is complete by the early 1990s, instead of the additional 5 years or so that it takes in the plot shown in Figure 2. These results suggest a decline in the persistence of inflation, though the left hand side variable is not what traditionally would be used in a regression meant to examine changes in persistence.[5]

The second column of Figure 2 shows what happens when the exercise above is repeated using the following equation:

$$E_t^M \pi_{t,t+4} = \beta_0 + \sum_{i=1}^{8} \beta_i \pi_{t-i,t-i+1} + \varepsilon_{2t} \qquad (2)$$

where $E_t^M \pi_{t,t+4}$ denotes the 1-year-ahead expected inflation from the Michigan survey.[6] The top panel on the right hand side shows that while the sum of lagged coefficients did decline in the mid-1990s (just as is the case in the panel on the left hand side), there is no evidence of any decline since the late 1990s. Instead, over the last few years the sum of lagged coefficients actually increased, so that one would be hard pressed to say that the sum of the coefficients at the end of the sample is any different from what it was in the beginning of the sample. Thus, it appears that---when forming expectations about inflation over the next year---households have continued to place a large weight upon recent inflation data till the very end of the sample. The adjusted-R^2 does show a decline towards the end of the sample, but is still substantial and quite a bit higher than what is obtained in the case of realized inflation.

More formal tests on the stability of the two equations provide consistent results. Table 1 shows how the coefficients in equations (1) and (2) change over the two halves of the 1978Q1-2009Q3 sample. Specifically, a dummy variable that is 0 until the end of 1993 and 1 afterwards is included both by itself and after being interacted with the inflation terms. In the first column (where realized inflation is the dependent variable), the sum of the coefficients on the inflation

[5]Regressing 1-quarter-ahead inflation on current and lagged quarterly inflation (which would be the traditional specification) leads to results that are very close to those shown in Chart 2.

[6] In view of the distinction between core and headline inflation that comes up when the SPF forecasts are examined below, it is worth noting that no evidence was found to suggest that household inflation expectations are more (or less) sensitive to oil or food prices than to other kinds of inflation. More specifically, terms representing increases in the price of oil, the price of food or the level of non-core inflation (as defined below) were almost always insignificant at the 10 percent level when included in equation (2).

terms interacted with the dummy (that is, the DP_i terms) is significantly negative, implying a significant decline in the inflation coefficients over the second half of the sample. Further, the null hypothesis that these terms can be dropped from the equation can be rejected at the 1 percent level. Notably, the hypothesis that the sum of the coefficients on all the inflation terms is zero in the second half of the sample cannot be rejected at the 5 percent level, though it can be rejected at the 10 percent level. The results for the expected inflation equation are quite different. Although the sum of the coefficients on inflation during the first half of the sample is close to that for the first equation, the results for the second half are very different. Most importantly, one cannot reject the hypothesis that there has been no change in the sum of these coefficients over the second half of the sample at conventional significance levels. And one can clearly reject the hypothesis that the sum of coefficients on all the inflation terms is zero in the second sample (that is, one can easily reject the hypothesis that $\Sigma P_i + \Sigma DP_i = 0$).

An alternative procedure to test the stability of the equations is to use the Bai-Perron tests, where both the dates of the breaks and the number of breaks are assumed to be unknown.[7] For the realized inflation regression, the WDMax test statistic is 130.2, compared to a 1 percent critical value of 24.8. Thus, the null of no break is decisively rejected. The sequential test finds four breaks at the five percent level: in 1981Q1, 1986Q1, 2004Q2 and 1990Q2 (in that order). In contrast, for the expected inflation regression, the value of the WDMax statistic is 15.8, compared to a 10 percent critical value of 18.1. And the sequential test finds no breaks at the 10 percent level. Thus, the results from the Bai-Perron tests reinforce the findings in Table 1.

A straightforward interpretation of the decline in the sum of the coefficients on lagged inflation shown in the first column of Figure 2 is that inflation is becoming less persistent over time, with a noticeable change in persistence having taken place in the early part of this decade. Given this result, the second column of charts suggests that in forming expectations about inflation over the next year, households are placing substantially more weight than they should on recent quarterly inflation data.[8]

---

[7] These tests are discussed in Bai and Perron (1998).

[8] Longer term consumer inflation expectations also appear to be excessively sensitive to recent inflation data. When expected inflation over the 5-to-10 year horizon is regressed on the current and 7 lags of quarterly inflation (for a 15-year rolling window whose right end point moves from 2004 to 2009), the sum of the coefficients is *positive* and always significantly different from zero. By contrast, in the regression for actual inflation over the same horizon, the sum of coefficients on quarterly inflation is always *negative* and significantly different from zero for the last 10 years or so. Because of data availability, the two sets of regressions do not span exactly the same period. Over the roughly 10-year overlapping sample period, the regression with realized inflation as the dependent variable has an

The recent decline in inflation persistence could well be part of a trend of declining persistence that has been in place since the 1980s. Among others, Taylor (2000), Cogley and Sargent (2005) and Levin and Piger (2004) have argued that the persistence of inflation has declined. Along the same lines, Blanchard and Gali (2007) and Mishkin (2007) argue that there has been a change in the way inflation responds to shocks. And, as noted above, both Leduc, Sill and Stark (2007) and Mehra and Herrington (2008) conclude that (roughly) since the 1980s, U.S. inflation has been a stationary process.

Many of these authors have suggested that the change in the inflation process represents a change in the conduct of policy. But there has been considerable debate about this. As pointed out by Sims (1999), what appears to be time variation in the estimated coefficients could really be the result of changes in the shocks hitting the system; this argument has been elaborated in Sims and Zha (2006). Using the Cogley-Sargent technology, Clark and Nakata conclude that a reduction in the size of the shocks hitting the economy is largely responsible for the reduction in the volatility of inflation and inflation expectations in recent years. In a similar vein, Pivetta and Reis (2007) argue that inflation persistence has not changed much over the postwar period because there has been little change in the size of the largest root in the inflation process since the 1960s. Stock and Watson (2007) provide a reconciliation of these findings in a model where inflation has both a permanent and a temporary component. In this model, a reduction in the variance of the innovation to the permanent component implies that a given change in inflation is more likely to be reversed than before, even though there has been no change in the largest root of the process.

The Stock and Watson (SW) model provides a characterization of the inflation process which is very different from the univariate autoregressions presented above. (According to SW, "The time-varying trend-cycle model is equivalent to a time-varying first-order integrated moving average (IMA(1,1)) model for inflation, in which the magnitude of the MA coefficient varies inversely with the ratio of the permanent to the transitory disturbance variance," p. 4) It is therefore interesting to see how their specification interprets recent changes in the inflation process. SW postulate a model in which inflation has two components: a stochastic permanent component and a serially uncorrelated temporary component. The variance of the disturbance

adjusted-R^2 of 0.13 with the sum of inflation coefficients equal to -0.14 while the regression for the Michigan survey data has an adjusted-R^2 of 0.92 and the sum of inflation coefficients is 0.48.

terms is allowed to change over time. Specifically, their (unobserved components-stochastic volatility) model is given by:

$$\pi_t = \tau_t + \eta_t, \qquad \eta_t = \sigma_{\eta,t}\zeta_{\eta,t}$$
$$\tau_t = \tau_{t-1} + \varepsilon_t, \qquad \varepsilon_t = \sigma_{\varepsilon,t}\zeta_{\varepsilon,t}$$
$$\ln\sigma^2_{\eta,t} = \ln\sigma^2_{\eta,t-1} + \upsilon_{\eta,t}$$
$$\ln\sigma^2_{\varepsilon,t} = \ln\sigma^2_{\varepsilon,t-1} + \upsilon_{\varepsilon,t}$$

where $\zeta_t = (\zeta_{\eta,t}, \zeta_{\varepsilon,t})$ is i.i.d. N(0,I_2) and $\upsilon_t = (\upsilon_{\eta,t}, \upsilon_{\varepsilon,t})$ is i.i.d. N(0,γ I_2). ζ_t and υ_t are independently distributed, and γ is a scalar parameter.

The estimates of $\sigma_{\eta,t}$ and $\sigma_{\varepsilon,t}$ (the standard deviations of the shocks to the temporary and permanent components) are plotted in Figure 3, together with an estimate of τ_t, the permanent component of inflation. As pointed out by SW, the standard deviation of the permanent component of inflation (shown in the middle panel) rose significantly from the 1960s to the early 1980s, before declining sharply over the remainder of that decade. It has moved very little since the mid-1990s. The standard deviation of the temporary component has moved in almost the opposite way; it did not move around very much prior to 2000, especially when compared to the permanent component. However, it has risen sharply since the beginning of this decade. By the end of the sample, it is more than six times as large as the contemporaneous standard deviation of the permanent component, and nearly twice as large as the maximum attained by the latter in the early 1980s.

As discussed by SW, the decline in the persistence of inflation after 1980 can be explained by the drop in the variance of the permanent component; as this variance declined over the second half of the 1980s and the early 1990s, movements in CPI inflation came to be dominated by the temporary component. Inflation became harder to forecast, even as the variance of inflation was falling. The relative importance of the permanent component has fallen even further in this decade—thus making inflation even harder to forecast—but that's happened because the variance of the temporary component has increased sharply.[9]

[9] The finding that the increase in variability is concentrated at the high frequencies does not hinge upon the functional form that is imposed upon the inflation process, but is evident in the raw inflation data itself. For instance, if a 10-year rolling window is used to calculate the variance of monthly inflation, there is a noticeable drop in the variance of inflation beginning (with samples that end) in the mid-1990s (which is the same time that the

Assume, now, that the only change that has taken place in the inflation process recently is an increase in the variance of the temporary component. Intuitively, this means that the current level of inflation has become a more noisy indicator of future inflation than before. The literature on inference suggests that when a signal becomes more noisy one should pay less attention to it.

To see how this intuition applies to the case at hand, suppose (first) that the inflation process is given by the Stock-Watson specification. For simplicity, also assume that households know the inflation process and the current period temporary shock. Expected inflation next period is then given by:

$$E_t \pi_{t+1} = \tau_t = \pi_t - \eta_t$$

Regressing this forecast of next period's inflation rate on today's inflation rate (an exercise similar to regression (2) above) leads to the following estimated coefficient:

$$1 - \frac{var(\eta)}{t\ var(\varepsilon) + var(\eta)}$$

where *t* is the sample size. For the fixed sample size used in the rolling regressions above, this coefficient will decline as the variance of the temporary component (η) increases. Alternatively, if one assumes that the long lived component of inflation is autoregressive of order one with a root 1/α that is close to, but not equal to, one, regressing the inflation rate expected to prevail next period on today's inflation rate leads to the coefficient

$$1 - \frac{\text{var}(\eta) + (1-\alpha)\,\text{var}(\tau)}{\text{var}(\eta) + \text{var}(\tau)}$$

which, again, will be smaller in a regime where the variance of the temporary shock is higher. Thus, even if the change in the inflation process is better modeled as an increase in the variance

regression coefficients change in the charts above), followed by an increase that begins in the early 2000s. This increase is much more obvious when one looks at the difference of inflation (which tends to emphasize higher frequency movements), and by the end of the sample the variance is slightly above the highs of the 1980s . When the same exercise is repeated at the annual frequency, there is only a very small increase in the variance of either inflation or the difference of inflation at the very end of the sample. The results at the quarterly frequency lie in between.

of the temporary component, the sum of coefficients plotted in the top panel on the right hand side of Figure 2 should decline over time.

3. The forecasts from the SPF survey

This section examines the forecasts from the survey of professional forecasters. The left hand column of charts in Figure 4 repeats –for the SPF forecasts--the exercise seen in Figure 2 above, that is, it shows what happens when the year-ahead SPF inflation forecasts are regressed on current and lagged inflation. The results turn out to be similar to those for actual inflation (see the left hand column in Figure 2). Thus, the forecasters in the SPF panel appear to be placing less weight on recent inflation data, and one could conclude that the professionals have recognized the change that has taken place in the inflation process, much as hypothesized by ABW.

However, a closer look at the data reveals that there is another dimension along which the forecasters' behavior looks quite different. The charts on the right hand side of Figure 4 show what happens when the SPF forecasts are projected on core CPI inflation data. If anything, the SPF forecasts have become more sensitive to core CPI data in recent years, though---given the size of the two-standard-error band--- one cannot reject the argument that there has been no change in their response to these data over the entire sample. Table 2 provides more direct evidence on these issues. The first column presents the estimates from a full sample regression of the SPF forecasts on quarterly CPI inflation, in which the constant and the coefficients on the inflation terms are allowed to change approximately midway through the sample (specifically, at the end of 1993, to allow comparison with the results in Table 1). The sum of the coefficients on the inflation terms interacted with the dummy variable is negative and significantly different from zero. As indicated at the bottom of the table, these variables cannot be excluded from the equation at the one percent level. And one cannot reject the hypothesis that during the second half of the sample, changes in CPI inflation have no permanent effect on the SPF forecasts. The results in the second column, where the SPF forecast is regressed on *core* CPI inflation, are quite different from those in the first. The coefficients on the inflation terms that have been interacted with the dummy are insignificantly different from zero and one can easily reject the hypothesis that the sum of all the coefficients on the core inflation terms is zero in the second half of the sample. Note that the fit of this equation is marginally better than the first, in which the forecasts are projected on headline inflation measures.

One way to reconcile these results is to argue that the SPF forecasters used to pay attention to both the core and non-core components of inflation until recently and now pay attention only to the former. An alternative argument is that the SPF forecasters always paid attention to the core CPI and not the headline, but this has only become obvious following the recent decline in the correlation between core and headline CPI inflation. I discuss each of these possibilities in turn. Before doing so, it is worth noting the results of an experiment meant to distinguish between the two. Specifically, the exercise in Figure 4 was repeated, except that the core and non-core[10] components of CPI inflation were entered separately. It turns out that while the noncore inflation terms did not account for much of the variation in the SPF forecast, they could not be excluded from the SPF regression for samples that end before 2002; after that, the evidence is mixed, with the noncore components significant for some samples and not others.

Why might the professional forecasters have reduced the attention they pay to the noncore component? One possible reason is that the increased noise identified earlier in the CPI is concentrated in this component, which would suggest that one should reduce the weight one attaches to the noncore component but continue to pay attention to the core inflation rate.

In order to see if the data are consistent with this hypothesis, Figure 5 shows the results obtained when the SW specification is imposed upon the core CPI inflation process. The decline in the variance of the permanent component is similar to that seen in the case of the headline CPI. Importantly, while the variance of the temporary component has been going up recently, the increase is nowhere near as marked as it was for headline CPI inflation.[11]

Thus, an argument can be made that because of recent changes in the inflation process, it is appropriate to pay less attention to the non-core component of CPI inflation.[12] But that does not justify the SPF forecasters practice of continuing to place a large weight on core inflation data when predicting headline inflation. The first column in Figure 6 demonstrates this point. It

---

[10] Non-core inflation is defined as the rate of headline inflation relative to core, following Stock and Watson (2008).

[11] In contrast to the headline CPI (see footnote 9), the raw data for core CPI do not provide clear evidence of changes in volatility.

[12] Even in this case, the practice of placing a zero weight on food and energy price data is hard to justify. First, it ignores the dynamics of the known temporary component, which is unlikely to be white noise. More problematic is the assumption of independence between the core and non-core components. Specifically, the non-core component will affect the core component (given the range of historical responses of monetary policy to such shocks); e.g., oil shocks are likely to affect other prices in the economy and so should be taken into account.

shows what happens when headline CPI inflation over the next year is regressed upon core CPI inflation, similar to the rolling regressions seen earlier. As can be seen, the relationship between headline and core inflation has deteriorated quite noticeably in recent years. For samples ending in the last four to five years, core CPI inflation data provide no information about future headline inflation, which is almost exactly the same result shown in Figure 2--where the right hand side variable was headline CPI inflation. This result could not be more different from that in Figure 4, where the SPF forecasts of headline inflation appear to have become more sensitive to core inflation data in recent years.

These results make one wonder whether it is only the relationship between headline and core inflation that has changed or if there has been a change in the behavior of core inflation as well. The panels on the right hand side of Figure 6 address this issue by regressing 1-year-ahead core CPI inflation on quarterly core CPI inflation. While the deterioration in the predictive power of the equation is not as great as when headline CPI inflation is regressed on itself (see Figure 2), the pattern is similar. The relationship begins to deteriorate by the mid-2000s; by the end of the sample period, the sum of the coefficients on inflation cannot be distinguished from zero and the adjusted-R^2 is below 0.1. The data suggest that these regressions could look worse as time goes by; when a 10-year rolling window is used (instead of the 15-year window used for the graphs), the sum of the coefficients on core inflation falls below zero in 2006 and continues to fall through the end of the sample.

As mentioned above, the other possibility is that the SPF forecasters have always relied on core inflation to forecast headline inflation, but this has only become obvious following the recent change in the relationship between the two. The SPF forecasters would not be unique in following such a procedure, if this is indeed what they were doing. For instance, Blinder and Reis (BR, 2005) argue that even if one is interested in headline inflation, it is better to generate forecasts of this variable by using data on core inflation. Based on a series of results for forecasting inflation at the 6, 12, 24 and 36 month forecasting horizons, they state that:

> "Every specification in the table points to the same conclusion: that recent core inflation is a better predictor of future *headline* inflation than is recent headline inflation…. Indeed, once you take core inflation into account, adding headline inflation has at best no effect on forecasting performance, and at most horizons makes forecasts even worse."

So, one could argue that the forecasters' decision to focus on core inflation was a reflection of prevailing opinion.[13] But, even if this was the right way to proceed in the past, should the forecasters have continued to do so in light of the evidence above?

Table 3 provides a comparison of different ways of forecasting CPI inflation since the beginning of 2003, which is just after the ABW (2007) sample ends and close to the time that the inflation process appears to have changed. The Michigan and SPF forecasts are compared to forecasts from several alternative specifications: A random walk, Stock and Watson's unobserved-component stochastic-volatility model, and three regression based specifications. Two of the regressions involve only inflation data (either headline or core) while the third is a Phillips curve specification, which adds the unemployment rate and noncore inflation. Each forecast from these equations is obtained by regressing inflation on a constant and the explanatory variable(s) over a 15 year sample that ends in the period prior to the forecast. The estimated equation is then used to forecast next period's inflation. The specification used here is the same as that used by BR (2005).

Timing issues become important when the regression based forecasts are compared to the survey forecasts. Both ABW and BR include the latest inflation data on the right hand side when estimating equations to predict future inflation.[14] As inflation data are released with a lag, this means that the forecasts obtained from the regressions will be based on more information than the survey respondents had when they made their forecasts. Note that the difference is less than one quarter, as survey respondents have access to monthly data but are being asked to forecast quarterly inflation once a quarter (SPF) or annual inflation every month (Michigan). More specifically, for the Michigan survey, this paper uses data from the final reading in the third month of the quarter, which is released at the end of that month. Since CPI data tend to be released around the middle of the month, the quarterly Michigan forecast used here is likely to be based on the monthly inflation rate for the first two months of the quarter. This is quite good, since the last month of data has a low weight in calculating the quarterly average inflation rate. Things are different for the SPF survey, as it is conducted once a quarter and is released in the middle of the middle month of the quarter. Thus, depending upon survey and data release dates, SPF survey respondents may or may not have information about CPI inflation in the first month

[13] However, such an opinion was not universally held. For an alternative, see Smith (2005).
[14] BR do not compare the equation based forecasts with the survey forecasts.

of the quarter. In order to avoid stacking the deck against the survey respondents, this paper excludes quarter *t* information from the right hand side of the estimated equations. In the second specification in Table 3A, for instance, inflation from quarter *t* to quarter *t+4* is predicted using inflation data through quarter *t-1*. Thus, the timing convention here differs from both ABW and BR.

The results turn out to be very different from these studies as well, though not because of the timing convention. Using lagged CPI inflation to predict headline inflation leads to the smallest root mean squared error in Table 3. According to the Diebold-Mariano-West (or DMW) test,[15] these forecasts are better than the SPF forecasts at the 1 percent level. Forecasts from the Phillips curve specification are better than the SPF forecasts at the 10 percent level. Forecasts based on lagged core inflation are the third most accurate (out of seven), though the associated RMSE is close to that for the SPF forecasts. This similarity is not surprising given the results above suggesting that the SPF forecasts can be well described as a linear combination of lagged core CPI inflation data. The Michigan survey forecasts, the forecasts from the unobserved components stochastic volatility model and the random walk specification bring up the rear, with RMSEs that are just above 2 percent.

For comparison, the lower panel of Table 3 presents results from a sample of the same size which ends at the end of 2002. Here we replicate the results found by ABW and BR. In particular, as pointed out by ABW, the SPF survey does best of all, and the Michigan survey has almost exactly the same RMSE. The forecasts from the SPF survey are better than those based on lagged core inflation at the 1 percent level and are better than those from the Phillips curve specification at the 5 percent level (DMW test again). And consistent with BR, forecasts of headline CPI inflation based on core CPI inflation are better than those based on headline CPI data.

The RMSEs in the post 2002 sample are noticeably larger than the earlier sample for every specification in Table 3. This reflects the increase in high frequency noise in CPI inflation over this period. Even so, the deterioration in the SPF forecast is surprising. Since 2003, professional forecasters are doing worse than a forecast based on headline inflation alone; here, it is worth pointing out that the latter specification has almost no ability to explain inflation within sample (in terms of the adjusted-R^2). This evidence makes one wonder whether the SPF

[15] See West (2006).

forecasters' apparent decision to pay little or no attention to the non-core inflation data in recent years and to continue to place a large weight on the core inflation data was motivated by a desire to predict headline inflation more accurately or in pursuit of some other objective.

BR (2005) provide an interesting rationale for what might be going on:

> "Another Greenspan innovation, which is rarely mentioned but is likely to prove durable, is the way he has focused both the Fed and the financial markets on core, rather than headline, inflation. This aspect of Federal Reserve monetary policy contrasts sharply with the concentration on headline inflation at the ECB and to the stated inflation targets of most other central banks, which are rarely core rates."

Perhaps there is more than one realizes to the Blinder and Reis argument that Chairman Greenspan turned public attention towards the core inflation data. Could it be that the SPF forecasters have followed the Fed and switched their attention to forecasting core CPI, even though they are being asked to forecast headline CPI?

Table 4 presents some evidence that is consistent with this hypothesis. It shows what happens (over the period since the beginning of 2003) when the SPF forecast of headline CPI inflation is treated as a forecast of core CPI inflation. For comparison, forecasts from the other specifications in Table 3 are also included, with the exception of the one that uses headline CPI on the right hand side. As can be seen, the SPF forecast of *headline* inflation turns out to be a pretty good predictor of *core* inflation. It is better than the random walk specification at the 1 percent level. And the RMSE of the SPF forecast is slightly better than that obtained when core CPI inflation is used to predict future core inflation, though the difference is nowhere near statistically significant. The Phillips curve specification turns out to be the best, though it outperforms the SPF forecast only at the 10 percent level. The Michigan survey is the worst by a wide margin. This is not a big surprise, as the respondents are being asked to predict headline – and not core--inflation; the surprise is that the SPF survey does so well.

Section 4: Conclusions

The evidence suggests that the inflation process has changed in recent years. The autoregressive representation of CPI inflation shows a noticeable decline in the sum of the coefficients on lagged inflation over time and by the end of the sample lagged inflation data have no predictive power for future inflation at all. If, instead, Stock and Watson's unobserved-component-stochastic-volatility specification is imposed on the data, the change shows up as a noticeable increase in the variance of the high frequency component. While it may be hard to determine the correct representation of the data, either kind of change should lead survey respondents to place less weight on recent inflation data when predicting future inflation.

Households do not appear to have learned about this change in the inflation process, as they do not appear to have changed the way in which they form expectations of inflation. Historically, households have placed a large weight on recent inflation data when forming inflation expectations, and they continue to do so now. The effects of this mistake show up in a marked deterioration in forecasting performance, as the Michigan forecasts have gone from being about the most accurate to the least accurate.

There is more reason to believe that professional forecasters have changed the way that they forecast inflation. They now seem to react very little, if at all, to noncore inflation data; at the same time, though, they do not appear to have changed the way they react to core inflation data. However, the changes in the inflation process documented above suggest that this strategy may be problematic, an assessment that is borne out by the noticeable deterioration in the relative forecasting performance of the professionals. These results suggest that professionals are placing too much weight on recent core inflation, just as households are placing too much weight on recent headline inflation.

The evidence presented above is also consistent with the interpretation that the professionals have stopped worrying about headline inflation and are now focusing on core CPI inflation. To the extent that this is a recent switch, and possibly one encouraged by then-Chairman Greenspan's advocacy of the core inflation rate, it suggests that analyses which use data from expectations surveys to determine how agents learn about the economy need to account for a wide variety of influences on agents.

References

Ang, Andrew, Geert Bekaert and Min Wei (2007) "Do macro variables, asset markets, or surveys forecast inflation better?" *Journal of Monetary Economics* 54, pp. 1163-1212.

Bai, Jushan and Pierre Perron. (1998) "Estimating and Testing Linear Models with Multiple Structural Changes," *Econometrica*, January, pp. 47-78.

Bernanke, Ben, Thomas Laubach, Frederic Mishkin and Adam Posen. (2001) "*Inflation Targeting: Lessons from the International Experience,*" Princeton University Press.

Blanchard, Olivier J and Jordi Gali. (2007) "The Macroeconomic Effects of Oil Shocks: Why Are the 2000s so different from the 1970s?" NBER Working Paper 13368.

Blinder, Alan S. and Ricardo Reis. (2005) "Understanding the Greenspan Standard," mimeo, Princeton University.

Carroll, Christopher. (2003) "Macroeconomic Expectations of Households and Professional Forecasters," *Quarterly Journal of Economics*, vol. 108, pp. 269-298.

Clark, Todd E. and Taisuke Nakata. (2008) "Has the Behavior of Inflation and Long-Term Inflation Expectations Changed?" *Economic Review*, Federal Reserve bank of Kansas City, First Quarter, pp. 17-50.

Cogley, Timothy and Thomas J. Sargent. (2005) "Drifts and volatilities: monetary policies and outcomes in the post WWII US," *Review of Economic Dynamics,* pp. 262-302.

Coibion, Olivier and Yuriy Gorodnichenko. (2008) "What can survey forecasts tell us about informational rigidities?" mimeo.

Croushore, Dean. (1998) "Evaluating Inflation Forecasts," Federal Reserve Bank of Philadelphia Working Paper No. 98-14.

Leduc, Sylvain, Keith Sill and Tom Stark. (2007) Self-fulfilling expectations and the inflation of the 1970s: Evidence from the Livingston Survey," *Journal of Monetary Economics*, vol 54, pp. 433-459.

Levin, Andrew and Jeremy Piger. (2004) "Is Inflation Persistence Intrinsic in Industrial Economies?" European Central Bank Working Paper 334, April.

Mehra, Yash P. and Christopher Herrington. (2008) "On the Sources of Movements in Inflation Expectations: A few Insights from a VAR Model," *Economic Quarterly*, 94(2), Spring, pp. 121-146.

Mankiw, N.Gregory, Ricardo Reiss and Justin Wolfers (2003) "Disagreement About Inflation Expectations," NBER Working Paper 9796.

Mishkin, Frederic S. (2007) "*Inflation Dynamics,"* NBER Working Paper 13147.

Mullineaux, Donald J. (1980) "Inflation Expectations and Money Growth in the United States," *The American Economic Review*, 70(1), pp. 149-161.

Orphanides, Athanasios and John C. Williams. (2005) "The decline of activist stabilization policy: natural rate misperceptions, learning, and expectations", *Journal of Economic Dynamics and Control*, 29(11), pp.1927-1950.

Pesando, James E. "A Note on the Rationality of the Livingston Price Expectations," *The Journal of Political Economy*, 83(4), pp. 849-858.

Pesaran, M. Hashem and Martin Weale. (2006) "Survey Expectations," pp. 746-775 in Handbook of Economic Forecasting, G Elliott, C.W.J. Granger and A. Timmerman (eds.), North Holland.

Pivetta, Frederic and Ricardo Reis. (2007) "The Persistence of Inflation in the United States," *Journal of Economic Dynamics and Control*, pp. 1326-1358.

Primiceri, Giorgio E., "Time Varying Structural Vector Autoregressions and Monetary Policy". *Review of Economic Studies*, Vol. 72, No. 3, pp. 821-852, July 2005

Sims, Christopher A. (1999) "Drifts and Breaks in Monetary Policy", unpublished.

Sims, Christopher A, and Tao Zha (2006) "Were There Regime Switches in Monetary Policy," *American Economic Review*, pp. 54-81.

Smith, Julie K. (2004) "Weighted Median Inflation: Is This Core Inflation?" *Journal of Money, Credit and Banking*, Vol 36(2), April, pp.253-263.

Stock, James and Mark Watson. (2007) "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money Credit and Banking*, Supplement to Vol. 39, no. 1, February, pp. 3-33.

---- (2008) "Phillips Curve Inflation Forecasts," mimeo, May.

Taylor, John. (2000) "Low inflation, pass-through and the pricing power of firms," *European Economic Review* 44, pp. 1389-1408.

West, Kenneth D. (2006) "Forecast Evaluation," pp. 99-134, *Handbook of Economic Forecasting*, Vol 1, G. Elliott, C. W. J. Granger and A. Timmerman (eds), Amsterdam: Elsevier.

Table 1. Projections of 1-year-ahead Realized and Expected CPI Inflation on Realized Inflation

Sample: 1978Q1-2009Q3

	Realized Inflation	Expected Inflation – Michigan Survey
Constant	1.44[1] (0.43)	1.01[1] (0.16)
D94	2.62[5] (0.87)	0.45[10] (0.24)
ΣPi	0.66[1] (0.08)	0.63[1] (0.03)
ΣDPi	-1.24[1] (0.33)	-0.09 (0.09)
$\overline{R}^2$	0.69	0.93
Exclude ΣDPi*	2.39[1]	1.29
ΣPi + ΣDPi = 0**	3.26[10]	44.68[1]

ΣPi is the sum of the coefficients on realized CPI inflation, ΣDPi is the sum of the coefficients on the realized inflation terms multiplied by the dummy D94, which equals 0 until the end of 1993 and 1 after that. HAC standard errors are reported in parentheses.

[1] denotes significant at 1 percent; [5] denotes significant at 5 percent; [10] denotes significant at 10 percent.

[*] F statistic for null that all 8 DPi terms can be excluded from the equation.

**Chi-square statistic for null that the sum of the coefficients on inflation is equal to zero in the second half of the sample.

Table 2. Projections of 1-year-ahead SPF Inflation Forecast
Sample: 1981Q3-2009Q3

	On Headline Inflation	On Core Inflation
Constant	2.22[1] (0.26)	1.72[1] (0.28)
D94	0.13 (0.39)	-0.30 (0.39)
ΣPi	0.50[1] (0.05)	0.55[1] (0.04)
ΣDPi	-0.46[1] (0.09)	-0.09 (0.15)
$\overline{R}^2$	0.88	0.91
Exclude ΣDPi*	4.49[1]	1.56
ΣPi + ΣDPi = 0**	0.32	10.53[1]

ΣPi is the sum of the coefficients on realized CPI inflation, ΣDPi is the sum of the coefficients on the realized inflation terms multiplied by the dummy D94, which equals 0 until the end of 1993 and 1 after that. HAC standard errors are reported in parentheses.

[1] denotes significant at 1 percent; [5] denotes significant at 5 percent; [10] denotes significant at 10 percent.

[*] F statistic for null that all 8 DPi terms can be excluded from the equation.

**Chi-square statistic for null that the sum of the coefficients on inflation is equal to zero in the second half of the sample.

Table 3. Predicting 1-year-ahead headline CPI Inflation

A. Sample: 2003:Q1 – 2009:Q3
(27 observations)

Using:	Random walk	Lagged headline inflation only	Lagged core inflation only	Phillips Curve	UC-SV Model	Michigan	SPF
Mean Error	-0.15	0.07	0.15	0.27	1.34	-0.52	0.30
Root Mean Square Error	2.05	1.41[1]	1.59	1.47[10]	2.03	2.04	1.62

Notes: [1] Better than the SPF forecast at 1 percent (Diebold-Mariano-West MSE test).
[10] Better than SPF forecast at 10 percent (DMW test).

B. Sample: 1996:Q2 – 2002:Q4
(27 observations)

Using:	Random walk	Lagged headline inflation only	Lagged core inflation only	Phillips Curve	UC-SV Model	Michigan	SPF
Mean Error	-0.21	-0.64	-0.33	-0.22	1.29	-0.37	-0.28
Root Mean Square Error	1.13	1.11	0.99	0.98	1.50	0.90	0.89[1,5]

Notes: [1] Better than lagged core inflation forecast at 1 percent
[2] Better than Phillips curve specification at 5 percent

Table 4. Predicting 1-year-ahead core CPI Inflation

Sample: 2003:Q1 – 2009:Q3
(27 observations)

Using:	Random walk	Lagged core inflation only	Phillips Curve	UC-SV Model	Michigan	SPF
Mean Error	-0.09	-0.09	-0.12	1.05	-1.08	-0.26
Root Mean Square Error	0.63	0.57	0.43[10]	1.06	1.30	0.51[1]

Notes: [1] Better than Random walk forecast at 1 percent (DMW MSE test)
[10] Better than SPF forecast at 10%

Figure 1: 1-year-ahead Expected Inflation

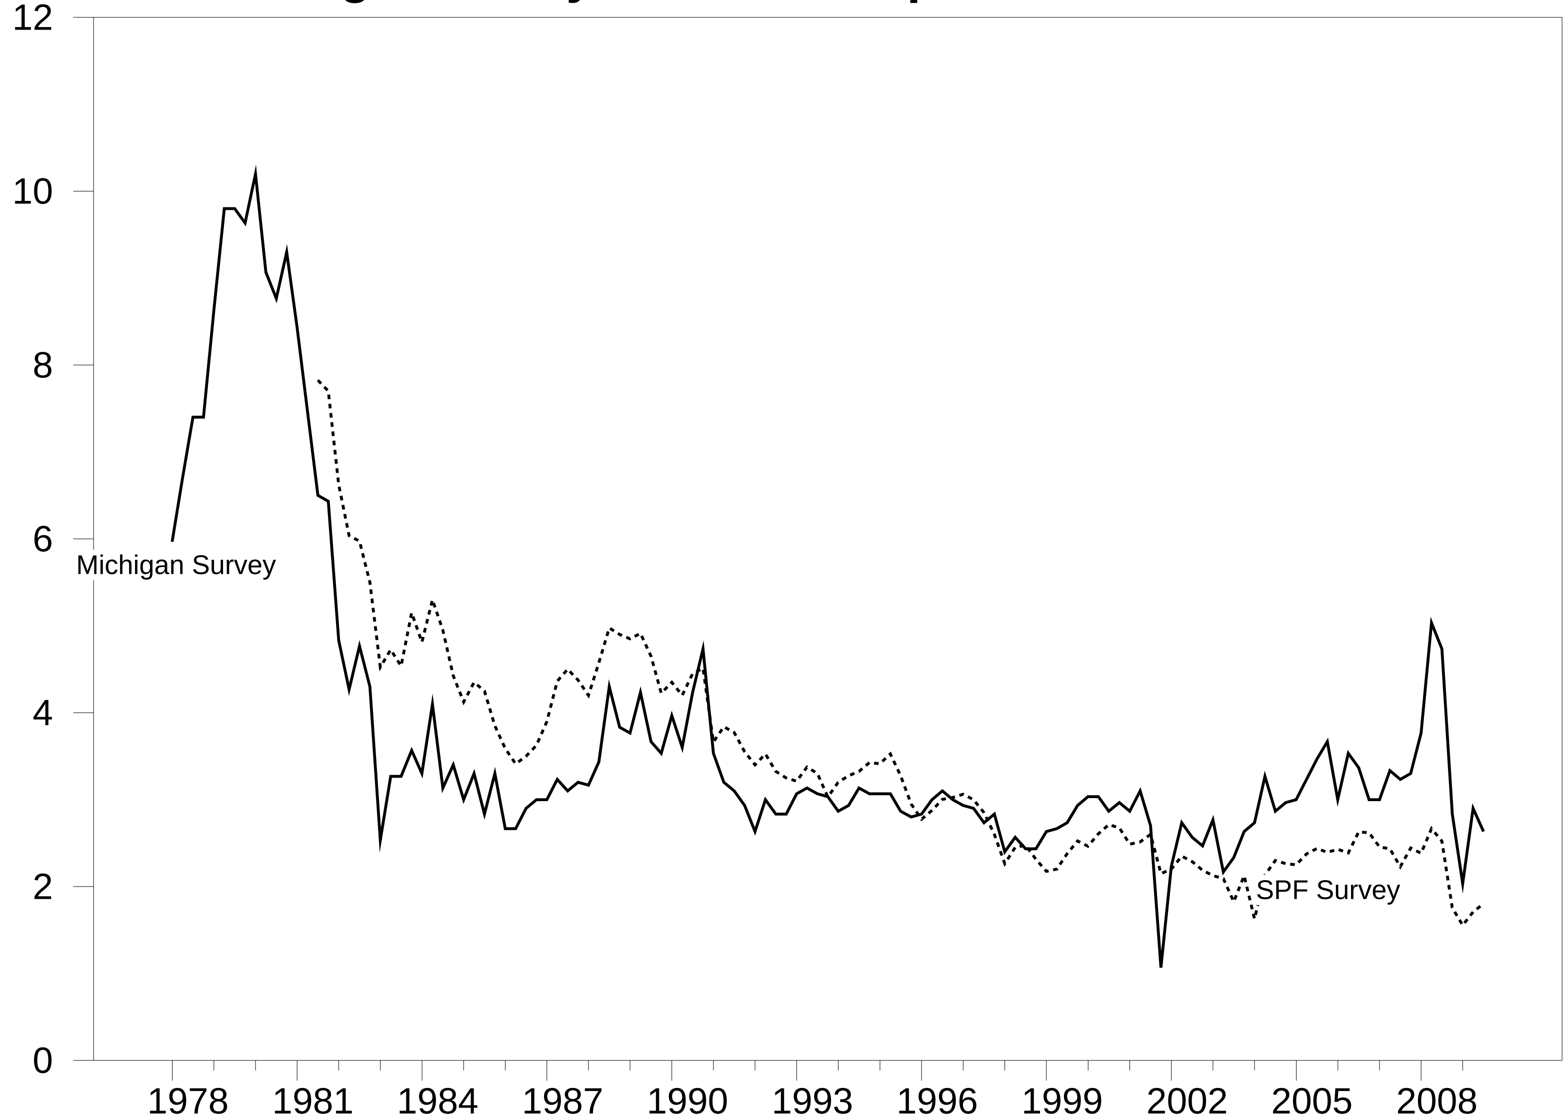

Figure 2: Projections on quarterly CPI inflation

15 year rolling sample

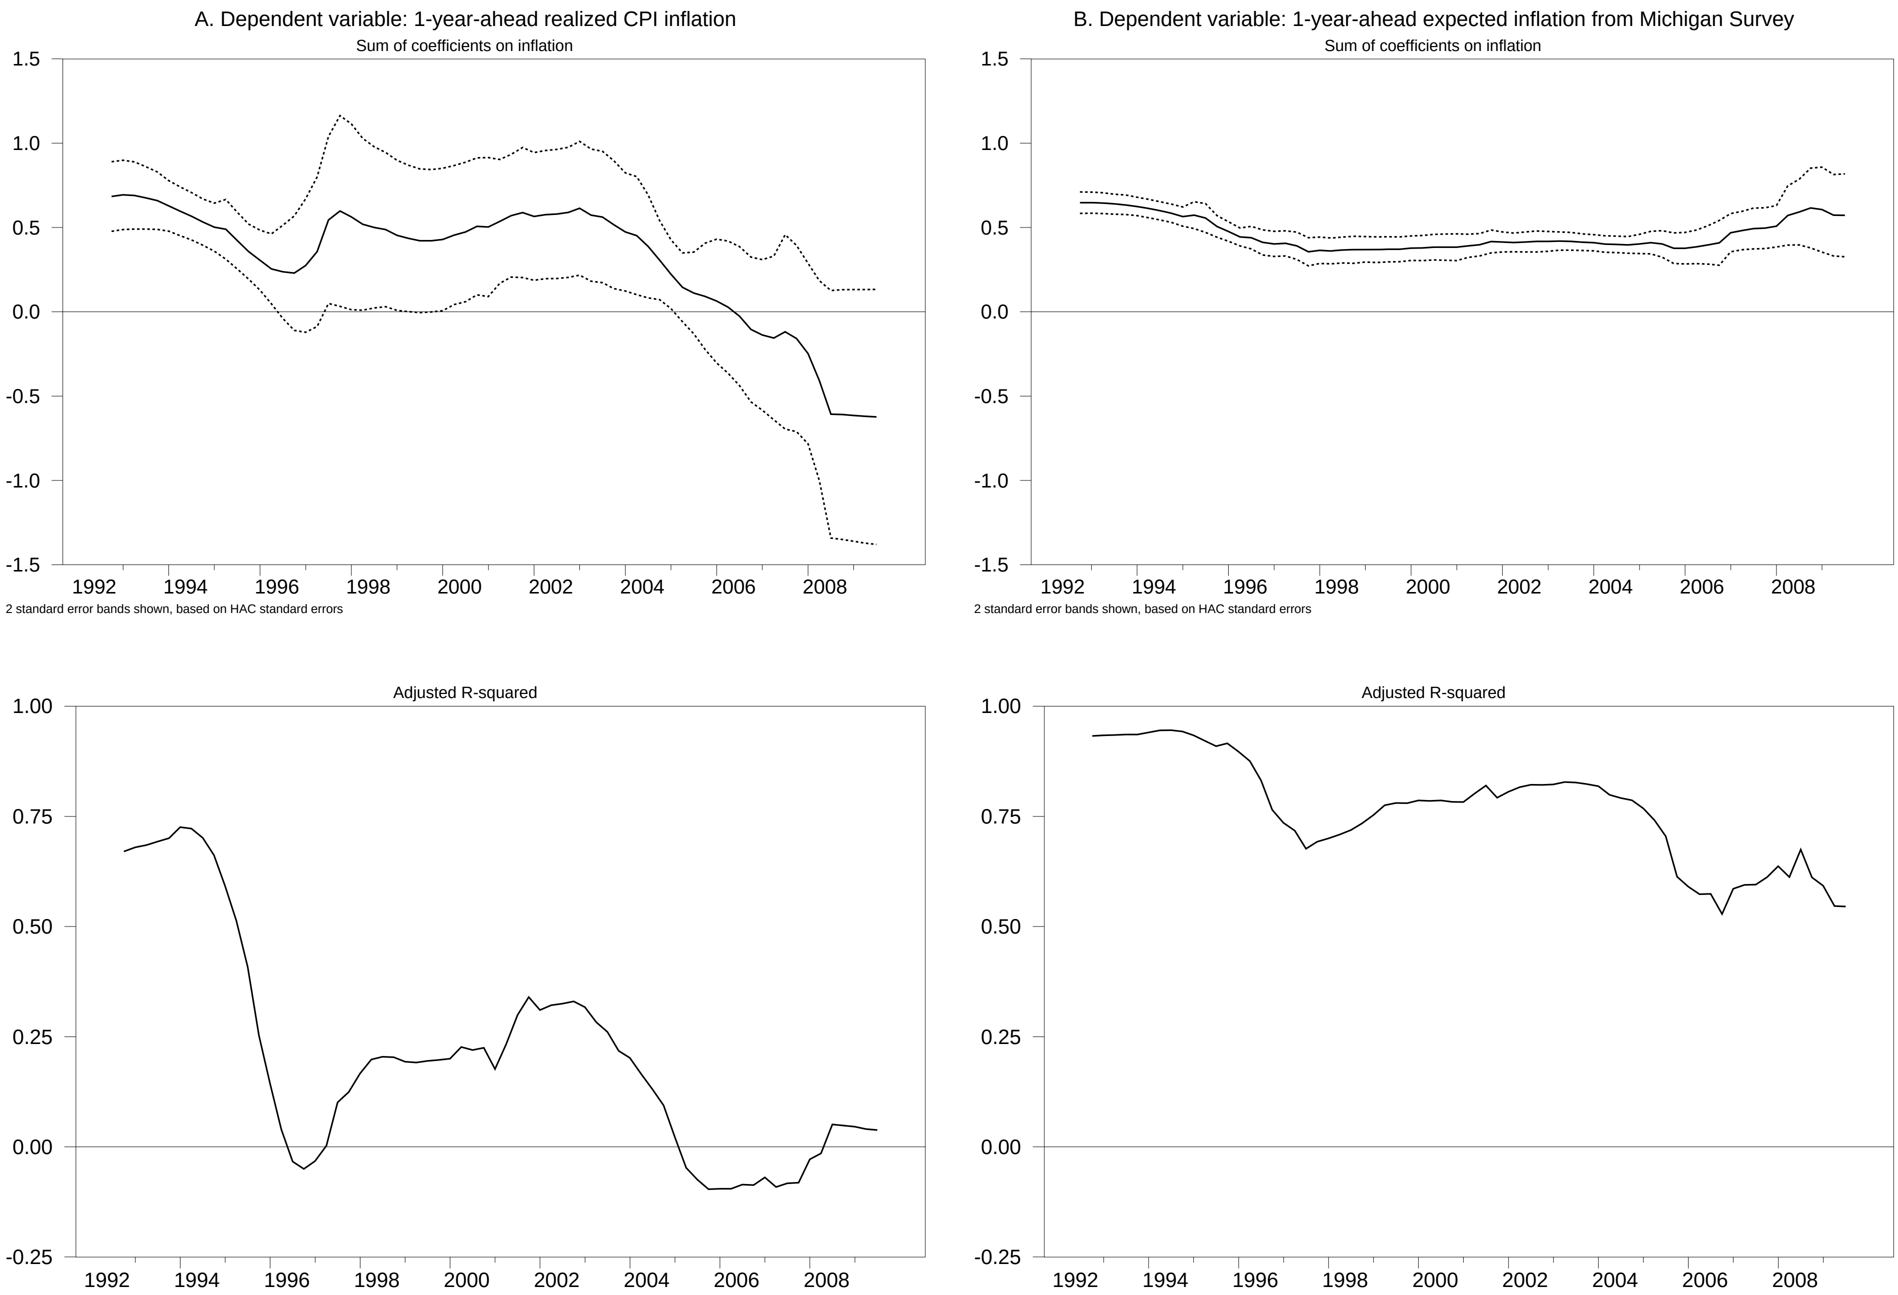

Figure 3: Estimates from unobserved component stochastic volatility model for CPI inflation

1960Q1-2009Q3

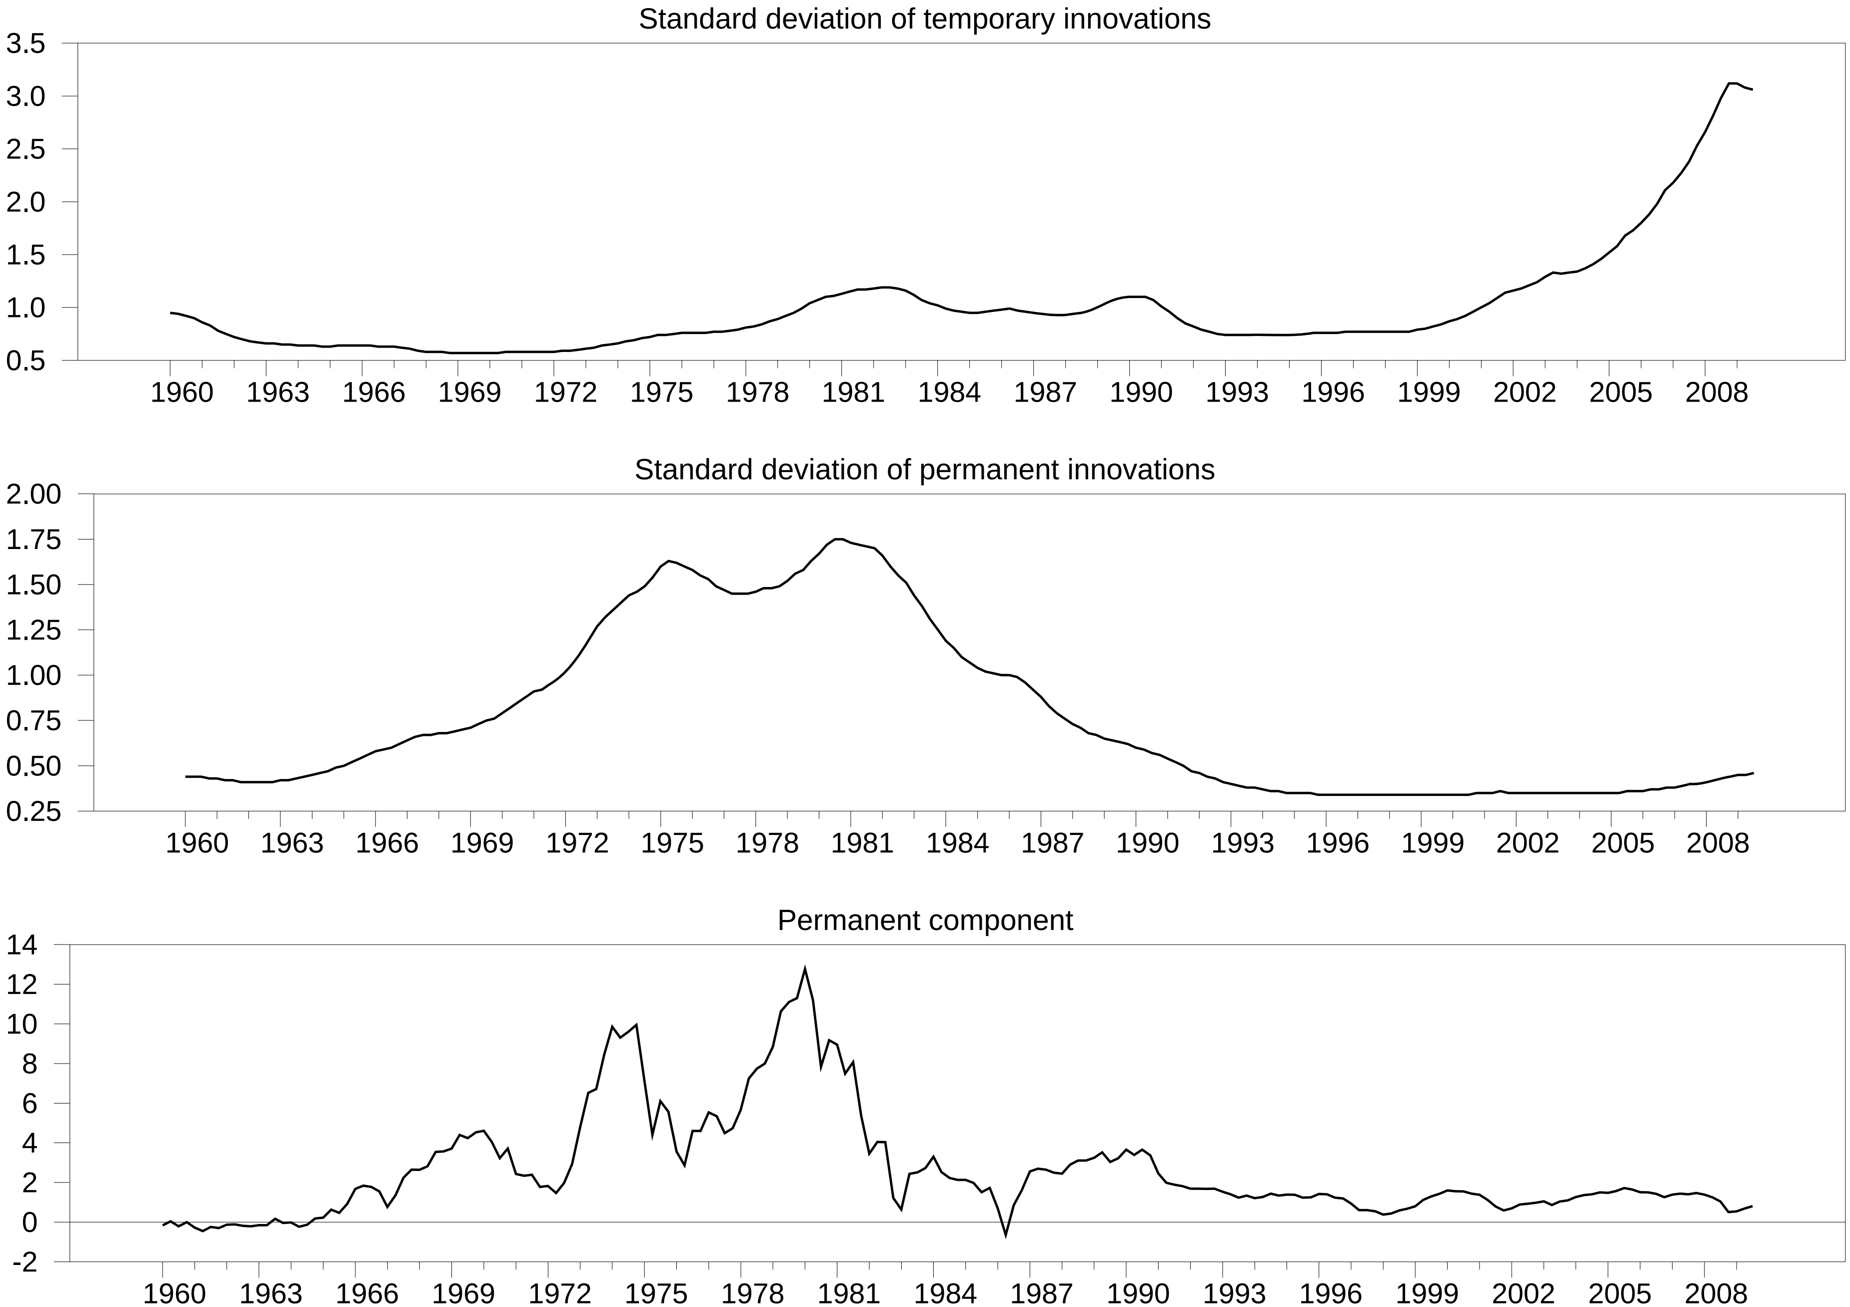

Figure 4: Projections of the 1-year-ahead SPF inflation forecast

15 year rolling samples

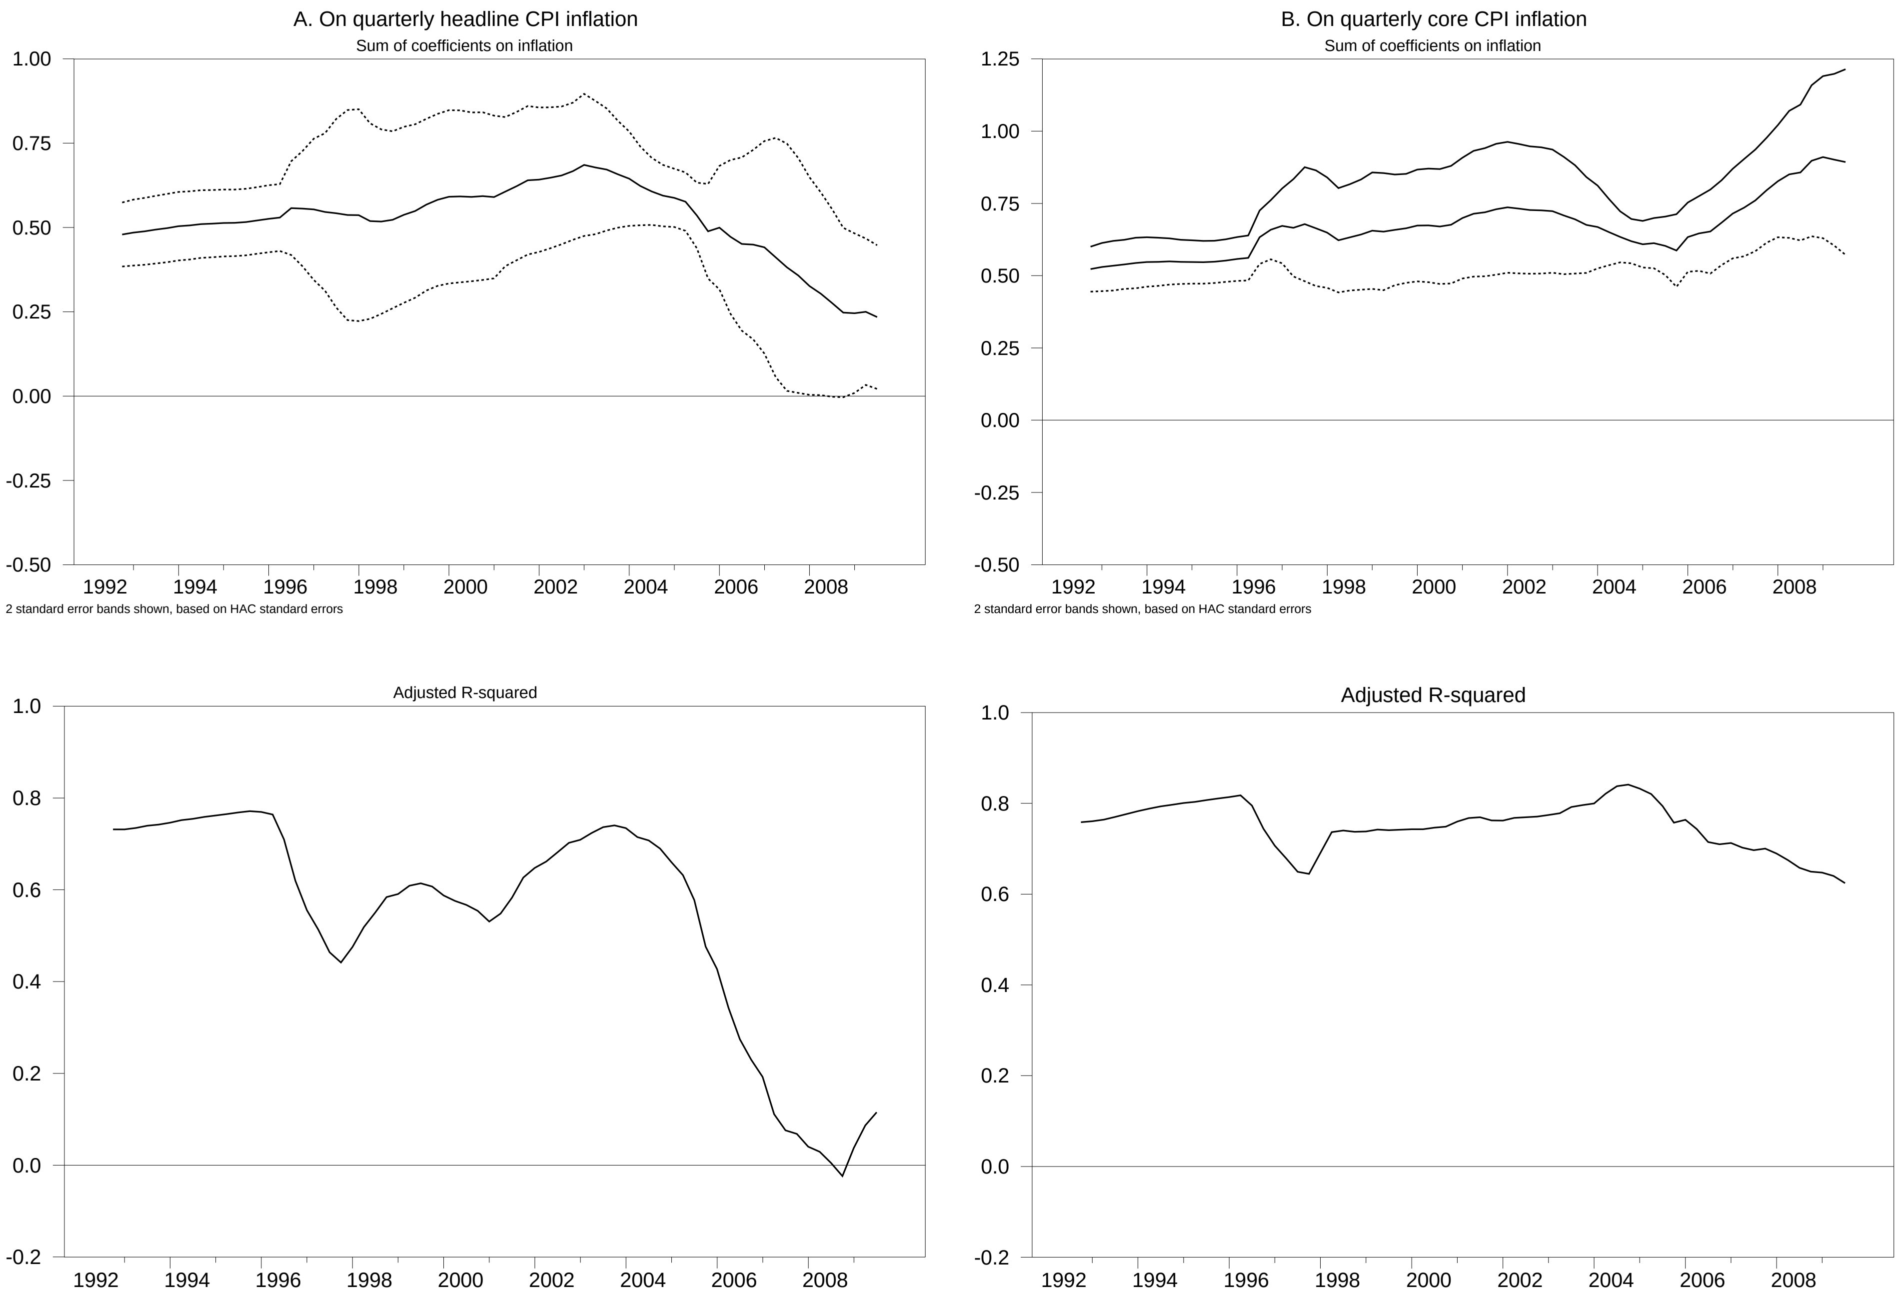

Figure 5: Estimates from unobserved component stochastic volatility model for core CPI inflation

1960Q1-2009Q3

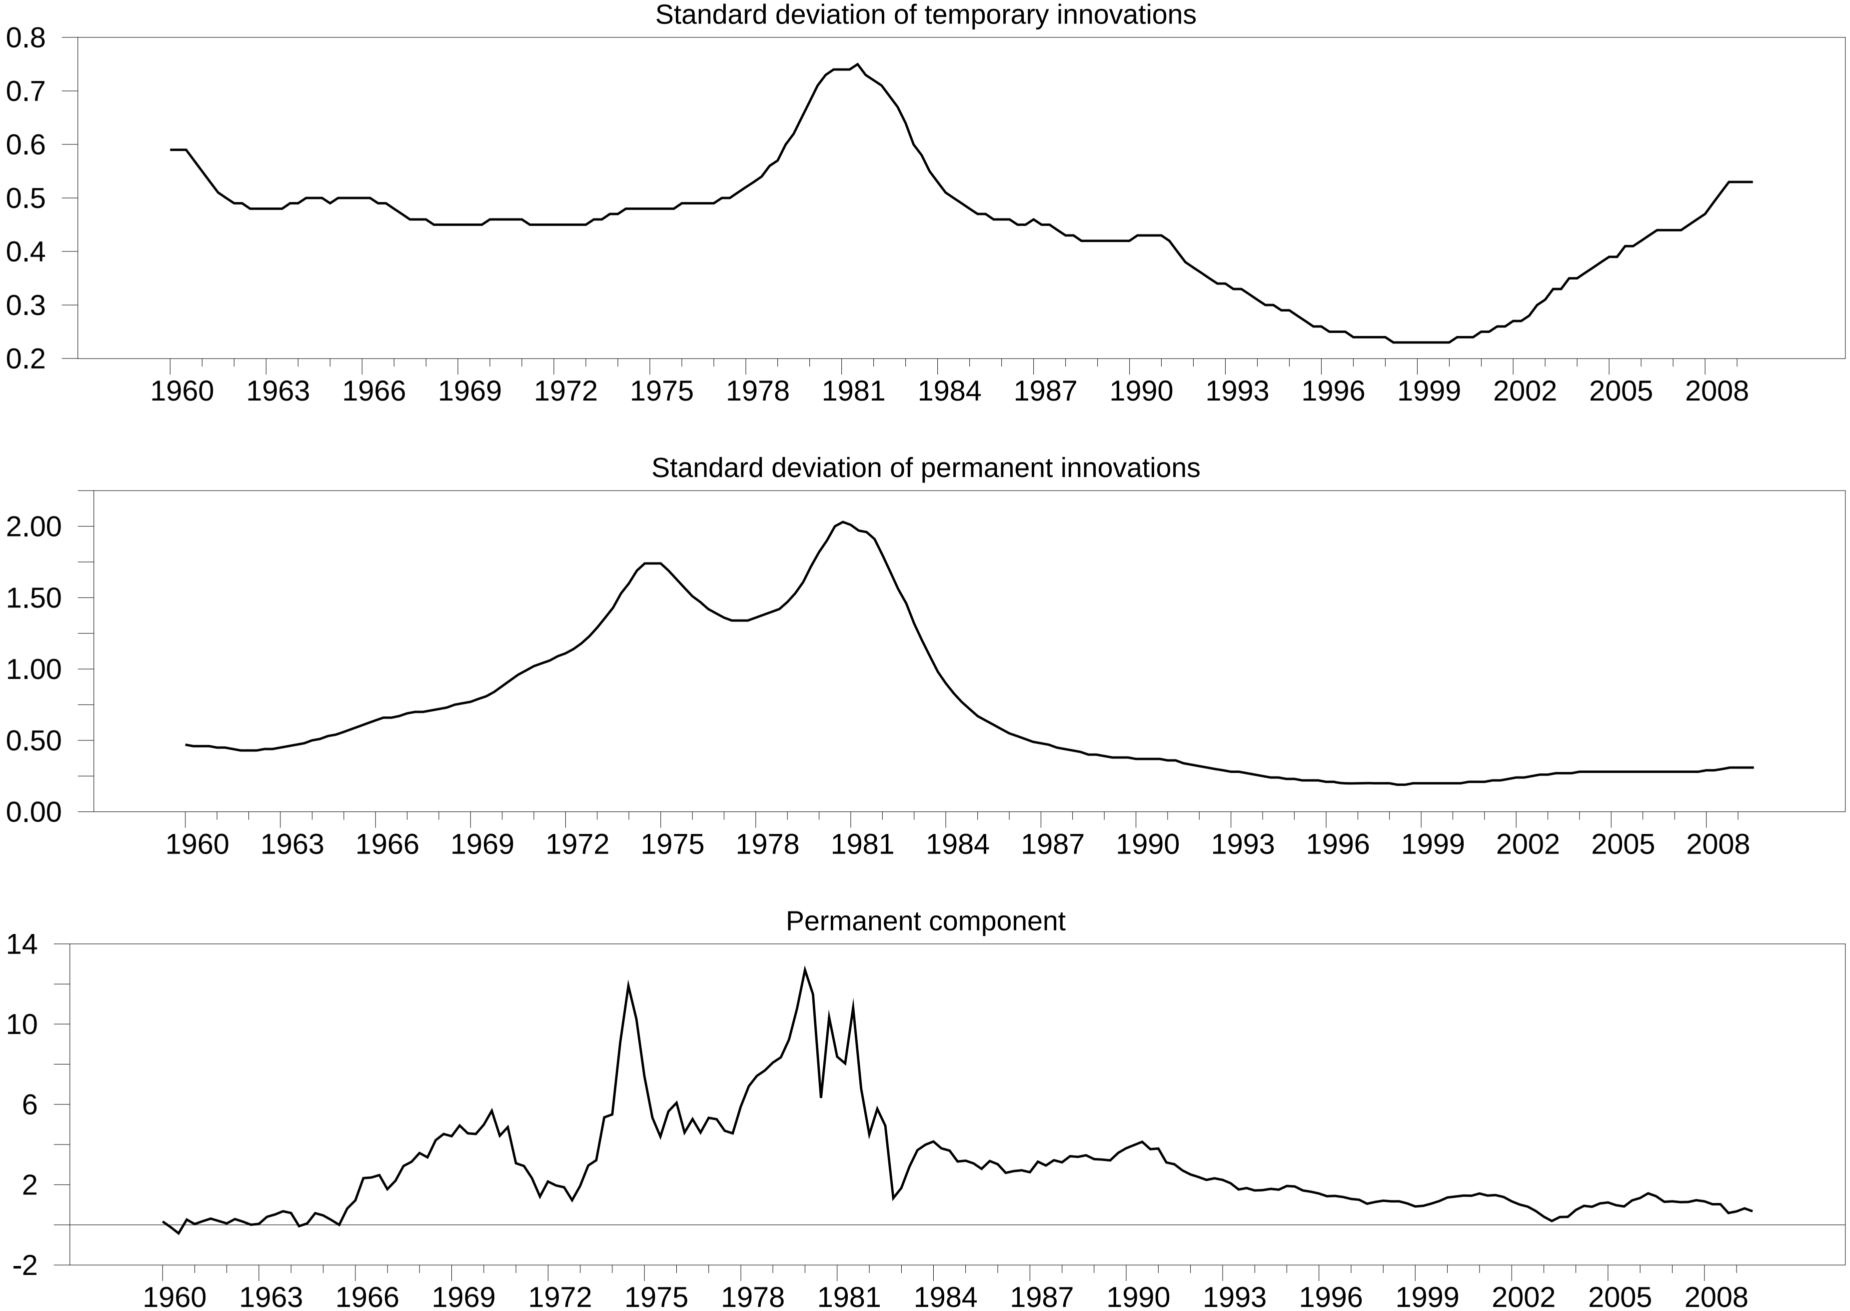

Figure 6: Projections on quarterly CORE CPI inflation

15 year rolling samples