A Simple Adaptive Measure of Core Inflation

Timothy Cogley*

November 1998

Abstract

This paper proposes a new measure of core inflation and compares it with several existing measures. The new measure is adaptive and is designed to track sudden and persistent movements in inflation, such as those arising from changes in monetary policy regimes. The adaptive measure is a superior predictor of (locally) mean reverting components of inflation, and appears to filter out transients more effectively than existing measures.

1 Introduction

In a series of papers, Cecchetti (1996), Bryan and Cecchetti (1993, 1994, and 1995), and Bryan, Cecchetti, and Wiggins (1997) have studied a variety of problems related to the measurement of inflation. Their work is motivated by the observation that as the focus of central banks shifts toward inflation targeting, it becomes increasingly important to have accurate measures of inflation.

One problem concerns the fact that conventional measures of inflation contain a great deal of high frequency noise. The presence of transient noise

*For comments and suggestions, I am grateful to Ken Kasa, Kevin Lansing, Bob Rasche, Glenn Rudebusch, and seminar participants at the Federal Reserve Bank of New York. I am also grateful to Niloofar Badie for research assistance. The main idea in this paper was motivated by discussions with Tom Sargent about policy regime changes and stochastic approximation algorithms. The opinions expressed herein are my own and do not necessarily reflect those of the Federal Reserve Bank of San Francisco or the Federal Reserve System.
Please address correspondence to the Research Department, 1130, Federal Reserve Bank of San Francisco, PO Box 7702, San Francisco, CA 94120, or to tim.cogley@sf.frb.org.

complicates the implementation of an inflation targeting rule, for it makes it hard to forecast changes in inflation one or two years ahead, and difficult to detect early warnings of sustained movements in inflation. In this paper, I propose a simple adaptive method for filtering incoming inflation data to remove the transient noise, and I demonstrate that with appropriately filtered data one can significantly improve inflation forecasts over multi-year horizons.

Various authors have proposed other methods for removing transient noise from inflation data. For example, the Bureau of Labor Statistics computes a "core inflation" measure that consists of a weighted average of all CPI components except for food and energy, which are removed because their prices tend to be more volatile than those of other CPI components. Bryan and Cecchetti also consider measures such as the median and trimmed mean price change among CPI components. These measures extend the BLS approach by automatically excluding large price changes from the CPI basket regardless of the sector in which they arise.

Both the BLS and the Bryan-Cecchetti measures are less volatile than CPI inflation itself, but they retain a substantial amount of high frequency variation. Perhaps as a consequence, they forecast only a small fraction of subsequent changes in inflation over horizons of 2 to 4 years (e.g. see Freeman 1998 and the results reported below). This is problematic, for the difference between actual and core inflation should predict subsequent changes in inflation: when actual inflation is above the core value, inflation should fall as the transients accounting for the high current level die out. The fact that existing measures of core inflation contain a great deal of high frequency variation, and that they do not forecast inflation reversals very well, suggest that there are important transients in the data other than those arising from volatile sectoral components.

This paper proposes a new measure of core inflation, which is motivated in part by Sargent's (1998) work on monetary policy regime changes. I view the problem as one of designing a core inflation measure that adapts quickly to occasional, sudden changes in mean inflation, such as those that occur when central banks adopt new decision rules. The measure that I propose can be interpreted as a "constant-gain" update of mean inflation, or alternatively as the output of a one-sided low-pass filter applied to current and past inflation.

The new measure of core inflation filters out transients more effectively than existing measures. It is smoother than existing measures, and it is a superior predictor of subsequent changes in inflation. Perhaps most importantly, the new measure unlocks much of the predictive power contained in conventional macroeconomic predictors of inflation. For example, in a simple forecasting model involving a measure of the output gap and a short-term real interest rate, the addition of the new core inflation measure raises R^2 statistics for medium-term forecasts from around 10 to 20 percent to between 40 and 60 percent. Thus, the new measure contains substantial, marginal predictive power for locally mean reverting components of inflation.

The remainder of the discussion is organized as follows. Section 2 defines core inflation, discusses the motivation for existing measures, and explains the rationale for the new measure. Section 3 evaluates the new measure and compares it with existing measures by investigating their predictive power for subsequent changes in inflation. The paper concludes with a summary.

2 Definition and Motivation

2.1 Defining Core Inflation

Bryan and Cecchetti (1994) define "core inflation" as "the component of price changes that is expected to persist over medium-run horizons of several years." In other words, core inflation is

$$\pi_{ct} = E_t \pi_{t+H}, \tag{1}$$

where $\pi_t = \ln(P_t/P_{t-1})$ is actual inflation and H is a suitably long forecast horizon. Their definition is analogous to Beveridge and Nelson's (1981) definition of a stochastic trend. Indeed, as the forecast horizon H tends to infinity, π_{ct} converges to the Beveridge-Nelson trend component. In what follows, I use the infinite horizon definition to discuss certain conceptual issues, but when it comes to empirical work I follow Bryan and Cecchetti and focus on the medium-run of several years.[1]

This definition fits naturally within Svensson's (1997) framework for thinking about inflation targeting regimes. Svensson demonstrates that strict inflation targeting implies inflation *forecast* targeting, provided that the central

[1] Quah and Vahey (1995) propose an alternative definition, based on the long run neutrality of inflation. They define core inflation as the "component of measured inflation that has no (medium- to) long-run impact on real output." While this definition is conceptually appealing, it requires a structural model for implementation, and I agree with Bryan and Cecchetti that such models are "difficult to formulate and easy to criticize." Hence I prefer Bryan and Cecchetti's forecast-based definition.

bank minimizes a quadratic loss function. Hence, core inflation as defined in equation (1) is the optimal intermediate target for monetary policy. Among other things, inflation forecast targeting simplifies both the implementation and monitoring of an inflation target rule. In terms of implementation, Svensson shows that the central bank should adopt a rule which adjusts the policy instrument so that inflation forecasts adjust toward the inflation target. In terms of monitoring, the availability of a transparent, easily replicated forecasting model would help the central bank communicate with the public, and help the public determine whether policy is on the right track. In this way, a simple, reliable forecasting model could enhance the public's confidence in an inflation targeting regime.

2.2 Motivating Existing Measures of Core Inflation

The best-known and most widely used measures of core inflation are the "ex food and energy" series produced by the U.S. Bureau of Labor Statistics and Bryan and Cecchetti's median core and trimmed mean series. At a basic level, the BLS and Bryan-Cecchetti measures share a common motivation. Both start by trying to identify sources of transient variation in individual price changes, and then attempt to remove transients from aggregate inflation by excluding sensitive items from the market basket.

The BLS constructs its core measure by removing food and energy from the market basket and averaging over the remaining commodities. The rationale for excluding these items is that food and energy prices are subject to more high frequency price variation than other elements of the basket, and that much of the variation is driven by non-monetary factors. For example, food prices are more sensitive to weather and seasonal factors than many other CPI components. The price of fresh fruit rises during the winter and falls during the summer, and a late frost or early rainy season may damage crops and temporarily drive up prices. Similarly, energy prices are also sensitive to weather and seasonal variation, as well as to factors that influence the ability of the OPEC cartel to enforce production quotas. By excluding food and energy from the market basket, the BLS can ensure that relative price shocks arising in these sectors do not affect measured inflation.

Bryan and Cecchetti take this idea one step further. They note that relative price shocks are not confined to the food and energy sectors but can occur in any sector of the economy, and that while the BLS measure is robust to transients arising in these particular sectors, it is contaminated by relative price shocks arising elsewhere. Furthermore, they point out that this problem cannot be easily remedied by expanding the list of excluded sectors, because at any given time it is difficult to predict the sectors in which relative price shocks are likely to occur.[2] Hence the strategy of excluding preselected sectors from the core index is bound to miss relative price shocks arising elsewhere in the economy. Instead, they recommend bounded influence estimators, such as the median or trimmed mean, to measure the central tendency of the cross section of price changes. These measures do not require an *a priori* selection of volatile sectors and exclude large price changes from the cross section regardless of the sector in which they arise.

Bryan and Cecchetti further motivate their recommendation by appealing to a model of Ball and Mankiw (1992). In that model, firms set nominal prices before observing the current state of the economy. After this information arrives, a firm has the option to pay a menu cost and reset its price immediately or to wait until the beginning of next period and reset its price for free. Firms experiencing large relative price shocks pay the menu cost and those experiencing small shocks choose to wait. In this model, the effect of relative price shocks on measured inflation depends on whether the cross section of price changes is skewed. If the cross section is symmetric, large relative price shocks in one direction are balanced by those in the opposite direction, and measured inflation is unaffected. If the cross section is asymmetric, the cross-sectional mean is pulled in the direction of the long tail, and relative price shocks do affect conventional measures of inflation. Bryan and Cecchetti point out that the sensitivity to relative price shocks can be mitigated by computing the median or trimmed mean, which are robust to outliers even when the cross section is asymmetric. Thus, the Bryan-Cecchetti measures filter out large relative price shocks regardless of the sector in which they originate.

Figure 1 plots quarterly data on CPI inflation, along with the BLS measure and Bryan and Cecchetti's median 15 percent trimmed-mean series. The sample period is 1967.Q1 to 1998.Q2.[3] While the core measures clearly succeed in removing some transients from CPI inflation (notably the 1986 fall in oil prices), they still contain a great deal of high frequency variation.

[2] Indeed, the motivation for excluding food and energy was based largely on the experience of the 1970s, but these sectors were selected after the fact. At the end of the 1960s, it would have been difficult to predict which sectors would experience the most severe relative price shocks in the next decade.

[3] The initial date is dictated by the availability of the Bryan-Cecchetti data.

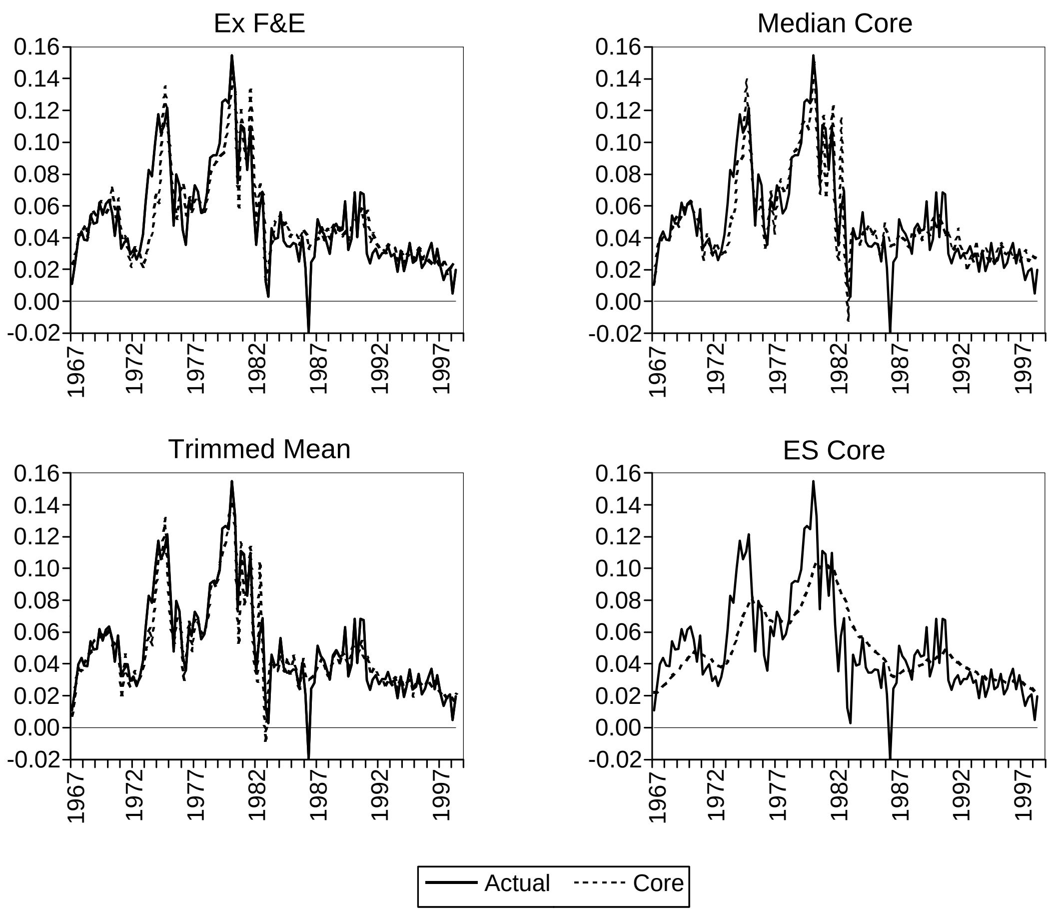

Figure 1: CPI Inflation & Core Inflation

This observation suggests that there may be other sources of transients in the data, and it motivates my interest in an alternative measure of core inflation.

2.3 Motivating a New Measure of Core Inflation

Instead of trying to identify sources of transient variation, I approach the problem from the opposite direction and consider the source of *persistent* movements in inflation. I start with the assumption that in the long run inflation is a monetary (or monetary-fiscal) phenomenon. Therefore, to identify the source of persistent variation, one needs to think about the choices central banks make and, in particular, about why they sometimes adopt new decision rules.

To make the discussion concrete, I find it useful to focus on a particular model, and for this purpose it is convenient to work with versions of the model that appears in chapters 7 and 8 of Sargent (1998). In this model, there are two players, the government and the private sector. Private sector behavior is summarized by a Phillips curve,

$$U_t - U^* = \theta(\pi_t - E_{t-1}\pi_t) + v_t, \tag{2}$$

where U_t represents the unemployment rate, U^* is the natural rate of unemployment, E_{t-1} is the conditional expectations operator, and v_t is an aggregate supply shock. Private agents forecast inflation and supply labor in accordance with this equation.

The government solves a version of Phelps' (1967) optimal policy problem. It fits an empirical Phillips curve model, and interprets it as an exploitable trade-off. Subject to the constraint on the evolution of (π_t, U_t) implied by that model, the central bank chooses a decision rule for inflation to minimize

$$V = E\sum_{t=0}^{\infty} \delta^t L(\pi_t, U_t), \tag{3}$$

where $L(\pi_t, U_t)$ is a quadratic loss function.

There are two versions of the model, which are distinguished by assumptions about how much agents know. In the first version, Sargent assumes that agents know the population moments that prevail in equilibrium and that they form parameter estimates from these moments. In the second version, he withdraws this knowledge and instead assumes that they must estimate parameters using adaptive or recursive methods.

For the first version of the model, in which agents know population moments, Sargent defines a generalized version of a rational expectations equilibrium that he calls a "self-confirming" equilibrium. The solution of the Phelps problem implies that the government's decision rule is conditioned on the parameters of its empirical approximating model. Private agents understand equations (2) and (3) as well as the government's approximating model and deduce the policy rule that is optimal given the government's beliefs. They forecast inflation rationally, in accordance with their understanding of the environment. The equilibrium is self-confirming if perceived and actual laws of motion coincide, so that the data evolving from a given government policy rule ratify the empirical approximating model from which that rule was derived. In that case, the government and private sector are both optimizing given their beliefs and constraints, the government has no incentive to modify its approximating model or policy rule, and the private sector has no incentive to alter its forecasting function.

Within a self-confirming equilibrium, inflation is a stationary random process around a fixed unconditional mean. Stationarity follows from three assumptions: first, that the natural rate of unemployment and other determinants of inflationary bias are constant;[4] second, that the government minimizes a quadratic loss function involving inflation; and third, that actual and perceived laws of motion coincide. The first assumption implies that mean inflation is constant (see Barro and Gordon 1983). The second implies certainty equivalence, which in conjunction with the government's optimization implies that policy will be set so that long run forecasts of inflation converge to the unconditional mean (see Svensson 1997). Among other things, this means that inflation cannot have a unit root, for in that case forecasts would not converge to the mean. The third assumption implies that the parameters of the government's approximating model and policy rule are time invariant. Indeed, time invariance is the hallmark of a self-confirming equilibrium. Agents believe they live in a stationary environment, and incoming data ratify this belief. There is no learning and no reason to alter decision rules.

Within such an equilibrium, measurement of core inflation is trivial. Because long run forecasts converge to the mean, core inflation (in the Beveridge-Nelson sense) is constant. The mean rate of inflation depends

[4] This could be weakened to allow for stationary variation in the determinants, but it rules out the kind of permanent variation assumed by Ireland (1998).

on details about the government's approximating model, its ability to pre-commit, and/or its reputation for preferring low inflation. But in Sargent's model, these features of the environment are constant within a self-confirming equilibrium.

To identify the source of persistent variation in inflation, we need to understand why central banks sometimes alter their decision rules. This requires changing some aspect of the model outlined above. The most natural modification is to relax the assumption that the government and private sector know the equilibrium population moments, and to replace it with an assumption that they use adaptive methods to learn about the world in which they live. In the second version of Sargent's model, the government re-estimates its approximating model each period as new data arrive. Then it solves for the new policy rule that is optimal for the updated model, and implements the first period action recommended by the new rule.

In particular, suppose the government's approximating model is

$$y_t = \gamma_t' x_t + \varepsilon_{1t}, \tag{4}$$

where the vectors x_t and y_t represent incoming data, and the vector γ_t lists the model's unknown parameters. In addition, suppose the government believes that γ_t evolves according to

$$\gamma_t = \gamma_{t-1} + \varepsilon_{2t}, \tag{5}$$

where ε_{1t} and ε_{2t} are Gaussian random vectors. As each new observation (x_t, y_t) arrives, the central bank uses Bayes' Law to update its beliefs about γ_t.

If the government believed that the parameters of its approximating model were time invariant (i.e., $var(\varepsilon_{2t}) = 0$), it would be optimal to update estimates by recursive least squares. Among other things, this means that new observations are given weight $(1/t)$ in the updating rule, and that the weight on new observations decreases as data accumulate. In the literature on stochastic approximation algorithms, this is known as a "decreasing gain" rule. The gain parameter refers to the weight on new data, which in this case is $(1/t)$.

On the other hand, if the government believed that the parameters of its approximating model were time varying ($var(\varepsilon_{2t}) > 0$), it would update parameter estimates by recursive *discounted* least squares. This updating rule assigns new data a constant weight and discounts observations in the distant past. Heuristically, when parameters drift through time, discounting is optimal because recent data are more informative about current parameters than data from the distant past. Learning rules that discount old data belong to a class known as "constant gain" algorithms.

Sargent demonstrates that when the government uses a decreasing gain algorithm to update its approximating model, the system converges to a self-confirming equilibrium. This version of the model provides little guidance about the source of persistent variation in inflation, except perhaps as a transitional phenomenon. But when the government employs a constant gain algorithm, the system does not converge to a self-confirming equilibrium, and instead gives rise to learning dynamics that look very much like monetary policy regime changes.

What are the nature of those changes? For our purposes, the most salient feature of Sargent's simulations is the occurrence of endogenous shifts in mean inflation as the economy shifts from the vicinity of one equilibrium to another. Loosely speaking, the sample paths resemble a state dependent process in which mean inflation shifts randomly from the basin of attraction surrounding the self-confirming equilibrium to the vicinity of zero inflation (the Ramsey outcome), and then back again.[5]

Regime shifts arising from changing beliefs strike me as a plausible way to model inflation persistence. In a linear context, the main alternative involves representations which assume that central banks allow inflation to drift. These seem less plausible, for conditional on beliefs and assuming that the central bank minimizes a loss function involving the variance of inflation, policy makers would not choose a rule that allowed inflation to drift, as this would result in an infinite variance. Therefore, conditional on remaining within a policy regime and assuming that policy is set optimally, long run

[5] In the model, shifts in mean inflation correspond to changes in the government's beliefs about the nature of the Phillips curve and the associated changes in its policy rule. The government pursues a high inflation policy when it suspects that the Phillips curve may be exploitable. But as data accumulate and it begins to suspect that the natural rate hypothesis is more plausible, it revises its approximating model and shifts toward a decision rule that attempts to maintain low inflation. Along the sample path, this accounts for the appearance of sudden disinflations. In the low inflation equilibrium, however, there is some backsliding because the data trace out an inverse relation between inflation and unemployment, essentially for the reasons given by Lucas (1972). Eventually this tempts the government to try once again to exploit the (apparent) Phillips tradeoff, and inflation drifts upward. This occurs very gradually and is repeatedly punctuated by subsequent disinflations.

forecasts should converge to a regime-specific constant.

On the other hand, if the central bank were to make a significant discovery about how the economy operated, it would alter its decision rule, and the change in policy rule would alter the regime-specific mean. If policy regimes last for years at a time (i.e. if significant discoveries occur only once in a while), learning dynamics generate persistent variation in inflation.

Representations with latent, random mean shifts and period-by-period drift both imply reduced form $I(1)$ representations for inflation, but a unit root arising from variation across regimes somehow seems more plausible than a unit root arising from variation within a regime.

If mean shifts arising from changes in policy rules are indeed the main source of inflation persistence, then perhaps we should consider measures of core inflation that are designed to adapt to such changes. Taking a cue from Sargent's agents, we might consider a constant gain algorithm for updating estimates of the mean,

$$\mu_t = \mu_{t-1} + g_0(\pi_t - \mu_{t-1}), \tag{6}$$

where μ_t is the period t estimate of mean inflation and g_0 is the gain parameter, which is assumed to lie between 0 and 1. This updating rule corresponds to simple exponential smoothing,

$$\mu_t = \left(\frac{g_0}{1-(1-g_0)L}\right)\pi_t, \tag{7}$$

which is of course a one-sided geometric distributed lag of past inflation,

$$\mu_t = g_0 \sum\nolimits_{j=0}^{\infty}(1-g_0)^j \pi_{t-j}. \tag{8}$$

Despite its simplicity, this measure has a number of appealing properties. First, the exponential smoother well approximates an ideal low-pass filter for suitable choices of g_0. For example, the top panel of figure 2 shows the mod-square of its transfer function for values of g_0 ranging from 0.075 to 0.15. As required, the filter removes high-frequency components of inflation and passes those at low frequencies. Similarly, the bottom panel illustrates the properties of the associated high-pass filter, which transforms inflation into the core deviation,

$$\pi_t - \mu_t = \left(\frac{(1-g_0)(1-L)}{1-(1-g_0)L}\right)\pi_t. \tag{9}$$

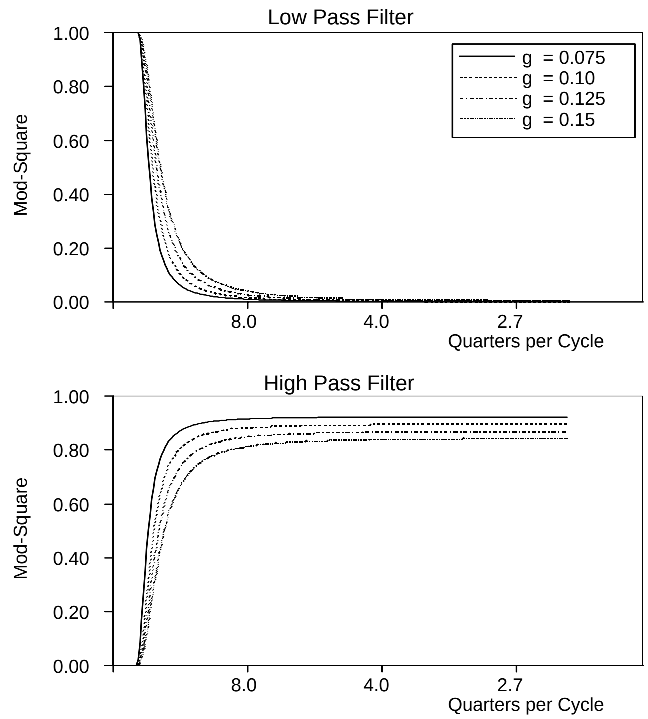

Figure 2: Properties of the Exponential Smoother

The core deviation contains the high frequency or short-lived components of inflation.

Second, unlike many other approximations to low pass filters, such as those advocated by Baxter and King (1995) or Hodrick and Prescott (1997), this filter is one-sided into the past and can be implemented in real time. As Bryan and Cecchetti note, two-sided low pass filters are of relatively little use to monetary policy makers, because they reduce the timeliness and relevance of incoming inflation data. In contrast, because this filter is one-sided into the past, its output would be available to policy makers as soon as new inflation data became available.

Third, this measure is absolutely trivial to compute and depends on only one free parameter, g_0, which can be calibrated *a priori*. Among other things, this means that the filter coefficients do not need to be re-estimated when new data arrive, and that the history of the core series need not be revised in response to changes in parameter estimates. To calibrate g_0, it is useful to work through the following thought experiment. Suppose the mean followed a step function, equaling 0 up to some unknown date and then shifting to 1. The parameter g_0 governs the rate at which the updating rule adapts to this shift. The half-life of the learning process is approximately $\ln(2)/g_0$, so to complete half the adjustment within 4 quarters one would set $g_0 = 0.17$, to complete half the adjustment within 8 quarters one would set $g_0 = 0.087$, and so on. In what follows, I set $g_0 = 0.125$, which implies a half life of 5.5 quarters, but parameter settings between 0.075 and 0.2 yield similar results.[6,7]

The bottom right panel of figure 1 illustrates the results of applying the exponential smoother to CPI inflation. The filtered measure, μ_t, is smoother than CPI inflation itself, and also smoother than the other measures of core inflation. Indeed, the exponential smoother removes nearly all the high frequency variation in the data. At first glance, this appears to be an attractive method for tracking persistent movements in inflation.

3 Evaluating the Various Measures

According to the Bryan-Cecchetti definition, a successful measure of core inflation should purge the transients from actual inflation. In this section, I consider which of the candidate measures does this most effectively. The first subsection investigates univariate properties of the proposed measures, the second examines which has the most weight in a combined measure, and the third considers their information content when combined with other macroeconomic predictors of inflation.

3.1 Univariate Properties

Let $\tilde{\pi}_{ct}$ represent a candidate measure of core inflation. For sufficiently large H, the core deviation, $\pi_t - \tilde{\pi}_{ct}$, should be inversely related to subsequent

[6] If one held a random walk prior for inflation, which implies a mean shift every quarter, one would set $g_0 = 1$, so that $\mu_t = \pi_t$.

[7] Of course, g_0 could also be estimated, for example by choosing the value that best forecasts future changes in inflation. In the empirical applications, it appears that the results are not too sensitive to the choice of g_0, which suggests that g_0 is not sharply identified. For what it's worth, it also appears that the calibrated value is close to the best fitting value.

changes in inflation, $\pi_{t+H} - \pi_t$. Moreover, in order for the candidate to satisfy equation (1), the coefficients in the regression

$$(\pi_{t+H} - \pi_t) = \alpha_H + \beta_H(\pi_t - \tilde{\pi}_{ct}) + u_{t+H} \tag{10}$$

should satisfy $\alpha_H = 0$ and $\beta_H = -1$. The restriction on α_H is trivial, and follows from the fact that $(\pi_{t+H} - \pi_t)$ and $(\pi_t - \tilde{\pi}_{ct})$ are (approximately) mean zero. The restriction on β_H indicates whether the core deviation correctly measures the magnitude of transient components. If β_H were negative but less than one in absolute value, the measured core deviation would overstate the magnitude of subsequent changes in inflation, and thus would also overstate the magnitude of current transients. Similarly, if β_H were negative but greater than one in absolute value, the measured core deviation would understate the magnitude of current transients.

Estimates of β_H are reported in figure 3. Solid lines represent point estimates, and dashed lines illustrate the boundaries of centered 90 percent confidence bands. The latter were computed from heteroskedasticity and autocorrelation consistent (HAC) estimates of the standard error for β_H. The sample period is 1967.Q1 to 1998.Q2 with an appropriate allowance at the end to adjust for the forecast horizon, H.

In each case, estimates of β_H are negative, as required. Hence, measured core deviations are negatively correlated with subsequent changes in inflation. For short forecast horizons, say from 1 to 6 quarters, most of the estimates are significantly less than 1 in absolute value. But the estimates decline as the forecast horizon lengthens, and for horizons greater than 8 to 10 quarters are not significantly different from -1. Thus, over horizons of 3 to 4 years, the measured transients have (more or less) the right magnitude. The exponentially smoothed series and the BLS measure are particularly accurate on this dimension, with point estimates close to -1 at horizons of 3 to 4 years.

Figure 4 plots the R^2 statistics from these regressions. In addition to knowing whether measured core deviations are negatively correlated with subsequent changes in inflation and whether they are approximately of the right magnitude, we also want to know which of the candidates has the highest R^2. Candidates that account for a greater percentage of subsequent changes in inflation filter out more transient variation and are preferred to those that account for less.

Over short horizons of 1 to 4 quarters, the Bryan-Cecchetti measures are most highly correlated with changes in inflation. For example, at the 1

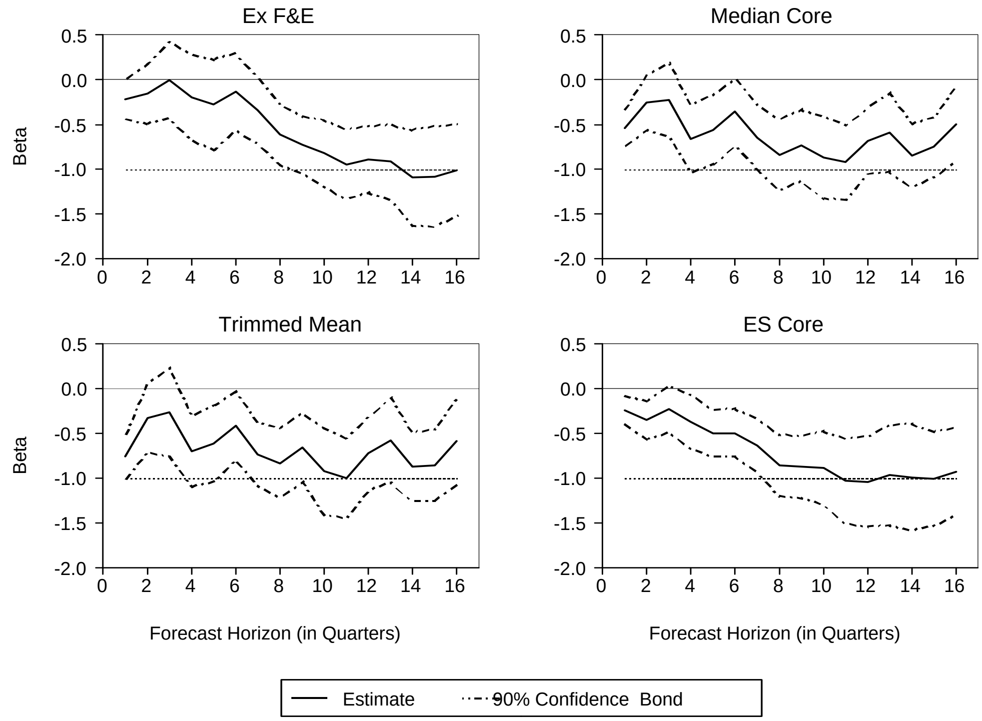

Figure 3: Coefficients in Univariate Regressions

quarter horizon, the median and trimmed mean account for roughly 20 to 25 percent of the variance of the change in inflation. At longer horizons, the exponentially smoothed measure appears to be more informative, with R^2 statistics between 25 and 30 percent. Thus, for the medium run of 2 to 4 years, the exponentially smoothed series is the single most informative measure of expected inflation.[8]

[8] In the case of the ES core measure, equation (10) represents a resitricted ARIMA model for forecasting equation (substitute equation 7 into 10 and re-arrange terms). This begs the question, "why not use the unrestricted version instead?" The answer is that the restrictions are essential for producing useful multi-year forecasts. To estimate the unrestricted analogs of (10) one would have to difference π_t, and the estimates would be dominated by the high frequency variation we seek to filter out. One can interpret the restrictions as a way to impose parsimony, or as a way to estimate the free parameters β_H from the low frequency information in the data. For formal arguments on this point, see Sims (1972) or Appendix C of Cochrane (1988).

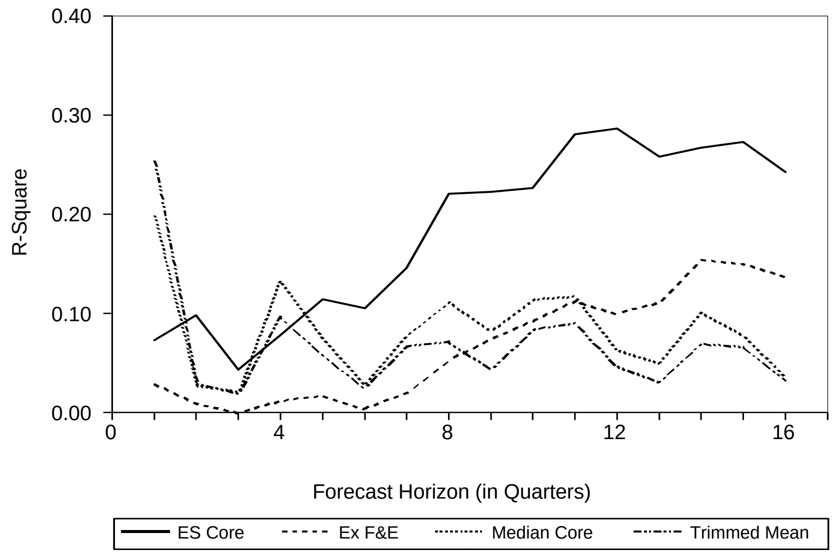

Figure 4: R^2 Statistics from Univariate Regressions

3.2 Optimal Combinations

There is of course no reason to confine one's attention to univariate measures of core inflation. One measure may contain information absent from the others, and even if it were not the most informative in a univariate context it would not be optimal to discard that information. If the various measures each contain information about different kinds of transients, it would be optimal to incorporate them all into a combined forecast, consisting of a weighted average of the individual measures. A simple way to choose the weights is by estimating the coefficients in the regression,

$$(\pi_{t+H} - \pi_t) = \alpha_H + \sum\nolimits_{j=1}^{n} \beta_{jH}(\pi_t - \tilde{\pi}_{jt}) + u_{t+H}, \tag{11}$$

where $\tilde{\pi}_{jt}$ is the jth candidate measure and β_{jH} is its weight in predicting the H-period change in inflation.

Figure 5 summarizes the marginal predictive power of the various measures by comparing R^2 statistics from univariate and multivariate regressions. For example, in the top left panel, the solid line graphs R^2 statistics from univariate regressions involving the exponentially smoothed series, the dotted line graphs R^2 statistics from univariate regressions involving the BLS

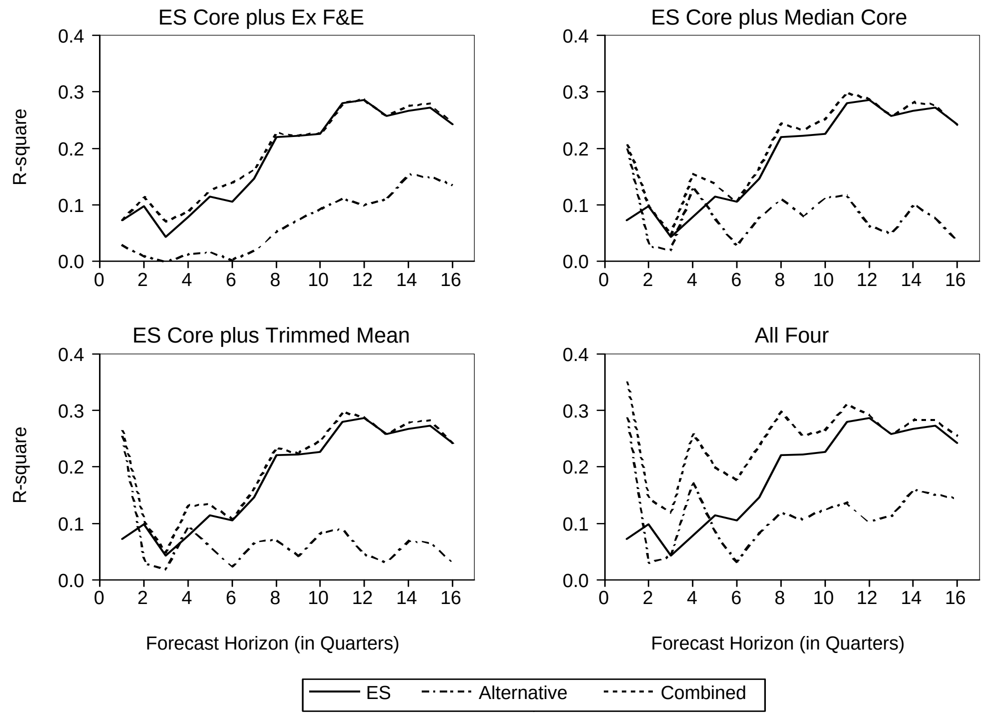

Figure 5: R^2 Statistics from Univariate & Multivariate Regressions

measure, and the dashed line graphs R^2 statistics from bivariate regressions involving both measures. As noted above, on a univariate basis, the exponentially smoothed series is a superior predictor, as it accounts for a greater fraction of the variance of subsequent changes in inflation. Hence, there is information in the exponentially smoothed series that is not in the BLS measure. Conversely, this panel shows that there is very little information in the BLS measure that is not in the exponentially smoothed series. The R^2 statistics from bivariate regressions are only slightly greater than those in univariate ES regressions, and the little additional information in the BLS measure is mostly relevant over short forecasting horizons. Over the medium run of 2 to 4 years, the exponentially smoothed measure contains all the information in the BLS measure, plus some information that is orthogonal to the BLS measure.

Similarly, the off-diagonal panels in figure 5 compare the exponentially smoothed measure with the median and trimmed mean series. At the 1 quarter horizon, the Bryan-Cecchetti measures subsume the information in the exponentially smoothed series. That is, in univariate regressions, they are more informative for 1-quarter changes in inflation, and combining them with the exponentially smoothed series yields only a negligible improvement in 1-quarter inflation forecasts. Over longer horizons, however, the outcome is reversed. At horizons of 2 to 4 years, the exponentially smoothed series is more informative in univariate regressions, and R^2 statistics rise only slightly when the exponentially smoothed series is combined with the median or trimmed mean. Again, over the medium run of a few years, the exponentially smoothed measure contains nearly all the information in the Bryan-Cecchetti measures, plus some information that is orthogonal to them.

The bottom right panel of the figure compares R^2 statistics from univariate ES regressions with those obtained from regressions including the other three candidates, as well as those obtained from regressions which combine all four measures. Relative to the exponentially smoothed series, the other three candidates are jointly superior at forecast horizons up to one year, but the univariate ES measure is a superior predictor at longer horizons. Similarly, when all 4 measures are combined, the other three have some marginal predictive power at forecast horizons up to 10 quarters, but at longer horizons univariate ES regressions fit almost as well as the four-variable model. Hence, for long horizon forecasts, the exponentially smoothed series contains most of the information in the other three, as well as some additional information not contained in them. Consequently, over the medium run of 3 to 4 years, there appears to be little scope for improving forecasts by combining the measures.

Figure 6 illustrates this result in a different way, plotting realizations of 4-year ahead inflation forecasts from the univariate ES regression and optimal combination regressions. The solid line in each panel is the fitted value of the univariate ES regression, and dashed lines illustrate the fitted values of bivariate or multivariate regressions. In each case, univariate and combination forecasts are essentially the same. Again, at the 4 year horizon, the other candidates have little information not contained in the exponentially smoothed series.

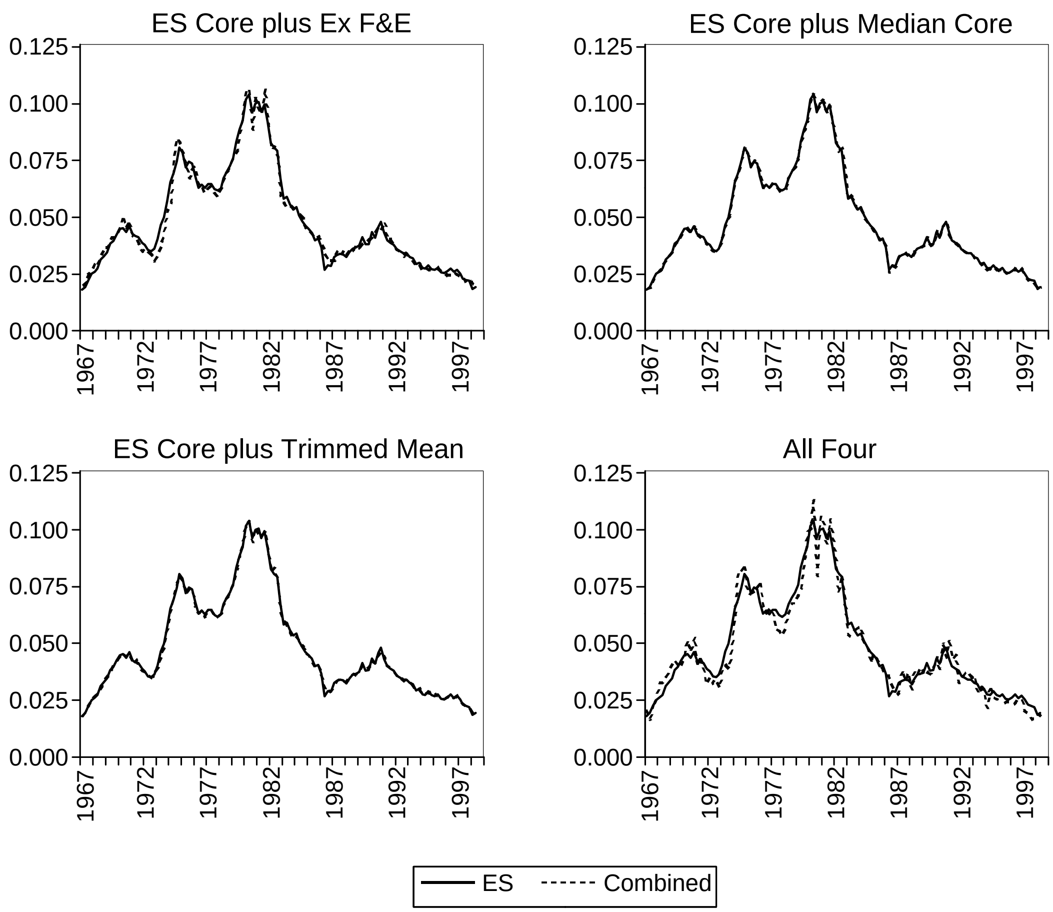

Figure 6: Expected Inflation Four Years Ahead

3.3 Marginal Predictive Power in Combination with Macroeconomic Variables

More generally, there is no reason to confine one's attention to measures of core inflation when forecasting movements in actual inflation. Other macroeconomic variables may contain information that is absent from the core measures, and it would be useful to incorporate this information as well. The simplest way to do so is by adding macroeconomic predictors to the forecast regressions,

$$(\pi_{t+H} - \pi_t) = \alpha_H + \beta_H(\pi_t - \tilde{\pi}_{ct}) + \gamma'_H x_t + u_{t+H}, \tag{12}$$

where x_t is a vector that lists additional variables and γ_H is a vector of free parameters.

In this subsection, I consider two questions. First, when combined with other macroeconomic predictors of inflation, which of the core measures is the most informative? Second, how much can inflation forecasts be improved by combining information in this manner?

To address these questions, I rounded up the usual suspects and experimented with a number of macroeconomic information sets, x_t. The list of suspects included, among others, a variety of measures of output and unemployment "gaps," short term real interest rates, long term yield spreads, a variety of measures of oil prices, and M2 growth. For the sake of brevity, I report results only for a parsimonious information set x_t consisting of ex post 3-month real interest rates and a measure of detrended output. Generally speaking, I found that it was hard to improve on this information set by including other macroeconomic predictors, especially in combination with the exponentially smoothed measure of core inflation.[9]

Real interest rates were measured by subtracting CPI inflation from the nominal yield on 3-month Treasury bills. Real GDP was detrended by applying the exponential smoother with $g_0 = 0.125$. The resulting measure of the output gap can be interpreted either as the deviation from an adaptive estimate of a deterministic trend or as the output of a high-pass filter applied to GDP. The measure is analogous to the one used by King and Watson (1994) in their study of the Phillips Curve, except that their filter was two-sided. Again, the restriction to one-sided filters is essential for forecasting.

[9] Glenn Rudebusch was not surprised by this, and reminded me that the output gap and short term real interest rate are the key variables in the Rudebusch-Svensson (1998) model.

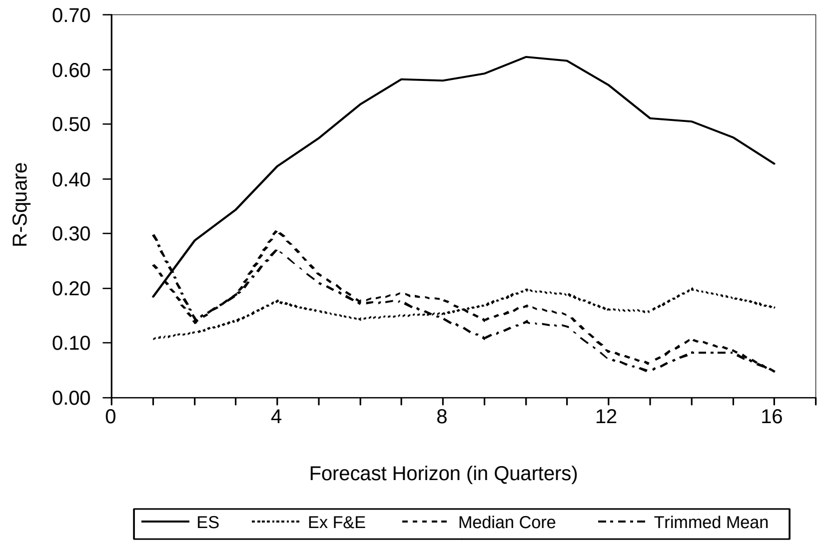

Figure 7: Relative Predictive Power when Combined with Macroeconomic Predictors

Figure 7 reports results on the predictive power of the various core measures in combination with the two macro variables. The figure plots R^2 statistics from versions of equation (12) over forecast horizons ranging from 1 quarter to 4 years. The results are similar to those reported above. The median and trimmed mean measures are most informative at the 1-quarter horizon, but the exponentially smoothed series is a superior predictor over longer horizons. Indeed, when combined with detrended output and the real bill rate, the exponentially smoothed inflation series accounts for approximately 40 to 60 percent of the variation in inflation over the 2 to 4 year horizon. In contrast, the other candidates account for only 10 to 20 percent of the variation at these horizons. Thus, over the medium run of several years, it again appears that the exponentially smoothed series is the single most informative measure of core inflation.

A comparison of the joint goodness-of-fit statistics in figure 7 and the univariate goodness-of-fit measures in figure 4 suggests that the macro variables contain much information about inflation that is not present in the ES core measure. But is there also some information in the latter that is not contained in the former? Figure 8 reports the results of some regressions

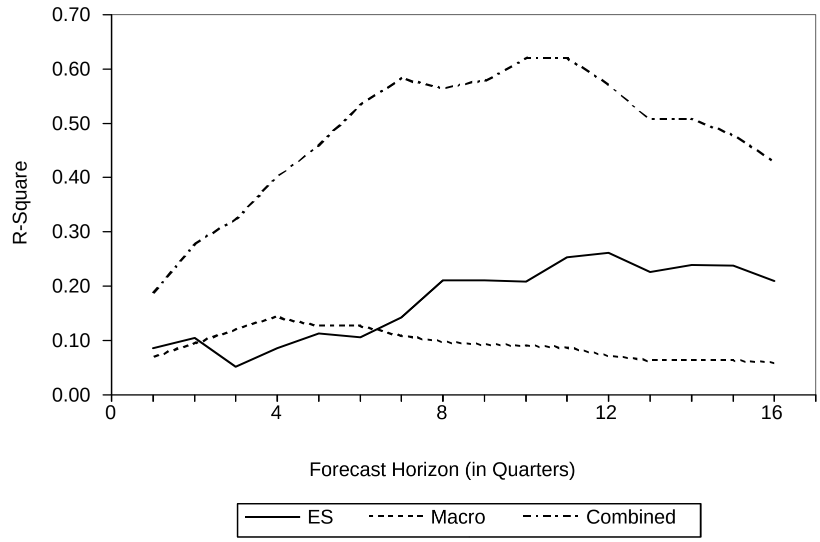

Figure 8: Marginal Predictive Power of ES Core Relative to the Macro Varibles

that decompose the information in the trivariate information set. In this figure, the solid line plots R^2 statistics for univariate regressions involving the core deviation, the dotted line plots R^2 statistics for bivariate regressions involving the two macro variables, and the dashed line plots R^2 statistics for trivariate regressions involving all three. Here the sample has been expanded to cover the period 1954.Q1 to 1998.Q2, with an appropriate allowance at the end for the forecast horizon H.[10]

By themselves, the macro variables account for only about 10 percent of the variation in $(\pi_{t+H} - \pi_t)$ over horizons of 2 to 4 years. Similarly, by itself the ES core measure accounts for about 20 to 25 percent of the variation over these horizons. The combination of all three is far superior to either, and accounts for roughly 40 to 60 percent of the medium-run variation. Hence, the ES core measure indeed has considerable incremental predictive power relative to the two macro variables.

Finally, the trivariate model compares quite favorably with other inflation forecasting models in the literature. For example, in contrast to the results

[10]From this point on, the Bryan-Cecchetti measures are no longer in play, and the constraint on their availability is no longer binding.

reported here, Cecchetti (1995) and Stock and Watson (1998) both report that inflation is difficult to forecast, especially over longer horizons. It seems that the key to improving inflation forecasts involves allowing for a shifting mean, in the manner described above. Macroeconomic predictors have substantial information about locally mean-reverting components of inflation, but it is harder to find globally mean-reverting components.

4 Conclusion

This paper proposes a new measure of core inflation and compares it with several existing measures. The new measure is adaptive and is designed to track sudden and persistent changes in inflation, such as those arising from changes in monetary policy decision rules. Formally, the measure can be interpreted either as a constant gain update of mean inflation or as the output of an approximate low-pass filter. Unlike many other approximations to low-pass filters, the filter used here is one-sided into the past, and the date t filtered measure can be computed as soon as date t inflation data become available.

I evaluate the new measure and compare it with existing measures by considering how well it filters transient movements out of the data. A measure that strips transients from actual inflation should be useful for predicting subsequent changes in inflation. Conversely, the measure that best predicts subsequent reversals in inflation is also the one that best filters current transients. Hence, various measures of core inflation can be evaluated by comparing their relative predictive power for future changes in inflation.

Over the medium run of 2 to 4 years, I find that the new measure is a better predictor than conventional measures such as the BLS ex food and energy series or Bryan and Cecchetti's median and trimmed mean series. The new measure subsumes the relevant information contained in existing measures and has incremental predictive power relative to them. The new measure is also a superior predictor in combination with other macroeconomic predictors of inflation, and it contains substantial incremental predictive power relative to the macro variables. For example, a simple forecasting model involving a short-term real interest rate and a measure of the output gap accounts for only about 10 to 20 percent of the variance of changes in inflation over horizons of 2 to 4 years. But when the new core measure is added to the information set, the augmented model accounts for approximately 40 to

60 percent of the variation.

Ideally, one would want to validate the forecasting results reported here by doing out-of-sample experiments. Unfortunately, with forecast horizons of 2 to 4 years and samples covering at most 44 years, there is not enough independent information in the data to do this. For example, suppose the sample were split in two, with the first 22 years used to initialize parameters and the second 22 years used to evaluate forecasts. For the 2-year ahead forecasting model, there are only 11 non-overlapping observations for initializing parameters and only 11 non-overlapping observations for evaluating out-od-sample forecasts. This leaves too few observations to do either job well. Thus, the ultimate proof of the new measure will have to come in practice.

5 References

Ball, Lawrence and N. Gregory Mankiw, 1992, "Relative Price Changes as Aggregate Supply Shocks," NBER Working Paper 4168.

Barro, Robert J. and David B. Gordon, 1983, "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy* 91, pp. 589-610.

Baxter, Marianne and Robert G. King, 1995, "Measuring Business Cycles: Approximate Bandpass Filters for Economic Time Series," NBER Working Paper 5022.

Beveridge, Stephen and Charles R. Nelson, 1981, "A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the 'Business Cycle,"' *Journal of Monetary Economics* 7, pp. 151-174.

Bryan, Michael F. and Stephen G. Cecchetti, 1993, "The Consumer Price Index as a Measure of Inflation," *Economic Review of the Federal Reserve Bank of Cleveland* 29, pp. 15-24.

—— and ——, 1994, "Measuring Core Inflation," in *Monetary Policy,* N. Gregory Mankiw (ed.), Chicago: University of Chicago Press for the National Bureau of Economic Research.

—— and ——, 1995, "The Seasonality of Inflation," *Economic Review of the Federal Reserve Bank of Cleveland* 31, pp. 12-23.

——, ——, and Rodney L. Wiggins, 1997, "Efficient Inflation Estimation," NBER Working Paper No. 6183.

Cecchetti, Stephen G., 1995, "Inflation Indicators and Inflation Policy," *NBER Macroeconomics Annual* 10, pp. 189-219.

——1996, "Measuring Short-Run Inflation for Central Bankers," NBER Working Paper No. 5786.

Cochrane, John H., 1988, "How Big is the Random Walk in GDP?" *Journal of Political Economy* 96, pp. 893-920.

Freeman, Donald G., 1998, "Do Core Inflation Measures Help Forecast Inflation," *Economics Letters* 58, pp. 143-147

Hodrick, Robert and Edward C. Prescott, 1997, "Post-War U.S. Business Cycles: An Empirical Investigation," *Journal of Money, Credit, and Banking* 29, pp. 1-16.

Ireland, Peter N., 1998, "Does the Time-Consistency Problem Explain the Behavior of Inflation in the United States?," Unpublished Manuscript, Department of Economics, Boston College.

King, Robert G. and Mark Watson, 1994, "The Post-War U.S. Phillips Curve: A Revisionist Econometric History," *Carnegie-Rochester Conference Series on Public Policy* 41, pp. 157-219.

Lucas, Robert E., Jr., 1972, "Expectations and the Neutrality of Money," *Journal of Economic Theory* 4, pp. 103-124.

Phelps, Edmund S., 1967, "Phillips Curves, Expectations of Inflation and Optimal Unemployment Over Time," *Economica* 2, pp: 22-44.

Quah, Danny and Shaun P. Vahey, 1995, "Measuring Core Inflation," *Economic Journal* 105, pp. 1130-1144.

Rudebusch, Glenn D. and Lars E.O. Svensson, 1998, "Policy Rules for Inflation Targeting," NBER Working Paper No. 6512.

Svensson, Lars E.O., 1997, "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets," *European Economic Review* 41, pp. 1111-1146.

Sargent, Thomas J., 1998, *The Conquest of American Inflation,* unpublished manuscript.

Sims, 1972, "Approximate Prior Restrictions on Distributed Lag Estimation," *Journal of the American Statistical Association* 67, pp. 169-175.

Stock, James H. and Mark W. Watson, 1998, "Forecasting Inflation," Unpublished Manuscript, Kennedy School of Government and Woodrow Wilson School of Public Policy.