An alternative measure of inflation

François R. Velde

Introduction and summary

Controlling inflation is a primary goal of monetary policy. In order to control inflation, central bankers need to be able to measure and forecast inflation as best they can. Forecasting is particularly important, given the fact that monetary policy operates with "long and variable lags," and therefore policymakers need to act well in advance of actual developments in inflation, on the basis of their forecasts.

Both the measurement and the forecasting of inflation have been subjects of ongoing debate and research in recent years. This article reports on my research on both aspects. Specifically, I develop a new measure of inflation, which can also be used to generate a forecast. The research is still very preliminary, but the first results are encouraging. In particular, it appears to provide some gains in forecasting compared with what remains the best and simplest forecasting model, namely, the *random walk* model of inflation.[1]

Motivation

For the U.S. Federal Reserve (and for many other central banks), price stability is a primary goal, mandated by law. This stability is usually interpreted to mean a low level of inflation (how low will not be debated here). So what is inflation and how do we measure it?

Inflation is generally defined as the rate of change of some price index: Well-known examples are the Consumer Price Index (CPI) and the Personal Consumption Expenditures (PCE) Price Index. Price indexes, generally speaking, result from the attempt to measure with a single number a change in a collection of prices p_i (for $i = 1, \ldots, n$) .

The simplest conceivable index is to take a straight average of the prices in each period, ignoring the quantities. But it seems more reasonable for many purposes to weight the prices. Movements in the price of an item that is of little importance relative to the others should not be given much weight. An item of little importance is one that does not represent a large share of expenditures, which naturally leads one to use expenditure shares in creating the weights (see box 1).

More generally, suppose we have observations on prices at which a given range of goods and services are bought and sold and also observations on the quantities bought and sold at those prices. We thus have a collection $\{p_i\}$ and a collection $\{q_i\}$. Suppose further that we have observations in two periods, 0 and 1. One period (either 0 or 1) is chosen as the reference period. The problem of constructing an index (for either prices or quantities) is that of devising a formula that takes the prices and quantities in both periods and yields a single number. The formula must be such that, if the prices and quantities are the same in the reference period and the other period, the number is 1. Note that, even if prices are unchanged between the two periods, a change in quantities will generally result in the index being different from 1. Even though prices are unchanged, the weighting of the prices, which is based in part on the quantities, will change the overall index.

From this brief overview, one can draw some general observations about price indexes. Most price indexes require information on quantities in order to weight the prices. For certain applications, the fact that quantities are measured less precisely, less easily, and less quickly than prices can be a problem. At a deeper level, there is an important connection between

François R. Velde is a senior economist at the Federal Reserve Bank of Chicago. *The author thanks Michael Kouparitsas for all his advice, David Kang for outstanding research assistance, and his colleagues at the Chicago Fed, particularly Craig Furfine and Marco Bassetto, as well as Lawrence Christiano and Lars Hansen, for their comments.*

BOX 1

Different kinds of indexes

A straightforward weighting scheme is to use the expenditure shares to weight the items. And, since the choice of the units to measure the quantities of goods, and therefore the prices per unit of goods, is arbitrary, the absolute level of an index is meaningless, and an index can only measure changes relative to a reference period. One thus arrives at the classic price indexes to measure changes between period 0 and period 1—the Laspeyres and Paasche indexes, depending on whether one chooses period 0 or period 1 as the reference period, and the Fisher index, which takes the geometric average of the two indexes and thus achieves a pleasing symmetry between the two periods.

The CPI is essentially a Paasche index, with weights based on a reference period that is changed from time to time. The current PCE index is a type of Fisher index, with no fixed reference period: The changes computed period by period are chained together to form an index series. The Fisher index has another nice feature—a quantity index that can be computed in exactly the same way (weighting quantities by expenditure shares); in any period the quantity index times the price index equals the total expenditure. Thus, the price index can be seen as a deflator of the nominal expenditures that yields an index of real quantities.

the prices and quantities. In much of index theory, the two collections are intimately related, and the index makes sense with respect to a particular set of quantities. For example, the CPI is based on weights representing the typical basket of goods and services consumed by an average urban consumer. The PCE deflator is based on the quantities of goods and services consumed in the economy as a whole. These indexes yield different inflation rates, partly because they are tied to different collections of goods and services, and the computation of the index depends on the collections themselves, as well as on the quantities.

More generally, price indexes are designed for a certain purpose and have optimal properties for that purpose, but they may not be well suited for others. Monetary policy needs measures of inflation, but it may well be that indexes designed to measure the value of a basket of consumer goods or to convert nominal consumption expenditures into real consumption expenditures are not perfectly suited for the goals of monetary policy.

In fact, policymakers have come to use variants of both the CPI and PCE index, the so-called core measures. The intuition behind these measures is that monetary policy is interested in broad and persistent movements in inflation and that certain price series, being too volatile, introduce noise and confusion in the measurement of these broad movements. Therefore, the troublesome series (typically food and energy-related items) are removed altogether from the price index.

The alternative approach I propose here extends the intuition behind the core measures of inflation. My premise is that inflation is a general movement in the price level or, put differently, a movement that is common to all individual price series. Once we posit the object to be measured (inflation) as a statistical series of its own, then the measurement problem can be seen in a different light, as a signal extraction problem. Constructing (weighted) averages is a way of measuring inflation that makes particular assumptions about the movements that are specific to each series: essentially, that they are a sort of observation noise that can be removed by taking averages and counting on the law of large numbers. But these movements specific to each price series can have a more complex structure than being just noise. As it happens, statistical tools are available to measure inflation and allow for more complex structures. The result of this approach is still, in a way, a weighting scheme, but it is a dynamic weighting scheme, and it is one that weights series not by their importance in a basket, but according to the information that they contain.

Method

The Kalman filter

The Kalman filter relies on a distinction made between what is observed and what is not. This is formalized by writing two equations, known as the *state* equation and the *observation* equation. The first equation posits the evolution over time of the hidden variables, gathered in a vector called the state vector. The specification is typically dynamic, meaning that current value taken by the state depends on past values. One of these hidden variables will be our general movement in the price level. The second equation describes the relation between the state and the observables.

I call the vector of observable variables y_t: That is, at every point in time t, $(y_{1t}, y_{2t}, y_{3t}, \ldots)$ represent the values of the series 1, 2, 3 at time t. There is another vector, made up of variables that are not observed: It is called the state vector, x_t. The state equation describes how this (unobserved) state changes over time. The general form is

$$x_t = A\, x_{t-1} + u_t,$$

where A is a matrix and u_t is a noise or error term.

The observation equation relates the observables and the state in the following way:

$$y_t = B\, x_t + v_t,$$

with v_t, another noise or error term, uncorrelated with u_t. Having specified this model of the relationships between the state and observables, I need to supply initial guesses about two things: the initial value taken by the state and the uncertainty surrounding that initial value. Typically, the initial value is assumed (in the absence of any other information) to be the long-run average of the state, and the uncertainty surrounding this value can be derived from the state equation.

I am now ready to apply the Kalman filter. It may seem a little magical to estimate the value of a variable (the state) that is never observed. The way it works is as follows: The method is recursive, meaning that at any point in time it takes the most recent guess and updates it in a systematic manner based on the newly available information. Given a guess as to the value of the state yesterday, and the uncertainty around it, the mechanical application of the state equation provides a best guess as to its value today (before I introduce any of today's information).

How do I represent today's information? The basic rule here is learning from one's mistakes. Since I have a guess of today's state, I can make a guess of today's observables, using the observation equation. Then I compare this guess with what actually happened: The difference between the two is the new information that is relevant to my model.

How do I incorporate today's information? I project the (unknown) value of today's state onto all of the information, which can be decomposed into the information available yesterday and the new information that became available today. A classic result of regression analysis tells us that this projection is the sum of two terms. The first is simply the best guess of today's state using the information up to yesterday. The second corrects this guess with the new information, but weights it according to two expectations: how correlated it is likely to be with today's state and how noisy it is. The more correlated this new information is with the state, the more weight I place on it; on the other hand, the noisier it is, the more I discount it. How these weights are determined depends on the particular values I have assigned to the coefficients of the state and observation equations.

This leads to a recursive formula: Today's guess is yesterday's guess updated with the (appropriately weighted) new information. Tomorrow, I will take today's guess of today's state, derive a guess of tomorrow's state, and repeat the procedure.

Of course, when tomorrow rolls around, I will correct my guess of tomorrow's state that was based on today's information. But I could also correct my guess of today's state based on today's information, or even yesterday's guess of yesterday's state. More generally, having proceeded recursively from the beginning to the end of the available sample, it is possible to go back and correct the guesses made for the value of the state in earlier periods based on the information of the whole sample. This procedure, which is also recursive but backwards (as it updates yesterday's guess based on today's error), is called the Kalman smoother.

This approach to measuring inflation, like any other, has costs and benefits. Some of the benefits are apparent if we think back to the initial motivation. Modeling inflation as a hidden variable allows me to bypass a number of the issues that arise for standard indexes. For example, the basic intuition behind core measures of inflation is fully extended. Price series are not ignored or deleted when they are volatile; rather, optimal use is made of the information that they contain. The approach helps me deal with the choice of optimal weights to apply to the price series because the Kalman filter algorithm itself chooses the weights that it applies recursively. But it doesn't choose them arbitrarily; rather, it tries to extract the information contained in the price series.[2] The choice of the series themselves is not eliminated, of course, but it is of less importance. There is no conceptual problem in mixing series of different origins (say, the PCE index and CPI) or choosing a subset of either collection of series. There is no "adding up" constraint; there is no need to fully represent a given basket or bundle of goods and services. The main consideration in adding another series to the collection we use should be: Is that additional series likely to provide information about inflation that was not contained in our collection already?

Finally, one major benefit of the approach is that it yields a forecasting tool at no additional cost, so to speak. My best guess of the value of the state at time $t+1$ based on information available at time t is simply my best guess of the state at t, projected forward one period using the state equation. The Kalman filter approach thus folds into one operation measurement and forecast.

There are disadvantages, however. One cost is of a technical nature, and another is more of a conceptual problem.

The technical difficulty becomes more apparent in the next section. Although index theory may rely on some assumptions about the economic process that generated the prices and quantities, the considerations that lead to the choice of an index are quite general and make few or no assumptions about the prices themselves. The Kalman filter approach requires that a modeling choice be made about the statistical processes that best represent the price series and the underlying inflation as well. That is, I have to take a stand on the structure of prices, their interdependence, the correlations of a series with its past values, as well as those between the series themselves, and so on. Fortunately, some statistical tests guide the choice of that structure, as I discuss in the next section.

The conceptual difficulty is the following: One might fairly argue that I am not so much measuring inflation as inventing a concept of inflation that I can measure. The underlying inflation, or "latent inflation," may be just a statistical artifact. My response is that, although it is indeed an invented concept, it is one that captures the intuition we have about inflation. But it would be better to think of the series I uncover as an index, perhaps not of inflation itself, but of the forces that affect inflation dynamics, at least in the short run. For lack of a simpler term, I choose to call this index latent inflation, but I need to show that, in practice, it can be closely related to more standard measures of inflation. I do this in the second part of the section that follows.

The model

General form

The general form of the model I use is relatively straightforward. Let Y_{it} denote the individual inflation series, with $i = 1, ..., N$. Let P_t be the latent inflation. I assume that the relation between them takes the form

$$Y_{it} = \lambda_i P_t + P_{it},$$

where λ_i is called the "loading factor." The term P_{it} represents the component specific to the individual inflation series i. I call it the "relative inflation" for the good or service i.

One would expect the loading factors to be close to 1. (As I explain later, one of them is normalized to be 1.) Indeed, it may be hard to think of a theory in which they would not be 1, since one would expect inflation to have the same impact on all series. To the extent that I do not find them to be 1, this can be reinterpreted as capturing any immediate dependence on the relative inflation from the general inflation, for example, the product of distortions generated by inflation on the pricing decisions in one sector. Formally, the equation can be rewritten as $P_{it} = P_t + P_{it}^{'}$ with $P_{it}^{'} = \left(\lambda_i - 1\right) P_t + P_{it}$. An alternative explanation is that loading factors different from 1 are picking up some model misspecification, such as nonlinear time trends.

Within this general framework, I consider a variety of statistical models for the relative inflation rates P_{it} and the latent inflation P_t.

Specific form

As I explained previously, part of the cost of the approach is that it involves many choices: not only a choice of series, but also a choice of the statistical model to apply to the series. Partly to avoid deciding, but mostly to explore the properties of the general model, I have experimented with a number of variations.

The PCE index or the CPI can be thought of as the apex of a pyramid. The general price index corresponds to the most aggregated level of observation. Immediately below, there is a first level of disaggregation. In the case of the PCE index, it contains three series: an index of durables prices, an index of nondurables prices, and an index of services prices. Further down is a second level of disaggregation, which contains 13 series, and a third level. At this stage of the research, I have experimented with a collection of three series of the PCE index (the first level of disaggregation), 13 series (the second level), and 52 series (selected from the third level).

The next decision is the choice of a statistical model for the individual series. For the sake of simplicity, I have imposed the same model on the latent and the relative inflation series, but I have varied the model. All models belong to the ARIMA (autoregressive integrated moving average) families of models, which I now explain.

The simplest statistical model one can think of is that a series is white noise; that is, it consists of realizations from uncorrelated, identically distributed statistical processes— each observation (at time t) is drawn from, say, a normal distribution with constant mean and variance. Obviously, this is not a good model for inflation, which is highly persistent, but it serves as a building block for other models. The next step is to allow for serial correlation, and imagine that inflation at time t can be decomposed into the sum of last period's realization multiplied by a factor ρ, and white noise e_t:

$$P_t = \rho\, P_{t-1} + e_t.$$

This introduces some persistence in the process. More generally, one can suppose that the process depends on more than one lag. The general form is then that of an AR(p), autoregressive process with p lags:

$$P_t = \rho_1 P_{t-1} + \rho_2 P_{t-2} + \ldots + \rho_p P_{t-p} + e_t.$$

Another step is to allow the innovation e_t to have effects that extend beyond the period when it occurs, without having as much persistence as the autoregressive part. That is, the innovation e_t affects not just P_t but also P_{t+1}:

$$P_t = \rho_1 P_{t-1} + \rho_2 P_{t-2} + \ldots + \rho_p P_{t-p} + e_t + \theta\, e_{t-1}.$$

This is a mixture of an autoregressive (AR) process with a moving average (MA) component: It is called an ARMA process. The moving average part can have q terms:

$$P_t = \rho_1 P_{t-1} + \rho_2 P_{t-2} + \ldots + \rho_p P_{t-p} + e_t + \theta_1 e_{t-1} + \ldots + \theta_q e_{t-q},$$

in which case the process is called ARMA(p,q).

Estimation is much simpler if the process is stationary; that is, its properties do not vary over time, and it tends to revert to its mean rather than drift away. This will be true if the sum of the autoregressive coefficients is less than 1 in absolute value. But this may not be a good assumption for inflation, which is so highly persistent that it can look like a random walk. One solution is to take the difference of inflation and to model that difference as an ARMA process; the original process is said to be integrated of order 1 if the first difference is stationary, and the process is called an ARIMA(p,1,q) process, where 1 denotes the fact that inflation needs to be differenced once. (I do not consider higher orders of integration.)

A final variant that I consider is to allow for feedback from the relative inflation to the latent inflation. This takes the following form: The relative inflation series are modeled as ARMA(p,q), and the latent inflation is modeled as

$$P_t = \rho_1 P_{t-1} + \rho_2 P_{t-2} + \ldots + \rho_p P_{t-p} + e_t + \theta_1 e_{t-1} + \ldots + \theta_q e_{t-q} + \Sigma\, \psi_t P_{it}.$$

I denote the model as ARIMA(p,i,q,ψ) if I allow for such feedback and ARIMA(p,i,q,~) otherwise. The same model is imposed on all series. (Further refinement of the analysis will involve imposing different models on the different series, eliminating terms that appear to be insignificant in the estimation.) Constants

BOX 2

Other methods and related literature

In addition to presenting my results, I want to say a few words about other methods.

The approach taken here is related to other work. The model I use can be seen as a special case of what are called "dynamic factor models." These models represent a given collection of variables $\{X_{1,t}, X_{2,t}, \ldots\}$ as being determined by a set of unobserved common factors $\{F_{1,t}, F_{2,t}, \ldots\}$ and their lags, to which observation noise is added. The general form of the equation modeling each variable would be

$$X_{i,t} = a_{i10} F_{1,t} + a_{i11} F_{1,t-1} + \ldots + a_{i20} F_{2,t} + a_{i21} F_{2,t-1} + \ldots + u_{i,t},$$

where the $u_{i,t}$ terms are not correlated over time and with each other (Sargent and Sims, 1977). Typically, the number of factors is kept small relative to the number of variables being modeled. More recently, researchers have found that the principal components of the collection X_i can be used to approximate the common factors F_i, an approximation that becomes valid as the number of variables X becomes large relative to the number of factors (Stock and Watson, 1998; Forni et al., 2000). These techniques are used by the Chicago Fed's National Activity Index (Evans, Liu, and Pham-Kanter, 2002). Cristadoro et al. (2002) use these methods to compute a measure of core inflation for the euro area using large numbers of economic series and extracting the slow-moving component of the common factor associated with inflation.

My approach is a particular form of a dynamic factor model, where the number of factors is the number of series plus one and estimation proceeds along the more traditional (and computer-intensive) line of maximum likelihood. Bryan and Cecchetti (1983) use this method with a small-scale model to estimate the degree of bias in the CPI (the bias being the difference between actual CPI and the estimated latent variable). They do not assess the properties of their estimated variable or its forecasting ability. Jain (1992, 2001) uses the state space approach with only price series to remove seasonal fluctuations from price series, but the focus is not on estimating latent inflation. Other uses of the state space model approach to estimate or predict inflation include Bomhoff (1982), who uses a small economic model to relate inflation to money and output; Hamilton (1985) and Burmeister, Wall, and Hamilton (1986), who estimate current expectations of inflation using variables such as interest rates; and Laubach and Williams (2003).

are included in all the ARMA models, allowing for potentially different trends in relative inflation.

Estimation method

To estimate each model, I use the so-called estimation-maximization (EM) algorithm detailed in Watson and Engle (1983). The problem is to find the values of the parameters of the model: the loading factors; the ρ, θ, ψ coefficients; and the variances of the innovations e_t for each relative inflation and for the latent inflation. The difficulty is that the Kalman filter and smoother formulas can compute estimates of the latent variable assuming that these parameters are known; however, they are not known, and they must themselves be estimated.

The EM algorithm uses the classic approach of assuming we know what we don't know. Specifically, one starts with a guess for the parameters, applies the Kalman smoother, and computes estimated series (the estimation step); then, pretending that these estimated series are observed data, one finds new estimates of the parameters, essentially by regressing the observed price series on the relative and latent inflation to compute the loading factors, as well as the hidden variables on their lags to compute the parameters of the ARMA model.[3] The main drawback of this algorithm is that it converges very slowly and makes computation time-intensive.

In box 2 (p. 59), I present a brief overview of other methods. In the following section, I discuss my results.

Results

The estimated latent inflation

As I have explained, the approach to modeling relative prices as well as latent inflation is somewhat agnostic: A variety of models have been estimated. Which does one choose? One criterion is how well the model fits the existing data (I use the sample period from 1959:Q1 to 2005:Q1). The estimation procedure tries to maximize the likelihood that the observed data were generated by the estimated model, and one can simply compare the resulting likelihood across models. Of course, models with more parameters will tend to do better, simply because they have more parameters, and ways have been devised to take this into account.[4]

Figures 1–3 show the estimated values of the latent variable over the sample, compared with the quarterly core inflation rate, for three models: two models chosen on the basis of fit—the ARIMA(2,0,0, ~) and the ARIMA(3,0,1, ψ)—and a model that will turn out to have good forecasting ability in the next section, the ARIMA(2,0,2, ~). The figures also show the forecasted path of the latent variable over the next 12 quarters. Note that this is not a forecast of core inflation, but only a forecast of the latent variable. I use this forecast of the latent inflation in the next section to forecast core inflation. Another important point is that neither the level nor the amplitude of the latent variable can be determined. The estimation procedure normalizes to 1 the first loading factor (in effect, I scale the latent variable so that its amplitude is comparable to that of the first price series), and in figures 1–3 I add the value of core inflation in 1959:Q2 to the level of the latent variable. Thus, the scale of the figure only applies to core inflation, and if the figures allow us to compare visually the two series, they should not be taken to mean that core inflation is higher or lower than the latent variable at any particular point in time.

Overall, the behavior of the latent variable is similar to that of core inflation. It's worth recalling that I did not remove food or energy from the series I used to estimate the latent variable.

Forecasting with the latent inflation measure

As I mentioned previously, one difficulty with the latent inflation approach is that the variable I am measuring is a construct. How can I be sure that it is

FIGURE 1

Estimated latent inflation, ARIMA(2,0,0,~) model

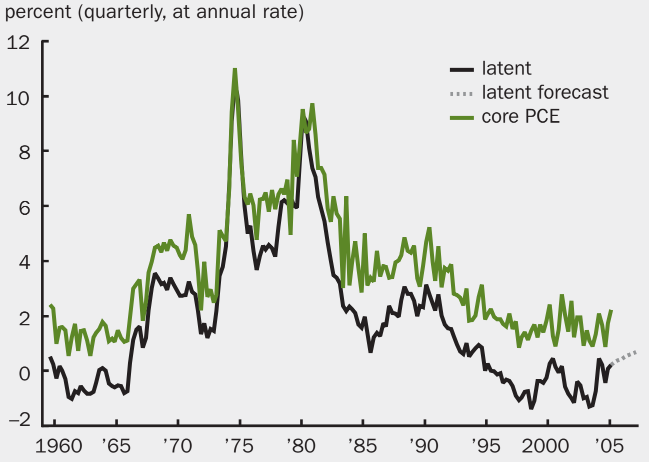

Notes: The scale of the figure only applies to core inflation. PCE means personal consumption expenditures.
Sources: Author's calculations and U.S. Bureau of Economic Analysis.

FIGURE 2

Estimated latent inflation, ARIMA(2,0,2,~) model

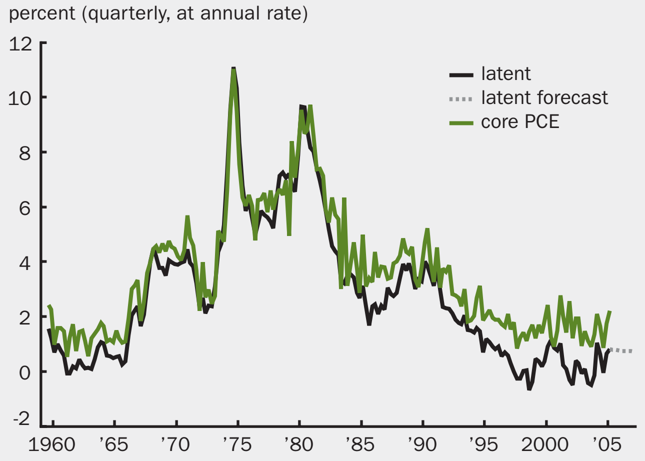

Notes: The scale of the figure only applies to core inflation. PCE means personal consumption expenditures.
Sources: Author's calculations and U.S. Bureau of Economic Analysis.

FIGURE 3

Estimated latent inflation, ARIMA(3,0,1,ψ) model

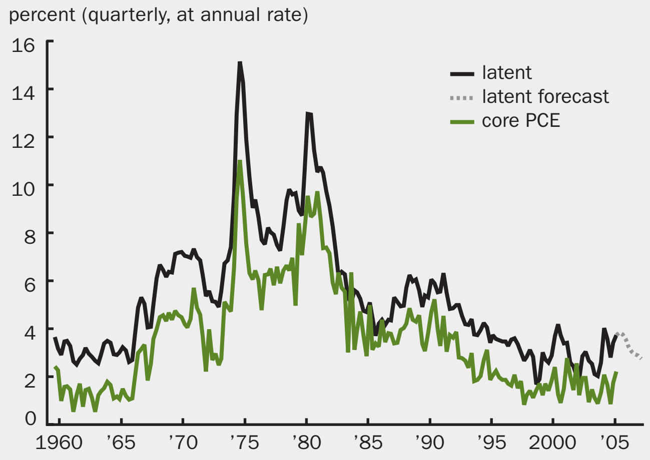

Notes: The scale of the figure only applies to core inflation. PCE means personal consumption expenditures.
Sources: Author's calculations and U.S. Bureau of Economic Analysis.

measuring what I think it might be measuring? Is it of any use or, more precisely, does it capture the latent inflationary pressures that are in play, at least in the short term?

One way to evaluate the latent inflation measure is to find out if it holds any predictive power for inflation as it is commonly measured. To find out, I carry out an out-of-sample forecasting exercise similar to Fisher, Liu, and Zhou (2002) and Brave and Fisher (2004). As this research has emphasized, the naive model of inflation, which predicts that inflation in the future will be what its most recent value was, is "the man to beat."

I proceed as follows. For each quarter T between 1984:Q2 and 2002:Q2, I take the sample ranging from the beginning of the series (1959:Q1) to the chosen quarter T. Using only the data in this sample, I estimate a family of ARIMA models. Then, I run various regressions of core inflation over various horizons (that is, core inflation from quarter $t-H$ to quarter t, where H ranges from 1 to 8) on the estimated measure of the latent inflation and current and lagged core inflation *within the sample*. Then, I construct a forecast of latent inflation over the horizon T to $T+H$ and use those forecasts as well as the values for current and lagged core inflation to project core inflation over the horizon T to $T+H$. Having done this for all quarters T between 1985:Q1 and 2002:Q2, I compute the root mean squared error (RMSE) of these forecasts. I compare this RMSE to the RMSE of the naive model, which simply predicts that core inflation over T to $T+H$ will be what core inflation was from $T-4$ to T.

Using the same notation as Brave and Fisher (2004), core inflation from $t-H$ to t is

$$\pi_t^H = \ln p_t - \ln p_{t-H},$$

while core inflation from $t-1$ to t is simply denoted $\pi_t = \ln p_t - \ln p_{t-1}$.

Note that

$$\pi_t^H = \pi_{t-H+1} + \pi_{t-H+2} + \ldots \pi_t.$$

Latent inflation x_t is calculated as the latent variable in the statistical model, and the latent variable over the T to $T+H$ horizon is simply

$$x_t^H = x_{t-H+1} + x_{t-H+2} + \ldots x_t.$$

The regression I run is

$$1)\quad \pi_t^H = \alpha x_t^H + \beta_0 \pi_{t-H} + \beta_1 \pi_{t-H-1} + \gamma.$$

The statistical model allows me to project $\hat{x}_{T+H}^H$ and then construct an estimate

$$\hat{\pi}_{T+H}^H = \alpha \hat{x}_{T+H}^H + \beta_0 \pi_T + \beta_1 \pi_{T-1} + \gamma.$$

Note that, in equation 1, latent inflation is included in the regression in addition to lagged inflation. Such an inclusion usually hurts the predictive power of the forecasting equation (out of sample). By contrast, if latent inflation helps significantly, this is a success. Note also that, although a lot of work goes into coming up with the series x_t^H and the forecast $\alpha \hat{x}_{T+H}^H$, the regression itself is simple and has only three regressors.

The results in terms of relative RMSE are shown in table 1. For each model and each horizon (one to eight quarters ahead), the table shows the model's RMSE relative to the naive model (a number lower than 1 indicates that the model performs better). The models are sorted by order of increasing likelihood.

The pattern of performance varies considerably across models. One group of models does substantially worse than the others: As it turns out, these are the models that allow feedback from lagged relative inflation to latent inflation. The models without feedback do markedly better than regressing core inflation on two quarters of inflation, the performance of which is given in the last row of table 1. In other words, the addition of the latent inflation to the regression substantially improves the forecasting performance. Which model performs best depends on the horizon: At the short horizon (one to three quarters), the ARIMA(2,1,0,~)

TABLE 1

Root mean squared error relative to naive model

	Forecast horizon one to eight quarters ahead							
ARIMA(2,1,0) without feedback	1.09	1.13	1.12	1.16	1.20	1.23	1.24	1.25
ARIMA(2,1,0) with feedback	1.34	1.70	1.71	1.82	1.89	1.91	1.93	1.94
ARIMA(3,1,0) with feedback	1.30	1.81	1.84	1.96	2.05	2.12	2.16	2.18
ARIMA(2,0,1) without feedback	1.11	1.16	1.15	1.17	1.20	1.23	1.24	1.25
ARIMA(2,0,2) without feedback	1.12	1.15	1.11	1.10	1.10	1.10	1.09	1.08
ARIMA(3,0,1) without feedback	1.10	1.15	1.11	1.11	1.12	1.11	1.11	1.10
ARIMA(2,0,0) without feedback	1.17	1.24	1.25	1.27	1.29	1.29	1.28	1.28
ARIMA(3,0,0) without feedback	1.19	1.25	1.22	1.21	1.22	1.22	1.21	1.20
ARIMA(2,0,0) with feedback	1.38	1.74	1.91	2.02	2.07	2.08	2.06	2.03
ARIMA(3,0,0) with feedback	1.36	1.74	1.91	2.00	2.07	2.09	2.08	2.05
ARIMA(3,0,1) with feedback	1.41	1.83	2.04	2.17	2.25	2.30	2.30	2.29
1 lag of inflation alone	1.08	1.17	1.24	1.31	1.37	1.44	1.50	1.54

TABLE 2

Performance of the moving average versions of the models' forecasts

	Two-quarter moving average							
ARIMA(2,1,0) without feedback	1.02	1.02	1.03	1.05	1.08	1.09	1.11	1.15
ARIMA(2,0,1) without feedback	1.04	1.01	0.97	0.97	0.99	1.01	1.05	1.08
ARIMA(2,0,2) without feedback	1.05	1.01	0.96	0.95	0.95	0.97	1.00	1.03
ARIMA(3,0,1) without feedback	1.03	1.00	0.95	0.94	0.95	0.97	1.01	1.04
ARIMA(2,0,0) without feedback	1.08	1.04	1.01	1.00	1.00	1.00	1.02	1.05
ARIMA(3,0,0) without feedback	1.08	1.03	0.98	0.96	0.96	0.97	1.00	1.03
	Three-quarter moving average							
ARIMA(2,1,0) without feedback	0.98	0.96	1.00	1.02	1.06	1.08	1.12	1.16
ARIMA(2,0,1) without feedback	0.98	0.95	0.94	0.96	0.99	1.03	1.08	1.12
ARIMA(2,0,2) without feedback	0.99	0.95	0.93	0.93	0.95	0.99	1.03	1.07
ARIMA(3,0,1) without feedback	0.97	0.94	0.92	0.93	0.96	1.00	1.04	1.08
ARIMA(2,0,0) without feedback	1.01	0.98	0.97	0.97	0.98	1.00	1.04	1.09
ARIMA(3,0,0) without feedback	0.99	0.96	0.93	0.93	0.95	0.98	1.02	1.06

does slightly better; for all other horizons, the winner is the ARIMA(2,0,2,~), with the ARIMA(3,0,1,~) not far behind. Neither one, however, manages to do any better than the naive model, though they come reasonably close, within 10 percent of the RMSE of the naive model.

FIGURE 4

Predictions of the naive model and ARIMA(2,0,2,~) model vs. actual inflation

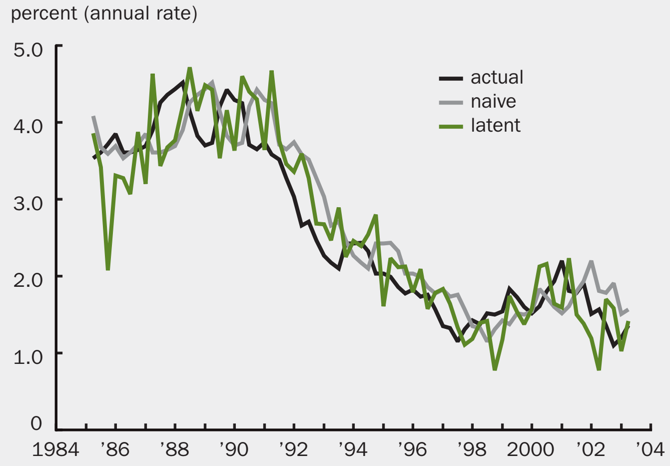

Note: See the text for further details.
Sources: Author's calculations and U.S. Bureau of Economic Analysis.

FIGURE 5

Predictions of the naive model and ARIMA(2,0,2,~) model (smoothed) vs. actual inflation

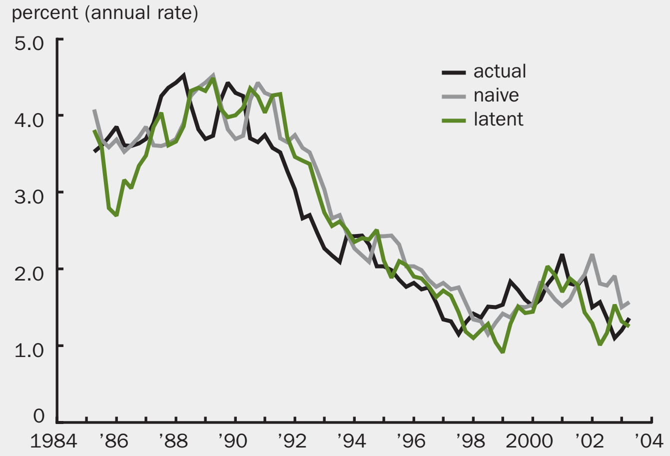

Note: See the text for further details.
Sources: Author's calculations and U.S. Bureau of Economic Analysis.

Figure 4 compares the forecasts of core inflation produced by the naive model (gray line) and the latent inflation model ARIMA(2,0,2,~) (green line) with actual core inflation (black line) at the two-year horizon. The date on the horizontal axis is the date at which the forecast is made. The gray line is the black line shifted by two years, since the naive model predicts that inflation two years hence will be the same as today. The predictions of the latent model are not substantially different from those of the naive model, and hence the latent model does not perform any better. But what is striking is how variable the green line is, relative to the gray line. The reason is as follows. The gray line averages actual inflation over the previous eight quarters and therefore smoothes out a lot of the quarter-to-quarter variability in inflation. The latent inflation model incorporates the new information that arrives in each quarter, and even though it weights it appropriately, the new information shifts the estimate of where latent inflation currently stands; this in turn shifts the whole projected path of latent inflation, and hence the forecast of core inflation. There is no smoothing mechanism here.

It is possible, of course, to add a smoothing mechanism.[5] For example, I have tried replacing the latent model's forecast with a two-quarter or three-quarter moving average of itself. This ad hoc procedure produces a smoother forecast. Its performance is shown in table 2, only for selected models.

The performance of both the ARIMA (2,0,2,~) and the ARIMA(3,0,1,~) models is improved markedly. Just taking a two-quarter moving average reduces the relative RMSE for the ARIMA(2,0,2,~) from 1.10 to 0.95. It becomes possible to beat the naive model, although not by a great amount. Figure 5 compares the predictions of this moving average: The green line is clearly smoother, and in some instances, it seems to do better in terms of predicting changes in inflation (for example, the downturn in the mid-1980s and the pick up in the early 1990s).

Conclusion

This article has presented recent research on measuring and forecasting inflation. The approach taken, that of state space modeling, consists of representing latent inflation as an unobserved variable affecting simultaneously a collection of individual price series, for example, the main components of an aggregate price index like the PCE deflator. The approach extends the intuition that lies behind the use of core measures of inflation in that it takes the individual price series to be noisy observations on true, underlying inflation, and filters out the noise in the individual price series. The resulting estimated latent inflation validates the use of core inflation, since the two series look very much alike. The latent inflation approach has the additional benefit of yielding a forecast of future inflation, and preliminary results indicate that some progress can be made in reducing out-of-sample forecasting error.

NOTES

[1]For an explanation of the random walk model, please refer to Brave and Fisher (2004) and Fisher, Liu, and Zhou (2002).

[2]Note that I am not fully escaping the use of weighted indexes, since the individual price series will, in practice, be indexes of their own.

[3]If the model has moving average components, the e_t series are treated as yet another unobserved variable.

[4]Two such criteria are commonly used, the Bayesian information criterion (BIC) and the Akaike information criterion (AIC), the former tending to be stricter than the latter. In my family of models, the BIC chooses the most parsimonious model, the ARIMA(2,0,0) with no feedback, while the AIC ranks almost equally the ARIMA (3,0,0) with feedback and the ARIMA(3,0,1) with feedback.

[5]I thank former Federal Reserve Board Chairman Alan Greenspan for this suggestion.

REFERENCES

Bomhoff, Eduard J., 1982, "Predicting the price level in a world that changes all the time," *Carnegie-Rochester Conference Series on Public Policy*, Vol. 17, pp. 7–56.

Brave, Scott, and Jonas D. M. Fisher, 2004, "In search of a robust inflation forecast," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 28, No. 4, Fourth Quarter, pp. 12–31.

Bryan, Michael F., and Stephen G. Cecchetti, 1983, "The Consumer Price Index as a measure of inflation," National Bureau of Economic Research, working paper, No. 4505, October.

Burmeister, Edwin, Kent D. Wall, and James D. Hamilton, 1986, "Estimation of unobserved expected monthly inflation using Kalman filtering," *Journal of Business & Economic Statistics,* Vol. 4, No. 2, April, pp. 147–160.

Cristadoro, Riccardo, Mario Forni, Lucrezia Reichlin, and Giovanni Veronese, 2002, "A core inflation index for the euro area," Center for Economic Policy Research, discussion paper, No. DP3097.

Evans, Charles L., Chin Te Liu, and Genevieve Pham-Kanter, 2002, "The 2001 recession and the Chicago Fed National Activity Index: Identifying business cycle turning points," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 26, No. 3, Third Quarter, pp. 26–43.

Fisher, Jonas D. M., Chin Te Liu, and Ruilin Zhou, 2002, "When can we forecast inflation?," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 26, No. 1, First Quarter, pp. 30–42.

Forni, Mario, Marc Hallin, Marco Lippi, and Lucrezia Reichlin, 2000, "The generalized dynamic factor model: Identification and estimation," *Review of Economics and Statistics*, Vol. 82, No. 4, November, pp. 540–554.

Hamilton, James D., 1985, "Uncovering financial market expectations of inflation," *Journal of Political Economy*, December, Vol. 93, No. 6, pp. 1224–1241.

Jain, Raj K., 2001, "A state space model-based method of seasonal adjustment," *Monthly Labor Review,* Vol. 124, No. 7, July, pp. 37–45.

__________, 1992, “A state space modelling approach to the seasonal adjustment of the Consumer Price and other BLS Indexes: Some empirical results,” U.S. Bureau of Labor Statistics, working paper, No. WP-229.

Laubach, Thomas, and John C. Williams, 2003, “Measuring the natural rate of interest,” *Review of Economics and Statistics*, Vol. 85, No. 4, November, pp. 1063–1070.

Sargent, Thomas J., and Christopher A. Sims, 1977, “Business cycle modeling without pretending to have too much a priori economic theory,” in *New Methods of Business Cycle Research*, Christopher A. Sims (ed.), Federal Reserve Bank of Minneapolis, pp. 45–109.

Stock, James H. and Mark W. Watson, 1998, “Diffusion indexes,” National Bureau of Economic Research, working paper, No. 6702, August.

Watson, Mark W., and Robert F. Engle, 1983, “Alternative algorithms for the estimation of dynamic factor, MIMIC, and varying coefficient regression models,” *Journal of Econometrics*, Vol. 23, No. 3, December, pp. 385–400.