

The Roles of Inflation Expectations, Core Inflation, and Slack in Real-Time Inflation Forecasting*

N. Kundan Kishor[†] Evan F. Koenig[‡]

Abstract

We extract information from surveys of long-term inflation expectations and multiple quarterly inflation series to undertake a real-time decomposition of headline PCE and GDP-deflator inflation rates into a common long-term trend, common cyclical component, and high-frequency noise components. We then explore alternative approaches to real-time forecasting of PCE inflation. We find that performance is enhanced if forecasting equations are estimated using inflation data that have been stripped of high-frequency noise. Performance can be further improved by including an unemployment-based measure of slack in the equations. The improvement is statistically significant relative to benchmark autoregressive models and also relative to professional forecasters at all but the shortest horizons. In contrast, introducing slack into models estimated using headline PCE inflation data or conventional core inflation data causes forecast performance to deteriorate. Finally, we demonstrate that forecasting models estimated using the Kishor-Koenig (2012) methodology—which mandates that each forecasting VAR be augmented with a flexible state-space model of data revisions—consistently outperform the corresponding conventionally estimated forecasting models.

Keywords: Inflation, Real-Time Forecasting, Unobserved Component Model, Slack

JEL Codes: E31, E37.

*The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Dallas or the Federal Reserve System. We have benefitted from comments received at the 2015 CIRANO Realtime Workshop (Montreal, October 2015), Midwest Econometrics Group Meeting (St. Louis, October 2015), the "day-ahead" meeting of the Federal Reserve System Committee on Macroeconomics (Cleveland, November 2015) and SNDE Symposium (Tuscaloosa, March 2016).

[†]Kishor: Associate Professor, Department of Economics, Box 413, Bolton Hall 822, University of Wisconsin-Milwaukee, Milwaukee, WI 53201. Email: kishor@uwm.edu.

[‡]Koenig: Senior Vice President and Principal Policy Advisor, Federal Reserve Bank of Dallas, 2200 N. Pearl Street, Dallas, TX 75201. Email: evan.f.koenig@dal.frb.org.

1 Introduction

Forecasting inflation is of great importance to policymakers, households and businesses, and the academic literature on forecasting inflation is vast. Nevertheless, there is no consensus on the relative usefulness of different inflation forecasting methodologies. This lack of agreement can be partly attributed to instability in U.S. inflation dynamics over time. Notably, Stock and Watson (2007) have shown that inflation in the U.S. became much less variable during the post-1983 "Great Moderation" period, and that changes in inflation simultaneously became much harder to predict. Mean-squared forecast errors from a variety of standard inflation models shrink during the Great Moderation, but it is very difficult to improve on the forecasts generated by a parsimonious autoregressive or random-walk model.

One strand of the inflation-forecasting literature attributes the apparent instability of the inflation process to changes in the behavior of long-run inflation expectations. Studies that explicitly model long-run inflation expectations include, for example, Kozicki and Tinsley (2001), Stock and Watson (2010), Cogley, Primiceri and Sargent (2010), and Mertens (2016). Other researchers (*e.g.*, Koenig and Atkinson, 2012; Faust and Wright, 2013; Clark and Doh, 2014) have used survey measures of expectations as an input into their forecasting models. Ang *et al.* (2007) show that surveys often do a better job of forecasting inflation than models that are based on macroeconomic and financial variables.

Properly controlling for changes in inflation's longer-run trend is one issue in the inflation-forecasting literature. Another is the link between inflation and economic slack. This connection has been studied going back at least to Phillips (1958), but remains controversial. In a widely cited paper, Atkeson and Ohanian (2001) show that slack adds nothing to the forecasting power of a simple random-walk inflation model. Researchers who focus on medium-frequency inflation movements, though, have found a strong and robust Phillips-curve relationship. Examples of such studies include Stock and Watson (2010), which looks at the impact of slack on core PCE inflation during and immediately following recessions, and Koenig and Atkinson (2012), which examines the link between the unemployment rate

and subsequent deviations of trimmed-mean PCE inflation from a survey measure of long-run inflation expectations.¹ Stock and Watson's forecasting exercises, though, do not use real-time data, and the Koenig-Atkinson study is constrained by the limited availability of real-time trimmed-mean inflation data, which extend back only to 2005.² So, the usefulness of economic slack in real-time inflation forecasting remains in doubt.

We view survey measures of long-run inflation expectations as informative for trend inflation, but not necessarily definitive.³ Similarly, we believe that the available evidence justifies skepticism about the usefulness of standard measures of core inflation in forecasting headline inflation.⁴ Accordingly, rather than rely on an off-the-shelf measure, we use a multivariate unobserved-components (UC) analysis to *infer* core inflation from a variety of published inflation and inflation-expectations series under two alternative assumptions about trend inflation.⁵ In our baseline specification, trend inflation is treated as unobserved and is estimated simultaneously with core inflation under the assumption that trend inflation became well anchored in the late 1990s.⁶ As a robustness exercise, however, we also consider a simple alternative specification in which trend inflation is obtained by applying a one-sided Hodrick-Prescott (HP) filter to long-run inflation expectations as captured by the Survey of

¹Dolmas (2005) describes the procedures used to calculate the trimmed-mean PCE inflation measure used by Koenig and Atkinson. Data are available on the Federal Reserve Bank of Dallas website.

²Stock and Watson undertake a "pseudo real-time" analysis, which means that they conduct a recursive forecasting exercise using latest-vintage data. The ex-food-and-energy PCE inflation data available on the Federal Reserve Bank of Philadelphia's website extend back to vintage 1996:Q1.

³See Chan, Clark and Koop (2015) for a similar take on the usefulness of survey measures.

⁴We take it for granted that the analyst's goal is to forecast headline inflation. For this purpose it's desirable that core inflation strip as much unforecastable noise as possible from headline inflation (Koenig, Dolmas, and Piger, 2003). However, other definitions of core inflation may have their justifications, too. For example, on theoretical grounds a central bank might want to stabilize sticky-price inflation rather than headline inflation. Then, a core inflation measure that excludes flexible-price goods and services would be of interest to policymakers. Similarly, in some models it's appropriate for monetary policy to react differently to supply-side shocks than to demand-side shocks. So, a core inflation measure that strips out the effects of supply shocks could be useful. For an early discussion of these issues, see Bryan and Cecchetti (1993). Smith (2004) and Crone *et al.* (2013) examine whether various core inflation measures are helpful for forecasting headline inflation.

⁵Our approach is in the spirit of Basistha and Nelson (2007) and Basistha and Startz (2008), who show that multivariate unobserved-component models provide more precise and economically meaningful estimates of the output gap and natural rate of unemployment.

⁶See Nalewaik (2015) and Koenig and Atkinson (2012) for evidence of a shift to an anchored inflation process in the 1990s.

Professional Forecasters. Next, we examine whether our trend and UC-filtered core inflation measures are helpful—either by themselves or in combination with labor-market slack—for forecasting headline inflation.⁷ The inference process and the forecasting exercises are all genuinely "real time": They use only data that would have been available in the quarter during which the forecast would have been prepared.

Key results are as follows: First, we confirm the Atkeson and Ohanian (2001) result that adding slack to an autoregressive model of headline inflation fails to improve forecasting performance, and the Crone *et al.* (2013) result that using conventional core inflation to predict headline inflation produces little improvement in forecasting performance. In the same vein, we find that professional forecasters perform no better—and often significantly worse—than a simple autoregressive model when looking out beyond the current quarter. Second, nevertheless, headline-inflation forecasts based on UC-filtered core inflation or, especially, UC-filtered core inflation together with slack substantially and significantly outperform forecasts based on headline inflation alone or headline inflation in combination with slack. These key results are robust to how we model trend inflation. Summarizing: Slack *is* of use in forecasting headline inflation, but only if inflation data are carefully filtered to exclude high-frequency noise prior to estimation of the forecasting equations.

To construct our real-time inflation forecasts we rely on the VAR methodology described in Kishor and Koenig (2012). This method augments each VAR forecasting model with a flexible state-space model of data revisions. In previous research (Kishor and Koenig, 2012, 2014), we've shown that this approach ("the KK method") typically produces forecasts that are more accurate than those produced by estimation which takes at face value the very latest data that would have been available in real time ("conventional" real-time estimation). Consistent with those results, the present analysis finds that real-time inflation forecasts obtained using the KK method consistently and often significantly outperform inflation forecasts from conventionally estimated models.

⁷Our benchmark measure of slack is the "unemployment recession gap", which is defined by Stock and Watson (2010) as the difference between the current unemployment rate and the minimum unemployment rate over the current and previous eleven quarters.

The remainder of the paper is organized as follows: Section 2 presents our baseline real-time, unobserved-components model of inflation; Section 3 provides a description of the data used in our empirical analysis; Section 4 reviews econometric issues in real-time forecasting and how we deal with them; Section 5 presents forecasting results; Section 6 examines the robustness of these results to our alternative model of trend inflation, and presents other robustness checks; and Section 7 concludes.

2 A Real-Time, Unobserved-Common-Component Model of Inflation

Survey measures of long-term expected inflation play an important role in policy deliberations and provide valuable information about future movements in inflation. Ang et al. (2007) have shown that survey measures of expected inflation do a better job of forecasting inflation than macro variables and asset prices. That survey measures of inflation expectations can be used to improve forecasting performance has been demonstrated by Faust and Wright (2013), Stock and Watson (2010), Koenig and Atkinson (2012), and Kozicki and Tinsley (2001), among others. Most of these papers use a single survey measure of long-run inflation expectations. We, instead, combine the information from two different survey measures: one from the Survey of Professional Forecasters (SPF) and the other from the Blue Chip. (Details are provided in Section 3, below.) Included in our analysis, also, are core CPI inflation and real-time headline PCE and GDP inflation rates. Our approach decomposes each of these five series into a common long-term trend, a cyclical component, and a combination of shared and idiosyncratic high-frequency noise⁸. The approach has several advantages. First, it overcomes constraints on the availability of real-time data that apply to many off-the-shelf measures of core inflation. Second, the definition of core inflation can be tailored to our intended purpose—which, here, is to help forecast headline inflation.⁹ Finally, the

⁸Stock and Watson (2016) use disaggregated data on sectoral inflation to estimate trend PCE inflation, whereas Kiley (2008) uses a bi-variate, common-trend framework to examine the role of food and energy prices in the dynamics of trend inflation for both PCE and CPI inflation.

⁹See the discussion in footnote #4, above.

extraction of different components from a multivariate model provides us with more-precise estimates of long-run inflation expectations and the various cyclical inflation components. The multivariate approach has been shown to be useful in a state-space setting by Clark (1989), Basistha and Nelson (2007), and Basistha and Startz (2008).¹⁰

Our unobserved-common-component model has the following structure:

$$\pi^{PCE} = \mu_1 + \tau_t + c_t + \lambda_t + \varepsilon_t + \eta_{1t} \quad (1)$$

$$\pi^{CCPI} = \mu_2 + \tau_t + c_t + \lambda_t + \eta_{2t} \quad (2)$$

$$\pi^{GDP} = \mu_3 + \tau_t + c_t + \lambda_t + \varepsilon_t + \eta_{3t} \quad (3)$$

$$\pi^{SPF} = \tau_t + c_t^{SPF} + \eta_{4t} \quad (4)$$

$$\pi^{BC} = \mu_4 + \tau_t + c_t^{BC} + \eta_{5t}. \quad (5)$$

Here π^{PCE} , π^{CCPI} and π^{GDP} are 1-quarter rates of headline PCE inflation, conventional-core (*i.e.*, ex-food-and-energy) CPI inflation, and GDP deflator inflation, respectively, while π^{SPF} and π^{BC} are long-term CPI inflation expectations from the SPF and the Blue-Chip survey. Each of the five inflation series incorporates the same unobserved long-term trend component, τ_t . In addition, the three short-term inflation rates have a shared cyclical component, c_t , and a shared white-noise term λ_t . To take into account a very high degree of correlation between GDP inflation and PCE inflation, we also introduce a white noise component ε_t that is common across these two short-term inflation rates. The two long-term survey measures have individual cyclical components (c_t^{SPF} and c_t^{BC}). Finally, each series has a noise

¹⁰An example of an analysis that takes the multivariate approach to infer longer-run trend inflation is Mertens (2016). Notably, Mertens finds that among different measures of realized inflation, trimmed-mean PCE inflation, because it filters out short-term noise, is particularly useful as a signal of trend inflation. Trimmed-mean PCE inflation is unsatisfactory for our purposes, however, because real-time estimates only become available in 2005.

component, η_{jt} , that is independently and identically distributed over time. Specifically, $\eta_{jt} \sim iidN(0, \sigma_{\eta_j}^2)$ for $j = 1, 2, 3, 4, 5$. By construction, all the noise components (common as well as idiosyncratic) are not forecastable. They are uncorrelated with each other. In addition, we also assume that shocks to trend and cycles are uncorrelated with each other. Idiosyncratic noise across equations are uncorrelated with the trend and also the cycles of each series.

UC-filtered core inflation is identical, up to a constant, for PCE, ex-food-and-energy core CPI, and GDP inflation. It includes trend inflation (which is common to π^{PCE} , π^{CCPI} , π^{GDP} , π^{SPF} and π^{BC}) and the cyclical component of inflation that is common to PCE, conventional-core CPI, and GDP inflation:

$$\pi^{FPCE} = \mu_1 + \tau_t + c_t \tag{6}$$

$$\pi^{FCPI} = \mu_2 + \tau_t + c_t. \tag{7}$$

$$\pi^{FGDP} = \mu_3 + \tau_t + c_t. \tag{8}$$

Importantly, given that the difference between headline inflation and UC-filtered core inflation, as we have defined the latter, is unforecastable noise, predictions of UC-filtered inflation also serve as predictions of headline inflation. Indeed, by stripping away unforecastable noise, one can expect to obtain more-precise coefficient estimates when estimating a forecasting equation for UC-filtered inflation than when estimating a forecasting equation for headline inflation. More-precise coefficient estimates mean more-accurate forecasts—including, particularly, more-accurate forecasts of headline inflation (Koenig, Dolmas, and Piger, 2003).

Our strategy is to construct real-time-vintage estimates of π^{FPCE} and π^{FGDP} ; estimate VAR forecasting models for these variables, augmented with equations describing the data-revisions process; and then look at how well these models forecast *headline* inflation, in real time, in comparison with simple autoregressions and in comparison with the headline

inflation forecasts of professional forecasters. Our primary focus is PCE inflation, as the Federal Reserve policymakers have defined price stability in terms of that inflation gauge and include it in their quarterly forecasting exercise (as reported in the Federal Reserve's Summary of Economic Projections). As a robustness check, however, we also present results for GDP inflation.

2.1 The Dynamics of Trend and Cycle

A significant amount of work that has undertaken to study the dynamics of inflation in the U.S. has shown the existence of a slow-moving trend.¹¹ Trend inflation gradually increased during the Great Inflation period that extended from the late 1960s through the 1970s, and then declined during the post-1983 Great Moderation period. This trend is often modeled as a random walk. However, there is a strong evidence of a break in the dynamics of trend inflation in the late 1990s, with long-run inflation expectations having become "well anchored" at that point (Koenig and Atkinson, 2012). To take into account the anchoring of inflation expectations that evidently took place in the late 1990s, we adopt the following dynamics for trend inflation:

$$\tau_t = \delta_S + \theta\tau_{t-1} + v_{st}, \quad (9)$$

where $v_{st} \sim iidN(0, \sigma_S^2)$, and where $\theta = 1, \delta_S = \delta_1$ and $\sigma_S^2 = \sigma_1^2$ if $t \leq 1997:Q4$, and $\theta = \theta_2, \delta_S = \delta_2$ and $\sigma_S^2 = \sigma_2^2$ if $t > 1997:Q4$. This specification takes into account the random-walk nature of long-run inflation expectations before 1997:Q4, and the mean-reverting nature of these expectations after 1997:Q4. A standard Bai and Perron (1998) stability test also identifies a clear break in the dynamics of long-run inflation expectations at the end of 1997.^{12,13}

¹¹See for example, Kozicki and Tinsley (2001, 2005), Cogley, Primiceri and Sargent (2010), Stock and Watson (2010), Faust and Wright (2013), and Wright (2012).

¹²Koenig and Atkinson (2012) also found a break in the dynamics of long-horizon inflation expectation in 1997:Q4. They performed Quandt-Andrews stability test on a regression of long-horizon inflation expectation on its own lag and lagged trimmed-mean inflation. We also perform Bai and Perron (1998) break test on the first difference of the HP trend of SPF's inflation expectation series and found a break in 1999:Q1. The break date assumed in our exercise (1997:Q4) falls within 95% confidence band of the break date estimated using Bai and Perron (1998) approach.

¹³Our assumption of a random walk trend in the earlier sample period and stationary expectation in the later sample period is also broadly consistent with recent work by Nalewaik (2015), which finds that the U.S.

An issue when doing real-time forecasting using our model is how to handle the 1997:Q4 break in the dynamics of long-run inflation expectations. Even if an analyst had known in 1998:01 that a break had taken place, she would have had only a single observation with which to estimate the new inflation process. Our assumption is that the analyst would not have taken the break into account until 2002:Q1. During the 4-year period from 1998:Q1 through 2001:Q4, we assume that the analyst would have continued to estimate a model with a random-walk inflation trend.¹⁴ Our handling of this issue is consistent with the policy discussion of the time. It was at the beginning of 2003, for example, that Fed-Governor Bernanke publicly noted that inflation had entered a new regime, stating that “Inflation breached the 2 percent barrier in the spring of 1996 and has remained consistently within the narrow range of 1.5 to 2 percent for the past six and a half years—for practical purposes, a good approximation to price stability” (Bernanke, 2003).

As previously noted, our three short-term inflation series share a common cyclical component, c_t . This component is assumed to follow an AR(2) process:

$$c_t = \phi_1 c_{t-1} + \phi_2 c_{t-2} + u_t, \quad (10)$$

where $u_t \sim iidN(0, \sigma_u^2)$. In addition, the two long-term survey measures have idiosyncratic cycles, c_t^{SPF} and c_t^{BC} , that are assumed to follow AR(1) processes. Note that the stationary component of τ_t is distinct from c_t (which is common across the three short-run inflation series). The trend component of inflation expectations is common across all equations, whereas c_t is common across only equations (1-3). This difference helps us identify the stationary portion of inflation expectations.

The above model can be put into state-space form and estimated using maximum likelihood via the Kalman filter.¹⁵ We do the estimation recursively, using real-time data, starting with a sample that runs from 1984:Q1 through 1992:Q1. The estimated inflation cycles and

economy entered a stable-mean-low-variance regime in the 1990s. An alternative approach—used by Stock and Watson (2007), for example—is to assume that trend inflation follows a random walk over the entire sample, but that the variance of trend-inflation innovations changes over time.

¹⁴Results don’t change if we allow the estimation break to take place one year earlier or one year later.

¹⁵Appendix A shows the state-space representation of the system. For the details on the estimation procedure, see Chapter 2 of Kim and Nelson (2000).

trend are revised with each extension of the sample period, both because the amount of data available for inference increases and because official PCE and GDP inflation data are revised.

3 Data Description

We use various vintages of official PCE and GDP deflator inflation data taken from the Federal Reserve Bank of Philadelphia's website. We use core CPI inflation—which is revised only when seasonal factors are reestimated—as another measure of inflation.¹⁶ Our inflation measures are calculated as quarterly percent changes in the price level and are annualized. The unemployment rate is the quarter-average civilian unemployment rate. We use two measures of long-run inflation expectations. The most straightforward of these comes from a Blue Chip survey published twice each year, in early March and early October: It is the average CPI inflation rate that respondents expect will prevail 6-to-10 years out. We use March survey results for Q1 and Q2 of each year and October survey results for Q3 and Q4.¹⁷ Our second measure of long-run inflation expectations is calculated from Survey of Professional Forecasters (SPF) 10-year and 1-year CPI inflation expectations: Specifically, it is defined as $(10 * cpi10 - cpi1) / 9$, where $cpi10$ and $cpi1$ are 10-year and 1-year median expected inflation rates, respectively. It captures the expectation for inflation 2-to-10 years out implicit in SPF respondents' 10-year and 1-year inflation forecasts.¹⁸ Importantly, both of our measures of long-run inflation expectations are *forward* rates—rates that exclude forecasters' expectations for the coming year. We don't want our own forecasts, which extend out as far as 4 quarters, to "piggyback" on the near-term forecasts of professionals. Indeed, we want to compare the accuracy of our inflation forecasts with the accuracy of professionals' forecasts. For that comparison, we splice together headline-PCE inflation forecasts from the SPF (first

¹⁶Revisions resulting from the reestimation of seasonal factors are typically small, and we ignore them.

¹⁷The timing is conservative. First-quarter PCE and GDP inflation releases are unavailable until late April, and third-quarter PCE and GDP inflation releases are unavailable until late October.

¹⁸SPF 10-year CPI inflation expectations are first available in 1991:Q4. Before then, we substitute Blue Chip 10-year CPI expectations.

available in 2000:Q1) and “Greenbook” forecasts prepared by the Board of Governors’ staff in advance of Federal Reserve policy meetings. It is important to note that their access to high-frequency data gives SPF and Greenbook forecasters a distinct information advantage when predicting current-quarter inflation.¹⁹

Our data run from 1984:Q1 through 2015:Q1. The sample is determined by the availability of survey measures of inflation expectations and professional forecasts of headline PCE inflation. However, our sample period also coincides with the period that has been found to be particularly challenging for inflation forecasters (Stock and Watson, 2010).

3.1 Results from the Unobserved-Common-Component Model

Figure 1 plots the inflation trend from the UC model (first revision) along with SPF 9-year, 1-year forward expected inflation and Blue Chip 5-year, 5-year forward expected inflation. The inflation trend from the UC model successfully filters out idiosyncratic short-term swings in the two survey measures, such as the dip in SPF expected inflation in the first quarter of 1999 and the temporary surge in Blue Chip expectations in early 2002. The sharp change in the behavior of SPF expected inflation and of our estimate of trend inflation starting in 1998 is clear.

First-release and first-revision estimates of UC-filtered core PCE inflation are plotted in Figure 2, along with end-of-sample-vintage headline PCE inflation. Similarly, Figure 3 plots first-release and first-revision estimates of UC-filtered core GDP inflation, along with end-of-sample-vintage headline GDP inflation. By construction, UC-filtered PCE and UC-filtered GDP inflation are very much alike. Indeed, they would parallel one another were it not for real-time updating of their respective constant terms (*c.f.* equations 6 and 7). The noise component of headline PCE inflation is clearly much larger than that of headline GDP inflation.

¹⁹If $P(t)$ is the log average price level in quarter t and $p_i(t)$ is the log price level in month i of quarter t , then to a close approximation $P(t) - P(t - 1) = [(p_3(t) - p_2(t)) + 2(p_2(t) - p_1(t)) + 3(p_1(t) - p_3(t - 1)) + 2(p_3(t - 1) - p_2(t - 1)) + (p_2(t - 1) - p_1(t - 1))]/9$. SPF participants have access to $(p_3(t - 1) - p_2(t - 1))$ and $(p_2(t - 1) - p_1(t - 1))$ at the time they prepare their forecasts, and may have some information pertaining to $(p_1(t) - p_3(t - 1))$, as well.

4 Forecasting Methodology

As we have seen, UC methods can be used to strip unforecastable noise from headline inflation. Our next step is to see whether the resultant "core" inflation measures can be used to improve real-time forecasts of headline inflation. Headline PCE and GDP inflation are subject to substantial revision, and so are our versions of core PCE and core GDP inflation.²⁰ We require a forecasting strategy that takes these revisions into account and handles them appropriately. The conventional approach to forecasting takes latest-available data at face value, which in practice means estimating equations using data that may have undergone many rounds of revision, and then substituting very recent releases into these equations to produce forecasts. Thus, it treats recently released and heavily revised data as interchangeable. Most *ex post* real-time analyses mimic this procedure, estimating equations and producing forecasts using the latest data that would have been available in real time. Because it mixes heavily revised with first-release and lightly revised data, this approach is unlikely to produce good forecasts (Koenig, Dolmas and Piger, 2003).

As an alternative to conventional real-time estimation and forecasting, we adopt the augmented VAR approach developed in Kishor and Koenig (2012). Kishor and Koenig (2012) assume that a VAR describes the evolution of "final-release" data. In practice, these data need not be truly final. It is only required that subsequent revisions be unforecastable. The VAR is augmented with a model of early data revisions that is flexible in its assumptions about how data releases evolve. The VAR and the revisions model are estimated together and the resultant equations are put into state-space form, which allows application of the Kalman filter. The filter projects what the most recent data will look like after revision, and it is this projection that is substituted into the VAR to produce forecasts. A more detailed description of the methodology follows.

Let x_t denote a vector of final (or, more generally, efficiently estimated) data, which become available after e revisions. It's assumed that the evolution of these data is governed

²⁰See Croushore and Stark (2001) for a general discussion of the importance of data revisions in macroeconomics. Croushore (2008) looks specifically at revisions to headline PCE inflation.

by a VAR of order L :

$$x_t = F_1 x_{t-1} + F_2 x_{t-2} + \dots + F_L x_{t-L} + v_0(t). \quad (11)$$

where $v_0(t)$ is white-noise error, so that $E(v_0(t)) = 0$ and $E(v_0(t)v_s(t)') = 0 \forall s \neq t$.

Rearranging terms:

$$z_t = F z_{t-1} + v_t \quad (12)$$

where $z_t = \begin{bmatrix} x_{t-e} \\ x_{t-e+1} \\ \dots \\ x_t \end{bmatrix}$, $v_t = \begin{bmatrix} 0 \\ 0 \\ \dots \\ v_0(t) \end{bmatrix}$,

and

$$F = \begin{bmatrix} 0 & I & 0 & \dots & 0 \\ 0 & 0 & I & \dots & 0 \\ \cdot & \cdot & \cdot & \dots & \cdot \\ \cdot & \cdot & \cdot & \dots & \cdot \\ 0 & 0 & F_L & \dots & F_1 \end{bmatrix}.$$

This VAR is augmented with equations that describe the data-revision process:

$$y_t = (I - G)F y_{t-1} + G z_t + \varepsilon_t \quad (13)$$

where $y_t = \begin{bmatrix} x_{t-e}^t \\ x_{t-e+1}^t \\ \dots \\ x_t^t \end{bmatrix}$, $\varepsilon_t = \begin{bmatrix} 0 \\ \varepsilon(e-1)_t \\ \dots \\ \varepsilon(0)_t \end{bmatrix}$

and

$$G = \begin{bmatrix} I & 0 & 0 & \dots & 0 \\ G_{e-1,e} & G_{e-1,e-1} & G_{e-1,e-2} & \dots & G_{e-1,0} \\ \cdot & \cdot & \cdot & \dots & \cdot \\ G_{1,e} & G_{1,e-1} & \cdot & \dots & G_{1,0} \\ G_{0,e} & G_{0,e-1} & \cdot & \dots & G_{0,0} \end{bmatrix}.$$

Here, x_{t-s}^t denotes the official estimate of x_{t-s} available at time t . Note that $x_{t-e}^t = x_{t-e}$, *i.e.* the official estimate of x_{t-e} available at time t is assumed to be efficient. It's assumed, additionally, that the transition-equation errors are uncorrelated with the observation-equation

errors at all leads and lags, and are serially uncorrelated. This specification can be shown to encompass several standard models of data revision (Kishor and Koenig, 2012).²¹

Because the Kishor-Koenig (KK) approach does not mix apples (heavily revised data) with oranges (first-release and lightly revised data), it avoids two problems that afflict conventional VAR estimation and forecasting. First, under the conventional approach, VAR coefficients are inefficiently estimated and possibly biased. Second, it is typically inappropriate to take end-of-sample data at face value and substitute them into a VAR that has been estimated with revised data. When the most recent data are out of line with what one would have expected given previously available information, the discrepancy can be exploited to predict how these data will later be revised. By predicting revisions to end-of-sample data, forecasts of the future path of the economy can be improved.

To summarize, our forecasting approach involves the following two steps:

Step 1: Estimate the state-space model outlined in equations 1 through 5 to generate the UC-filtered trend and cyclical components of PCE, GDP and core-CPI inflation. The initial sample for the recursive estimation is 1984:Q1-1992:Q1. The final sample is 1984:Q1-2013:Q4. The sequence of estimations yields real-time first-release and first-revision estimates of inflation's trend and cyclical components running from 1992:Q1 through 2013:Q4. It is those estimates that are used in the KK forecasting approach. The state-space estimations also yield a sequence of end-of-sample-vintage ("smoothed") estimates of inflation's trend and cyclical components such as would have been used by an analyst taking the conventional approach to forecasting.²²

Step 2: Using the filtered first-release and first-revision estimates of trend and cyclical inflation from step 1, we apply the KK forecasting approach outlined in equations 10 through 12 to generate 0-to-4-step-ahead forecasts. The first set of forecasts covers the pe-

²¹Appendix A provides the estimation details of this model for the special case of one revision and one lag in the VAR system.

²²For example, the UC model for the sample period 1984:Q1-2001:Q4 provides us a smoothed estimate of trend and cycle for 1984:Q1 through 2001:Q4. We perform this exercise for each iteration and the final smoothed series for the conventional real-time estimation runs from 1984:Q1-2013:Q4. To have the same sample size for all of the recursive forecasting exercises, the sample period for conventional approach starts in 1992:Q1, just as it does for the KK approach.

riod 2002:Q1-2003:Q1 and is based on filtered inflation estimates spanning 1992:Q1-2001:Q4. The final set of forecasts covers the period 2014:Q1-2015:Q1 and is based on filtered inflation estimates from 1992:Q1 through 2013:Q4. For conventional estimation, we estimate a VAR for each end-of-sample vintage of smoothed inflation estimates starting with 2001:Q4.

5 Estimation Results

5.1 In-Sample Estimation Results

For illustrative purposes and also for motivating our real-time forecasting exercise, we first undertake in-sample regressions of various inflation measures on lagged labor-market slack. Our baseline slack measure is the "unemployment-recession gap." As defined in Stock and Watson (2010), this gap is the difference between the current unemployment rate and the minimum unemployment rate over the current and previous eleven quarters. Unlike the gap between output and potential output or the gap between the unemployment rate and the natural rate of unemployment, the unemployment-recession gap does not depend on an unobservable and is not subject to revision. Yet the unemployment-recession gap captures the sharp upward spikes in unemployment that are associated with recessions.

Results for headline PCE inflation are consistent with the existing literature (Table 1): Slack appears to be of no use in explaining inflation movements. Even lagged inflation is barely significant, and the equation explains only 6 percent of the variation in headline PCE inflation. However, slack *is* significant in explaining one-quarter-ahead movements in conventional (ex-food-and-energy) core and UC-filtered PCE inflation, and in explaining movements in headline GDP inflation.²³ Those results suggest that high-frequency, unforecastable noise obscures the relationship between slack and future inflation—especially PCE inflation.

²³Results are similar using alternative measures of slack, such as the change in unemployment and the deviation of unemployment from its 4-quarter moving average.

5.2 Forecasting Inflation in Real-Time

In this section, we apply the KK real-time forecasting methodology to first-release and first-revision UC estimates of PCE inflation’s trend and cycle. We forecast trend and cycle separately because they have distinct dynamics (*c.f.* Equations 8 and 9).²⁴ For comparison, we also do conventional real-time VAR estimation and forecasting using the most up-to-date UC estimates of trend and cycle that would have been available in real time. Each round of forecasts extends from the current quarter ($h = 0$) to 4 quarters hence ($h = 4$). The first round of forecasts covers the period 2002:Q1-2003:Q1, and the final round covers the period 2014:Q1-2015:Q1. The starting point for the exercise is late enough that an analyst likely would have taken into account the 1998:Q1 shift in the dynamics of trend inflation (*c.f.* the discussion following Equation 8). That is our assumption as we forecast trend inflation.

As an alternative to pre-filtering PCE inflation using our UC model, we estimate several real-time forecasting models in conventional-core PCE inflation. Finally, to facilitate comparison with the existing literature, we estimate real-time forecasting models in headline PCE inflation.

We start with univariate forecasting models, and then consider models that include the unemployment-recession gap. In our model of UC-filtered inflation, we include the unemployment-recession gap only in the forecasting equation for the cyclical component of inflation as we do not expect the trend component to be related to slack.

We compare forecasting performance across models and also with spliced SPF/Greenbook inflation forecasts. We evaluate each model—including those models estimated using UC-filtered or conventional-core inflation—on its ability to predict the latest-vintage official *headline* PCE inflation releases that have yet to undergo a comprehensive revision.²⁵

²⁴We take this approach rather than forecast the gap between inflation and trend inflation and then add the latest estimate of trend inflation back in to generate an inflation forecast (Stock and Watson, 2010 and Faust and Wright, 2013, among others). The gap approach does well when trend inflation is well approximated by a random walk, but that is not the case over our forecast period. The results are shown in Appendix B.

²⁵Comprehensive revisions to the national income and product accounts are idiosyncratic and wide reaching. Realistically, their effects cannot be predicted.

5.3 Univariate Forecasts

Table 2 reports the forecasting results for headline PCE inflation obtained using univariate models. We report mean-squared forecast errors (MSEs) for specific horizons ($h = 0, 1, \dots, 4$) as well as over two multi-quarter horizons ($h = 0 - 4$ and $h = 1 - 4$) that are relevant for policy. The first column of results (labeled "AR1") is the performance benchmark in the literature on inflation forecasting: It shows the real-time forecast performance of a first-order autoregressive (AR1) model in headline inflation.²⁶ The second column of results (labeled "SPF") is another possible performance benchmark: It shows MSEs from the SPF/Greenbook. Asterisks (*) mark when an alternative forecast performs significantly better than the corresponding headline AR1 model according to the Diebold and Mariano (1995) and West (1996) non-nested forecast-comparison test (DMW test). Pound signs (#) mark when an alternative forecast performs significantly better than the corresponding Greenbook/SPF forecast according to the DMW test.

Three simple alternatives to the benchmark AR1 model are examined: (1) a conventionally estimated AR model in ex-food-and-energy core PCE inflation ("Core"), (2) a conventionally estimated AR model in UC-filtered inflation ("Conv"), and (3) an AR model in UC-filtered inflation, estimated in the manner recommended by Kishor and Koenig (2012) ("KK").²⁷ In all cases it is against realized headline PCE inflation that forecasts are compared when calculating MSEs.

The main message from Table 2—a message consistent with much of the recent literature on inflation forecasting—is that it is difficult to substantially improve on the forecasts generated by a simple autoregressive model in headline inflation. Professional forecasters outperform the simple AR1 model only in their current-quarter predictions, where their

²⁶As an alternative benchmark, we considered an IMA(1,1) model similar to the UC-stochastic volatility model favored by Stock and Watson (2007). Over our forecast period, however, the simple AR(1) benchmark model consistently outperforms the IMA(1,1) model.

²⁷In the KK model, first-revision estimates of filtered inflation are treated as final for estimation purposes. Forecasts of first-revision filtered inflation are used as forecasts of headline inflation for forecast-evaluation purposes.

access to high-frequency data gives them a natural advantage (*cf.* footnote #19). Using ex-food-and-energy core inflation to estimate the forecasting equation fails to produce significant improvement at any horizon. Applying conventional real-time estimation to UC-filtered trend and cyclical inflation produces statistically significant improvement at only one forecast horizon ($h = 1$ in the column labeled "Conv"). Applying KK estimation to UC-filtered trend and cyclical inflation shows greater promise, producing statistically significant forecast improvement at $h = 0$, $h = 1$, and $h = 0 - 4$. The combination of UC filtering and KK estimation and forecasting also fares well against the "Core" and "Conv" models: It produces lower mean-square errors at both the $h = 0 - 4$ and $h = 1 - 4$ multi-quarter forecast horizons and at most of the single-quarter horizons. The multi-quarter differences are statistically significant at the 10-percent level according to the DMW test.²⁸

5.4 Is Slack Helpful for Forecasting Inflation?

The answer is, "It depends." Specifically, it depends on whether and how inflation is filtered before the forecasting equation is estimated. Adding Stock and Watson's unemployment-recession gap to a forecasting equation estimated in headline inflation or estimated in ex-food-and-energy core inflation is counterproductive. In contrast, adding labor-market slack to forecasting equations estimated in UC-filtered core inflation often yields a statistically significant improvement in real-time performance.

Results are in Table 3, which displays four pairs of MSEs at each forecast horizon. Within each pair, the first entry shows the MSE obtained from a model without slack. These duplicate results reported in Table 2. The second entry within each pair is new. It shows the MSE obtained from an otherwise-identical model with the unemployment-recession gap included. The presence of an inequality sign ($>$) signals that the difference between the first and second entries is statistically significant according to the Clark-West (2007) nested forecast-comparison test (CW test). In addition, asterisks (*) signal a statistically significant

²⁸The KK model's advantage over the Core model at the $h = 0$ horizon is also statistically significant at the 10-percent level.

advantage over the baseline AR1-model forecast, according to the DMW test.

Starting with the first two columns in Table 3, adding slack to a forecasting model in headline PCE inflation causes real-time forecast performance to *deteriorate* at every forecast horizon—although in no case is the deterioration statistically significant by the CW test. Similarly, according to the results in columns 3 and 4, adding slack to a forecasting model estimated in ex-food-and-energy core PCE inflation causes forecast performance to deteriorate at all but the very shortest horizon ($h = 0$). The story is very different when the unemployment-recession gap is included in a forecasting model estimated with UC-filtered inflation. When estimation and forecasting are done conventionally, as in columns 5 and 6, forecast performance improves at every horizon but one ($h = 1$). Moreover, the improvement is statistically significant in nearly every case ($h = 0$, $h = 3$, $h = 4$, $h = 0 - 4$, and $h = 1 - 4$), and is large enough that the bivariate model in UC-filtered inflation significantly outperforms the benchmark AR1 headline inflation model at $h = 0$, $h = 4$, and $h = 0 - 4$. It also significantly outperforms the SPF/Greenbook at every horizon except $h = 0$ and $h = 1$. (Results available on request.) Results are even stronger when the Kishor-Koenig methodology is applied to UC-filtered inflation, as in columns 7 and 8. Forecast performance now improves at *every* horizon as a result of the inclusion of slack. The improvement is almost always statistically significant, and it is large enough that the KK model with slack produces forecasts that significantly outperform those of the benchmark AR1 model at every horizon except $h = 2$ and $h = 3$. Additionally, the forecasts of the KK model with slack ("KK+S") significantly outperform those of the conventionally estimated model with slack ("Conv+S") at $h = 1$, $h = 3$, $h = 0 - 4$, and $h = 1 - 4$; and they significantly outperform SPF/Greenbook forecasts at every horizon except $h = 0$ and $h = 1$.

Summarizing: UC filtering of PCE inflation is, by itself, only modestly useful for improving forecast performance. Adding slack to models of headline PCE inflation or a model of ex-food-and-energy PCE inflation is counterproductive. Yet, UC-filtering combined with slack produces forecasts of headline PCE inflation that are clearly superior to those obtained by forecasting headline inflation directly, and which are superior to the forecasts of sophisti-

cated professionals at the horizons which are of greatest interest to policymakers²⁹. Results are strongest when revisions to inflation data are properly accounted for during estimation and forecasting.

6 Robustness Exercises

In this section, we confirm that the results displayed in Table 3 are not sensitive to changes in how we measure inflation's longer-run trend, to changes in the inflation measure being forecasted, or to changes in the measure of cyclical variation in real activity.

6.1 An Alternative Measure of Trend Inflation

First, we examine the robustness of our out-of-sample PCE forecasting results to a simple alternative measure of long-run trend inflation that makes no assumption about the anchoring of longer-run inflation expectations. The alternative trend is obtained by applying a one-sided Hodrick-Prescott (HP) filter to SPF long-forward inflation expectations. It uses only information that would have been available in real time and it is not subject to revision. Inflation's cyclical component is estimated by applying the Kalman filter to Equation 9 and Equations 1 - 3 after subtracting the HP trend from each of the three measures of inflation. As before, cyclical inflation is forecasted separately from trend inflation, and our period- t forecast of headline inflation h periods hence is the sum of our forecasts of cycle and trend. The forecast of cycle is obtained using the Kishor-Koenig (2012) method, and we simply use the actual estimate of the trend at time t as a forecast of trend for $t + h$.

Results are reported in Table 4, in the column labeled "KK(HP)" for the univariate version of the model, and the column labeled "KK(HP)+S" for the model that includes the unemployment-recession gap as a measure of slack. For comparison purposes, Table 6 also displays mean-square forecast errors from our benchmark AR1 model, from the

²⁹To verify that the performance improvement from the inclusion of slack is not due to a handful of data points, we examined the behavior of the estimated coefficient on lagged slack in equation (10) over time. Results are reported in Appendix C.

SPF/Greenbook, and from the KK and KK+S models described in Section 5 (and reported on in Table 3), which rely on an estimate of trend inflation obtained within our UC model.

MSEs for the KK(HP) model are more often than not lower than those for the benchmark AR1 model, but the difference is statistically significant only at $h = 0$ and $h = 0 - 4$. Greater improvement is obtained when UC filtering is combined with slack, as in the KK(HP)+S model.³⁰ Now, MSE differences relative to the AR1 benchmark are statistically significant at $h = 0$, $h = 0 - 4$, and $h = 1 - 4$. Moreover, forecasts from the KK(HP)+S model significantly outperform those from the SPF/Greenbook at $h = 2$, $h = 3$, $h = 0 - 4$, and $h = 1 - 4$. These results are completely consistent with our earlier conclusions that (1) UC filtering of PCE inflation is, by itself, modestly useful for improving forecast performance; while (2) UC-filtering combined with slack produces forecasts of headline PCE inflation that are clearly superior to those obtained by forecasting headline inflation directly, and which are superior to the forecasts of sophisticated professionals at the horizons which are of greatest interest to policymakers. However, comparing forecast errors from the HP-trend model with slack (KK(HP)+S) to errors from the corresponding model with UC-filtered trend (KK+S) shows that the latter model has a lower MSE at all but one forecast horizon, and that the differences are statistically significant at $h = 1$, $h = 2$, $h = 0 - 4$, and $h = 1 - 4$. So, although there is a payoff from filtering headline inflation even if we use an off-the-shelf measure of inflation's long-run trend, the payoff is larger if the trend is estimated within the UC model.

6.2 GDP Inflation as an Alternative to PCE Inflation

Forecast comparison results for GDP inflation (Table 5) are similar to those for PCE inflation (Table 3). In particular, including slack in a model of headline inflation *worsens* forecasting performance at all horizons (*c.f.* the columns labeled "AR1" and "GDP+S" in Table 5), although the deterioration is never statistically significant. In contrast, including slack in models of UC-filtered GDP inflation significantly *improves* forecasting performance at every

³⁰As documented in Table 4, the MSEs produced by the KK(HP)+S model are significantly lower than those produced by the KK(HP) model at all but the shortest forecast horizons.

horizon (KK method) or nearly every horizon (conventional method). For the UC-filtered inflation model estimated in the manner recommended by Kishor and Koenig (2012), the forecast improvement is large enough that the KK model with slack outperforms each alternative model at most forecast horizons. The advantage of the KK model with slack over the benchmark AR1 model is statistically significant at every horizon except $h = 3$ and $h = 4$. The advantage of the KK model with slack over the conventional model with slack is statistically significant at $h = 0$ and $h = 1$. (Results available on request.)

6.3 The Change in the Unemployment Rate as an Alternative to the Unemployment-Recession Gap

Motivated by estimation results reported in Koenig and Atkinson (2012), which suggest that the change in labor-market slack may be more important than the level of slack in PCE inflation forecasting, we also examine the effects of replacing the unemployment-recession gap with the quarterly change in unemployment rate. The new results are reported in Table 6, which is identical in format to Table 3. Our main conclusions carry through. Thus, including the change in the unemployment rate in a model of headline PCE inflation causes forecasting performance to deteriorate at every horizon, albeit not significantly. (Compare the columns headed "AR1" and "PCE+S" in Table 6.) Similarly, including the change in the unemployment rate in a model of ex-food-and-energy core inflation produces significant improvement in forecasting performance only at the very shortest horizon. (Compare the columns headed "Core" and "Core+S" in Tables 6.) The payoff to including the unemployment change is noticeably greater when inflation is filtered using our unobserved components model—especially if the filtered inflation data are handled as recommended by Kishor-Koenig (2012). The mean-square forecast error falls at nearly every forecast horizon—by enough that the bivariate KK model significantly outperforms the AR1 model at every forecast horizon except $h = 2$ and $h = 3$. DMW tests (not reported in the table, but available on request) show that the bivariate KK model significantly outperforms the same model estimated conventionally ("Conv+S") at every forecast horizon except $h = 4$, and significantly outperforms

SPF/Greenbook forecasts at the $h = 2$, $h = 3$, $h = 0 - 4$, and $h = 1 - 4$ horizons.

To summarize, the overall message of this paper, that substantial improvement in real-time performance can be achieved if one uses UC-filtered inflation together with real activity to forecast headline inflation, is robust to alternative measures of real activity.

7 Conclusion

Is it possible to utilize the information from multiple inflation series and multiple surveys of long-run inflation expectations to improve inflation forecasts in real-time? Does slack matter in forecasting inflation? We find that the answers to these questions are intertwined. It is the *combination* of filtered inflation and slack that produces significantly improved inflation forecasts.

We propose an unobserved-components model that draws on information from SPF and Blue Chip long-term inflation expectations and from multiple quarterly inflation series. We use this model to strip out the noise from quarterly real-time headline PCE and GDP inflation and propose filtered measures of inflation that capture both medium-term cyclical inflation movements and changes in long-term inflation expectations. As an alternative to conventional real-time estimation and forecasting, which inappropriately mixes heavily revised and lightly revised data, we use the methodology proposed by Kishor and Koenig (2012). The KK method augments a standard VAR in revised data with a state-space model of data revisions. It takes data revisions into account without imposing restrictive assumptions on the revisions process.

We find that there is a payoff to carefully filtering high-frequency, volatile movements out of headline PCE inflation prior to estimating forecasting equations, even when what one is ultimately interested in is forecasting headline inflation. That payoff is most evident when a measure of slack is also included in the model and estimation and forecasting are undertaken using the KK method. Introducing slack into a model of UC-filtered inflation produces forecasts of headline inflation that dominate those obtained by modeling headline inflation directly, with or without slack, and which “meet or beat” the inflation forecasting

performance of the Greenbook and SPF at all but the very shortest horizons. Ex-food-and-energy "core" inflation is *not* a good substitute for UC-filtered inflation. Forecasting models estimated using conventional core inflation do not perform significantly better than an AR1 model of headline inflation.

We conclude that three ingredients are required for superior inflation forecasting performance: (1) careful filtering to remove unforecastable high-frequency variation, (2) proper allowance for data revisions, and (3) provision for the influence of economic slack.

References

- [1] Ang, Andrew, Geert Bekaert and Min Wei (2007), “Do Macro Variables, Asset Markets or Surveys Forecast Inflation Better?,” *Journal of Monetary Economics*, 54, pp. 1163-1212.
- [2] Atkeson, Atkeson and Lee E. Ohanian (2001), “Are Phillips Curves Useful for Forecasting Inflation?,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 25, pp. 2-11.
- [3] Bernanke, Ben (2003), “Constrained Discretion,’ and Monetary Policy,” Remarks Before the Money Marketeters of New York University, New York, New York, February 3, 2003.
- [4] Bai, J., & Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, 47-78.
- [5] Basistha, A. and C.R. Nelson (2007), “New measures of the output gap based on the forward-looking new Keynesian Phillips curve,” *Journal of Monetary Economics*, 54, pp. 498-511
- [6] Basistha, A. and R. Startz (2008), “Measuring the NAIRU with Reduced Uncertainty: A Multiple-Indicator Common-Cycle Approach,” *The Review of Economics and Statistics*, 90, pp. 805-811.
- [7] Bryan, M. F., and Cecchetti, S. G. (1993), “Measuring core inflation,” *NBER Working Paper Series*, No. 4303.
- [8] Chan, Joshua C.C., Todd.E. Clark and Gary Koop (2015), “A New Model of Inflation, Trend inflation, and Long-run Inflation Expectations,” Federal Reserve Bank of Cleveland Working Paper.
- [9] Clark, Peter K. (1989), “Trend Reversion in Real Output and Unemployment,” *Journal of Econometrics*, 1, pp. 15-32.

- [10] Clark, Todd E., and Taeyoung Doh (2014), "Evaluating Alternative Models of Trend Inflation." *International Journal of Forecasting* 30(3), pp. 426-448.
- [11] Clark, Todd E. and Kenneth D. West (2007), "Approximately Normal Tests for Equal Predictive Accuracy in Nested Models," *Journal of Econometrics*, 138, pp. 291-311.
- [12] Cogley, Timothy, Giorgio Primiceri, and Thomas J. Sargent (2010), "Inflation-Gap Persistence in the U.S.," *American Economic Journal: Macroeconomics*, 2, 43-69.
- [13] Crone, Theodore, N. Neil K. Khettry, Loretta J. Mester and Jason A. Novak (2013), "Core Measures of Inflation as Predictors of Total Inflation," *Journal of Money, Credit and Banking*, 45, pp. 505-519.
- [14] Croushore, Dean and Tom Stark (2001), "A Real-Time Data Set for Macroeconomists," *Journal of Econometrics*, 105, pp. 111-130.
- [15] Croushore, Dean (2008), "Revisions to PCE Inflation Measures: Implications for Monetary Policy," *Federal Reserve Bank of Philadelphia Working Paper*.
- [16] Dolmas, Jim (2005), "Trimmed mean PCE inflation" Federal Reserve Bank of Dallas Working Paper 506.
- [17] Diebold, F.X., and R.S. Mariano (1995), "Comparing Predictive Accuracy," *Journal of Business and Economic Statistics*, 13, pp. 253-263.
- [18] Faust, Jon and Jonathan Wright (2009), "Comparing Greenbook and Reduced Form Forecasts Using a Large Real-Time Dataset," *Journal of Business and Economic Statistics*, pp. 468-479.
- [19] Faust, Jon and Jonathan Wright (2013), "Forecasting Inflation," *Handbook of Economic Forecasting*, edited by Graham Elliott and Alan Timmermann, Vol 2, Elsevier.

- [20] Gurkaynak, Refet S., Brian P. Sack and Jonathan H. Wright (2010), “The TIPS Yield Curve and Inflation Compensation,” *American Economic Journal: Macroeconomics*, 2, pp. 70-92.
- [21] Kiley, M. T. (2008), “Estimating the common trend rate of inflation for consumer prices and consumer prices excluding food and energy prices,” Divisions of Research & Statistics and Monetary Affairs, Federal Reserve Board.
- [22] Kim C-J., and C. R. Nelson (1998), “Business Cycle Turning Points, A New Coincident Index, and Tests of Duration Dependence Based on a Dynamic Factor Model With Regime Switching,” *Review of Economics and Statistics*, 80(2), 188-201.
- [23] Kim C-J., and C. R. Nelson (2000), *State-Space Models with Regime Switching*, MIT Press, Cambridge.
- [24] Kishor, N. Kundan and Koenig, Evan F. (2012), “VAR Estimation and Forecasting When Data Are Subject to Revision,” *Journal of Business and Economic Statistics*, 30, pp. 181-190.
- [25] Kishor, N. Kundan and Koenig, Evan F. (2014), “Credit Indicators as Predictors of Economic Activity: A Real-Time VAR Analysis,” *Journal of Money, Credit and Banking*, 46, pp. 545-564.
- [26] Koenig, Evan F., Sheila Dolmas and Jeremy Piger (2003), “The Use and Abuse of Real-Time Data in Economic Forecasting,” *Review of Economics and Statistics*, 85, pp. 618-628.
- [27] Koenig, Evan F. and Tyler Atkinson (2012), “Inflation, Slack and Fed Credibility,” Dallas Fed *Staff Paper* 16.
- [28] Kozicki, Sharon and Peter A. Tinsley (2001), “Shifting Endpoints in the Term Structure of Interest Rates,” *Journal of Monetary Economics*, 47, pp. 613-652.

- [29] Kozicki, Sharon and Peter A. Tinsley (2005), "What do you expect? Imperfect policy credibility and tests of the expectations hypothesis," *Journal of Monetary Economics*, 52, pp. 421-447.
- [30] Mertens, E. (2016), "Measuring the Level and Uncertainty of Trend Inflation," *Review of Economics and Statistics*, 98(5), 950-967.
- [31] Nalewaik, Jeremy (2015), "Regime Switching Models for Estimating Inflation Uncertainty," Working Paper, Federal Reserve Board.
- [32] Phillips, A. William (1958), "The Relationship between Unemployment and the Rate of Change of Money Wages in the United Kingdom 1861-1957," *Economica*, 25, pp. 283-299.
- [33] Sims, Christopher A. (2002), "The Role of Models and Probabilities in the Monetary Policy Process," *Brookings Papers on Economic Activity*, 2, pp. 1-40.
- [34] Smith, Julie K. (2004), "Weighted Median Inflation: Is this Core Inflation?" *Journal of Money, Credit and Banking*, 36(2): 253-263.
- [35] Stock, James H. and Mark W. Watson (1999), "Forecasting Inflation," *Journal of Monetary Economics*, 41, pp. 788-829.
- [36] Stock, James H. and Mark W. Watson (2007), "Why Has Inflation Become Harder to Forecast?," *Journal of Money, Credit and Banking*, 39, pp. 3-34.
- [37] Stock, James H. and Mark W. Watson (2009), "Phillips Curve Inflation Forecasts," in *Understanding Inflation and the Implications for Monetary Policy*, Jeffrey Fuhrer, Yolanda Kodrzycki, Jane Little, and Giovanni Olivei (eds), MIT Press, Cambridge.
- [38] Stock, James H. and Mark W. Watson (2010), "Modeling Inflation After the Crisis," in *Macroeconomic Challenges: The Decade Ahead*, Federal Reserve Bank of Kansas City Economic Policy Symposium.

- [39] Stock, James H. and Mark W. Watson (2016), “Core inflation and trend inflation,” *Review of Economics and Statistics*, 98(4), 770-784.
- [40] West, K.D. (1996), “Asymptotic Inference about Predictive Ability,” *Econometrica* 64, pp.1067-1084.

Table 1. In-sample Forecasting Regressions

Inflation Measure	Lagged Inflation	Slack	R ²
Headline PCE	0.219 (0.03)	-0.137 (0.48)	0.061
Conventional-core PCE	0.625 (0.00)	-0.200 (0.05)	0.452
UC-filtered PCE	0.642 (0.00)	-0.090 (0.00)	0.541
GDP Deflator	0.423 (0.00)	-0.466 (0.00)	0.352

Notes:

The sample period is 1992:Q1-2013:Q4. Newey-West P-values are in parentheses.

Slack is the unemployment-recession gap, defined as the current unemployment rate less the minimum rate over the current and previous 11 quarters.

UC-filtered inflation is first-revision filtered inflation from our UC model.

Table 2. Headline PCE Inflation MSEs: Models without Slack

Horizon (h)	AR1	SPF	Core	Conv	KK
0	3.108	0.872**	2.917	2.890	2.859*
1	2.986	2.967	2.855	2.807*	2.793*
2	2.837##	3.339	2.813##	2.791##	2.780##
3	2.842##	3.343	2.832##	2.967	2.887#
4	2.942	3.199	2.905	2.805	2.791#
0 - 4	0.702	0.861	0.651#	0.640#	0.602***
1 - 4	0.871#	1.237	0.838##	0.830##	0.796##

Notes:

The first set of forecasts is for 2002:Q1-2003:Q1; the final set is for 2014:Q1-2015:Q1.

* Outperforms the benchmark AR1 model at the 10-percent level.

** Outperforms the benchmark AR1 model at the 5-percent level.

Outperforms SPF/Greenbook forecasts at the 10-percent level.

Outperforms SPF/Greenbook forecasts at the 5-percent level.

The smallest MSE in each row is bolded.

Table 3. Headline PCE Inflation MSEs: Models with and without Slack

Horizon (h)	AR1	PCE+S	Core	Core+S	Conv	Conv+S	KK	KK+S
0	3.108	3.469	2.917 >	2.825	2.890 >>	2.783**	2.859* >	2.822**
1	2.986	4.089	2.855	2.877	2.807*	2.833	2.793* >	2.772*
2	2.837	4.004	2.813	2.981	2.791	2.784	2.780	2.772
3	2.842	3.536	2.832	2.955	2.967 >	2.903	2.887 >>	2.798
4	2.942	3.179	2.905	2.940	2.805 >	2.748*	2.791 >>	2.755*
0 - 4	0.702	1.395	0.651	0.719	0.640 >>	0.607*	0.602* >>	0.553**
1 - 4	0.871	1.642	0.838	0.940	0.830 >	0.817	0.796 >>	0.753*

Notes:

The first set of forecasts is for 2002:Q1-2003:Q1; the final set is for 2014:Q1-2015:Q1.

" +S" indicates that the unemployment-recession gap is included in the model.

*Outperforms the benchmark AR1 model at the 10-percent level.

** Outperforms the benchmark AR1 model at the 5-percent level.

> The MSE difference between adjacent entries is significant at the 10-percent level.

>> The MSE difference between adjacent entries is significant at the 5-percent level.

The smallest MSE in each row is bolded.

Table 4. Headline PCE Inflation MSEs: Alternate Measure of Trend

Horizon (h)	AR1	SPF	KK		KK(HP)		KK(HP)+S		KK+S
0	3.108	0.872**	2.859*		2.847*		2.857*		2.822**
1	2.986	2.967	2.793*		2.837		2.844	>	2.772*
2	2.837##	3.339	2.780##	<	2.866#	>	2.849#	>	2.772##
3	2.842##	3.343	2.887#		2.852#	>>	2.786#		2.798##
4	2.942	3.199	2.791		2.849	>>	2.785		2.755*#
0 - 4	0.702	0.861	0.602*##	<	0.638*##	>>	0.598*##	>	0.553*##
1 - 4	0.871#	1.237	0.796##		0.837##	>>	0.795*##	>	0.753*##

Notes:

The first set of forecasts is for 2002:Q1-2003:Q1; the final set is for 2014:Q1-2015:Q1.

"+S" indicates that the unemployment-recession gap is included in the model.

"(HP)" indicates that trend inflation is estimated using the one-sided HP filter.

*Outperforms the benchmark AR1 model at the 10-percent level.

** Outperforms the benchmark AR1 model at the 5-percent level.

Outperforms SPF/Greenbook forecasts at the 10-percent level.

Outperforms SPF/Greenbook forecasts at the 5-percent level.

> The MSE difference between adjacent entries is significant at the 10-percent level.

>> The MSE difference between adjacent entries is significant at the 5-percent level.

The smallest MSE in each row is bolded.

Table 5. GDP Inflation MSEs: Models with and without Slack

Horizon (h)	AR1	GDP+S	Conv	Conv+S	KK	KK+S
0	1.053	1.091	0.842** >>	0.787**	0.783** >>	0.754**
1	0.966	1.014	0.817* >	0.809*	0.794* >>	0.744**
2	0.956	1.129	0.891 >>	0.831*	0.889 >	0.850*
3	0.985	1.139	1.013	1.024	1.038 >	1.009
4	1.076	1.322	1.130 >	1.106	1.111 >	1.085
0 - 4	0.411	0.504	0.390 >>	0.334	0.363 >>	0.326*
1 - 4	0.457	0.608	0.431 >	0.399	0.428 >	0.389*

Notes:

The first set of forecasts is for 2002:Q1-2003:Q1; the final set is for 2014:Q1-2015:Q1.

"+S" indicates that the unemployment-recession gap is included in the model.

*Outperforms the benchmark AR1 model at the 10-percent level.

** Outperforms the benchmark AR1 model at the 5-percent level.

> The MSE difference between adjacent entries is significant at the 10-percent level.

>> The MSE difference between adjacent entries is significant at the 5-percent level.

The smallest MSE in each row is bolded.

Table 6. Headline PCE Inflation MSEs: An Alternative Measure of Slack

Horizon (<i>h</i>)	AR1	PCE+S	Core	Core+S	Conv	Conv+S	KK	KK+S
0	3.108	3.148	2.917 >>	2.787*	2.890 >>	2.863*	2.859* >>	2.788*
1	2.986	3.434	2.855	2.872	2.807* <	2.831	2.793* >>	2.716*
2	2.837	3.258	2.813	2.872	2.791	2.841	2.780 >	2.772
3	2.842	3.037	2.832	2.869	2.967 >	2.942	2.887 >	2.858
4	2.942	2.905	2.905	2.873	2.805 >	2.762	2.791	2.813*
0 - 4	0.702	0.941	0.651	0.650	0.640 >	0.628*	0.602* >	0.560**
1 - 4	0.871	1.121	0.838	0.857	0.830	0.830	0.796 >	0.769*

Notes:

The first set of forecasts is for 2002:Q1-2003:Q1; the final set is for 2014:Q1-2015:Q1.

" +S" indicates that the change in the unemployment rate is included in the model.

*Outperforms the benchmark AR1 model at the 10-percent level.

** Outperforms the benchmark AR1 model at the 5-percent level.

> The MSE difference between adjacent entries is significant at the 10-percent level.

>> The MSE difference between adjacent entries is significant at the 5-percent level.

The smallest MSE in each row is bolded.

Figure 1: Measures of Inflationary Expectations

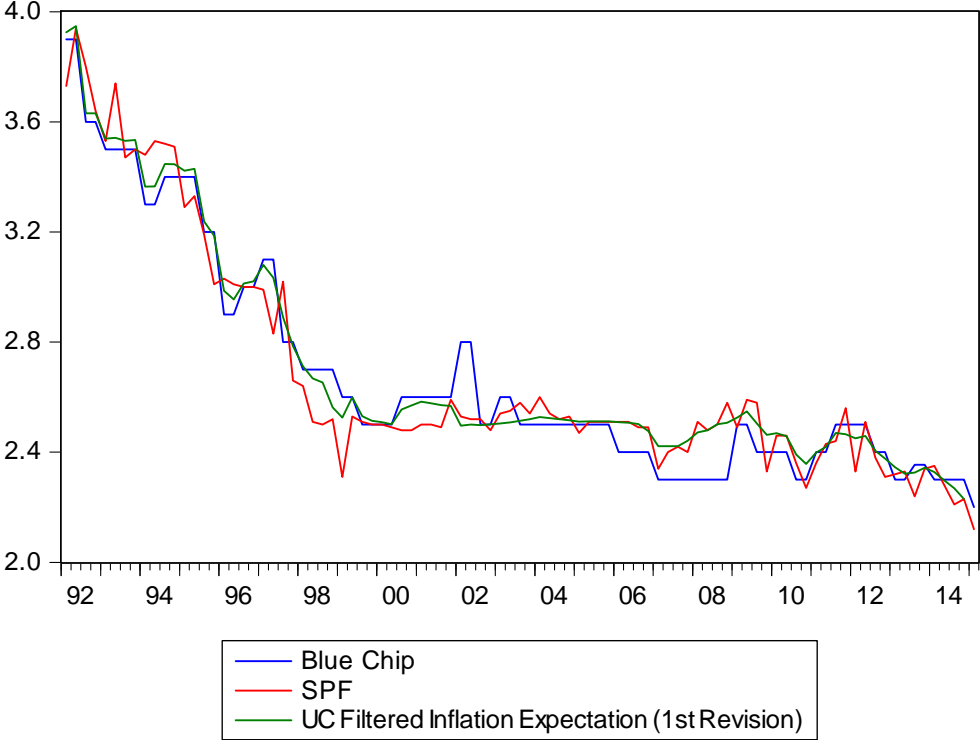


Figure 2: Headline PCE Inflation and UC Filtered Inflation


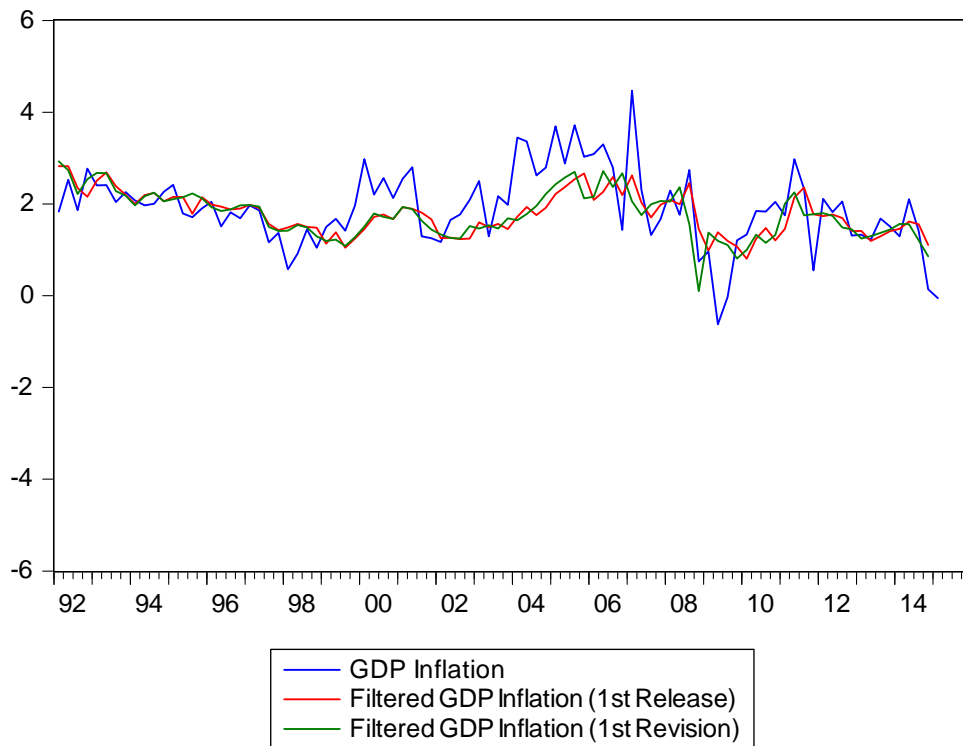


Figure 3: GDP Inflation and UC Filtered Inflation



Appendix A: Model Estimation and Forecasting

The measurement equation of the state-space system represented by equations (1) through (5) and equations (8) through (9) can be represented as:

$$\begin{bmatrix} \pi^{PCE} \\ \pi^{CORE} \\ \pi^{GDP} \\ \pi^{SPF} \\ \pi^{BC} \end{bmatrix} = \begin{bmatrix} \mu_1 \\ \mu_2 \\ \mu_3 \\ 0 \\ \mu_4 \end{bmatrix} + \begin{bmatrix} 1 & 1 & 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \tau_t \\ c_t \\ c_{t-1} \\ \lambda_t \\ \varepsilon_t \\ c_t^{SPF} \\ c_t^{BC} \end{bmatrix} + \begin{bmatrix} \eta_{1t} \\ \eta_{2t} \\ \eta_{3t} \\ \eta_{4t} \\ \eta_{5t} \end{bmatrix}$$

In matrix form, the measurement equation can be written as

$$y_t = \alpha + H\beta_t + e_t, e_t \sim iidN(0, R)$$

The transition equation can be written as:

$$\begin{bmatrix} \tau_t \\ c_t \\ c_{t-1} \\ \lambda_t \\ \varepsilon_t \\ c_t^{SPF} \\ c_t^{BC} \end{bmatrix} = \begin{bmatrix} \delta_s \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} \theta_s & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \phi_1 & \phi_2 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \phi^{SPF} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \phi^{BC} \end{bmatrix} \begin{bmatrix} \tau_{t-1} \\ c_{t-1} \\ c_{t-2} \\ \lambda_{t-1} \\ \varepsilon_{t-1} \\ c_{t-1}^{SPF} \\ c_{t-1}^{BC} \end{bmatrix} + \begin{bmatrix} v_{s,t} \\ u_t \\ 0 \\ \lambda_t \\ \varepsilon_t \\ u_t^{SPF} \\ u_t^{BC} \end{bmatrix}$$

The transition equation can be represented as:

$$\beta_t = \mu_s + F_s\beta_{t-1} + w_{st}, w_{st} \sim iidN(0, Q_s), s = 1, 2$$

$$\text{where } w_1 = \begin{bmatrix} v_{1,t} \\ u_t \\ 0 \\ \lambda_t \\ \varepsilon_t \\ u_t^{SPF} \\ u_t^{BC} \end{bmatrix}, \text{ and } F_1 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \phi_1 & \phi_2 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \phi^{SPF} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \phi^{BC} \end{bmatrix}$$

$$\text{and } w_2 = \begin{bmatrix} v_{2,t} \\ u_t \\ 0 \\ \lambda_t \\ \varepsilon_t \\ u_t^{SPF} \\ u_t^{BC} \end{bmatrix}, \text{ and } F_2 = \begin{bmatrix} \theta & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \phi_1 & \phi_2 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \phi^{SPF} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \phi^{BC} \end{bmatrix}$$

The above model can be put into state-space form and estimated using classical maximum likelihood via the Kalman filter. The unconditional mean and unconditional covariance matrix have been used as the initial values of the state vector and state variance for the stationary part of the transition equation. For the non-stationary component, we assign large uncertainty. The prediction and the updating equation in the Kalman filter recursion takes into account the break in 1997:Q4. We estimate the above model recursively, using real-time data, starting with a sample that runs from 1984:Q1 through 1992:Q1. The estimated inflation cycles and trend are revised with each extension of the sample period, both because the amount of data available for inference increases and because official PCE and GDP inflation data are sometimes revised.

Once we estimate the real-time filtered real-time-vintage estimates of π^{FPCE} and π^{FGDP} , we generate forecasts of headline PCE inflation and GDP inflation using different methods. In particular, we summarize the Kishor-Koenig (2012,2014) real-time forecasting methodology as outlined in section 4. In our estimation, we consider first revision of the filtered inflation and a slack measure that is not subject to revision.

Let z_{1t} and z_{2t} are the first-revised estimate of quarter-t inflation and the actual slack. We assume that

$$z_t = Fz_{t-1} + v_t$$

where $z'_t = [z_{1t} \ z_{2t}]$, $v'_t = [v_{1t} \ v_{2t}]$ is vector white noise. Let y_{1t} be the first-release of the filtered inflation in our exercise. As outlined in Kishor and Koenig (2012, 2014) we augment the above equation by supplementing the above equation with the equation

$$z_{1t} - y_{1t} = k(z_{1t-1} - y_{1t-1}) + u_t$$

where k is a parameter and where u_t has mean zero and is serially uncorrelated, but is allowed to be contemporaneously correlated with both v_{1t} and v_{2t} . Letting $z_{3t} = z_{1t} - y_{1t}$. We can convert the above model in a state-space form, with the following state and measurement equations

$$\begin{bmatrix} z_{1t} \\ z_{2t} \\ z_{3t} \end{bmatrix} = \begin{bmatrix} \theta_s & 0 & 0 \\ 0 & \phi_1 & \phi_2 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} z_{1t-1} \\ z_{2t-1} \\ z_{3t-1} \end{bmatrix} + \begin{bmatrix} v_{1,t} \\ v_{1,t} \\ u_t \end{bmatrix}$$

and

$$y_{1t} = \begin{bmatrix} 1 & 0 & -1 \end{bmatrix} \begin{bmatrix} z_{1t} \\ z_{2t} \\ z_{3t} \end{bmatrix}$$

The above equations are estimated using maximum likelihood via the Kalman filter. The state equation is then iterated forward to produce forecasts of future filtered inflation, future slack, and future inflation revisions.

Appendix B: Forecasting Results from Inflation-Gap and IMA(1,1) Models

Horizon (h)	AR1	KK	Inflation Gap	IMA(1,1)
0	3.108	2.859	3.170	3.541
1	2.986	2.793	3.225	3.491
2	2.837	2.780	2.967	3.192
3	2.842	2.887	2.920	3.247
4	2.942	2.791	2.967	3.262
0 - 4	0.702	0.602	0.818	1.106
1 - 4	0.871	0.796	0.980	1.286

Inflation-gap models forecast the gap between inflation and trend inflation and then add the most recent estimate of trend inflation back in to generate an inflation forecast [Stock and Watson (2010), Faust and Wright (2013)]. This approach performs well when trend inflation is well-approximated by a random walk. It does much less well over a forecast period like ours, when longer-run inflation expectations are well anchored. An IMA(1,1) model, similar to the UC-stochastic volatility model favored by Stock and Watson (2007), also does poorly over our forecast period.

Appendix C: Recursive Estimate of Slack Coefficient

The accompanying plot shows the recursively estimated coefficient on lagged slack in a VAR model of UC-filtered inflation and slack (Equation 10), where slack is measured by the unemployment-recession gap. Although the coefficient varies considerably over our forecasting period, it is consistently negative, suggesting a robustly downward-sloping Phillips Curve relationship.

Estimated Slack Coefficient

