Modeling Inflation After the Crisis

James H. Stock and Mark W. Watson

I. Introduction

The past five decades have seen tremendous changes in inflation dynamics in the United States. Some of the changes arguably stem from transformations in the U.S. economy. Energy is a smaller share of expenditures than it was during the oil price shocks of the 1970s, labor union membership has declined sharply over the past 40 years, and there has been a shift from production of goods to production of services. Monetary policy, too, has undergone dramatic transformations: The stance against inflation has become more aggressive, there have been discussions of formal or informal inflation targets, and there has been a recognition of the importance of expectations—and of expectations management—in determining the path of inflation.

These changes have created major headaches for inflation forecasters. Research over the past decade has documented considerable instability in inflation forecasting models, see, for example, Cogley and Sargent (2002, 2005), Cogley, Primiceri and Sargent (2010), Levin and Piger (2004), and Stock and Watson (2007); the literature on this instability is surveyed in Stock and Watson (2009). Given this instability, inflation forecasters have a dearth of reliable multivariate models for forecasting inflation. In fact, it is exceedingly difficult to improve systematically

The complete set of charts from this paper can be found at: http://www.economics.harvard.edu/faculty/stock/files/w16488_rev.pdf.

upon simple univariate forecasting models, such as the Atkeson-Ohanian (2001) random walk model (although that model seems to have broken down in the 2000s) or the time-varying unobserved components model in Stock and Watson (2007).

Yet this picture of the instability and unreliability of multivariate forecasting models conflicts with the broad historical regularity that the major postwar U.S. disinflations have all occurred during or just following recessions. Chart 1 plots the paths of the unemployment rate and the four-quarter rate of inflation[1] (π_t^4) in the core personal consumption expenditure (PCE) price index over the eight National Bureau of Economic Research (NBER)-dated recessions from 1960 to 2010. Because the 1980Q1 recession was only six quarters peak-to-peak, Chart 1 combines the 1980Q1 and 1981Q3 recessions into a single episode, so the eight recessions and their aftermath are presented as seven recessionary episodes. The plotted series are deviated from their values at the date of the NBER peak. For example, in the recession beginning in 1960Q2, the unemployment rate rose from 5.2 percent in 1960Q2 to 7.0 four quarters later (1961Q2), an increase of 1.8 percentage points. Over those four quarters, the four-quarter rate of core PCE inflation fell from 1.9% to 1.2%, a decline of 0.7 percentage points; these changes, relative to 1960Q2, are plotted in the first panel of Chart 1. In five of the seven recessionary episodes since 1960, inflation fell through the date at which the unemployment rate reached its peak, and then either plateaued or continued to fall for at least several more quarters. The most notable exception is the 1973Q4 recession, which was accompanied by sharp oil price increases and, as discussed below, much higher oil price pass-through to core than is currently observed.

One way to see the commonality among these episodes is to superimpose the panels of Chart 1. This is done in Chart 2, where the data for each episode have been scaled so that the unemployment rate increases by one unit between the NBER peak (time 0) and the unemployment peak (time 1).[2] Chart 2 also plots the mean of these scaled unemployment and inflation rates, along with one-standard error bands. The 1973Q4 recession is omitted from Chart 2—but not from our econometrics—because of the atypical sequence of

Chart 1
U.S. Rates of Unemployment and Inflation in All post-1960 Recessions.

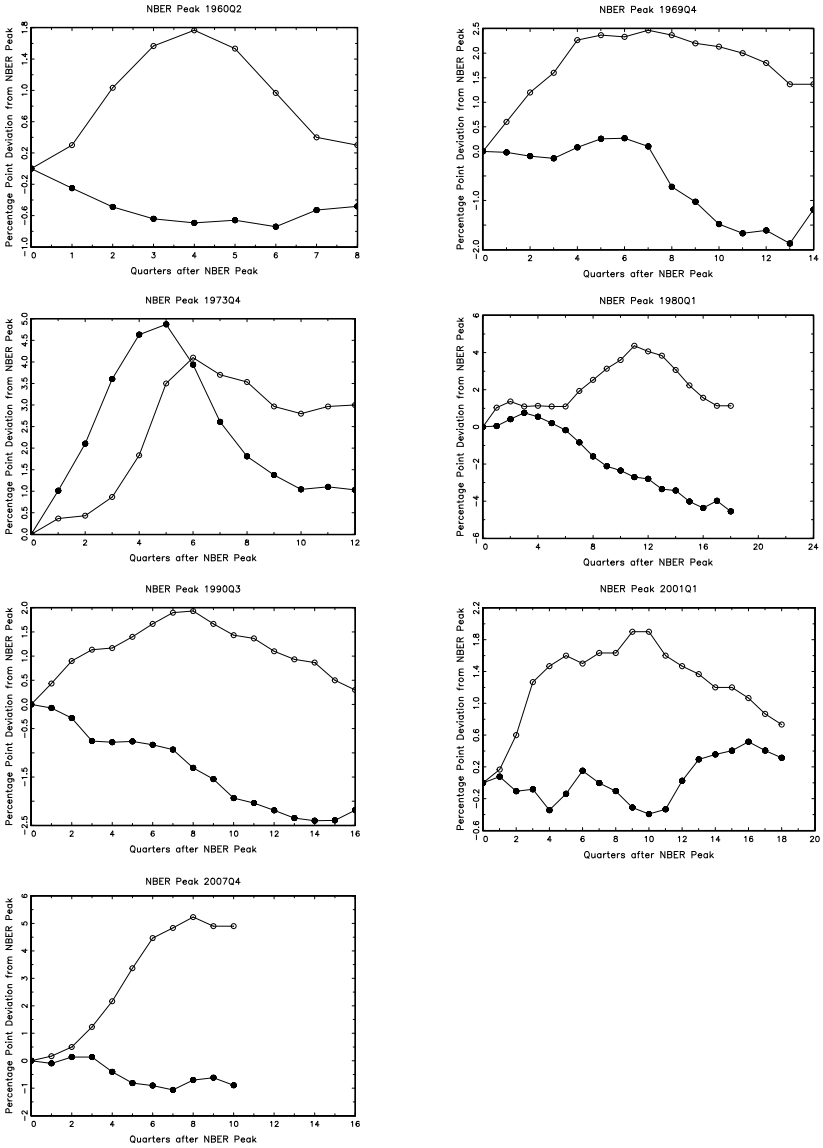

Unemployment rate (open circles) and four-quarter rate of core PCE inflation (solid circles) during the eight U.S. recessions since 1960 (the 1980 and 1981 recessions are merged). The series are plotted as deviations from their values at the NBER peak.

Chart 2
U.S. Rates of Unemployment and Inflation in Six post-1960 Recessions, Combined

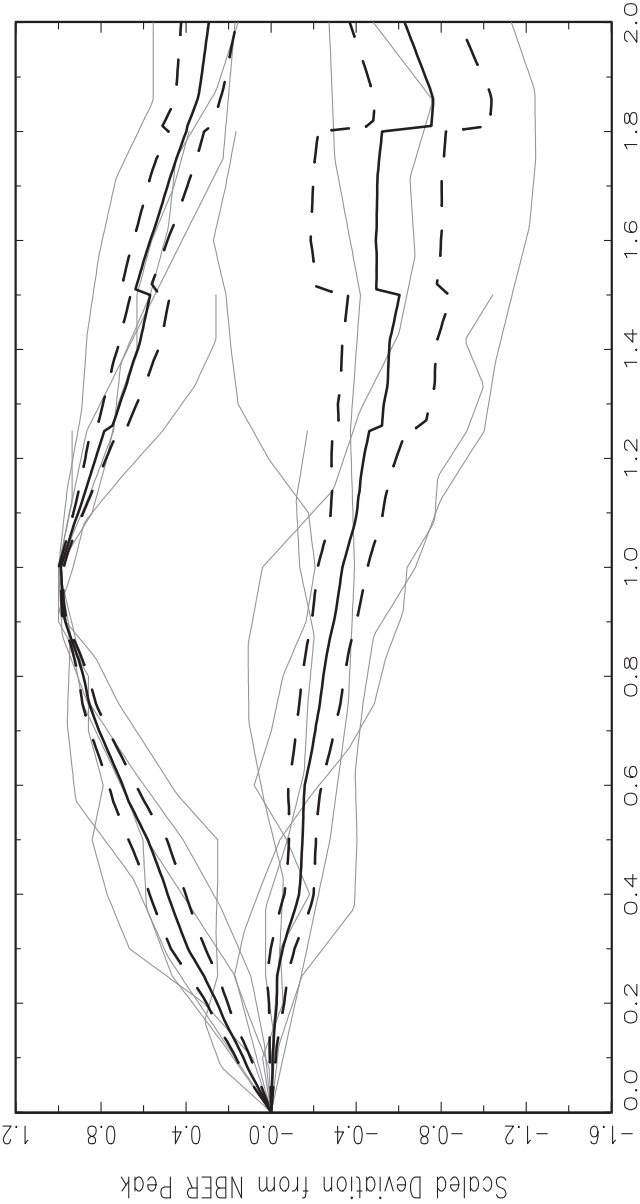

Unemployment rate (upper lines) and 4-quarter rate of core PCE inflation (lower lines) over six U.S. recessions from 1960 to 2010, including the mean (dark solid) and ± 1 standard error bands (dashed). The series are plotted as deviations from their values at the NBER peak, scaled so that the unemployment rate reaches a maximum of 1 at date 1. 1973 is omitted, and the 1980 and 1981 recessions are merged.

energy price increases through the first six months of the recession. Averaged over the six episodes in Chart 2, by the time that the unemployment rate peaks, the four-quarter rate of core PCE inflation has fallen by 0.37 percentage points (standard error = 0.13) for each percentage point rise in the unemployment rate. By the time that the episode is 50 percent beyond the peak unemployment rate (that is, at time scale 1.5 in Chart 2), the four-quarter rate of core PCE has fallen by 0.59 percentage points (SE = 0.23) for each percentage point peak increase in the rate of unemployment.

Two of the episodes in Chart 2 are of particular interest. The first is 2001Q1. Inflation fell through the first 10 quarters of this episode: by the second quarter of 2003, four-quarter core PCE inflation had fallen to 1.5% and there was increasing concern about deflation (e.g., Bernanke [2003]). In 2004, however, inflation deviated from the historical pattern by increasing. The second episode of interest is the recession that began in 2007Q4. Based on currently available data, the path of core PCE inflation in this episode is only slightly above the post-1960 average. We return to both of these episodes below.

Chart 2 captures the essential empirical content of the Phillips curve: Inflation declines during periods of economic weakness. On average over these recessionary episodes, inflation at first falls slowly, then more rapidly as the unemployment rate increases. At some point after the unemployment rate peaks, the inflation rate stabilizes at a lower level. With only six episodes, the standard errors are fairly large and increase with the time after the NBER peak, so these dynamics are estimated imprecisely.

The goal of this paper is to reconcile the apparent contradiction between the instability of Phillips curve forecasting models (and multivariate inflation forecasting models more generally) and the empirical regularity in Chart 2. We do so by drawing upon four sets of evidence. First, we provide nonparametric and parametric evidence of a stable linear relationship between inflation and a new gap measure, which we term a recession gap. The unemployment recession gap is the difference between the current unemployment rate and the minimum unemployment rate over the current and previous 11 quarters. This new gap is designed to turn the empirical regularity

in Chart 2 into a variable that can be used in a regression. Second, we provide nonparametric evidence of nonlinearities in the relation between four-quarter inflation and traditional unemployment and output gap measures; this evidence is consistent with the nonlinear parametric specification found by Barnes and Olivei (2003). Third, we conduct a pseudo out-of-sample forecasting exercise using the unemployment recession gap along with other activity measures, including both parametric and nonparametric forecasts; we find that simple linear models using the unemployment recession gap provide episodic improvements over univariate forecasts of four-quarter inflation, where the forecasting improvements occur during economic downturns. These episodic improvements are consistent with, but sharper than, those noted in Stock and Watson (2009). Fourth, we conduct a dynamic simulation of inflation using the recession gap model and find a good match between the actual and predicted inflation paths, given the unemployment path, over the five downturns of Chart 2.

The econometrics in this paper consider a multivariate forecasting model in which a candidate variable, say x_t, is used to predict the forecast errors from a univariate forecast of inflation over the next four quarters, π^4_{t+4}. The univariate model we adopt is the unobserved components model of inflation proposed in Stock and Watson (2007), in which the rate of inflation is represented as the sum of a stochastic trend, τ_t, and a transitory component, where the volatility of the two components varies over time. In this model, the forecast of future inflation using date t information is the best estimate of the trend at date t, $\tau_{t|t}$, so the forecast error for four-quarter ahead inflation is $\pi^4_{t+4} - \tau_{t|t}$. Cogley, Primiceri and Sargent (2010) refer to the deviation of inflation from τ_t as the inflation gap, and like them we focus on predictability of this gap. Specifically, the multivariate forecasting models we consider have the form,

$$\pi^4_{t+4} = \tau_{t|t} + \gamma_4 x_t + e^4_{t+4}, \tag{1}$$

where γ_4 is an unknown coefficient and e^4_{t+4} is an error term, and where the subscript/superscript "4" indicates that (1) applies to the four-quarter inflation rate.

Our primary focus is on the unemployment recession gap as the predictor variable x_t in (1). However, we also estimate (1) using other predictors x_t, in particular other measures of economic activity, survey expectations of inflation, and measures of the money supply. The findings using other activity variables are consistent with those using the unemployment recession gap: Activity variables provide episodic improvements over the univariate model, which are sharpest if the activity variable is a recession gap. In contrast to the findings in Ang, Bekaert, and Wei (2007), we find that, on average over our sample period, augmenting activity variable forecasts with survey measures of inflation expectations tends to make little difference, relative to using only the activity measure. Consistent with the literature, monetary variables produce forecasts of inflation that are less accurate out of sample than univariate forecasts, both on average over the full sample and episodically.

Before turning to our analysis, we make several remarks about the interpretation of our forecasting model and our results. First, the recession gap is not a standard gap measure, in the sense that it measures only the severity and timing of economic contractions. This paper focuses on only one part of the Phillips curve—what happens during downturns—and is silent about the behavior of inflation in booms.

Second, we think of the estimated trend in (1), $\tau_{t|t}$, as capturing long-term inflation expectations. The extent to which these expectations, as captured by $\tau_{t|t}$, are "anchored" or "resilient" is allowed to change over time. We show in Section III that our trend measure closely tracks inflation expectations as reported by the Survey of Professional Forecasters. In a sense, this should not be surprising: It is very difficult to beat univariate inflation forecasting models, and $\tau_{t|t}$ is computed from a competitive univariate forecasting model that allows for time variation in the resilience of trend inflation, so it makes sense that the forecasts from this model would line up with professional forecasts. Because our trend is derived as a uni-

variate long-run forecast, conceptually $\tau_{t|t}$ differs from private-sector inflation expectations, although as a practical matter this difference seems to be slight. Our interpretation of $\tau_{t|t}$ as long-term expected inflation also accords with Cogley, Primiceri and Sargent's (2010) interpretation of τ_t as the Fed's perceived inflation target.

Third, our analysis focuses on backwards-looking models, in which expectations are in effect estimated by a reduced-form time series model. To the extent that $\tau_{t|t}$ captures inflationary expectations, (1) can be thought of as a New Keynesian Phillips Curve in which observed expectations are used for estimation. An alternative approach is to use model-based expectations in conjunction with a New Keynesian Phillips curve. Fuhrer and Olivei (2010) provide simulations using this latter approach in the context of the current recession and those simulations complement the forecasting approach in this paper.

There are several other recent papers related to ours. Liu and Rudebusch (2010) provide different evidence that the behavior of inflation in the current downturn is consistent with the historical U.S. Phillips curve, and Meier (2010) provides international evidence that recessions are associated with declines in inflation. Williams (2009) provides Phillips-curve forecasts of the decline in inflation during this recession, in which he emphasizes the importance of the substantial increase in expectations anchoring in muting the disinflationary pressures of the currently large gaps. Giannone, Lenza, Momferatou, and Onorante (2010), using quite different methods, also provide evidence of a Euro-zone Phillips curve during the current episode.

Section II of this paper shows that the pattern in Chart 2 also holds for core CPI, the GDP price index, headline PCE, and headline CPI. Section III presents our econometric analysis of using the unemployment recession gap and other unemployment rate gaps. Section IV extends this analysis to other predictors. Section V discusses implications for the current recession, and Section VI concludes.

Data note: All the data used in this paper are quarterly from 1959Q1 – 2010Q2. The values of monthly series are averaged over the quarter. The data are the most recent revised data as of August 26, 2010. All

Chart 3
U.S. Rates of Unemployment and Inflation in post-1960 Recessions: Other Price Indexes

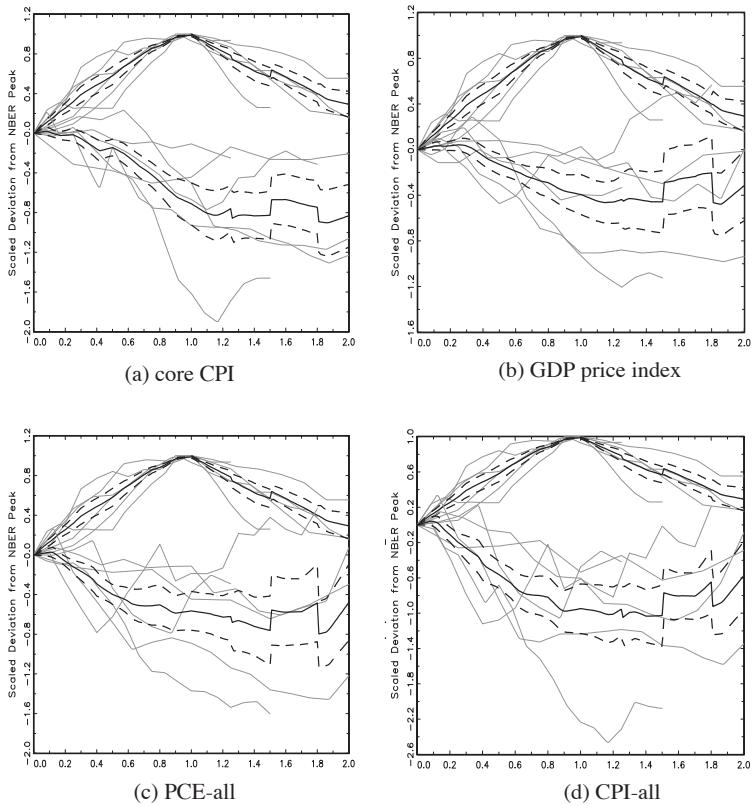

Unemployment rate (upper lines) and four-quarter rates of inflation (lower lines) over six U.S. recessions from 1960 to 2010, including the mean and ± 1 standard error bands, for four price indexes. Construction and line schemes are the same as in Chart 2.

predictors x_t are constructed to be one-sided using revised data; we do not consider issues raised by data revisions. Gaps and trend inflation are computed using pre-1959 data for initial conditions when available. Except for Section II, we focus on inflation as measured by the PCE price index less food and energy (core PCE) because it is methodologically consistent and because it eliminates the noise from energy price fluctuations, which have recently been very large (e.g., Hamilton [2009]); results for other inflation measures can be computed using the replication files that are available for this paper.

II. Price Inflation During Recessions, 1960 – 2010: Other Price Indexes

In addition to core PCE inflation, other measures of price inflation also fall during periods of economic weakness. Chart 3 plots the recession behavior of four-quarter inflation computed using four other price indexes: core Consumer Price Index (CPI), the chain-weighted GDP price index, the headline PCE price index, and the headline CPI. The construction of Chart 3 is the same as Chart 2, except for the price index used.

The pattern of inflation for the four price indexes in Chart 3 is similar to that seen using core PCE in Chart 2. The magnitudes of the decline in inflation depend on the price index. By the time that the episode is 50% beyond the peak unemployment rate (a value of 1.5 on the time scale in Chart 3), four-quarter core CPI inflation has fallen by 0.83 percentage points (SE = 0.25), inflation measured by the GDP price index has fallen by 0.45 percentage points (SE = 0.27), and headline PCE and headline CPI have respectively declined by 0.74 (SE = 0.33) and 1.02 (SE = 0.33) percentage points. The standard errors of the mean declines for headline inflation are larger than for core because of movements in energy and food prices that differ from one recession to the next. Nevertheless, the basic pattern remains the same.[3]

Because the behavior of the four inflation measures in Chart 3 matches the overall pattern observed for core PCE inflation in Chart 2, for the rest of this paper we focus solely on core PCE inflation.

III. Price Inflation During Recessions, 1960 – 2010: Econometrics

The graphical evidence of the previous section is suggestive but informal, so we now turn to an econometric investigation of price inflation during recessions. In this section, we continue to focus on unemployment-based measures of activity. We begin with additional details about our forecasting model (1), including our measure of trend inflation, the implications of time variation in our trend estimate for the long-run effect of a change in x_t on inflation, and unemployment gaps including our new unemployment recession gap. We then report the results of four complementary econometric investigations. First, we examine nonlinearities in the Phillips curve as suggested by recent work by Barnes and Olivei (2003), Stock and Watson (2009), and Fuhrer and Olivei (2010); we confirm that there is evidence of Barnes-Olivei (2003) nonlinearities using a standard gap measure, but not using the recession gap. Second, we estimate parametric (linear) Phillips curve models and find that models with the recession gaps exhibit less instability than models with conventional gaps. Third, we conduct a pseudo out-of-sample forecasting study that compares various unemployment-based forecasts; all the unemployment gap measures exhibit the "episodic" improvements (during recessions) discussed in Stock and Watson (2009), but those improvements are sharpest for the recession gap measure. Finally, we conduct a dynamic simulation using a full-sample, one-quarter ahead forecasting model based on the recession gap and find that, given the unemployment path, the predicted inflation path matches the actual path of inflation in each of the six episodes plotted in Chart 2. This model contains only two estimated coefficients, a time-varying moving average parameter and a single (stable) short-run Phillips curve slope coefficient. Thus this model provides a parsimonious parametric summary of Chart 2.

III.A. Measures of Trend Inflation and Real-Time Gaps

Trend inflation. Implementation of (1) as a forecasting equation requires a measure of trend inflation computed using contemporaneous and past, but not future, data—that is, a one-sided measure of trend inflation. The trend measure we use here is derived from the univariate

Chart 4
UCSV Model of Core PCE Inflation: Estimated Time-Varying Standard Deviations of the Trend and Transitory Components (panels (a) and (b)) and the Implied Time-Varying Moving Average Coefficient.

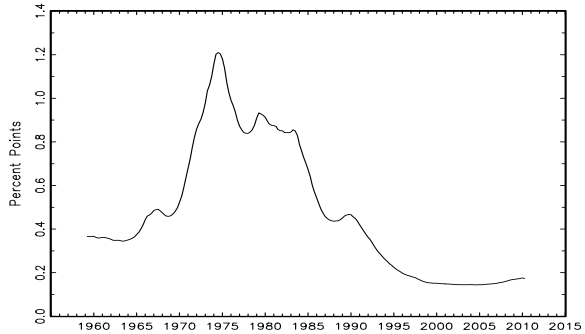

(a) Standard deviation of the change in trend ($\sigma_{\varepsilon,t}$)

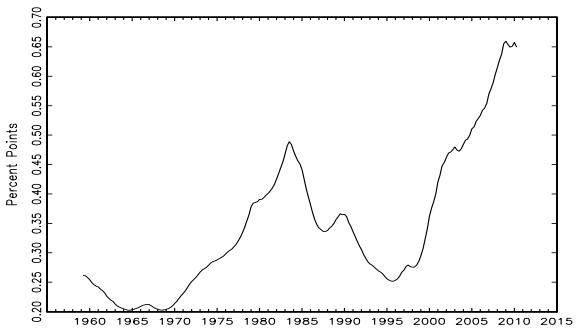

(b) Standard deviation of transitory component ($\sigma_{\eta,t}$)

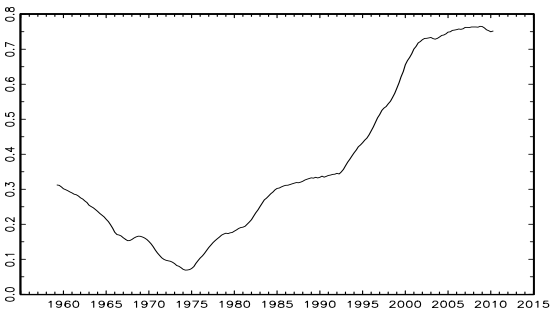

(c) Moving average coefficient (θ_t)

Chart 5
The Estimated Trend in Core PCE ($\tau_{\tau|t}$) and the Five-year Ahead Median Inflation Forecast From the Survey of Professional Forecasters

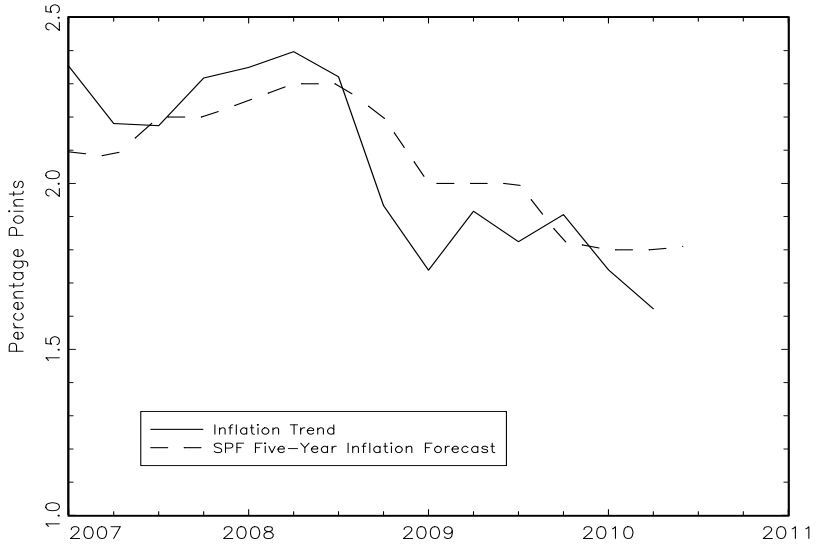

time series model of inflation developed in Stock and Watson (2007), in which the rate of inflation is represented as the sum of two unobserved components, a trend τ_t and a transitory disturbance η_t, where the variances of these two disturbances can change over time:

$$\pi_\tau = \tau_t + \eta_t, \quad E\eta_t = 0, \text{ var}(\eta_t) = \sigma^2_{\eta,t} \tag{2}$$

$$\tau_t = \tau_{t-1} + \varepsilon_t, E\varepsilon_t = 0, \text{ var}(\varepsilon_t) = \sigma^2_{\varepsilon,t}, \text{cov}(\eta_t, \varepsilon_t) = 0. \tag{3}$$

The time-varying variances are modeled as evolving as random walks in logarithms. This so-called unobserved components-stochastic volatility (UC-SV) model is estimated using nonlinear filtering methods, for details see Stock and Watson (2007). The estimate of trend inflation ($\tau_{t|t}$) which we use to estimate (1) is the one-sided (that is, filtered) estimate of τ_t obtained from the UC-SV model.

The UC-SV model implies that inflation has a time-varying moving average representation in first differences (a time-varying IMA(1,1) representation),

$$\Delta\pi_t = a_t - \theta_t a_{t-1}, \qquad Ea_t = 0, \ \text{var}(a_t) = \sigma^2_{a,t}, \tag{4}$$

where θ_t and $\sigma^2_{a,t}$ are functions of $\sigma^2_{\eta,t}$ and $\sigma^2_{\varepsilon,t}$.

From the perspective of inflation forecasting, the key feature of the UC-SV model is that, conditional on $\sigma^2_{\varepsilon,t}$ and $\sigma^2_{\eta,t}$, it results in a linear forecast of inflation with potentially long lags where the lag structure is time-varying but parsimoniously parameterized by only two parameters. The variances $\sigma^2_{\varepsilon,t}$ and $\sigma^2_{\eta,t}$ determine the variability of the trend and transitory components. Allowing these innovation variances to change over time produces time variation in the resilience of the trend. In particular, a regime shift in monetary policy that induces a change in the extent to which expectations are anchored will be captured by a decrease in the variance of the trend innovation and an increase in the resilience of the estimated trend.

Chart 4 presents the standard deviations $\sigma_{\eta,t}$ and $\sigma_{\varepsilon,t}$ and the implied time-varying moving average coefficient θ_t, for core PCE inflation. Over the past decade, the volatility of the trend ($\sigma_{\varepsilon,t}$) has been at historic lows, and the persistence of inflation forecasts, as measured by θ_t, has been at historic highs. During the 2000s, inflation tended to revert to a stable trend, whereas in the 1970s and 80s the trend moved to track inflation.

Chart 5 plots the estimated trend $\tau_{t|t}$ from the UC-SV model along with the median five-year ahead forecast that has been reported in the Survey of Professional Forecasters since 2007. The two series move together very closely. Although the time span is very short, Chart 5 suggests that the trend $\tau_{t|t}$ can be thought of as a substitute measure of long-term inflation expectations.

The equivalence of the unobserved components and IMA(1,1) representations allows a useful link between the value of θ and the resilience of the trend. Setting aside time variation for the

moment, the filtered trend can be expressed as a distributed lag of past inflation, specifically,

$$\tau_{t|t}=(1-\theta)\sum_{i=0}^{\infty}\theta^{i}\pi_{t-i} \ . \tag{5}$$

The weights in this expression sum to one, and the smaller is θ, the more weight is placed on recent observations and the more volatile is the trend. In the limit that θ approaches one, the estimated trend is simply the sample average of past inflation.

From (1) to a backwards-looking Phillips curve with time-varying parameters. In the UC-SV model, $\tau_{t|t}$ is the optimal univariate time-t forecast of π_{t+h} for all $h \geq 1$, so that $a_{t+1} = \pi_{t+1} - \tau_{t|t}$, where a_t is the forecast error in (4), the MA(1) version of the univariate model. We consider the possibility that this univariate forecast error is predictable using some variable x_t, so that $a_{t+1} = \gamma_1 x_t + e^1_{t+1}$, where the subscript/superscript "1" indicates that γ_1 and e^1_{t+1} apply to this one-step ahead projection. This yields the one-step ahead model,

$$\pi_{t+1} = \tau_{t|t} + \gamma_1 x_t + e^1_{t+1} . \tag{6}$$

If we continue to ignore time variation in θ, then substituting (5) into (6) and rearranging yields the autoregressive-distributed lag model,

$$\Delta\pi_{t+1} = -\theta\sum_{i=0}^{\infty}\theta^{i}\Delta\pi_{t-i} + \gamma_1 x_t + e^1_{t+1} , \tag{7}$$

Equation (7) is just a tightly parameterized backwards-looking Phillips curve forecasting model with potentially long lags in the tradition of Gordon (1982, 1990, 1998) and Brayton, Roberts, and Williams (1999), without the dummy variables and supply shock variables found in the Gordon (1990) "triangle" model.

Equation (7) provides a useful framework for understanding two implications of time variation in θ (with time variation, (7) is an approximation which holds for slow time variation). First, time variation in θ implies time variation in the Phillips curve coefficients on lagged inflation. Second, time variation in θ implies time variation in the long-run effect on inflation of a change in x_t. Spe-

cifically, the long-run effect on inflation of a unit exogenous change in x_t is $(1 - \theta)\gamma_1$. Thus, even if γ_1 is constant (we provide evidence below that this is so, when x_t is the unemployment recession gap), the long-run effect on inflation varies over time because θ varies over time. In particular, when θ is large (close to one), then the long-run multiplier is less than when θ is small. Said differently, when the innovations to trend inflation are relatively small—that is, when inflation expectations are well-anchored—then θ is near one. Even if the one-quarter ahead effect of a change in x_t on inflation is constant over time, the anchoring of expectations means that the long-run impact of a change in x_t is less than if expectations were less well anchored.

Iterating (4) forward four quarters yields $\pi^4_{t+4} = \tau_{t|t} + b_{t+4}$ or, equivalently, $\pi^4_{t+4} - \pi_t = -\theta\sum_{i=0}^{\infty}\theta^i\Delta\pi_{t-i} + b_{t+4}$ where $b_{t+4} = [a_{t+4} + (2-\theta)a_{t+3} + (3-2\theta)a_{t+2} + (4-3\theta)a_{t+1}]/4$. As in the one-step forecast, suppose that future univariate forecast errors a_{t+1}, a_{t+2}, a_{t+3}, a_{t+4} (and thus b_{t+4}) are predictable using x_t, so that $b_{t+4} = \gamma_4 x_t + e^4_{t+4}$. Thus we have that $\pi^4_{t+4} = \tau_{t|t} + \gamma_4 x_t + e^4_{t+4}$, which is (1), or equivalently,

$$\pi^4_{t+4} - \pi_t = -\theta\sum_{i=0}^{\infty}\theta^i\Delta\pi_{t-i} + \gamma_4 x_t + e^4_{t+4}. \qquad (8)$$

When derived in this way the coefficient γ_4 in (8) is seen to depend on θ because b_{t+4} is a function of θ.[4] Thus, time variation in θ may lead to time variation in γ_4 even if γ_1 is time invariant.

Real-time gaps. A challenge in forecasting inflation using activity variables is constructing reliable one-sided measures of activity gaps, which can differ substantially from two-sided gaps estimated with the benefit of subsequent data. Here, we consider two one-sided gaps, one standard in the literature and one new, plus a "differences" transformation of activity.

The new one-sided gap measure, which we refer to as a "recession gap," focuses attention on economic downturns by computing the gap as the deviation of unemployment from its minimum over the current and previous 11 quarters. That is, the unemployment recession gap is,

$$\textit{unemployment recession gap}_t = u_t - \min(u_t, \ldots, u_{t-11}), \qquad (9)$$

Chart 6
The Unemployment Rate (panel (a)) and Three Activity Measures based on the Unemployment Rate (panel (b)): the One-Sided Bandpass Gap, the Four-Quarter Difference, and the 12-quarter Unemployment Recession Gap

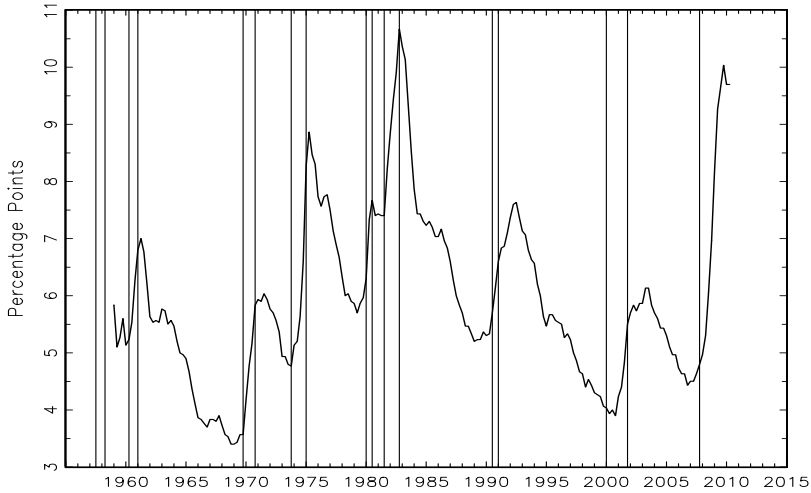

(a) Civilian unemployment rate

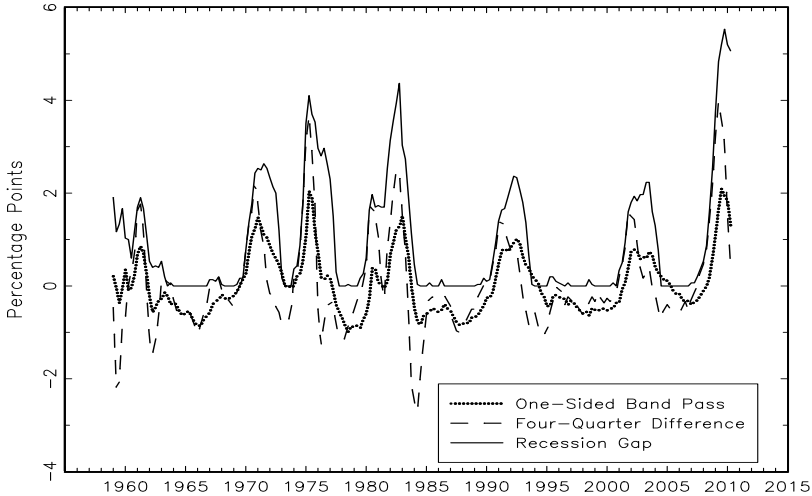

(b) Derived unemployment activity measures

Chart 7
Nonparametric Phillips Curves

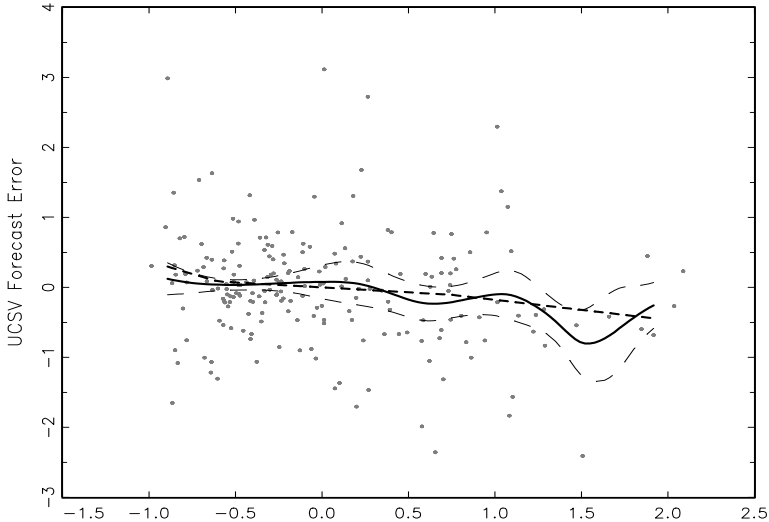

(a) unemployment gap: one-sided bandpass filtered gap

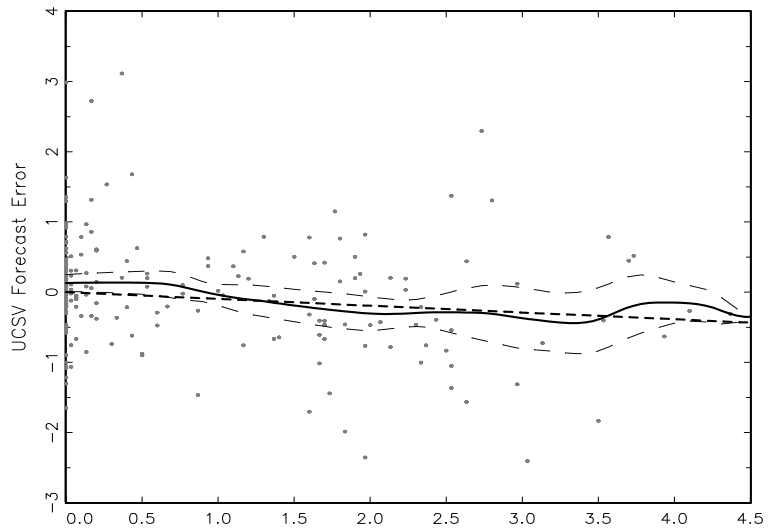

(b) unemployment gap: 12-quarter unemployment recession gap

Scatterplot of UCSV 4-quarter ahead forecast error ($\pi^4_{t+4} - \tau_{tlt}$) vs. real-time (one-sided) unemployment gaps, for two measures of the gap: (a) one-sided bandpass filtered, and (b) 12-quarter recession gap. Kernel nonparametric regression functions and one standard error bands (dashed) are shown in black. Parametric regression functions (dashed) are shown in gray: in panel (a), a Barnes-Olivei (2003)-type piecewise linear regression function, in panel (b), a linear regression function.

where u_t denotes the unemployment rate at date t. In effect, the unemployment recession gap takes on the value of the unemployment rate in Chart 2 during downturns and is zero otherwise. Thus, the unemployment recession gap translates Chart 2 into something that can be analyzed using linear regression.[5]

We also examine a conventional one-sided gap computed using a one-sided bandpass filter. Following Stock and Watson (2007), one-sided bandpass gaps are computed as the deviation of the series augmented with univariate forecasts of future values from a symmetric two-sided MA(80) approximation to the optimal lowpass filter with pass band corresponding to periodicities of at least 60 quarters.

The third unemployment-based predictor we consider is a difference (or changes) transformation, in which the predictor is the four-quarter change in the unemployment rate, $u_t - u_{t-4}$.

Chart 6 plots the unemployment rate and these three unemployment-based measures. The three measures have broad similarities but important differences. Most notably, the bandpass and differences measures vary during economic expansions, whereas the recession gap essentially varies only during downturns.

III.B. Nonlinearities in the Phillips Curve

Does the Phillips curve slope depend on the size of the gap?

Chart 7 provides scatterplots of $\pi^4_{t+4} - \tau_{t|t}$ against the one-sided bandpass gap (upper panel) and the unemployment recession gap (lower panel). Both panels also show a nonparametric kernel regression line (with 95 percent confidence bands) and a parametric regression function. Barnes and Olivei (2003) found evidence supporting a piecewise linear Phillips curve, so for the one-sided bandpass regression the parametric regression is a piecewise linear function, with the thresholds chosen so that 70 percent of the observations fall in the middle section and 15 percent in each outer section. The parametric regression in the recession gap scatterplot is linear.

Chart 7 provides support for the Barnes-Olivei (2003) specification applied to the one-sided bandpass gap: the Barnes-Olivei (2003) type piecewise linear function is remarkably close to the nonpara-

Chart 8
Dependance of γ_4 on the Level of Inflation

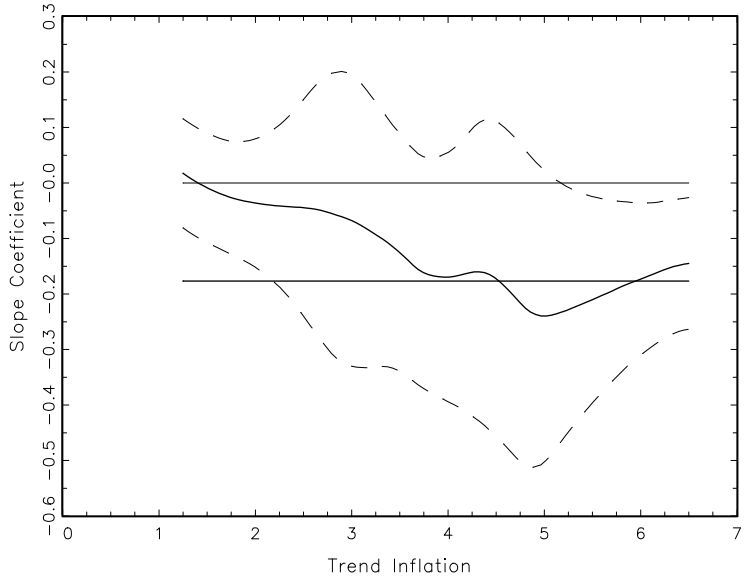

Nonparametric regression (gray solid) and 95% confidence bands (gray dashed) of the slope coefficient γ_4 as a function of the value of trend inflation at date t (τ_{it}), using the unemployment recession gap. Black solid line is the parametric estimate (-0.18, SE = 0.06). Parametric and nonparametric regressions are full-sample.

metric regression function. There is a large central region—normal times of moderate and small gaps—in which the Phillips relation is essentially flat, but in periods of large (bandpass) gaps, the curve steepens.[6] In the pseudo out-of-sample forecasting exercise reported below we therefore consider both linear and nonlinear (nonparametric) specifications for the bandpass gap.

In contrast, there is little evidence of nonlinearities in the Phillips curve using the recession gap, so the work below adopts a linear specification as a function of the recession gap.

Does the Phillips curve slope depend on the level of inflation? The possibility that the Phillips curve flattens at low levels of inflation has long been an element of the literature, see for example Ball, Mankiw, and Romer (1988), Akerlof, Dickens, and Perry (1996) (on downward wage rigidity), and, for a recent empirical treatment, Aron and Muellbauer (2010). We investigated this type of nonlinearity, in

Table 1
Estimated 1- and 4-Quarter Ahead Forecasting Regressions Using Unemployment Gaps

	1959Q2 – 2009Q2		1959Q2 – 1983Q4		1984Q1 – 2009Q2		t-test for break in 1984Q1	QLR stability test *p*-value
1-quarter ahead	γ_1	R^2	γ_1	R^2	γ_1	R^2		
Recession gap	–0.10 (0.04)	0.035	–0.11 (0.06)	0.037	–0.07 (0.03)	0.033	0.55	0.96
1-sided bandpass gap	–0.20 (0.09)	0.028	–0.30 (0.15)	0.045	–0.09 (0.08)	0.010	1.23	0.75
Fourth difference	–0.13 (0.05)	0.033	–0.18 (0.07)	0.046	-0.08 (0.06)	0.018	1.12	0.63
4-quarter ahead	γ_4	R^2	γ_4	R^2	γ_4	R^2		
Recession gap	-0.18 (0.06)	0.077	-0.21 (0.07)	0.084	-0.11 (0.07)	0.066	1.00	0.32
1-sided bandpass gap	-0.41 (0.10)	0.079	-0.60 (0.12)	0.120	-0.11 (0.13)	0.017	2.76**	0.02 (1983Q1)
Fourth difference	-0.29 (0.09)	0.111	-0.42 (0.11)	0.166	-0.08 (0.07)	0.028	2.66**	0.08 (1983Q1)

Notes: The one-quarter ahead regressions are $\pi_{t+1} - \tau_{tlt} = \gamma_1 x_t + e^1_{t+1}$, and the four-quarter ahead regressions are $\pi^1_{t+1} - \tau_{tlt} = \gamma^4 x_t$, `where x_t is a predictor known at date t. The first six numeric columns present the estimates of γ_1 (or γ_4, as appropriate), its standard error (in parentheses), and the regression R^2 for the row predictor and column sample. Standard errors are heteroskedasticity-robust for one-quarter ahead regressions and are Newey-West standard errors (6 lags) for four-quarter ahead regressions. The QLR (sup-Wald) statistic was computed using symmetric 15% trimming. If the QLR test rejects stability, the estimated break date appears in parentheses. The t-statistic in the second to last column is significant at the *5% **1% significance level.

which the slope of the Phillips curve, specifically γ_4 in (1), depends on the level of inflation; here, we focus on the recession gap Phillips curve.

Chart 8 presents a nonparametric estimate of the slope γ_4 (the coefficient on the unemployment recession gap) as a function of the current estimate of trend inflation (τ_{tlt}).[7] The estimated slope is clearly less in absolute value for small values of trend inflation than for large values, however the 95% confidence bands are wide and the full-sample linear regression estimate of -0.18 is contained within the confidence band for almost all values of trend inflation. Parametric models incorporating this nonlinearity do not seem to be particularly robust, with the statistical significance of the nonlinearity depending on the details of the specification. One reason for this imprecision and apparent lack of robustness is that there is limited historical experience at very low levels of inflation, so the evidence we have essentially rests on two historical episodes: the early 1960s and the

early 2000s. This imprecision and lack of robustness is underscored by pseudo out-of-sample forecasting experiments (not reported) in which specifications in which the slope depends on $\tau_{t|t}$ were found to exhibit instability.

Because the time series evidence is limited, we also consider evidence from the micro literature on price setting. One argument for a flattening of the Phillips curve at low levels of inflation is that there is resistance to reducing nominal wages and prices. The micro literature, however, presents little evidence of a price change floor at zero. For example Nakamura and Steinsson (2008) find that one-third of price changes for the same goods are negative, a finding consistent with Klenow and Kryvtsov (2008). Some additional evidence on whether the distribution of price changes truncates or piles up at zero is provided in the Appendix, which examines annual price changes for 233 disaggregated components of the PCE price index. Price changes at this level of disaggregation accord with the micro finding of little price resistance at zero. While the absence of resistance to price declines does not imply an absence of resistance to wage declines, this micro and subaggregate evidence does not on its face suggest that a price Phillips curve would flatten at low levels of inflation.

Given the limited evidence in the time series data and the lack of evident price resistance at zero in the micro and subaggregate data, for the rest of this paper we adopt specifications in which the Phillips curve slope does not depend on the level of inflation. This said, the hint of nonlinearity in Chart 8 remains an intriguing topic for further research.

III.C. Gap Models: Estimates and Stability

Table 1 reports various regression statistics for estimates of γ_1 in (6) (one-quarter ahead) and γ_4 in (1) (four-quarter ahead) using the three unemployment gaps. All R^2s are low, underscoring that inflation is difficult to forecast. The final two columns report statistics testing for stability of the slope coefficient, first by testing for a break in 1984Q1 (a common choice for the Great Moderation break) and second using the Quandt Likelihood Ratio (QLR) statistic (also

Chart 9
Pseudo Out-of-Sample Forecasts of Four-Quarter Core PCE Inflation Using Various Unemployment Gaps

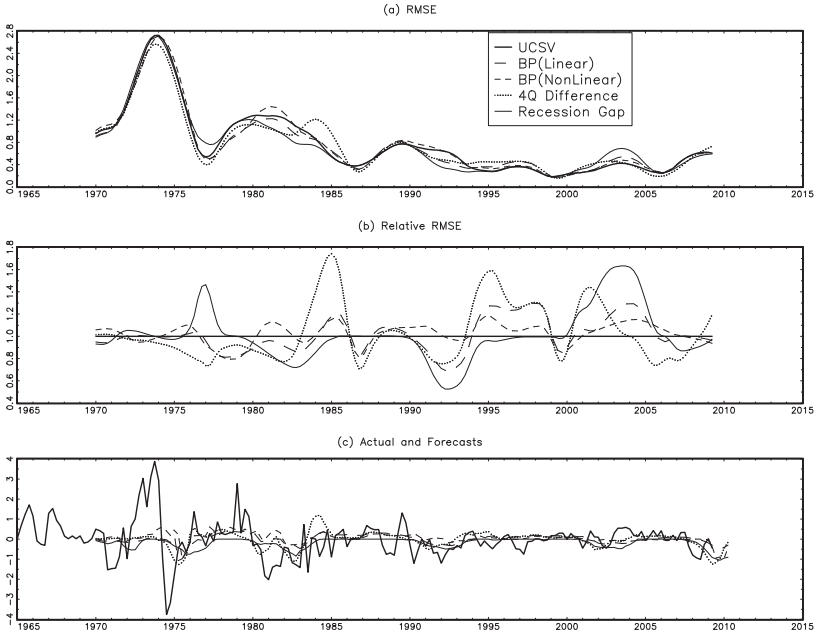

Plots are rolling root mean squared errors (top panel), rolling RMSEs relative to the UCSV model (middle panel), and forecasts (bottom panel). Forecasts are one-sided bandpass gap, nonlinear one-sided bandpass gap, four-quarter change in unemployment, and recession gap. In panel (c), actual values of $\pi^4_{t+4} - \tau_{t/t}$ are in black.

known as the sup-Wald statistic) testing for a single break at an unknown time. For the one-step ahead regressions, these test statistics fail to reject the null hypothesis of coefficient stability. The one-step ahead point estimates of γ_1 for the unemployment recession gap and the R^2 also indicate stability of this predictive regression. In contrast, the coefficients on the one-sided bandpass gap and the fourth difference change by a large factor across the subsamples, as do the regression R^2s, suggesting less stability than the unemployment recession gap regression despite the failure of the stability tests to reject for any of the one-step ahead specifications. Consistent with the discussion in Section 3.1, the estimates of γ_4 in the four-step ahead regression appear less stable than in the one-step ahead specification. Indeed, both stability tests reject for the four-quarter ahead bandpass gap and

fourth-differences specifications, and the estimated coefficients and R^2s change dramatically for these two measures from the pre-84 to post-84 parts of the sample. In contrast, the hypothesis of stability is not rejected for the recession gap coefficient in the four-quarter ahead specification, the magnitude of its change is small relative to the other variables, and its R^2 is stable across the two samples.

III.D. Pseudo Out-of-Sample Forecasting Results

The pseudo out-of-sample forecasting method. This section examines the forecasting performance of the three unemployment variables, relative to the univariate UC-SV benchmark, in a pseudo out-of-sample forecast experiment. At a given date t, forecasts of π^4_{t+4} using each model are made using data only available through date t. For the exercise here, the first forecast date is the later of 1970Q1 or the date necessary for the shortest regression to have 40 observations, and the final forecast date is four quarters before the end of the sample.

A useful statistic is the centered rolling root mean forecast error (RMSE). This is the square root of a weighted moving average of the squared pseudo out-of-sample forecast error, centered so that the moving average extends seven quarters on either side.[8]

Chart 9 presents rolling RMSEs, the rolling RMSEs relative to the rolling RMSE of the UC-SV model, and the pseudo out-of-sample forecasts for the three unemployment gap models. Because of the possible nonlinearity in the Phillips curve using the bandpass gap, for that gap forecasts were computed using both a linear model and a nonparametric nonlinear forecast (the predicted value is read off the recursively estimated nonparametric regression curve).

Five findings are apparent in Chart 9. As is documented in the next section, these results are robust to using other activity measures and including other variables, so we spend some time discussing them here.

1. Consistent with findings elsewhere in the literature, there is considerable variation over time in the predictability of inflation. In 2006, the rolling RMSEs were near historic lows, but they have recently crept up to levels of the early 1990s.

2. The forecasting improvements made by Phillips curve forecasts are episodic, and the greatest improvements are evident in downturns. This finding is similar to that in Stock and Watson (2009).

3. The recession gap model improves upon the UC-SV model during the disinflations of the early 1980s, the early 1990s, and (by a smaller margin) during the current recession. The only two periods in which the recession gap model does relatively poorly is during 1976-7 and 2004. Both of these failures correspond to the unusual periods observed in Chart 1: the increase in inflation following the 1973Q4 recession and the increase in inflation during 2004, which (as can be seen in Chart 2) was atypical for this stage of the business cycle.

4. The fourth-difference forecasts substantially improve upon the recession gap forecasts only during 2004-2005, when the slow decline of unemployment led to forecasts of increasing inflation, whereas the recession gap forecasts had inflation falling.

5. The nonparametric nonlinear forecasts made using the one-sided bandpass gap, which appeared promising based on the analysis of Section III.B, end up differing little from the UC-SV forecasts. The linear bandpass gap forecasts provide smaller improvements during downturns than the recession gap forecasts, and provide essentially no improvements over the UC-SV model during the current downturn. The reason for this is that the one-sided gap estimate at the end of the sample heavily weights the current unemployment rate, so by this measure the unemployment gap has been small (less than two percentage points) throughout this recession, see Chart 6(b).

III.E. Parametric Dynamic Simulations

We now turn to the question of whether the unemployment gap model is quantitatively consistent with the paths of inflation in Charts 1 and 2, given the observed path in unemployment. To address this question we conduct a dynamic simulation using the one-quarter ahead regression (6) in which x_t is the unemployment

Chart 10
Dynamic Simulations of Four-Quarter Core PCE Inflation in Five Downturns, Computed Using the Unemployment Recession Gap Model

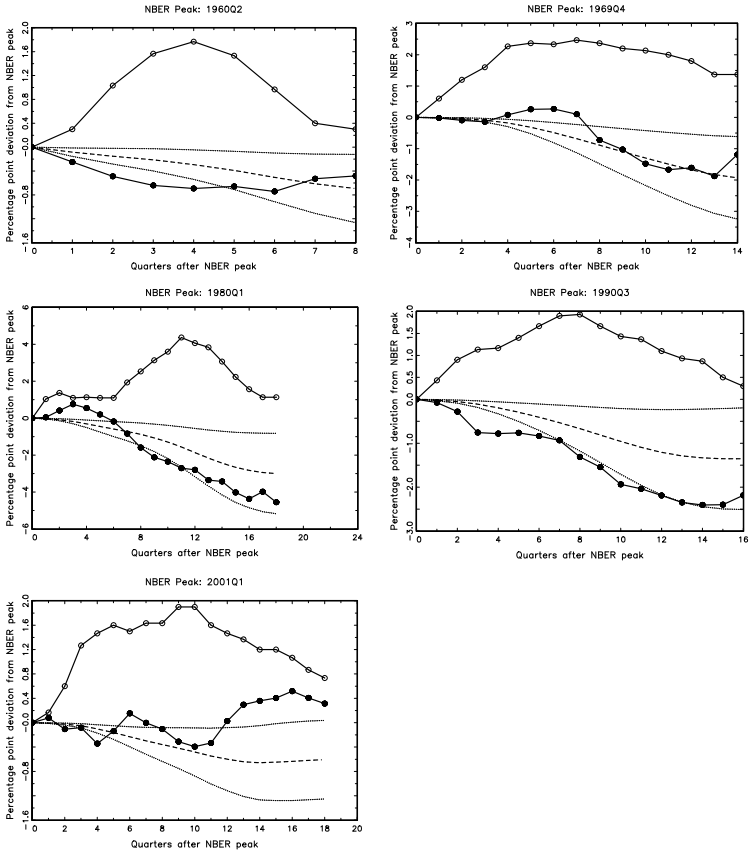

All series are plotted as percentage point deviations from their values at the NBER peak. Dashes are predicted values given the unemployment path, dots are 90% confidence bands.

Table 2
Relative Root Mean Squared Error of Activity-Based Pseudo Out-of-Sample Forecasts of Four-Quarter Core PCE Inflation, Relative to UC-SV Model

Series	**1970Q1 – 2010Q1**	**1970Q1 – 1979Q4**	**1980Q1 – 1989Q4**	**1990Q1 –1999Q4**	**2000Q1 –2010Q1**
UC-SV (RMSE)	0.97 (158)	1.61 (40)	0.86 (40)	0.45 (40)	0.39 (38)
A. Recession gaps					
Unemployment	0.98 (158)	1.01 (40)	0.85 (40)	0.79 (40)	1.20 (38)
Capacity utilization	0.94 (127)	1.03 (9)	0.85 (40)	0.87 (40)	1.22 (38)
GDP	0.98 (158)	1.03 (40)	0.82 (40)	0.78 (40)	1.18 (38)
Industrial production	0.98 (158)	1.01 (40)	0.86 (40)	0.82 (38)	1.24 (38)
Employment	0.99 (158)	1.00 (40)	0.83 (40)	0.84 (40)	1.55 (38)
CFNAI	1.01 (94)	(0)	1.04 (16)	0.77 (40)	1.24 (38)
Median recession gap	0.97 (158)	1.01 (40)	0.84 (40)	0.78 (40)	1.22 (38)
B. One-sided bandpass gaps					
Unemployment	0.98 (158)	0.98 (40)	0.97 (40)	0.96 (40)	1.08 (38)
Capacity utilization	0.93 (127)	0.80 (9)	0.93 (40)	1.00 (40)	1.10 (38)
GDP	0.95 (158)	0.96 (40)	0.90 (40)	0.87 (40)	1.11 (38)
Industrial production	0.97 (158)	0.97 (40)	0.91 (40)	0.97 (40)	1.19 (38)
Employment	0.99 (158)	0.99 (40)	0.98 (40)	0.95 (40)	1.13 (38)
CFNAI	0.90 (126)	0.81 (8)	0.88 (40)	0.92 (40)	1.14 (38)
Median BP gap	0.96 (158)	0.97 (40)	0.91 (40)	0.94 (40)	1.11 (38)
C. Four-quarter differences					
Unemployment	0.96 (158)	0.94 (40)	0.99 (40)	0.99 (40)	1.06 (38)
Capacity utilization	1.05 (123)	0.94 (5)	1.06 (40)	1.08 (40)	1.16 (38)
GDP	0.96 (158)	0.95 (40)	0.97 (40)	0.96 (40)	1.12 (38)
Industrial production	0.97 (158)	0.96 (40)	0.93 (40)	1.03 (40)	1.19 (38)
Employment	0.94 (158)	0.91 (40)	0.90 (40)	0.89 (40)	1.52 (38)
CFNAI	1.05 (122)	0.87 (4)	1.02 (40)	1.02 (40)	1.48 (38)
Median recession gap	0.96(158)	0.94 (40)	0.97 (40)	0.96 (40)	1.18 (38)
Overall median – all activity	0.95 (158)	0.97 (40)	0.86 (40)	0.86 (40)	1.12 (38)

Notes: The first line reports the standard deviation of the UC-SV forecast errors over the column sample period; the remaining lines report the ratio of the row forecast RMSE to the US-SV RMSE over the column sample. Numbers of observations used in the computation are given in parentheses. CFNAI denotes the Chicago Fed National Activity Index.

Chart 11
Pseudo Out-of-Sample Forecasts of Four-Quarter Core PCE Inflation Using Six Activity Measures
(unemployment rate, capacity utilization, GDP, industrial production, employment, and the CFNAI).

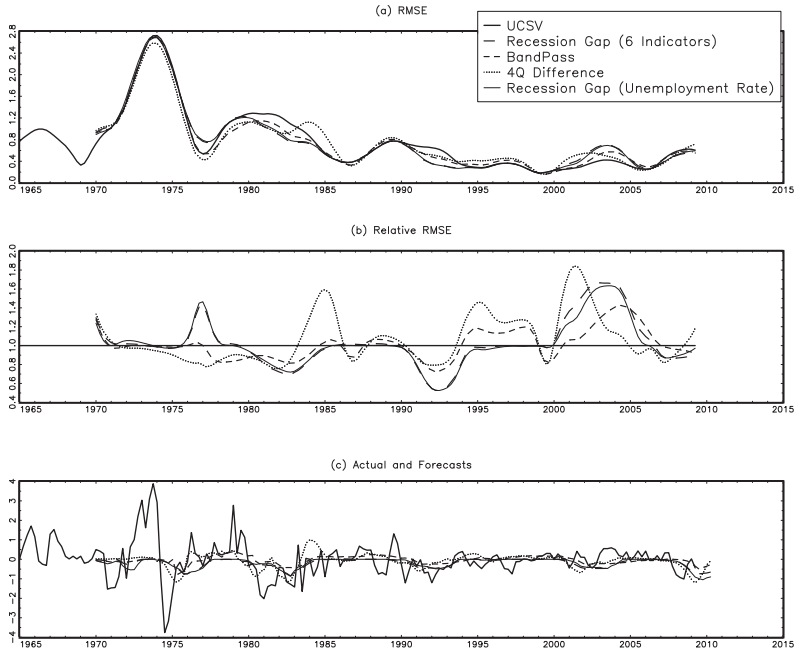

Panels are rolling root mean squared errors (top), rolling RMSEs relative to the UCSV model (middle), and recursive forecasts (bottom). Forecasts are median recession gap, median one-sided bandpass gap, median four-quarter difference, and the unemployment recession gap, where the median forecasts are across the six activity variables.

recession gap, using the full-sample estimate of γ_1 reported in Table 1. The simulation allows θ to vary across episodes by using the estimated value of θ_t at each episode's NBER peak date. We conduct the dynamic simulation by computing the value of π_t for the months over the recessionary episode plotted in Chart 2, given the path of unemployment.[9] Note that, except for initialization at the NBER peak, no actual values of inflation enter the simulation.

The dynamic simulation paths differ by episode both because the unemployment paths differ and because θ varies over time. When θ is large, there will be more inertia in trend inflation so that while a given value of x_t has a constant effect on one-quarter inflation, four-quarter inflation will fall by less than it would were θ smaller.

The dynamic simulation results, along with one standard error confidence bands, are presented in Chart 10. Two conclusions are evident. First, the predicted paths of inflation are similar to actual inflation in the 1960Q2, 1969Q4, 1980Q1, and 1990Q3 episodes. Second, the inflation path also is fairly close to its predicted value during the 2001Q1 episode through the peak of unemployment, but thereafter drifts upwards and away from the predicted continued disinflation. By 2004Q4, the dynamic simulation predicts the four-quarter inflation rate to have fallen since 2001Q1 by 0.6 percentage points, when in fact it rose by 0.5 percentage points. The standard error band for this episode is wide, but the increase in inflation falls outside that band.

IV. Other Predictors

This section examines the pseudo out-of-sample forecasting performance of other activity variables, activity variables augmented by survey expectations, and monetary variables. In many cases we focus on performance of the median forecast within a category (e.g. recession gap activity variables) both to streamline presentation and because of the well-known virtues of forecast pooling.

IV.A. Other Activity Variables

Table 2 summarizes the pseudo out-of-sample forecasting performance of six activity variables (the unemployment rate, the capacity

utilization rate, real GDP, the index of industrial production, employment, and the Chicago Fed National Activity Index [CFNAI]), each subject to three gap or changes transformations (recession gaps, one-sided bandpass gaps, and fourth differences).[10] Chart 11 plots the rolling RMSEs and forecasts of the median combined forecast, by gap transformation.

Table 2 and Chart 11 largely confirm the findings based on the analysis of the unemployment rate discussed in Section III.D. The forecasts based on the various activity variables tend to move together (for a given gap transformation). On average, the Phillips curve forecasts offer little improvement over the UC-SV benchmark, but they do offer improvements in recessionary episodes. The exception, again, is 2004, in which all the activity variable forecasts perform poorly relative to the UC-SV benchmark.

IV.B. Expectations

The models analyzed so far are variants of backwards-looking Phillips curves. Although we have interpreted τ_{tlt} as reflecting expectations (τ_{tlt} is the optimal long-term forecast of inflation from the UC-SV model), the empirical models do not explicitly incorporate forward-looking expectations. Expectations can be incorporated into Phillips curve forecasts either as model-based expectations (forecasts obtained using a model that includes a New Keynesian Phillips curve) or by using survey- or market-based expectations. Here, we consider the effect on the activity-based forecasts of Section IV.A. of adding survey expectations as an additional predictor in (1). Although market-based expectations are an appealing alternative approach, we do not use them here because of complications introduced by liquidity effects in the TIPS market.

We consider five real-time survey measures of inflation expectations: the Survey of Professional Forecasters (SPF) forecasts of GDP inflation one year ahead; the SPF forecast of CPI inflation one and 10 years ahead; and the University of Michigan survey forecast of inflation expectations one-year-ahead and five to 10 years ahead. Because these series are persistent, we analyze them as expectation gaps,

Table 3
Relative Root Mean Squared Error of Expectations-Augmented Activity Pseudo Out-of-Sample Forecasts of Four-Quarter Core PCE Inflation, Relative to Activity Variables Alone

Series	1970Q1 – 2010Q1	1970Q1 – 1979Q4	1980Q1 – 1989Q4	1990Q1 –1999Q4	2000Q1 –2010Q1
A. Recession gaps					
Unemployment	1.06 (113)	(0)	1.10 (35)	0.97 (40)	1.04 (38)
Capacity utilization	1.06 (113)	(0)	1.10 (35)	1.00 (40)	1.02 (38)
GDP	1.06 (113)	(0)	1.09 (35)	0.98 (40)	1.03 (38)
Industrial production	1.05 (113)	(0)	1.08 (35)	0.96 (40)	1.07 (38)
Employment	1.00 (113)	(0)	1.03 (35)	1.01 (40)	0.96 (38)
CFNAI	1.01 (94)	(0)	0.97 (16)	1.01 (40)	1.04 (38)
Median recession gap	1.05 (113)	(0)	1.08 (35)	0.99 (40)	1.02 (38)
B. One-sided bandpass gaps					
Unemployment	1.08 (113)	(0)	1.12 (35)	1.03 (40)	0.99 (38)
Capacity utilization	1.08 (113)	(0)	1.14 (35)	1.03 (40)	0.97 (38)
GDP	1.01 (113)	(0)	1.03 (35)	1.02 (40)	0.96 (38)
Industrial production	1.03 (113)	(0)	1.07 (35)	1.05 (40)	0.94 (38)
Employment	1.06 (113)	(0)	1.09 (35)	1.01 (40)	1.03 (38)
CFNAI	1.08 (113)	(0)	1.14 (35)	1.05 (40)	0.96 (38)
Median BP gap	1.06 (113)	(0)	1.10 (35)	1.04 (40)	0.98 (38)
C. Four-quarter differences					
Unemployment	1.02 (113)	(0)	1.08 (35)	0.95 (40)	0.87 (38)
Capacity utilization	0.96 (113)	(0)	1.00 (35)	0.93 (40)	0.82 (38)
GDP	1.04 (113)	(0)	1.11 (35)	0.97 (40)	0.86 (38)
Industrial production	1.03 (113)	(0)	1.09 (35)	0.97 (40)	0.88 (38)
Employment	1.02 (113)	(0)	1.15 (35)	1.02 (40)	0.78 (38)
CFNAI	0.94 (113)	(0)	1.04 (35)	0.93 (40)	0.69 (38)
Median Four-quarter difference	1.00 (113)	(0)	1.06 (35)	0.97 (40)	0.85 (38)
Overall median–all activity	1.04 (113)	(0)	1.09 (35)	1.01 (40)	0.96 (38)

Notes: Numbers of observations used in the computation are given in parentheses. The inflation expectations are SPF one-year CPI and GDP price index, SPF 10-year CPI, and University of Michigan one- and five-to-10-year inflation surveys.

Chart 12
Pseudo Out-of-Sample Forecasts of Four-Quarter Core PCE Inflation Using the Unemployment Recession Gap Augmented with Various Survey Measures of Inflation Expectations

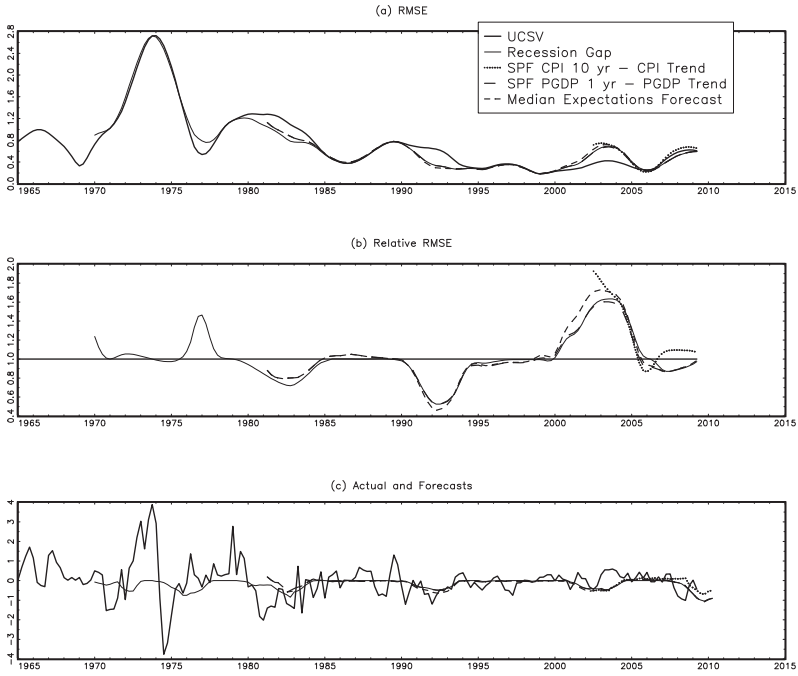

Panels are rolling root mean squared errors (top), rolling RMSEs relative to the UCSV model (middle), and recursive forecasts (bottom). Forecasts are the unemployment recession gap, not augmented, and augmented with: the SPF 10-year core CPI forecast, the SPF 1-year GDP price index forecast, and the median forecast using the five expectations measures in Table 3.

Table 4
Relative Root Mean Squared Error of Money-Based Pseudo Out-of-Sample Forecasts of Four-Quarter Core PCE Inflation, Relative to UC-SV Model

Series	1970Q1 – 2010Q1	1970Q1 – 1979 Q4	1980Q1 – 1989Q4	1990Q1 –1999Q4	2000Q1 –2010Q1
Forecasts of inflation over next four quarters					
UCSV (RMSE)	0.97 (158)	1.61 (40)	0.86 (40)	0.45 (40)	0.39 (38)
Monetary base (growth rate)	1.10 (155)	1.04 (37)	1.16 (40)	1.12 (40)	1.61 (38)
Monetary base (change in growth rate)	1.12 (154)	1.01 (36)	1.10 (40)	1.18 (40)	2.33 (38)
M2 (growth rate)	1.09 (155)	1.07 (37)	1.15 (40)	1.08 (40)	1.16 (38)
M2 (change in growth rate)	1.06 (154)	1.02 (36)	1.15 (40)	1.17 (40)	1.21 (38)
M3 (growth rate)	1.08 (141)	1.07 (37)	1.16 (40)	0.93 (40)	1.16 (24)
M3 (change in growth rate)	1.04 (140)	1.02 (36)	1.11 (40)	1.05 (40)	1.05 (24)
MZM (growth rate)	1.06 (155)	1.05 (37)	1.12 (40)	1.04 (40)	0.97 (38)
MZM (change in growth rate)	1.20 (154)	0.99 (36)	1.77 (40)	1.08 (40)	1.11 (38)
Median – four quarter ahead	1.05 (155)	1.03 (37)	1.10 (40)	1.05 (40)	1.04 (38)
Forecasts of inflation over next eight quarters					
UCSV (RMSE)	1.21 (154)	1.98 (40)	1.13 (40)	0.55 (40)	0.35 (34)
Monetary base (growth rate)	1.05 (147)	0.96 (33)	1.24 (40)	1.17 (40)	1.01 (34)
Monetary base (change in growth rate)	1.05 (146)	0.98 (32)	1.19 (40)	1.17 (40)	1.07 (34)
M2 (growth rate)	1.08 (147)	0.97 (33)	1.33 (40)	1.00 (40)	1.31 (34)
M2 (change in growth rate)	1.05 (146)	0.98 (32)	1.21 (40)	1.12 (40)	1.02 (34)
M3 (growth rate)	1.04 (137)	0.95 (33)	1.29 (40)	0.83 (40)	1.50 (24)
M3 (change in growth rate)	1.04 (136)	0.97 (32)	1.19 (40)	1.09 (40)	0.99 (24)
MZM (growth rate)	1.27 (147)	0.97 (33)	1.92 (40)	1.08 (40)	0.91 (34)
MZM (change in growth rate)	1.21 (146)	0.97 (32)	1.77 (40)	1.09 (40)	0.99 (34)
Median – eight quarter ahead	1.03 (147)	0.96 (33)	1.23 (40)	1.06 (40)	1.00 (34)

Notes: The first line in each block reports the standard deviation of the UC-SV forecast errors over the column sample period; the remaining lines report the ratio of the row forecast RMSE to the US-SV RMSE over the column sample. Numbers of observations used in the computation are given in parentheses. MZM denotes St. Louis Fed zero-maturity money.

Chart 13
Pseudo Out-of-Sample Forecasting of Four-Quarter Core PCE Inflation Using Monetary Variables

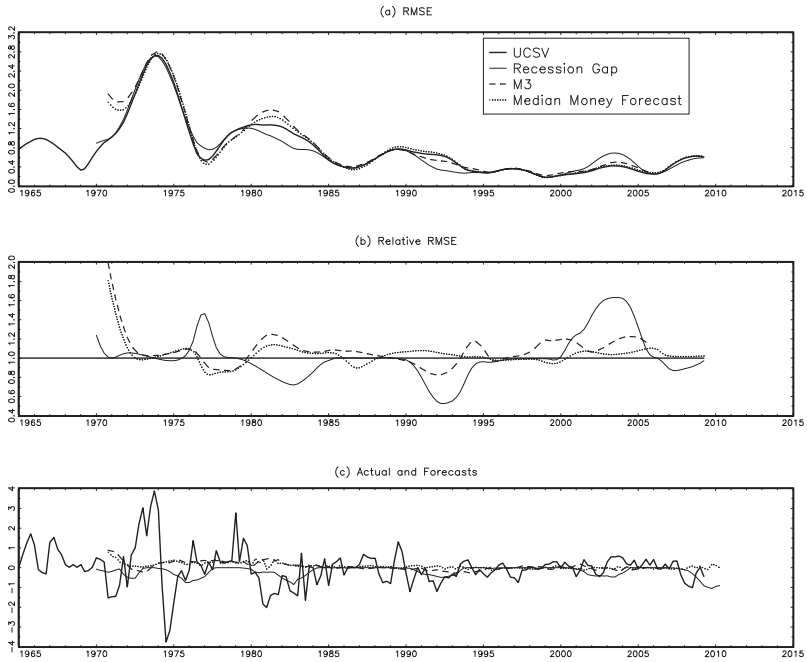

Forecasts are the unemployment recession gap (alone), and the forecast using the growth rate of M3 (alone), and the median forecast based on the eight measures of money in Table 4. Panels are rolling root mean squared errors (top), rolling RMSEs relative to the UCSV model (middle), and recursive forecasts (bottom).

that is, deviations from UC-SV trend CPI inflation (the SFP GDP inflation forecast is deviated from the GDP inflation trend).[11]

The results, presented as median combination forecasts across the various expectations measures, are summarized in Table 3. Chart 12 examines forecasts based on the unemployment recession gap augmented by individual survey forecasts, as well as the median survey-augmented unemployment recession gap forecast. The results in Chart 12 are striking and typical. Throughout almost all of the sample, the survey measures introduce negligible changes to the recession gap forecast.

IV.C. Monetary Aggregates

Table 4 and Chart 13 summarize the results of pseudo out-of-sample forecasts based on four monetary aggregates, with four transformations each. Unlike all previous models, these specifications include lagged values of x_t (here, money) as well as current values; lag lengths were chosen recursively by the Akaike Information Criterion. The recursive forecasts exhibit instability and have greater RMSEs by decade than the UC-SV model or the activity-based models. There are no episodes in which the monetary predictors outperform the UC-SV model. Because the coefficients are estimated to be small, the median combination forecasts are essentially the UC-SV forecast, with noise added. As the second part of Table 4 indicates, the qualitative findings are the same for two-year-ahead forecasts of inflation as for one-year-ahead forecasts. This negative assessment and indications of instability are consistent with the studies of monetary models of inflation at longer horizon by Sargent and Surico (2008) and Benati (2010). This is not to say that monetary expansions and inflation are unrelated, rather, the evidence here is that the predictive relationship between money and inflation is weak and unstable at short to medium horizons.

V. The 2007-2009 Recession

V.A. Energy and Housing

Before turning to the implications of this analysis for the current recession, we briefly consider the implications of energy and housing prices for core PCE inflation over the past several years.

Oil price pass-through. As discussed in the introduction, the volatility of oil prices since 2007 is an important reason that we have focused on core inflation in this paper. The question remains, however, about the extent to which energy price increases are passed through to core inflation. Hooker (2002) provides evidence that oil price increases led to increases in core inflation during the 1970s, but that after 1981 the extent of pass-through declined significantly. Hooker (2002) focused on the oil coefficients in triangle-type Phillips curve specifications, with a full-sample estimate of the NAIRU.

Table 5
Predicted Change in PCE Inflation Resulting From a 1 Percentage Point Increase in PCE-Energy Inflation, *q* Quarters Earlier

Entries in the first block are cumulative dynamic multipliers estimated over the indicated time period using the distributed lag regression,

$$\pi_t = \sum_{i=0}^{8} \beta_i \Delta\pi_{t-i}^{PCE-energy} + error_t$$

A. Pass-through to headline PCE

q	Cumulative dynamic multiplier after *q* quarters		
	1962Q1 – 1982Q4	**1983Q1 – 1994Q4**	**1995Q1 – 2006Q4**
0	0.075 (.008)	0.058 (.009)	0.048 (.003)
1	0.081 (.020)	0.057 (.014)	0.052 (.007)
2	0.144 (.025)	0.053 (.020)	0.052 (.011)
3	0.169 (.035)	0.064 (.022)	0.053 (.014)
4	0.189 (.039)	0.074 (.025)	0.056 (.015)
5	0.179 (.041)	0.083 (.030)	0.063 (.017)
6	0.213 (.042)	0.096 (.034)	0.065 (.019)
7	0.234 (.041)	0.110 (.034)	0.065 (.020)
8	0.199 (.045)	0.116 (.036)	0.059 (.022)
Oil share in PCE (final year of subsample)	**8.7%**	**5.2%**	**5.8%**

B. Pass-through to core PCE

q	Cumulative dynamic multiplier after *q* quarters		
	1962Q1 – 1982Q4	**1983Q1 – 1994Q4**	**1995Q1 – 2006Q4**
0	0.020 (.006)	-0.009 (.010)	0.000 (.004)
1	0.034 (.017)	-0.007 (.018)	0.004 (.007)
2	0.100 (.030)	-0.014 (.025)	0.003 (.011)
3	0.132 (.035)	-0.007 (.028)	0.000 (.014)
4	0.153 (.042)	0.009 (.032)	0.002 (.016)
5	0.155 (.044)	0.015 (.036)	0.007 (.017)
6	0.176 (.043)	0.028 (.040)	0.010 (.020)
7	0.199 (.045)	0.040 (.041)	0.007 (.022)
8	0.185 (.042)	0.045 (.046)	0.000 (.025)

Notes: Newey-West standard errors are in parentheses.

We reexamine the extent of energy price pass-through using a simpler specification than Hooker (2002), which does not involve a NAIRU assumption. Table 5 reports the cumulative coefficients in a distributed lag regression of quarterly inflation in headline PCE (panel A) and core PCE (panel B) on current and eight quarterly lagged values of PCE energy inflation. During the 1970s, the pass-through of energy prices to headline PCE was approximately twice energy's share. Unlike Hooker, we find that the effect of energy prices on headline inflation is twice energy's share during the 1980s and early 1990s, although the cumulative pass-through to core is not statistically significant. During the past 15 years, however, the pass-through of energy to headline inflation has occurred in the initial quarter and equals energy's (declining) share, and the dynamic pass-through to core is precisely estimated to be zero.[12] Thus, although the methods and samples are different, the results in Table 5 largely confirm Hooker's (2002) conclusion, although perhaps the reduction in pass-through occurred more gradually through the 1980s and early 1990s than Hooker (2002) estimates. Concerning the current recession, we therefore proceed to focus on core PCE inflation without special concern that the results are being distorted by energy prices, despite their recent large fluctuations.

Housing. Housing prices have fallen dramatically and, with a lag, so have the rents and owner-equivalent rents that enter PCE inflation. This raises the possibility that the collapse of housing prices, an important feature of this recession, might be driving measured declines in inflation. Hobijn, Eusepi, and Tambalotti (2010) examined the extent to which movements in the housing component of core PCE is exceptional over the past several years. They find that while the housing component has dropped, so have the other components of core PCE, and the differences between core PCE and core PCE excluding housing are negligible over 2008 and 2009. We therefore continue to focus on core PCE, with no special treatment of housing prices.

V.B. Forecasts and Dynamic Simulations

The dynamic simulation for the current recession is presented in Chart 14. The dynamic simulation uses the August 2010 SPF forecasted path of unemployment for quarters 2010Q3-2011Q3. Currently, the path

Chart 14
Dynamic Simulation of Four-Quarter Core PCE Inflation From 2007Q4 to 2011Q3 Computed Using the Unemployment Recession Gap Model

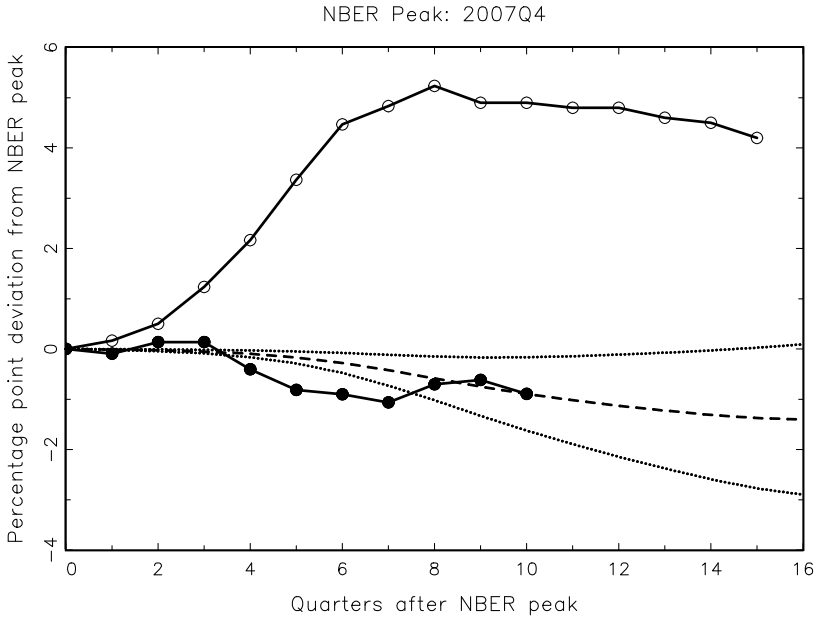

Unemployment values from 2010Q3 through 2011Q3 are SPF median forecasts. All series are plotted as percentage point deviations from their values at the NBER peak. Dashes are mean predicted values, dots are 90% confidence bands.

of inflation is on the conditional mean of the dynamic simulation, after initially dropping more sharply than the simulation path then increasing slightly. Based on the SPF forecasted path of unemployment, by 2011Q3 the four-quarter rate of core PCE inflation is expected to drop another 0.5 percentage points from its 2010Q2 value of 1.5%.

The four-quarter ahead forecasts using the estimated regression (1) and the unemployment recession gap (or the activity recession gaps) have generally tracked the downward movement of inflation over this recession, although the forecasts did not match the timing. The sharpest falls in inflation in this recession occurred from 2008Q3 to 2009Q3, and the pseudo out-of-sample forecasts of four-quarter inflation over this period (made four quarters prior to this decline) missed the decline and forecasted the decline to occur later because unemployment did not start to rise substantially until 2008Q3.

The projections based on the dynamic simulation are consistent with direct four-quarter ahead forecasts using the estimates of (1) reported in the previous sections. The four-quarter ahead unemployment recession gap model, both alone and the median expectations-augmented forecast, forecast a decline in the rate of four-quarter core PCE inflation of 0.8 percentage points from 2010Q2-2011Q2. The median four-quarter ahead forecast over all recession gap activity variables indicates a somewhat smaller decline, by 0.6 percentage points over this period.

We stress that there is considerable uncertainty surrounding these point estimates of further declines in inflation. The standard error bands in Chart 14 are consistent with declines that are considerably more modest or much more severe. One important source of uncertainty is how variable trend inflation currently is. For example, according to the dynamic simulation, if θ were to take on a value one standard error below its estimated value in 2007Q4, then the predicted decline in the rate of four-quarter PCE inflation from 2010Q2 to 2011Q3 would be 0.9 percentage points instead of 0.5 percentage points. The decline along the lower confidence band in Chart 14 (which reflects uncertainty in both γ_1 and θ) from 2010Q2 to 2011Q3 is steeper, 1.2 percentage points. The standard error bands in Chart 14 understate the uncertainty because they only incorporate estimation uncertainty and ignore future shock uncertainty.

The range of these declines in inflation is similar to that reported in Fuhrer and Olivei (2010) based on the entirely different and complementary approach of solving for model-based expectations with inflation determined by a New Keynesian Phillips Curve.

VI. Discussion and Conclusion

We have suggested that the empirical regularity of Chart 2—that U.S. recessions are associated with declines in inflation—can be captured by a simple model in which the deviation of core inflation from a long-run (statistical) trend is predicted by a new measure, the unemployment recession gap. Regressions using this gap measure appear to be stable based on standard statistical tests, although we point this out with trepidation because the history of the inflation forecast-

ing literature is one of apparently stable relationships falling apart upon publication. As the dynamic simulations in Charts 9 and 13 show, this model does a reasonable job of matching inflation dynamics given only the path of unemployment over a recessionary episode.

The results in this paper need to be understood in the context of three important sources of uncertainty. The first pertains to uncertainty within our econometric model. In that model, the long-term movement of inflation in response to a given short-term decline in activity depends on the volatility of the trend component of inflation: If the trend is resilient then inflation reverts to the trend, but if the trend is volatile then the trend tracks inflation. Both regimes are present in the record over the past 50 years. The volatility of trend inflation is currently at historically low levels, although at the very end of the sample it is inching up. An increase in that volatility, holding constant the volatility of the transitory component, would make the inflation path in Chart 14 steeper.

Second, making projections in the current recession requires extrapolating to rates of inflation at the edge of or outside the range of the data. There are some hints that the slope parameter γ_4 might be smaller in absolute value at low levels of inflation (Chart 8), but these hints are not robustly confirmed by statistical tests. Moreover, inflation dynamics could change in the region in which conventional monetary policy becomes ineffective and the parametric model could be ill-equipped to handle this.

Third, there is a key episode that does not match the historical regularity, the increase in inflation in late 2003 through 2004. This increase occurred despite the "jobless recovery" in which the unemployment rate lingered for quarters near its peak. Because the unemployment rate remained high, the recession gap model predicted falling inflation over 2004 when in fact the four-quarter rate of inflation increased by 0.7 percentage points from 1.5% in 2003Q4 to 2.2% in 2004Q4. As late as August 2003, the FOMC minutes record concern about continuing declines in inflation.[13] There are two leading explanations for this increase. The first is that the increase stemmed from special factors and price increases external to the U.S. economy, which in turn passed through to core inflation. This expla-

nation is the one put forward at contemporaneous FOMC meetings. According to the minutes from the spring and summer of 2004, the increase in core inflation during early 2004 was largely attributed to energy costs (which had risen sharply) and to a depreciation of the dollar.[14] The energy price part of this explanation, however, does not square with the econometric evidence that the pass-through from oil prices to core was zero on average from 1995 to 2006 (Table 5). Although housing prices were increasing sharply during 2004, the housing component of PCE did not start to increase substantially until the end of 2005, well after the unexplained rise in inflation in 2004.

The second leading explanation of 2004 is that this was an episode of successful expectations management by the Fed, in which market participants expected core inflation to increase (perhaps because of a perceived increase in the Fed's inflation target), and this expectation led to an actual increase in inflation. Interestingly, we can find no allusion to this mechanism in the FOMC minutes from the period. Throughout this episode, 10-year inflationary expectations remained steady, and incorporating inflationary expectations improves upon the unemployment recession gap forecasts for these years; however, we are reluctant to read too much into this improvement because including inflationary expectations produced worse forecasts on average for the decade and on average over the full sample. At the moment, we are sympathetic to the "special factors" explanation, in large part because that is the explanation offered by contemporary observers, including the FOMC, but ambiguity about this episode remains. Absent an explanation for the rise in inflation in 2004, we cannot rule out a similar fortuitous rise in the remaining quarters of the current episode, but neither can we be confident it will happen again.

Authors' Note: We thank Larry Ball, Ben Bernanke, Richard Berner, Roberto Billi, Jeffrey Fuhrer, Bob Gordon, Bart Hobijn, Peter Hooper, Michael Kiley, Mickey Levy, Emi Nakamura, Athanasios Orphanides, Glenn Rudebush, Frank Smets, Doug Staiger, and John Williams for helpful comments and suggestions.

Appendix A

Chart A-1
Empirical cumulative distributions of annual percentage price changes of 233 components of PCE inflation at five years between 1990 and 2007

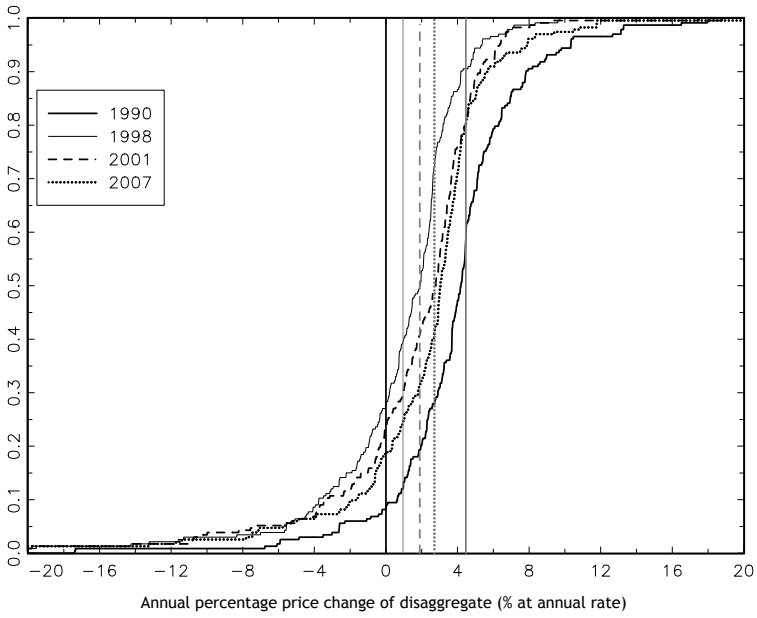

The vertical lines are monthly PCE (all) inflation rates for that month, for the date corresponding to the cumulative distributions. Discontinuities at zero indicate a "pile-up" of zero price changes. The infrequency of pile-ups at zero and the smooth shifting of the cdf through zero are consistent with the micro (individual-good) evidence in Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008) that price declines are commonplace.

Source: BEA Personal Income Web site.

Endnotes

[1]Throughout this paper we compute the four-quarter rate of price inflation as $\pi^4_t = 100\ln(P_t/P_{t-4})$, where P_t is the quarterly value of the price index. The one-quarter rate of inflation, at an annual rate, is computed as $\pi_t = 400\ln(P_t/P_{t-1})$, so $\pi^4_t = (\pi_t + \pi_{t-1} + \pi_{t-2} + \pi_{t-3})/4$. If the original price index is monthly, P_t is the average value of the price index for the months in the quarter.

[2]For example, in the 1960Q2 recession the quarterly unemployment rate rose by 1.8 percentage points from 1960Q2 to its peak in 1961Q2. Chart 2 thus plots $[u(s) - u_{1960Q2}]/1.8$, where the time scale s is set so that $s = 0$ is 1960Q2 and $s = 1$ is 1961Q2. The four-quarter rate of inflation is plotted in the same way, that is, as $[\pi^4(s) - \pi^4_{1960Q2}]/1.8$, on the same time scale as unemployment. When unemployment peaked in 1962Q2, four-quarter core PCE inflation had fallen by 0.7 percentage points, so the value plotted for inflation for this episode at $s = 1$ is -0.7/1.8 = -0.4.

[3]This pattern of inflation declines over recessions is robust to treating the 1980Q1 and 1981Q3 recessions separately instead of treating them as a single episode; for example, for the core PCE 4-quarter inflation decline at time scale 1.5, the mean decline and standard error are unchanged to two decimal points if these two recessions are treated separately. The mean declines are even robust to including the 1973Q4 recession, even though its special circumstances make it less relevant. With 1973Q4 included, at time 1.5 the mean decline in four-quarter core PCE is 0.43 (SE = 0.15), the mean decline in core CPI is 0.62 (SE = 0.30), in GDP price index inflation is 0.39 (SE = 0.23), in headline PCE is 0.66 (SE = 0.28), and in headline CPI is 0.93 (SE = 0.29).

[4]Let $\xi_j = Ea_{t+j}x_t/E$. From the expression for b_{t+4}, $\gamma_4 = [\xi_4 + (2-\theta)\xi_3 + (3-2\theta)\xi_2 + (4-3\theta)\xi_1]/4$.

[5]Our unemployment recession gap is related to two distinct literatures. The first concerns hysteresis in the Phillips curve and in the labor market (for a recent contribution and references, see Ball (2009)). If our recession gap is interpreted as a conventional gap, then the NAIRU is modeled as the running three-year minimum of the unemployment rate, which tends to drift down slowly in expansions and to adjust upwards after long contractions. We take no stand on such an interpretation and prefer to think of the unemployment recession gap as simply translating Chart 2 into a variable that can be studied using regressions. The second related literature is on nonlinearities in the univariate output process. Specifically, Beaudry and Koop (1993) propose a variable, the "current depth of recession," which is the difference between current real GNP and its historical maximum value. Our implementation differs from theirs in several ways. First, we consider the unemployment rate, not real GNP, and when we construct recession gaps using output measures we use local detrending (see Section IV.A.). Second, our unemployment rate minimum is for a three-year rolling window for instead all of history. Third,

Beaudry and Koop (1993) apply their measure to nonlinearities in the univariate output process, whereas our focus is models of inflation.

[6]The evidence for a piecewise nonlinear Phillips curve is stronger using a nonforecasting specification in which the unemployment gap dating overlaps with the dating of the dependent variable.

[7]The slope was estimated by local linear regression of (1) using a biweight kernel as a function of $\tau_{t|t} - \tau$, where τ appears on the horizontal axis of Chart 8, with a bandwidth of 1.3.

[8]Specifically, following Stock and Watson (2009), the rolling RMSE is computed as *rolling RMSE*(t) =

$$\left[\sum_{s=t-7}^{t+7} K(|s-t|)\left(\pi_{s+4}^{4} - \pi_{s+4|s}^{4}\right)^2 \Big/ \sum_{s=t-7}^{t+7} K(|s-t|)\right]^{1/2},$$

where K is the biweight kernel, $K(x) = (15/16)(1 - x2)^2\mathbf{1}(|x|\leq 1)$.

[9]This is equivalent to using a VAR to compute the response of inflation to a sequence of unemployment shocks chosen to match the episode-specific path of unemployment in chart 1, under the restriction that lagged inflation does not enter the unemployment equation, and using the nonlinear recession gap transformation to link the unemployment path and the inflation path. The restriction that lagged inflation does not enter the unemployment equation is not rejected at the 10% level.

[10]GDP, industrial production, employment, and the CFNAI were initially transformed by taking logarithms. The capacity utilization rate recession gap was computed as the negative of the deviation of the capacity utilization rate from its maximum value over the current and previous 11 quarters. The recession gaps for the remaining four series were computed by first computing a recursive locally detrended series, then setting the recession gap to be the negative of the deviation of the detrended series from its maximum value over the current and previous 11 quarters.

[11]Using the notation of (1), the regression estimated is $\pi_{t+4}^{4} - \tau_{t|t} = \gamma x_t + \delta\left(\pi_t^e - \tau_{t|t}^e\right) + e_{t+4}^e$, where $\tau_{t|t}^e$ is the trend used to detrend the inflation expectation π_t^e. Were the two trends the same so $\tau_{t|t} = \tau_{t|t}^e$, this regression would simplify to $\pi_{t+4}^{4} = \delta\pi_t^e + (1-\delta)\tau_{t|t} + \gamma x_t + e_{t+4}^e$, which is a Mincer-Zarnowitz forecast comparison regression comparing the survey forecast $\tau_{t|t}^e$ and the UC-SV model forecast $\tau_{t|t}$, augmented with x_t. Because π_{t+4}^{4} is core PCE inflation and the survey detrending uses either CPI or the GDP price index, the two trends are not the same, but this algebra suggests that the regression can still be given a Mincer-Zarnowitz forecast combination interpretation.

[12]The dynamic multipliers in the final column of Table 5B imply that the cumulative effect of energy price changes from 2007Q4 to 2010Q1 on core PCE inflation is a net reduction of core PCE inflation by 0.02 percentage points (2 basis points).

[13]"Committee members generally perceived the upside and downside risks to the attainment of sustainable growth for the next few quarters as roughly equal; however, they viewed the probability, though minor, of a substantial and unwelcome fall in inflation as exceeding that of a pickup in inflation from its already low level. On balance, the Committee believed that the concern about appreciable disinflation was likely to predominate for the foreseeable future." FOMC minutes, August 12, 2003. Billi (2009) argues that, in hindsight, the inflation "scare" of 2003 was overstated: Real-time estimates of core PCE inflation indicated inflation falling through the end of 2003 to under 1% per year, while the fully revised estimates now available show that core PCE inflation started to increase in mid-2003 and reached 2% by early 2004. Our interest here is not the confusion caused by these large data revisions, but the puzzling increase in core inflation in late 2003 and early 2004 in the fully revised data. Also see Dokko, et. al. (2009) for an analysis of real-time monetary policy during this period.

[14]Dokko, et al. (2009) document that the FOMC projection for 2004 inflation (headline PCE) in the *Monetary Report to Congress* at the start of 2004 was 2 percentage points below realized 2004 inflation, and states that this "miss is partly explained by an unexpected jump in the price of oil that year" (p. 14).

References

Aron, Janine and John Muellbauer (2010). "New Methods for Forecasting Inflation, Applied to the U.S.," manuscript, Nuffield College Oxford.

Akerlof, George A., William T. Dickens, and George L. Perry (1996). "The Macroeconomics of Low Inflation," *Brookings Papers on Economic Activity* 1996:1, 1-76.

Ang, A., G. Bekaert, and M. Wei (2007), "Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?" *Journal of Monetary Economics* 54, 1163-1212.

Atkeson, A. and L.E. Ohanian (2001), "Are Phillips Curves Useful for Forecasting Inflation?" *Federal Reserve Bank of Minneapolis Quarterly Review* 25(1):2-11.

Ball, Laurence, N. Gregory Mankiw, and David Romer (1988). "The New Keynesian Economics and the Output-Inflation Trade-off," *Brookings Papers on Economic Activity*, 1988:1, pp. 1-65.

Ball, Laurence M. (2009), "Hysteresis in Unemployment: Old and New Evidence." Ch. 7 in *Understanding Inflation and the Implications for Monetary Policy* (2009), Jeffrey Fuhrer, Yolanda Kodrzycki, Jane Little, and Giovanni Olivei (eds). Cambridge: MIT Press: 361-382.

Barnes, Michelle L. and Giovanni P. Olivei (2003). "Inside and Outside the Bounds: Threshold Estimates of the Phillips Curve," Federal Reserve Bank of Boston, *New England Economic Review*, 3-18.

Beaudry, Paul and Gary Koop (1993). "Do Recessions Permanently Change Output?" *Journal of Monetary Economics* 31, 149-163.

Benati, Luca (2010), "Long Run Evidence on Money Growth and Inflation," ECB DP 1027.

Bernanke, Ben S. (2003), "An Unwelcome Fall in Inflation?" speech, the Economics Roundtable, University of California – San Diego, July 23, 2003, at *http://www.federalreserve.gov/boarddocs/speeches/2003/20030723/.*

Billi, Roberto M. (2009). "Was Monetary Policy Optimal During Past Deflation Scares?" *Economic Review of the Federal Reserve Bank of Kansas City* 2009, 3rd Quarter, 67-98.

Brayton, Flint, John. M. Roberts, and John C. Williams (1999). "What's Happened to the Phillips Curve?" FEDS Discussion Paper 99-49.

Cogley, Timothy, Giorgio E. Primiceri, Thomas J. Sargent (2010). "Inflation-Gap Persistence in the U.S." *American Economic Journal: Macroeconomics* 2,43-69.

Cogley, Timothy and Thomas J. Sargent (2002). "Evolving Post-World War II U.S.

Inflation Dynamics," *NBER Macroeconomics Annual* 2001, MIT Press, Cambridge.

Cogley, Timothy and Thomas J. Sargent (2005). "Drifts and Volatilities: Monetary Policies and Outcomes in the Post World War II U.S." *Review of Economic Dynamics* 8, 262-302.

Cogley, Timothy and Argia Sbordone (2008). "Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Philips Curve." *American Economic Review* 98(5): 2101–2126.

Dokko, Jane, Brian Doyle, Michael T. Kiley, Jinill Kim, Shane Sherlund, Jae Sim, and Skander van der Heuvel (2009), "Monetary Policy and the Housing Bubble," FEDS 2009-49.

Fuhrer, Jeffrey C. and Giovanni P. Olivei (2010). "The Role of Expectations and Output in the Inflation Process: An Empirical Assessment," Public Policy Brief no. 10-2, Federal Reserve Bank of Boston.

Giannone, Domenico, Michele Lenza, Daphne Momferatou, and Luca Onorante (2010), "Short-Term Inflation Projections: A Bayesian Vector Autoregressive Approach," manuscript, European Central Bank.

Gordon, Robert J. (1982). "Inflation, Flexible Exchange Rates, and the Natural Rate of Unemployment." in *Workers, Jobs and Inflation*, Martin N. Baily (ed). Washington, D.C.: The Brookings Institution, 89-158.

Gordon, Robert J. (1990). "U.S. Inflation, Labor's Share, and the Natural Rate of Unemployment." *In Economics of Wage Determination* (Heinz Konig, ed.). Berlin: Springer-Verlag.

Gordon, Robert J. (1998). Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU," *Brookings Papers on Economic Activity* 1998:2, 297-333.

Hamilton, James D. (2009). "Causes and Consequences of the Oil Shock of 2007-08," *Brookings Papers on Economic Activity* 2009:1, 215-261.

Hobijn, Bart, Stefano Eusepi, and Andrea Tambalotti (2010). "The Housing Drag on Core Inflation." *Federal Reserve Bank of San Francisco Economic Letter* 2010-11, April 5. *http://www.frbsf.org/publications/economics/letter/2010/el2010-11.html.*

Hooker, Mark A. (2002). "Are Oil Shocks Inflationary? Asymmetric and Nonlinear Specifications versus Changes in Regime," *Journal of Money, Credit, and Banking* 34(2), 540-561.

Klenow, Peter J. and Oleksiy Kryvtsov (2008). "State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics* 123(3), 863-904.

Levin, Andrew and Jeremy Piger (2004). "Is Inflation Persistence Intrinsic in In-

dustrial Economies?" *ECB Working Paper* 334.

Liu, Zheng and Glenn D. Rudebusch (2010), "Inflation: Mind the Gap," Federal Reserve Bank of San Francisco, *Economic Letter* 2010-02.

Meier, André (2010), "Still Minding the Gap – Inflation Dynamics during Episodes of Persistent Large Output Gaps," IMF Working Paper WP/10/189.

Nakamura, Emi and Jón Steinsson (2008). "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics* 123(4), 1415-1464.

Roberts, John M. (2004), "Monetary Policy and Inflation Dynamics," FEDS Discussion Paper 2004-62, Federal Reserve Board.

Sargent, Thomas J. and Paolo Surico (2008), Monetary Policies and Low-Frequency Manifestations of the Quantity Theory," manuscript, NYU.

Staiger, Douglas O., James H. Stock, and Mark W. Watson (2001). "Prices, Wages and the U.S. NAIRU in the 1990s," Ch. 1 in *The Roaring Nineties*, A. Krueger and R. Solow (eds.), Russell Sage Foundation/The Century Fund: New York (2001), 3 – 60.

Stock, J.H., and M.W. Watson (2007), "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit, and Banking* 39, 3-34.

Stock, J.H., and M.W. Watson (2009), "Phillips Curve Inflation Forecasts," Ch. 3 in *Understanding Inflation and the Implications for Monetary Policy* (2009), Jeffrey Fuhrer, Yolanda Kodrzycki, Jane Little, and Giovanni Olivei (eds). Cambridge: MIT Press. (with M.W. Watson): 101-186.

Williams, John (2009), "The Risk of Deflation," Federal Reserve Bank of San Francisco, *Economic Letter* 2009-12.