No. 16
January 2012

StaffPAPERS
FEDERAL RESERVE BANK OF DALLAS

Inflation, Slack, and Fed Credibility

Evan F. Koenig
and
Tyler Atkinson

StaffPAPERS *is published by the Federal Reserve Bank of Dallas. The views expressed are those of the authors and should not be attributed to the Federal Reserve Bank of Dallas or the Federal Reserve System. Articles may be reprinted if the source is credited and the Federal Reserve Bank of Dallas is provided a copy of the publication or a URL of the website with the material. For permission to reprint or post an article, email the Public Affairs Department at dal.webmaster@dal.frb.org.*

Staff Papers *is available free of charge by writing the Public Affairs Department, Federal Reserve Bank of Dallas, P.O. Box 655906, Dallas, TX 75265-5906; by fax at 214-922-5268; or by phone at 214-922-5254. This publication is available on the Dallas Fed website, www.dallasfed.org.*

No. 16, January 2012

StaffPAPERS Federal Reserve Bank of Dallas

Inflation, Slack, and Fed Credibility

Evan F. Koenig
Vice President and Senior Policy Advisor
Federal Reserve Bank of Dallas

Tyler Atkinson
Senior Research Analyst
Federal Reserve Bank of Dallas

Abstract

It is generally agreed that slack has some impact on inflation. There is much less agreement on what form the relationship takes and whether it is stable enough to reliably help predict inflation. This analysis focuses on the Great Moderation period. We find that slack (as measured by the unemployment rate) and changes in slack are negatively correlated with changes in inflation and also deviations of inflation from long-forward inflation expectations. These relationships could have been exploited to produce forecasts of trimmed mean PCE inflation more accurate than rule-of-thumb forecasts. Forecasts of trimmed mean PCE inflation also serve well as predictions of GDP inflation and headline PCE inflation. Our analysis suggests that currently high levels of slack should hold inflation below two percent over 2012.

JEL codes: E31, E37
Keywords: Inflation, slack, forecasting

As of this writing (November 2011), the four-quarter headline PCE inflation rate is 2.9 percent, up sharply from 1.5 percent four quarters earlier. Does the increase signal the beginning of a potentially dangerous upward trend? Alternatively, is it likely to be reversed in coming quarters? Advocates of the latter view cite persistent real resource slack in the economy. The unemployment rate, for example, averaged 9.1 percent in third quarter 2011—over 3 percentage points above its long-run (20-year) average and down only modestly from the 9.6-percent rate recorded a year earlier. In this paper, we examine whether slack or changes in slack are systematically related to inflation and, more pertinently, whether slack or changes in slack are useful for *predicting* inflation.

The relationship between slack and inflation has long been a topic of interest to economists, of course. The original Phillips curve was a negative empirical relationship between the unemployment rate and (wage) inflation. When this relationship proved to be unstable, economists shifted to thinking of slack as impacting inflation measured relative to some baseline inflation rate. The baseline might be lagged inflation, for example, in which case an unemployment below normal would be associated with *rising* inflation, rather than with a high constant level of inflation. This paper starts by reviewing the empirical implications of several theoretical inflation models. Then it looks for an empirical relationship between the change in inflation and lagged slack, and between the deviation of inflation from trend inflation and lagged slack. The sample used runs from first quarter 1984 to first quarter 2011. It is dominated by the Great Moderation—a period of generally low and stable inflation during which naive forecasts are known to perform well. The final section undertakes a forecasting exercise that uses only data that would have been available when each forecast was prepared.

Our main conclusions are as follows. First, we find that slack does matter for inflation. However, discerning slack's effects is sometimes difficult because they can be obscured by quarter-to-quarter inflation fluctuations that are driven by other factors. Trimmed mean inflation measures remove many of these transitory influences, allowing the effects of slack to come to the fore. Second, while slack is an important influence on inflation, it is not the only important influence. Expected future monetary policy, as reflected in the public's perception of the Fed's long-term inflation objective, also has immediate and direct inflation effects. Expected future monetary policy determines the long-term trend in inflation, while slack and changes in slack help explain near-term deviations away from that trend. Third, even if it is the stability of headline inflation that ultimately matters to policymakers, adjusting policy in response to realized headline inflation is a mug's game. Because headline inflation has a large transitory component, reacting to it is like chasing a will-o'-the-wisp: You'll end up where you don't want to go. It's better to react to *forecasted* headline inflation or to trimmed mean inflation. Finally, forecasts of headline and trimmed mean inflation are essentially identical at the horizons relevant for policy. Indeed, forecasting equations fitted to trimmed mean inflation do a better job of predicting headline inflation than forecasting equations fitted directly to headline inflation. Intuitively, because noise is stripped away from the regression, forecasting equations fitted using trimmed mean inflation have more precise coefficient estimates than those fitted with headline inflation.

1. INFLATION MODELS

There is fairly general agreement that slack ought to matter for inflation: Inflation should tend to be low or to decrease when slack is high, once one controls appropriately for other influences. The question is, "Be low or decrease relative to what?" That is, "What is the appropriate baseline against which to compare current inflation?" The baseline for comparison is of considerable importance for how inflation responds to monetary policy and for how useful slack is, in practice, for inflation forecasting. The following is a quick rundown on some of the more important inflation models:

The NAIRU Model (a.k.a. the Accelerationist Phillips Curve)

In NAIRU (non-accelerating inflation rate of unemployment) inflation models, current inflation is compared with lagged inflation: High slack means one should expect falling inflation, absent (mostly transitory) cost-push shocks (Friedman 1968). NAIRU models once dominated both the theoretical and the empirical inflation literatures, and NAIRU thinking continues to influence policy discussions today. A key implication of the model is that changes in the expected future conduct of monetary policy have no current effects on inflation except insofar as they impact current slack. A corollary is that if you want to bring inflation down, you must be willing to put up with high unemployment for a time. NAIRU models appear to perform pretty well, empirically, but that may be partly because measures of slack are revised, ex post, to fit the model. Considerable effort is devoted to inferring movements in the NAIRU from the observed behavior of inflation. (Inflation didn't fall as expected? The NAIRU must have increased [Gordon 1997].) NAIRU models have the (to many, unrealistic) implication that the Fed can keep the unemployment rate low for however long it is willing to tolerate rising inflation.

New-Keynesian Phillips Curve (a.k.a. the Calvo Pricing Model)

In New-Keynesian inflation models, the baseline is expected future inflation rather than lagged inflation. The New-Keynesian Phillips Curve is now the dominant theoretical model. New-Keynesian firms adjust their prices the Dos Equis' way: "I don't always change my price, but when I do, I change it with an eye toward the future conduct of monetary policy" (Calvo 1983). The result is that inflation is a "jumping variable": Changes in the expected future conduct of policy have immediate effects on inflation. According to some, the predicted effects are so large as to be difficult to reconcile with the data (Fuhrer and Moore 1995). Because the inflation impact of policy changes is front loaded, New-Keynesian models say that one ought to expect inflation to rise over time when the unemployment rate is high.

Hybrid (NAIRU and New-Keynesian) Phillips Curve

The baseline for the hybrid model is a (typically 50–50) weighted average of lagged actual and expected future inflation rates (Gali and Gertler 1999). The hybrid model is more appealing than the plain-vanilla NAIRU model because inflation is at least somewhat forward looking, and does better empirically than the plain-vanilla New Keynesian model because the presence of lagged inflation in the baseline gives the inflation process some inertia. Attempts to justify the hybrid model involve some hand waving and talk of "rule-of-thumb" price setters. It is the dominant empirical model these days. A variant compares current inflation with a weighted average of expected future inflation and some measure of trend inflation (rather than lagged inflation).

P-Bar Inflation Model

In Bennett McCallum's P-bar model, the baseline against which inflation is compared is the flexible-price or market-clearing inflation rate—that is, the inflation rate you'd see in an otherwise identical economy without price rigidities (McCallum 1994). Roughly, the trend in inflation is determined by long-term prospective growth in the money supply relative to long-term prospective growth in potential output, while deviations around that trend are determined by slack. We are not aware that this approach has ever been subject to careful empirical evaluation (probably because the market-clearing inflation rate is not directly observable).

Fischerian or Sticky-Information Phillips Curve

Under this framework, price changes, per se, are costless and continuous. It's re-optimizing planned price paths that's subject to frictions (Koenig 1996 and Mankiw and Reis 2002). Here, the inflation baseline depends on the time horizon. Near-term inflation forecasts are tied to lagged expectations of current inflation (which, in practice, usually means they are tied to lagged trend inflation). At longer horizons, though, predicted inflation is tied to today's expectation of future market-clearing inflation, much as in the P-bar model. So, the response of near-term inflation forecasts to a shock is primarily influenced by slack, but the longer-horizon predicted response is directly influenced by policy expectations. Intuitively, inflation becomes more and more sensitive to changes in anticipated policy as new information about future policy diffuses to a larger and larger fraction of firms. This alternative to the hybrid Phillips curve has not received a great deal of attention—probably because empirical implementation is even more complicated than for the P-bar model.

Atkeson and Ohanian

Atkeson and Ohanian didn't put forward a theory of inflation, but in an influential and controversial article, they pointed out that during the Great Moderation period it has, in practice, proven to be very difficult to predict inflation changes using measures of slack (Atkeson and Ohanian 2001). Since Atkeson and Ohanian published their findings, a key question when evaluating any inflation forecasting model is "Can it beat lagged inflation?"

2. EMPIRICAL RESULTS: IS THERE EVIDENCE THAT SLACK HELPS EXPLAIN INFLATION?

As noted above, inflation models typically assume that deviations of inflation from some baseline are related to slack, but different models use different baselines. Is there evidence that slack is, in fact, important for understanding inflation? We look at whether deviations of inflation from lagged inflation are explained by slack and at whether deviations of inflation from long-run trend inflation are explained by slack. Both GDP and trimmed mean PCE inflation measures are considered. (For precise definitions of the variables used in this paper, see Table 1.) These series have broad coverage yet are less affected by transitory cost-push shocks than headline CPI or headline PCE inflation.[1] We use the unemployment rate to measure slack. Unlike the output gap, the unemployment rate is directly observed and is not subject to revision.[2] The analysis is entirely "in sample," using latest-available inflation data. (A later out-of-sample forecasting exercise uses real-time inflation data to the extent possible.) The sample covers the Great Moderation period over which inflation is difficult to predict.

Table 1: Variable Definitions

Variable	Definition
GDP inflation	$100*(\frac{P_t}{P_{t-4}} - 1)$, where P is the GDP chain-type price index
PCE inflation	$100*(\frac{P_t}{P_{t-4}} - 1)$, where P is the PCE chain-type price index
Trimmed mean PCE inflation	Quarterly average of trimmed mean PCE 12-month inflation rate
Unemployment rate	Quarterly average of civilian unemployment rate
9-year, 1-year-forward expected CPI	$\frac{10*cpi10y - cpi1y}{9}$, where $cpi10y$ and $cpi1y$ are 10-year and 1-year CPI inflation expectations from the Survey of Professional Forecasters

Changes in Inflation

In NAIRU models, slack impacts changes in the inflation rate. Accordingly, we estimated equations of the form:

$$\pi_t - \pi_{t-4} = \gamma(u_{t-4} - NAIRU) + \delta(u_{t-4} - u_{t-5}) \quad \text{(A)}$$

where π is the (end-of-sample-vintage estimate of the) inflation measure of interest and u is the unemployment rate. Results are displayed in Table 2. The lesson from Table 2 is that changes in GDP inflation are not well explained by the rate of unemployment. Instead, changes in the unemployment rate appear to be helpful for understanding movements in GDP inflation. For trimmed mean PCE inflation, the level and the change in the unemployment rate both matter. Each 1 percentage point increase in the unemployment rate lowers subsequent GDP inflation by about 0.8 percentage points and subsequent trimmed mean PCE inflation by about 0.7 percentage points. In addition, for each percentage point that unemployment exceeds the NAIRU, subsequent trimmed mean PCE inflation falls by 11 basis points. In 2010, GDP inflation rose from its recessionary lows while trimmed mean PCE continued to drift down. This behavior is consistent

[1] Trimmed mean PCE inflation strips out the prices of those items that increased or decreased the most in each month. The percent trimmed is calibrated to best capture the medium-term trend in inflation (Dolmas 2005). Trimmed mean PCE inflation is published each month by the Federal Reserve Bank of Dallas.

[2] Changes in demographics, unemployment insurance and other factors potentially affect the equilibrium rate of unemployment. We will ignore such variation. Consequently, our empirical results may understate the strength of the connections between slack and inflation.

with the results in Table 2, as the level of slack held down trimmed mean inflation and the change in slack boosted GDP inflation.

The NAIRU implied by equation A using trimmed mean PCE inflation is 4.9 percent, with a standard error of 0.5 percentage points. The GDP inflation equation does not yield a useful NAIRU estimate.

Figures 1 and 2 show simple scatter plots of the unemployment rate lagged four quarters with four-quarter changes in GDP inflation and trimmed mean PCE inflation, respectively. The correlations are –0.04 and –0.41 from first quarter 1984 to first quarter 2011. (The correlations strengthen to –0.35 and –0.65 when the unemployment series is not lagged.) It is striking how little information is conveyed by the unemployment rate on whether inflation is likely to rise or to fall over the next four quarters. It is not uncommon for inflation to rise when the unemployment rate exceeds 8 percent or to fall when the unemployment rate is below 4.5 percent.

Deviations of Inflation from Trend

In P-bar and Fischerian (sticky-information) models, slack impacts inflation by temporarily pushing it away from its long-run trend. Equation B relates deviations of inflation from trend to slack and lagged deviations of inflation from trend:

$$\pi_t - \pi_t^{LF} = \beta(\pi_{t-4} - \pi_{t-4}^{LF}) + \gamma(u_{t-4} - NR). \qquad \text{(B)}$$

Here, NR is the equilibrium or "natural" rate of unemployment and π^{LF} is the long-run trend measured using the nine-year, one-year-forward inflation expectation from the Survey of Professional Forecasters (SPF). This expected inflation measure is not published by the Federal Reserve Bank of Philadelphia, but calculated from the ten-year and one-year CPI inflation forecasts. It is the inflation rate expected over the nine years starting one year from the present.[3] It is meant to capture what the public perceives to be policymakers' long-term inflation goal, after short-term influences wash out. The long-forward expected inflation measure is for CPI inflation, and CPI inflation tends to run above GDP inflation and trimmed mean PCE inflation. To prevent these differences from biasing natural-rate estimates, the sample average of the difference between CPI inflation and the inflation measure of interest is subtracted from long-forward CPI inflation expectations when calculating π^{LF}. Specifically, π^{LF} equals long-forward CPI expected inflation less 0.5 percentage points when π is GDP inflation and equals long-forward CPI expected inflation less 0.3 percentage points when π is trimmed mean PCE inflation.

According to Table 3, slack is useful for explaining subsequent deviations of GDP inflation and trimmed mean PCE inflation from their long-run trends. Inflation tends to be low relative to long-forward expected inflation when the unemployment rate is high and to be high relative to long-forward expected inflation when the unemployment rate is low. (For both measures of inflation, the change in unemployment is insignificant when included in the regression.) The implied natural rates, 5.10 and 5.16, are close to the NAIRU implied by Equation A. For each 1 percentage point that unemployment exceeds the natural rate, GDP inflation runs 0.14 percentage points below trend in the short run (over the following four quarters) and runs 0.28 percentage points below trend in the long run. For each 1 percentage point that unemployment exceeds its natural rate, trimmed mean PCE inflation subsequently runs 0.27 percentage points below trend.

[3] Until 1992, the ten-year CPI inflation forecasts were only collected in the first and fourth quarter. The second and third quarter nine-year, one-year-forward expectations up until 1992 are estimated with a regression on the three-year, seven-year-forward Treasury yield: $cpi9y1yfwd_t = 0.22 + 0.13 * r3y7yfwd_t + 0.88 * (cpi9y1yfwd_{t-1} - 0.13 * r3y7yfwd_{t-1})$ where $r3y7yfwd_t$ is the three-year, seven-year-forward Treasury yield.

Table 2: Are Changes in Inflation Related to Slack?

$$\pi_t - \pi_{t-4} = \gamma(u_{t-4} - NAIRU) + \delta(u_{t-4} - u_{t-5})$$

Sample: 1984:Q1–2011:Q1

Inflation Measure	γ	δ	$NAIRU$	**Adj. R^2/S.E.**
GDP inflation	–0.0194	–	0.8618	–0.008
	(0.0686)		(18.6362)	0.683
GDP inflation	–0.0158	–0.7917**	–0.7642	0.116
	(0.0855)	(0.2765)	(35.6360)	0.639
GDP inflation	–	–0.7831*	–	0.107
		(0.3152)		0.643

Comments: It is not the level but the change in the unemployment rate that is helpful for predicting changes in GDP inflation. A 1-percentage point increase in the unemployment rate reduces the forecasted change in GDP inflation by 0.8 percentage points.

Inflation Measure	γ	δ	$NAIRU$	**Adj. R^2/S.E.**
Trimmed mean	–0.1108**	–	4.9826**	0.160
PCE inflation	(0.0336)		(0.5903)	0.361
Trimmed mean	–0.1076**	–0.7043**	4.8923**	0.468
PCE inflation	(0.0221)	(0.0831)	(0.5010)	0.287

Comments: Both the level and the change in the unemployment rate are helpful for predicting changes in trimmed mean PCE inflation. For each 1 percentage point that unemployment exceeds the NAIRU, forecasted inflation falls by 0.1 percentage points. Also, each 1-percentage-point increase in the unemployment rate reduces the forecasted change in inflation by 0.7 percentage points.

Standard errors of the estimated coefficients, in parentheses, are Newey-West corrected.

* Significant at 5-percent level.

** Significant at 1-percent level.

Table 3: Are Deviations of Inflation from Trend Related to Slack?

$$\pi_t - \pi_t^{LF} = \beta(\pi_{t-4} - \pi_{t-4}^{LF}) + \gamma(u_{t-4} - NR)$$

Sample: 1984:Q1–2011:Q1

Inflation Measure	β	γ	**NR**	**Adj. R^2/S.E.**
GDP inflation	–	–0.2811**	5.0990**	0.264
		(0.0648)	(0.5739)	0.674
GDP inflation	0.5058**	–0.1375*	5.1040**	0.448
	(0.1545)	(0.0647)	(0.8229)	0.584

Comments: The estimates indicate that deviations in inflation from long-run trend inflation are systematically related to the unemployment rate. For each 1 percentage point that unemployment exceeds its natural rate, inflation runs 0.14 percentage points below trend in the short run and hypothetically runs 0.28 percentage points below trend in the long run. The change in the unemployment rate is insignificant when included in the regression.

Inflation Measure	β	γ	**NR**	**Adj. R^2/S.E.**
Trimmed mean	–	–0.2657**	5.0710**	0.566
PCE inflation		(0.0333)	(0.3049)	0.340
Trimmed mean	0.2652	–0.2187**	5.1603**	0.608
PCE inflation	(0.1340)	(0.0341)	(0.3023)	0.322

Comments: Here, again, deviations in inflation from long-run trend inflation are systematically related to the unemployment rate. For each 1 percentage point that unemployment exceeds its natural rate, inflation runs 0.27 percentage points below trend. The change in the unemployment rate is insignificant when included in the regression.

Standard errors of the estimated coefficients, in parentheses, are Newey-West corrected.

* Significant at 5-percent level.

** Significant at 1-percent level.

Figures 3 and 4 show the connection between the four-quarter lagged level of the unemployment rate and deviations of GDP inflation and trimmed mean PCE inflation, respectively, from SPF long-forward inflation expectations (corrected for sample-average differences with CPI inflation). Slack has a tighter relationship with trimmed mean PCE inflation than GDP inflation, but both have a clear, negative relationship. (The correlations in the two charts are –0.52 and –0.75.) Notably, an unemployment rate above 6.1 percent virtually guarantees that inflation will fall short of trend over the coming year.

Figure 1: The Unemployment Rate Is Uncorrelated with Subsequent Changes in GDP Inflation

SOURCES: Bureau of Economic Analysis; Bureau of Labor Statistics.

Figure 2: The Unemployment Rate Is Negatively Correlated with Subsequent Changes in Trimmed Mean PCE Inflation

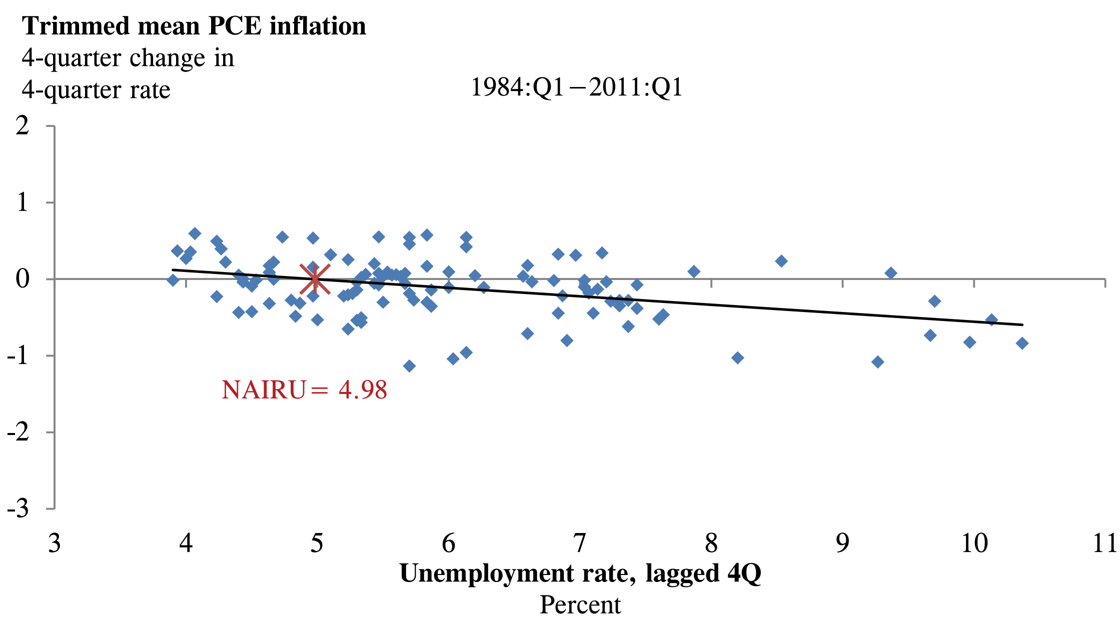

SOURCES: Federal Reserve Bank of Dallas; Bureau of Labor Statistics.

Figure 3: The Unemployment Rate Is Negatively Correlated with Detrended GDP Inflation

*GDP inflation − SPF 9-year, 1-year-forward CPI inflation expectations + 0.5.
SOURCES: Bureau of Economic Analysis; Bureau of Labor Statistics.

Figure 4: The Unemployment Rate Is Negatively Correlated with Detrended Trimmed Mean PCE Inflation

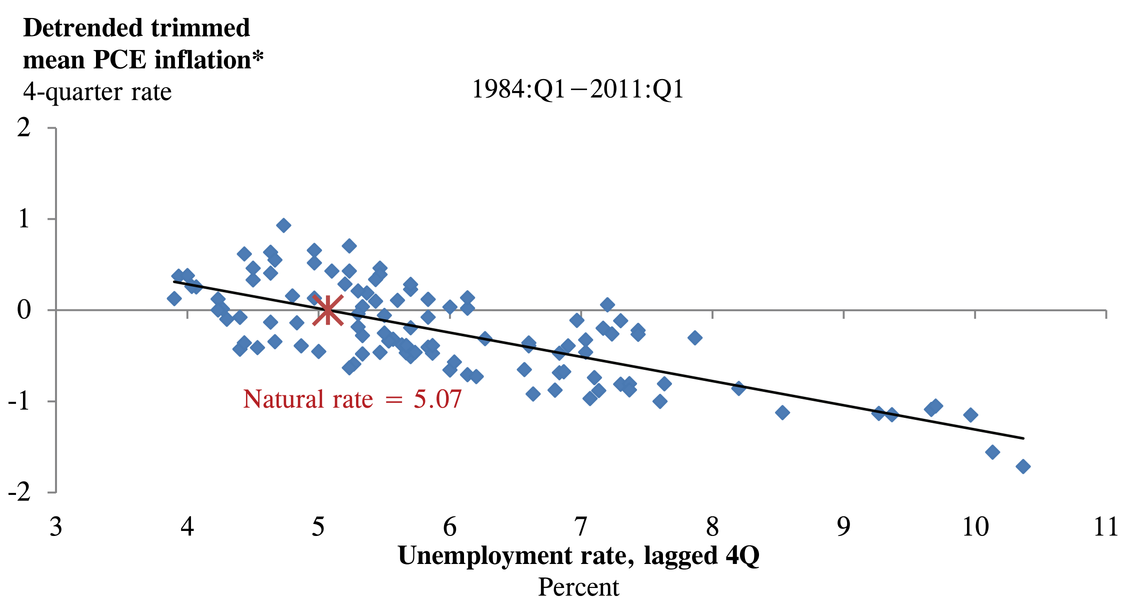

*Trimmed mean PCE inflation − SPF 9-year, 1-year-forward CPI inflation expectations + 0.3.
SOURCES: Federal Reserve Bank of Dallas; Bureau of Labor Statistics.

3. IS THERE EVIDENCE THAT SLACK HELPS FORECAST INFLATION?

An in-sample statistical relationship between inflation and slack as these variables appear today does not necessarily mean that slack helps to forecast inflation in real time, for three reasons. First, slack's influence may be statistically significant but practically unimportant. Second, relationships estimated using today's data can be very different from those estimated using data that would have been available at the time—which are the relationships that matter for real-time forecasting. Third, the links between slack and inflation may be unstable over time. To determine whether or not slack is likely to be of practical help in predicting inflation, we estimate two different Phillips-curve-style forecasting equations recursively, using only information that would have been available at the time, and we compare the resultant real-time inflation predictions to simple rule-of-thumb alternatives, in the spirit of Atkeson and Ohanian (2001). Results confirm that slack and changes in slack are useful for predicting trimmed mean PCE inflation measured either as a deviation from lagged trimmed mean inflation or as a deviation from the public's long-forward inflation expectations. Interestingly, there's no need to develop or estimate separate forecasting equations for GDP inflation or headline PCE inflation: The best forecasts of these inflation measures come from the trimmed mean PCE forecast equation. Since slack and changes in slack help predict trimmed mean inflation, they are indirectly also helpful for predicting headline inflation. The Phillips-curve specification that measures trimmed mean inflation as a deviation from long-forward inflation expectations performs particularly well, suggesting that current inflation adjusts quickly and one-for-one with changes in the public's perception of the Fed's long-run inflation objective. From 1998 onward, this perceived objective has been firmly anchored at 2.5 percent, as measured by CPI inflation.

Background

We look at the real-time performance of two different inflation forecasting equations for each of three different measures of inflation over the past ten years. The three inflation measures are GDP inflation, headline PCE inflation, and trimmed mean PCE inflation. GDP and headline PCE inflation are both of direct interest to policymakers. (The original Taylor rule is based on GDP inflation. The Federal Open Market Committee has decided to define long-term price stability in terms of headline PCE inflation.) Even though it may not be of direct policy interest, the fact that trimmed mean inflation strips high-frequency noise out of headline inflation makes it potentially useful as an indicator of headline trends and as a forecasting tool. Every effort is made here to restrict the information used for forecasting to that which would have been available in real time. For example, slack is measured using unfiltered unemployment-rate data: We do not try to construct or obtain real-time, time-varying NAIRU or natural rate estimates.[4] Any lagged inflation rates that appear on the right-hand side of a forecasting equation are second-release estimates whenever possible. (Unfortunately, real-time estimates of trimmed mean PCE inflation don't begin until 2005.) Left-hand-side inflation rates are always the latest vintage that would have been available at the time of the estimation. Each sample begins in first quarter 1984—the beginning of the Great Moderation period. The first sample ends in first quarter 2000, and the estimated equation is used to forecast inflation over the four quarters ending first quarter 2001. The sample is then extended by one quarter, coefficients are updated, and a new forecast prepared. The final inflation observation covers the four quarters ending first quarter 2011.

Forecasting Equation #1

In our first forecasting equation, the baseline with which inflation is compared is a weighted average of lagged inflation, lagged SPF long-forward inflation expectations, and lagged trend inflation as captured by the trimmed mean PCE inflation rate:

$$\pi_t = \alpha + \beta_1 \pi_{t-4}^{RT} + \beta_2 \pi_{t-4}^{LF} + \beta_3 \pi_{t-4}^{TM} + \gamma u_{t-4} + \delta(u_{t-4} - u_{t-5}), \tag{1}$$

where π is the (end-of-sample-vintage estimate of the) inflation measure of interest, π^{RT} is the second-release "real-time" estimate of the inflation measure of interest, π^{LF} is nine-year, one-year "long-forward" expected CPI inflation from the SPF, π^{TM} is the second-release estimate of "trimmed mean" PCE inflation, and u is the unemployment rate. The beta coefficients are constrained to sum to 1. For trimmed mean inflation, π^{RT} and π^{TM} are exactly the same thing, so we set $\beta_1 = 0$. For all three inflation measures, the hypothesis that $\beta_1 = \beta_2 = 0$ cannot be rejected at standard significance levels. So, in practice, Equation 1 reduces to:

$$\pi_t - \pi_{t-4}^{TM} = \alpha + \gamma u_{t-4} + \delta(u_{t-4} - u_{t-5}). \tag{1'}$$

[4] We looked at the Stock-Watson "unemployment recession gap" measure of slack but did not find it to be useful. See Stock and Watson (2010).

The deviation of inflation from lagged trimmed mean PCE inflation depends on the lagged unemployment rate and the lagged change in the unemployment rate. Equation $1'$ is a NAIRU-style model: Changes in the expected future conduct of monetary policy have no impact on inflation except through the unemployment rate.

Estimates of Equation $1'$ are presented in Table 4A for each of three measures of inflation over each of two sample periods. Results for trimmed mean inflation are clear cut. They show that trimmed mean inflation is lower by 0.1 percentage points for each percentage point that the unemployment rate exceeds its mean and is lower by approximately 0.8 percentage points for each percentage point that the unemployment rate increases. GDP-inflation and headline-PCE-inflation responses are quite similar, but estimated less precisely. Still, the hypothesis that the equations all have the same coefficients is rejected.

Forecasting Equation #2

Our second forecasting equation is in the spirit of the P-bar inflation model or the New-Keynesian Phillips curve, in that it builds in the potential for forward-looking price-setting behavior. Here, the baseline against which current inflation is compared is current long-forward inflation expectations. In its most general form:

$$\pi_t - \pi_t^{LF} = \alpha + \beta_1(\pi_{t-4}^{RT} - \pi_{t-4}^{LF}) + \beta_2(\pi_{t-4}^{TM} - \pi_{t-4}^{LF}) + \gamma u_{t-4} + \delta(u_{t-4} - u_{t-5}), \quad (2)$$

where all variables are defined as in Equation 1, but coefficients may take on different values. In particular, the beta coefficients are not constrained to sum to 1.

For trimmed mean inflation, π^{RT} and π^{TM} are exactly the same, so we set $\beta_1 = 0$. For the other inflation measures, the hypothesis that $\beta_1 = 0$ cannot be rejected at standard significance levels. So, in practice, Equation 2 reduces to:

$$\pi_t - \pi_t^{LF} = \alpha + \beta(\pi_{t-4}^{TM} - \pi_{t-4}^{LF}) + \gamma u_{t-4} + \delta(u_{t-4} - u_{t-5}). \quad (2')$$

The deviation of inflation from long-forward expected inflation depends on the lagged deviation of trimmed mean inflation from long-forward inflation expectations, the unemployment rate, and the lagged change in the unemployment rate. *Note that according to Equation $2'$ any change in the FOMC's perceived inflation objective has an immediate one-for-one impact on current inflation, independent of the amount of slack.* Estimates of Equation $2'$ are presented in Table 4B for each of three measures of inflation, over the same samples as in Table 4A. The results show that trimmed mean inflation is lower by about 0.2 percentage points for each percentage point that the unemployment rate exceeds its mean, and is lower by approximately 0.4 percentage points for each percentage point that the unemployment rate increases. GDP inflation results are very similar, and coefficient estimates are about equally precise. Precision suffers in the headline PCE equation, but coefficient estimates are in the same ballpark. Again, though, the hypothesis that the equations are identical is rejected.

Table 4A: Real-Time Inflation Forecasting Equations: Equation 1′

$$\pi_t - \pi^{TM}_{t-4} = \alpha + \gamma u_{t-4} + \delta(u_{t-4} - u_{t-5})$$

Inflation Measure	α	γ	δ	Adj. R^2/S.E.
		1984:Q1–2000:Q1		
Trimmed mean	0.411	–0.097**	–0.822**	0.318
PCE inflation	(0.230)	(0.037)	(0.196)	0.287
GDP inflation	0.039	–0.082	–0.745**	0.165
	(0.352)	(0.055)	(0.186)	0.378
PCE inflation	0.401	–0.092	–0.388	0.015
	(0.586)	(0.085)	(0.290)	0.596
		1984:Q1–2011:Q1		
Trimmed mean	0.627**	–0.110**	–0.795**	0.480
PCE inflation	(0.159)	(0.024)	(0.098)	0.304
GDP inflation	0.502	–0.124	–1.097**	0.243
	(0.426)	(0.066)	(0.228)	0.663
PCE inflation	0.263	–0.062	–0.811*	0.080
	(0.517)	(0.076)	(0.348)	0.823

Standard errors of the estimated coefficients, in parentheses, are Newey-West corrected.
* Significant at 5-percent level.
** Significant at 1-percent level.

Table 4B: Real-Time Inflation Forecasting Equations: Equation 2′

$$\pi_t - \pi_t^{LF} = \alpha + \beta(\pi_{t-4}^{TM} - \pi_{t-4}^{LF}) + \gamma u_{t-4} + \delta(u_{t-4} - u_{t-5})$$

Inflation Measure	α	β	γ	δ	Adj. R^2/S.E.
			1984:Q1–2000:Q1		
Trimmed mean	0.687	0.401*	–0.184**	–0.381*	0.503
PCE inflation	(0.359)	(0.189)	(0.061)	(0.178)	0.325
GDP inflation	0.303	0.548*	–0.150*	–0.441*	0.435
	(0.384)	(0.217)	(0.057)	(0.186)	0.387
PCE inflation	0.639	0.882**	–0.118	–0.396	0.388
	(0.591)	(0.294)	(0.091)	(0.310)	0.566
			1984:Q1–2011:Q1		
Trimmed mean	0.866**	0.403*	–0.194**	–0.402**	0.622
PCE inflation	(0.231)	(0.168)	(0.046)	(0.152)	0.316
GDP inflation	0.648	0.604*	–0.172*	–0.871**	0.336
	(0.447)	(0.231)	(0.079)	(0.284)	0.641
PCE inflation	0.384	0.659*	–0.100	–0.630	0.169
	(0.512)	(0.292)	(0.084)	(0.379)	0.800

Standard errors of the estimated coefficients, in parentheses, are Newey-West corrected.
* Significant at 5-percent level.
** Significant at 1-percent level.

Public Perceptions of the Fed's Inflation Goal

Equation 2′ produces an inflation forecast that is conditioned on long-forward inflation expectations—that is, the forecast is conditioned on the perceived long-term-inflation policy objective. How do these perceptions evolve? In the empirical inflation literature it has been common to assume that the FOMC's inflation objective follows a random walk and that public perceptions of this objective follow a persistent process that is sensitive to realized past inflation. The idea is that the credibility of any announced objective is gradually eroded if inflation stays too high (or too low) for too long. (Take this line of reasoning very far and you travel full circle and end up back at the NAIRU model.) Fuhrer and Olivei (2010), for example, assume that the Fed's perceived inflation goal evolves according to

$$\pi_t^{LF} = \rho\pi_{t-4}^{LF} + (1-\rho)\pi_{t-4}^{AVG} + noise,$$

where π^{AVG} is an average of recent realized inflation rates. We estimated a somewhat more general equation, with lagged trimmed mean PCE inflation taking the place of π_{t-4}^{AVG}:

$$\pi_t^{LF} = \rho_0 + \rho_1\pi_{t-4}^{LF} + \rho_2\pi_{t-4}^{TM} + noise. \tag{3}$$

A standard (Quandt-Andrews) stability test identifies a clear break in this relationship at the end of 1997. Before 1998, the equation simplifies to Fuhrer-Olivei with $\rho = \frac{3}{4}$:

$$\pi_t^{LF} = 0.755\pi_{t-4}^{LF} + \underset{(0.095)}{0.245}\,\pi_{t-4}^{TM} + noise \qquad \text{S.E.} = 0.336.$$

After 1997, however, Equation 3 collapses to a very different specification:

$$\pi_t^{LF} = \underset{(0.010)}{2.498} + noise \qquad \text{S.E.} = 0.075. \tag{3'}$$

The "*noise*" term in Equation 3′ is uncorrelated with inflation and unemployment information available at $t-4$. The implication is that the best four-quarter-ahead forecast of SPF long-forward CPI inflation expectations is a constant: 2.5 percent. From 1998 on, long-term inflation expectations have been extremely "well anchored." Figure 5 shows a plot of the SPF long-forward inflation expectation, with a vertical line marking the late-1990s regime shift. Toward the end of the sample, the chart also shows the five-year, five-year-forward inflation expectations implied by TIPS yields. The TIPS-implied rate is higher frequency and more volatile, with a shorter sample. The two series match up well, outside of the Lehman-collapse aftermath.

When using Equation 2′ to forecast inflation, we condition on $\pi_t^{LF} = 2.5$, consistent with Equation 3′. The analyst who believes that inflation expectations are becoming unanchored should adjust these inflation forecasts accordingly. For example, the analyst who believes that SPF long-forward inflation expectations will drift upward to 2.75 percent between now and third quarter 2012, should add 25 basis points to the third quarter 2012 inflation forecast we report.

Figure 5: Behavior of Long-Forward SPF Inflation Expectations Shifts in 1998

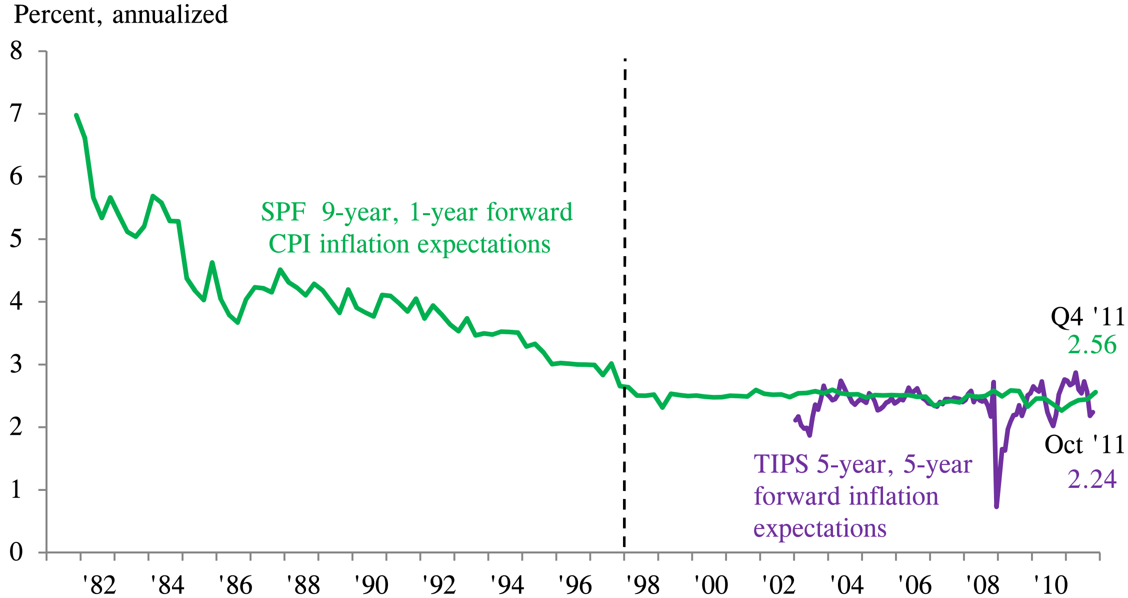

SOURCES: Federal Reserve Bank of Philadelphia; Federal Reserve Board; authors' calculations.

Results

Tables 5A, B, and C report root-mean-square forecast errors (RMSEs) achieved by Equations 1′ and 2′ for trimmed mean PCE, GDP, and headline PCE inflation, respectively. The smaller the RMSE, the better the forecast performance. In addition, the tables report how well one would have done by simply using lagged inflation to predict future inflation, by simply using lagged trimmed mean PCE inflation to predict future inflation, and by simply using lagged long-forward SPF inflation expectations to predict future inflation.[5] These "rule-of-thumb" forecasts ignore slack, very much in the spirit of Atkeson and Ohanian. Finally, the tables report how one would have done by setting forecasted GDP inflation and forecasted headline PCE inflation equal to forecasted trimmed mean PCE inflation, where forecasts of trimmed mean PCE inflation come either from Equation 1′ or Equation 2′. Trimming strips unforecastable noise from headline inflation. As a result, trimmed mean inflation is relatively easy to predict, and forecasting-equation coefficients are more precisely estimated. Perhaps a trimmed mean inflation forecast can do double or triple duty, by also usefully serving as a forecast of GDP or headline PCE inflation.

In each row of the tables, the best-performing forecasts are identified by asterisks, and the RMSE of the worst-performing forecast is put in parentheses. Forecasting performance is assessed for the periods first quarter 2006 to first quarter 2011 and first quarter 2001 to first quarter 2011 because real-time vintage trimmed mean PCE inflation data are not available until 2005. The forecasts for first quarter 2006 to first quarter 2011 are as they would have appeared at the time.

The message from Table 5A is that unemployment and changes in unemployment are useful for predicting trimmed mean PCE inflation. In particular, Equation 2′ produces lower RMSEs than all other forecasting methods considered. Equation 1′ comes in second. Figure 6 compares the forecasts from Equations 1′ and 2′ with actual trimmed mean inflation. The forecasts did remarkably well in predicting the steady decline in inflation after 2007. From 2003 to 2006, the forecasts were consistently below realized inflation as currently estimated. (The gap is mostly due to upward revisions to early estimates of trimmed mean inflation.)

Table 5B shows that the forecasts of trimmed mean PCE inflation from Equation 2′ are also the best forecasts of GDP inflation. The trimmed mean forecasts consistently do a better job of predicting GDP inflation than either Equation 1′ or Equation 2′ applied directly to GDP inflation and also do a better job than the rule-of-thumb forecasts that simply extrapolate from lagged inflation. Figure 7 compares the

[5] 0.3 percentage points are subtracted from the SPF long-forward inflation expectations in Table 5A and C, approximating the usual differential between CPI inflation and trimmed mean PCE inflation. In Table 5B, 0.5 percentage points are subtracted from the SPF long-forward inflation expectations.

trimmed mean forecasts from Equation 2′ with actual GDP inflation and with SPF GDP-inflation forecasts.[6] Again, the forecasts were consistently low from 2003 to 2006 but capture the recessionary drop fairly well.

According to Table 5C, SPF long-forward inflation expectations and forecasts of trimmed mean PCE inflation from Equation 2′ perform about equally well as forecasts of headline PCE inflation over 2001–11, and generally dominate alternative forecasts. The clear loser in the forecasting horse race is lagged headline PCE inflation: There is simply a great deal of variation in headline inflation that is uninformative for future headline inflation. Much of this variation is unrelated to long-forward inflation expectations and to the level or change in slack, too. Thus, even our best headline-PCE forecasts have root-mean-square errors in excess of 100 basis points over 2001–11, as compared with about 80 basis points for GDP inflation and about 30 basis points for trimmed mean PCE inflation. The difficulty of forecasting headline PCE inflation over the past 10 years is quite evident in Figure 8.

Table 5: Root-Mean-Square Errors of Alternative Real-Time Forecasts

A. Trimmed Mean PCE Inflation Forecasts

	Trim Mean Forecast		Lagged Inflation Measures	
Interval	Eq. 1′	Eq. 2′	Trim Mean	Long-Fwd
’06-’11	0.324*	0.323**	0.577	(0.714)
’01-’11	0.375*	0.327**	0.512	(0.543)

B. GDP Inflation Forecasts

	GDP Forecast		Lagged Inflation Measures			Trim Mean Forecast	
Interval	Eq. 1′	Eq. 2′	GDP	Long-Fwd	Trim Mean	Eq. 1′	Eq. 2′
’06-’11	0.904	0.881	0.926	1.036	(1.092)	0.783*	0.781**
’01-’11	1.017	0.967	0.961	0.909	(1.022)	0.862*	0.823**

C. Headline PCE Inflation Forecasts

	PCE Forecast		Lagged Inflation Measures			Trim Mean Forecast	
Interval	Eq. 1′	Eq. 2′	PCE	Long-Fwd	Trim Mean	Eq. 1′	Eq. 2′
’06-’11	1.340	1.303	(1.950)	1.307	1.338	1.271*	1.229**
’01-’11	1.178	1.152	(1.538)	1.046**	1.148	1.103	1.057*

Notes:
** best-performing forecast.
* second-best-performing forecast.
() worst-performing forecast.

[6] The SPF survey and the trimmed mean PCE inflation forecasts from Equation 2′ are in a virtual dead heat when it comes to forecasting GDP inflation: SPF RMSEs are 0.784 and 0.824 over 2006:Q1-2011:Q1 and 2001:Q1-2011:Q1, respectively.

Figure 6: Best-Performing Forecasts of Trimmed Mean PCE Inflation

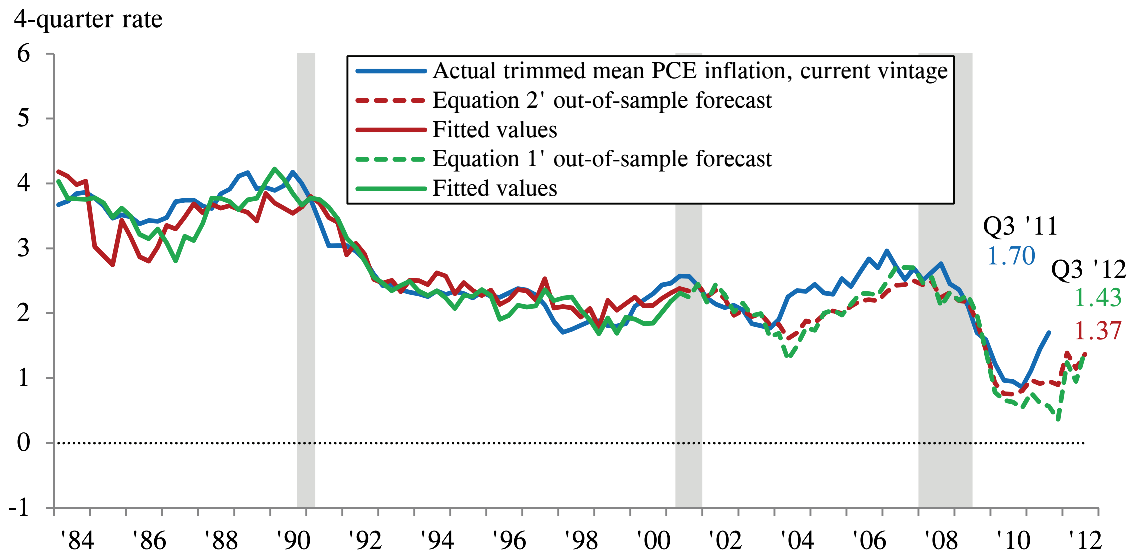

NOTE: Gray bars represent recession.
SOURCES: Federal Reserve Bank of Dallas; authors' calculations.

Figure 7: Best-Performing Forecasts of GDP Inflation

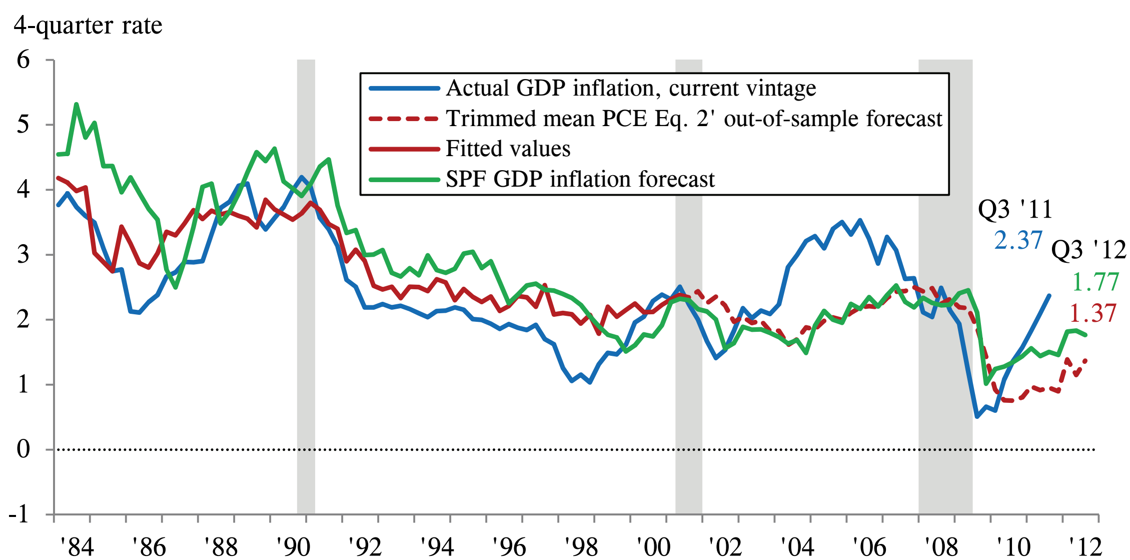

NOTE: Gray bars represent recession.
SOURCES: Bureau of Economic Analysis; Federal Reserve Bank of Philadelphia; authors' calculations.

Figure 8: Best-Performing Forecasts of Headline PCE Inflation

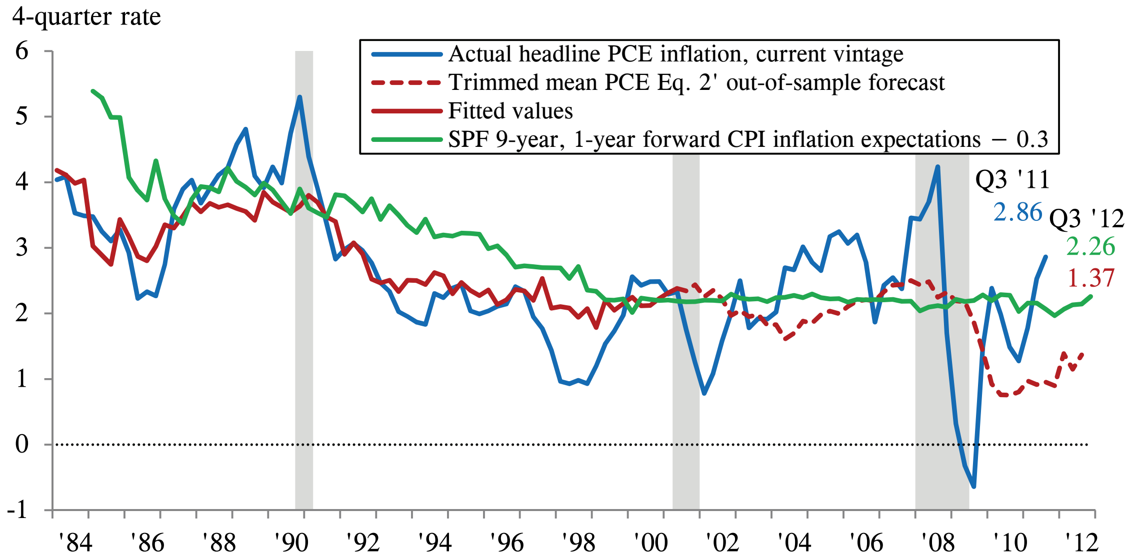

NOTE: Gray bars represent recession.
SOURCES: Bureau of Economic Analysis; Federal Reserve Bank of Philadelphia; authors' calculations.

4. CONCLUSION

There's no need for multiple inflation forecasting equations. It is enough to forecast trimmed mean PCE inflation because forecasts of trimmed mean inflation are also superior forecasts of GDP inflation and headline PCE inflation. When forecasting trimmed mean inflation, slack matters. For each percentage point that the unemployment rate exceeds its average, trimmed mean inflation can be expected to come in 0.2 percentage points below its long-run trend. Changes in slack matter, too. Trimmed mean inflation can be expected to come in 0.1 percentage points below trend for each quarter-point increase in the unemployment rate. Historically, changes in public perceptions of the Fed's long-run inflation objective have also had a strong influence on trimmed mean inflation. However, since the late 1990s, movements in these perceptions have been small and transitory—hence unimportant for forecasting. Our analysis suggests that trimmed mean PCE inflation, GDP inflation, and headline PCE inflation will all come in at a bit under 1.5 percent over the four quarters from third quarter 2011 to third quarter 2012. Confidence bands are wide—especially for GDP and headline PCE inflation.

REFERENCES

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

Calvo, Guillermo A. (1983), "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics* 12 (3): 383–98.

Dolmas, Jim (2005), "A Fitter, Trimmer Core Inflation Measure," Federal Reserve Bank of Dallas *Southwest Economy*, no. 3.

Friedman, Milton (1968), "The Role of Monetary Policy," *The American Economic Review* 58 (1): 1–17.

Fuhrer, Jeffrey C., and George Moore (1995), "Inflation Persistence," *The Quarterly Journal of Economics* 110 (1): 127.

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

Gali, Jordi, and Mark Gertler (1999), "Inflation Dynamics: A Structural Econometric Approach," *Journal of Monetary Economics* 44 (2): 195–222.

Gordon, Robert J. (1997), "The Time-Varying NAIRU and Its Implications for Economic Policy," *Journal of Economic Perspectives* 11 (1): 11–32.

Koenig, Evan F. (1996), "Aggregate Price Adjustment: The Fischerian Alternative," Federal Reserve Bank of Dallas, Research Department Working Paper no. 96-15 (December).

Mankiw, N. Gregory, and Ricardo Reis (2002), "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics* 117 (4): 1295–1328.

McCallum, Bennett T. (1994), "A Semi-Classical Model of Price-Level Adjustment," in *Carnegie-Rochester Conference Series on Public Policy*, 41, Elsevier, 251–84.

Stock, James H., and Mark W. Watson (2010), "Modeling Inflation After the Crisis," NBER Working Paper no. 16488 (Cambridge, Mass., National Bureau of Economic Research).

FEDERAL RESERVE BANK OF DALLAS
P.O. BOX 655906
DALLAS, TX 75265-5906