What's Up with Inflation?

Jeff Fuhrer[1]

January 18, 2018

In answer to the titular question, not enough. In recent years, the behavior of inflation has been somewhat puzzling to many. The most recent puzzle centers on the downturn of inflation in March of 2017, despite widely-accepted estimates of a closed or near-closed unemployment gap. But perhaps more puzzling was the tepid response of inflation to the economic disruption during the Great Recession. Most models at the time would have envisioned a more pronounced decline in inflation in the wake of ten percent unemployment, even with "anchored" expectations.[2] Well-anchored expectations decrease the response of inflation to shocks, but they do not imply that inflation should remain immobile in the wake of historic shocks.

A representation of the wedge between inflation outcomes and Committee members' expectations is depicted in Figure 1 below. The figure shows the realized four-quarter PCE inflation rate, the blue line, and the SEP 2-year-ahead forecast for four-quarter PCE, plotted in the period in which that inflation is expected to occur. As the chart suggests, the SEP submissions have consistently expected core PCE to rise toward the two percent inflation goal as output and employment gaps have closed, but that has yet to occur.[3]

Figure 1

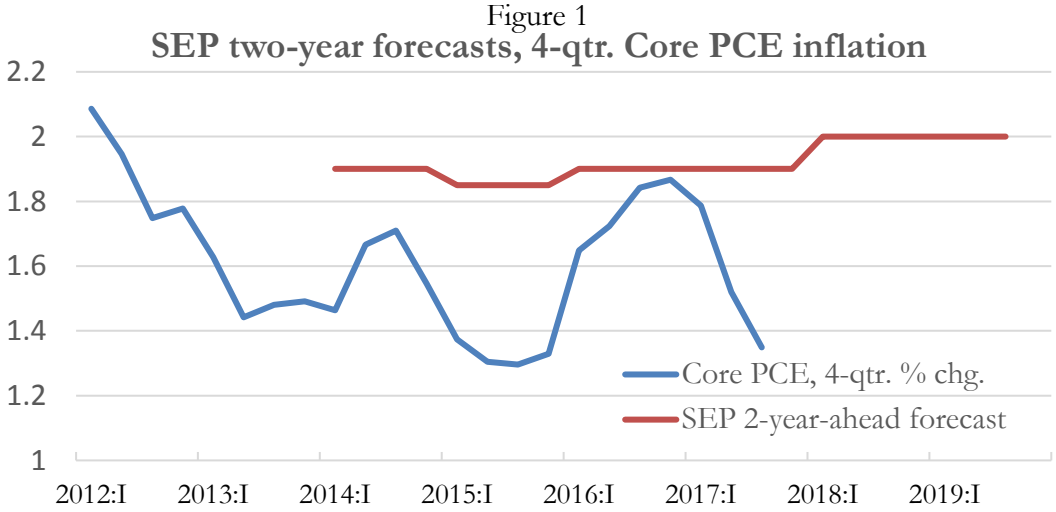

The puzzlement has arisen in the context of fading support for older "accelerationist" theories of inflation that suggest that the change in inflation is related to activity gaps and crude proxies for inflation expectations, and in the context of new empirical models that place primary weight on the anchoring effect of long-run expectations. The former have faced difficulty explaining inflation since the Great Recession, while the latter have less well-established theoretical backing,

---

[1] Federal Reserve Bank of Boston. Thanks to Todd Clark, David Lebow, Giovanni Olivei, Ekaterina Peneva, Jeremy Rudd and Bill Wascher for helpful comments.

[2] An exception is the DSGE exercise by Del Negro, Giannoni and Schorfheide (2015). They achieve their success in part by assuming a rapid return of expected real marginal cost to its pre-recession equilibrium, a development that did not occur.

[3] The most recent quarterly data available for core PCE, the blue line, are through the third quarter of 2017.

and are supported by little rigorous empirical testing. Many economists assume that long-run expectations matter, although exactly how they figure into firms' price-setting decisions is less clear. Earlier, it was widely assumed that short-run expectations were central, but some current empirical central bank models simply drop this influence. Most economists believe that activity gaps drive medium-term deviations of inflation from central bank inflation targets, but it is difficult to estimate the influence of such gaps with precision, especially in recent years. Earlier work added the influence of inflation "persistence" to the mix, suggesting that in addition to the aforementioned, inflation simply moved sluggishly in response to most other determinants. That influence appears to have diminished in importance in recent years.

In recent public pronouncements, policymakers have noted that inflation has been below its target for at least five years. While this observation is accurate, not all of this period implies a problem with our models of inflation: During most of that period, output was below potential and unemployment above its estimated natural rate, so inflation would be expected to be below its target. More recently, Fed discussions have mixed underlying models of inflation, in the following sense. One canonical model suggests that deviations of inflation from target are proportional to current and expected output gaps. That model implies that as the gap closes, other things equal, inflation should move to target (or perhaps to long-run expectations, which we hope remain "anchored" on the central bank target). Thus in some discussions, we are puzzled because inflation has not converged to target as activity gaps have closed. But in other discussions, FOMC participants have said that the unemployment rate has only recently dipped below most estimates of the natural rate, so we should not expect to see much upward pressure on inflation. This logic harkens back to an earlier, "accelerationist" model that does not imply that inflation converges to target when activity gaps close. Instead, activity gaps determine the change in inflation, so that inflation is expected to rise when unemployment dips below the natural rate, and fall when unemployment is above the natural rate. The level to which inflation falls depends on the precise sequence of gaps over time. That sequence of gaps is presumed to be manipulated by the central bank so as to achieve the inflation goal.

In some venues, policymakers have discussed the possibility that activity gaps no longer are important in affecting inflation. Instead, it's expectations that matter. The problem with this argument is that one then needs to take a stand on where expectations come from. In the earliest models of inflation, expectations were proxied by lags of inflation. That was an imperfect, but semi-explicit way of thinking about expectations—the models implied that expectations were formed via a moving average of recent inflation observations. In more recent models, expectations of inflation are implicitly dependent on expectations of future output gaps (or real marginal cost). This linkage is precisely the way in which the central bank controls inflation in the model: By affecting real interest rates, the central bank could affect the current and expected path of activity gaps (or real marginal cost), thus moving inflation toward its target.

But if activity gaps and the like no longer affect inflation, this not only complicates our understanding of inflation expectations, it also implies that the standard channel by which the central bank influences inflation is gone. Thus the statement that inflation now depends only on expectations is not just an interesting empirical observation, it unravels the underpinnings of expectations formation and monetary policy transmission. Absent a transmission mechanism, there

is no way to anchor expectations, or to move inflation, either up or down. Inflation is therefore decoupled from the rest of the economy. Houston, we have a problem.[4]

This memo will discuss what we know and don't know about the issues outlined above, and come to some tentative conclusions about what is likely happening with inflation. The key word to blend into this discussion, always and everywhere, is "humility." There is much that is not understood about how inflation evolves in a modern low-inflation economy, so one should be careful in imposing too rigidly one's favorite theories on real-time policy decisions. But I will argue that our biggest fears—that inflation will spiral downward, stay always below our target—are probably unfounded, even though we cannot categorically rule them out.

1. Historical perspective on recent inflation readings

It may be helpful to put our current troubles in historical perspective. In an earlier day, had we said that we averaged a 0.5 percentage point error in achieving our inflation goal (had we had an explicit goal), the Board room would have been filled with cheers rather than wailing and gnashing of teeth. Due to forecast errors, model uncertainty, measurement error and the ever-present long and variable lags, it would never have been thought feasible to achieve tighter control of inflation than the margin we have recently attained. One-year-ahead forecast errors for inflation were typically 1-1.5 percentage points.[5] To be sure, a series of one-sided forecast (or control) errors should be more worrisome, but again, given that a sizable portion of the period in question entailed significant activity gaps, one-sided control errors (deviations of inflation from target, not from forecast) should not have been unexpected. This observation is simply meant to provide some historical context: Things have been much worse in prior periods. This is a puzzle, and an important one, but not on the scale of the Great Inflation, the Great Disinflation, or the Great Recession. It's just not that Great.

A second means of judging whether recent inflation outcomes are unusually large is to compute confidence intervals for inflation outcomes from the perspective of a standard model. Here we estimate a model similar to the Board staff's baseline specification,[6]

$$\pi_t = a\pi_{t-1} + (1-a)\pi_t^{LR} + \sum_{i=1}^{2} b_i (U_{t-i} - U_{t-i}^{*}) + \sum_{i=0}^{1} c_i \Delta RPI_{t-i} \ , \qquad (1.1)$$

[4] This may not imply that the monetary authority loses control of inflation, but instead it suggests a dichotomous economy in which the real sector and the monetary sector are completely separate, as in the Real Business Cycle models of Kydland and Prescott (1982) and Long and Plosser (1983). In this case, the monetary authority might gain control of inflation by setting a fixed money growth target, relying on the long-run stability of the quantity equation to yield desirable inflation outcomes. But such a framework is very far from the heart of most all macroeconomic models today.

[5] These numbers correspond to the root mean-squared forecast errors (RMSEs) for the Greenbook/Tealbook forecasts for the inflation measures used from 1970-1989 by the Board staff, using current-vintage data. RMSE's dropped to below 1 percentage point starting in 1990, and have been even lower over the past ten years. Similar numbers arise if we use real-time actual data, but these numbers sidestep the debates over which vintage of real-time data forecasters have in mind at the time of forecast.

[6] The Board specification does not use the CBO natural rate, and includes a term for energy prices, neither of which alter the key inferences drawn from the specification estimated in equation (1.1).

where the inflation measure is the one-quarter annualized percentage change in the core PCE, long-run expectations are measured by the FRB/US PTR variable,[7] the unemployment gap is computed relative to the CBO's estimate of the natural rate, and the "RPI" term is the quarterly percentage change in core import prices relative to the lagged four-quarter percentage change in the core PCE.[8] The model is estimated from 2000:Q1 through 2014:Q4, yielding coefficients for $[a, b, c]$ of $[0.11, -0.16, 0.080]$, with all but the lagged inflation term $[a]$ precisely estimated. Taking the paths for PTR, the unemployment gap, and core relative import prices as given, we then simulate the model 10,000 times, taking random normal draws for the coefficient vector that match the estimated variance-covariance matrix of the parameter estimates for equation (1.1), and drawing a set of normal random *iid* shocks with the same variance as that estimated in equation (1.1). The resulting median simulated values (the red line) and 70 and 90% confidence intervals (the shaded areas) are displayed in Figure 2 below, along with the actual quarterly percent changes in the core PCE, in blue.

Figure 2

Stochastic simulation of model 1.1, with 70 and 90% confidence intervals

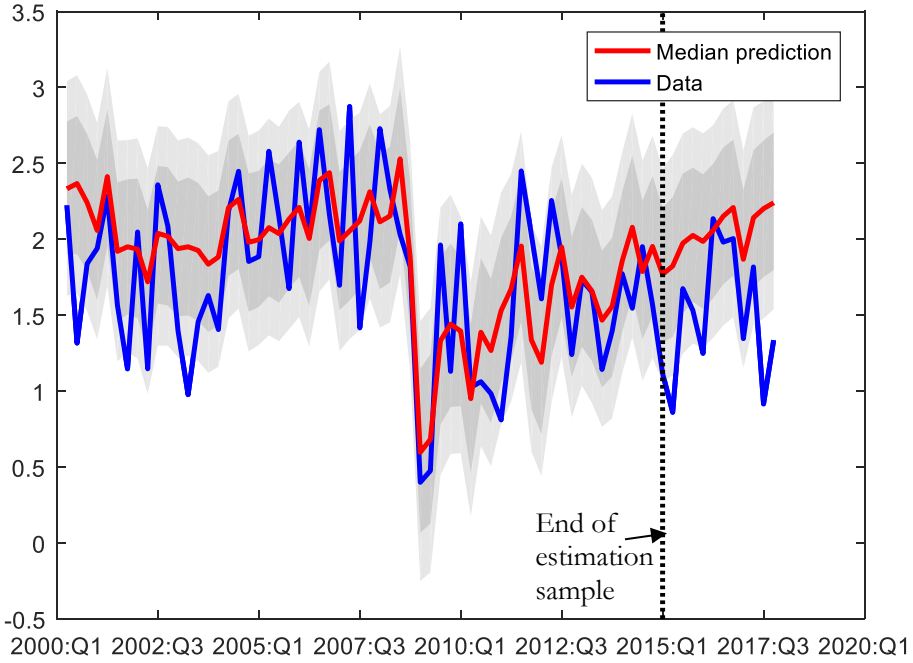

[7] The PTR variable used in the FRB/US model is derived for this part of the sample from the Survey of Professional Forecasters (SPF) 10-year average inflation expectations for the CPI. It is adjusted to match the average level of the PCE until 2007, when the SPF begins to measure the 10-year average PCE expectation directly.

[8] While this measure is often weighted by the import share, I do not do so here, as the differences between the two measures, weighted and unweighted, over the post-1999 sample are trivial. The correlation between the two measures over this sample is 0.994.

The model over-predicts somewhat after the estimation sample (that is, starting in 2015:Q1). Note that the model also overpredicted for quite a few quarters at the beginning of the sample, and this is of course within the estimation period. A key question is whether these over-predictions represent a significant change in inflation dynamics, or ordinary noise in the inflation data. Over the period presented in the figure, seven of the 71 observations fall outside the 90% confidence interval implied by the coefficient and residual uncertainty estimated for the model.[9] However, two of those seven observations fall in 2015, and two are adjacent observations in 2017, which might suggest that the behavior of inflation has been a bit unexpected recently. However, given the full cloud of uncertainty around the predictions for inflation—taking account of all the uncertainty, not just parameter and residual uncertainty—it might be premature to call these extraordinary observations. Were the simulations to take into account uncertainty in the trajectories for unemployment, long-run expectations, or import prices, the confidence intervals would of course expand further.[10]

Because the out-of-sample over-prediction suggests the possibility of a break in the relationship, we test for a breakpoint in the coefficients of the model at some point prior to the last few quarters. Reserving ten percent of the sample (seven observations on either end) for endpoint considerations, the Quandt-Andrews unknown breakpoint test cannot reject the hypothesis of zero breakpoints in favor of one, developing a *p*-value for the maximum likelihood ratio and Wald tests of 0.43. Of course, if the breakpoint occurred in the last few quarters, there would be no way to determine this statistically.[11]

Finally, a word on the precision of our inflation control in low-inflation environments, such as the one we have enjoyed over the past twenty-five years. When inflation is low and stable, small deviations of inflation from the Fed's target are likely of second-order concern to wage- and price-setters. A consequence of our success in achieving a low-inflation environment is that we are now expecting a more precise degree of control over inflation than we have achieved in previous eras.[12] In earlier times, such as when inflation was high and variable prior to the 1990s, and thus more powerfully affected nonfinancial and financial decisions, we necessarily focused on getting inflation

---

[9] One would of course expect about one-tenth of the observations to fall outside the 90% confidence interval, as is the case in this exercise. Note that six of these seven observations fall on the low side of the confidence interval, including the last two quarters.

[10] Todd Clark's background memo provides additional evidence suggesting that recent prediction errors are not extraordinary by historical standards.

[11] Extending the breakpoint test sample back to 1990 reveals significant evidence of a single break in 2000, consistent with our use of that sample above.

[12] Of course, "rational inattention" models (Sims 2003, Maćkowiak and Wiederholt 2009) explore this line of thought, although such models are generally not used in empirical form by central banks. See also Boivin, Giannoni and Mihov (2009) for an empirical exercise that attempts to distill the common component among a large number of firm-specific prices, and identifies a link between monetary policy shocks and this common component.

right to the first digit to the left of the decimal place. That was a degree of accuracy we could almost surely attain. The degree of accuracy we aspire to today is arguably a higher hurdle.[13]

2. Theory

a. The role of expectations, short-run and long-run

In words, the simplest model of inflation suggests that inflation depends on expected inflation, an activity gap, and (perhaps) lagged inflation (the so-called persistence effect). In earlier versions of the model, lagged inflation was used as a proxy for expected inflation and (perhaps) other frictions that made inflation slow to adjust to changes in the gap. More recent versions posit "rational" expectations for inflation. Such models of course make expectations explicit, rather than an implicit function of lagged inflation. Thus in these models lagged inflation represents frictions, rather than proxying for expectations.

The expected inflation term implies dependence of inflation on future expected activity gaps (if today's inflation depends on the expected inflation next period and the activity gap, then next period's inflation depends on expected inflation in the following period and next period's activity gap. Lather, rinse, repeat). This dependence is motivated in turn by the notion that most prices remain fixed for more than one period, and so the economic conditions in future periods should be taken into account in setting prices today.[14]

Note that in this description, the relevant expectations are short-run expectations, not long-run. In most models, the expectation that the central bank will ultimately steer inflation to its target means that longer-run expectations of inflation will be anchored to the central bank's target, and thus short- and long-run expectations will generally move towards that target over time, as the economy equilibrates.

A more recent forecasting model drops the short-run expectations linkage, and instead suggests that inflation is equal to long-run expectations (perhaps the target), except when output deviates from equilibrium output.[15] The theoretical motivation for dropping short-run expectations

---

[13] It is possible that, had we specified our goal in terms of a range rather than as a point, our communication about policy and our expectations about the desired degree of control might have been different.

[14] Cogley and Sbordone (2008) posit a model in which inflation fluctuates around "trend inflation," a concept that essentially captures time-variation in the central bank's inflation goal. In earlier samples, prior to the Fed's announcement of its two percent inflation goal, the implicit goal was (a) unknown, and (b) likely varied over time, although the precise amount of variation is the subject of debate. In more recent years, the inflation goal has not varied, although there is a reasonable concern that long-run expectations, which should be tied to the inflation goal, may have deviated from the announced goal, perhaps because economic agents believe that the Fed has lowered its effective medium-term inflation goal, or because agents believe the Fed will have a hard time attaining its long-run goal in the foreseeable future.

[15] While this is not how the model was developed, one could begin with a very simple version of Cogley and Sbordone's trend inflation model $\pi_t - \pi_t^T = \beta[E_t\pi_{t+1} - \pi_{t+1}^T] + \gamma s_t$, where π^T is trend inflation and s_t is marginal cost. Iterating this equation forward, one obtains $(\pi_t - \pi_t^T) = \gamma E_t[s_t + \beta s_{t+1} + \beta^2 s_{t+2} + ...]$. Then, if one assumes (e.g.)

is usually not explicit, but as an empirical description of recent inflation, the model works reasonably well. Inflation has fluctuated around an average of a bit below two percent, with larger deviations from two percent when the output/unemployment gap has opened up. Intuitively, one should worry if long-run inflation expectations do not center on the central bank's inflation goal. Interestingly, in work that surveyed hundreds of private firms about their price-setting behavior, Truman Bewley (2016) reports not a single mention of inflation expectations or the central bank's inflation goal as a determinant of their price-setting decisions.[16] This is not proof that long-run expectations enter nowhere (perhaps through wage negotiations), but it does suggest that this key anchor looms not quite so large in firms' minds as it does in those of central bankers'.[17]

b. The role of activity gaps

At the heart of the Phillips curve, as suggested above, are two inflation determinants: The level of the activity gap (or marginal cost), and expectations of next period's inflation.[18] As suggested above, next period's expected inflation enters as a compact representation of the notion that today's inflation depends upon a *sequence* of future activity gaps. Thus the gap or cost measure really stands at the center of most inflation models. It is tempting to say that expectations are really at the center, but the logic of the frameworks makes it clear that such a view ultimately devolves toward placing activity gaps at the center. How expectations of future activity gaps or costs are formed—are they nearly "rational," is the relevant horizon long or somewhat shorter—is of interest, but again these differences translate into the way in which gaps, current and future, affect inflation.

A common empirical finding in recent academic and Fed research is that the effect of activity gaps on inflation appears to have diminished in recent years. Put differently, the Phillips

---

a random walk process for s_t , one could collapse this into a simple equation $\pi_t = \pi_t^T + \Gamma s_t$, where the coefficient Γ would be a function of the underlying parameters $[\beta, \gamma]$ ..

[16] See point (5) on slide 11 of Bewley's presentation here https://www.bostonfed.org/great-recovery2016/agenda/. Bewley reports that "Nobody I interviewed about the pricing of goods mentioned expectations about inflation or future Federal Reserve policy as a factor, and questions along these lines provoked ridicule." Blinder, Canetti, Lebow and Rudd's (1998) earlier volume develops somewhat similar findings.

[17] One might assume that this empirical model is just another way to write the trend inflation model of Cogley and Sbordone (2008), in which all inflation terms enter as deviations from trend inflation π_t^T :

$\pi_t - \pi_t^T = a(\pi_{t-1} - \pi_t^T) + (1-a)(E\pi_{t+1} - \pi_t^T) + bGap_t$. But that model implies that the coefficient on trend inflation in a regression of inflation on lagged inflation and trend inflation should be zero.

$\pi_t = a\pi_{t-1} + (1-a)E\pi_{t+1} + \pi_t^T(a-1-a+1) + bGap_t = a\pi_{t-1} + (1-a)E\pi_{t+1} + \pi_t^T(0) + bGap_t$.

Instead, this model is better thought of as a model in which trend inflation is a linear combination of lagged inflation and the time-varying inflation goal π_t^* : $\pi_t^T = a\pi_{t-1} + (1-a)\pi_t^*$, and then inflation is equal to trend inflation plus a gap term.

[18] While current structural models most often make marginal cost the driving variable, in what follows I will primarily use gap variables, as they are more commonly used by the Fed and many other central banks. In theory-based models, under some simplifying assumptions, marginal cost will be proportional to the output gap.

curve has "flattened."[19] Such a development could arise for a number of reasons, not least of which is the possibility that price-setters believe that the Fed will systematically and promptly act to return inflation to its goal, in contrast to earlier historical eras during which price-setters were less confident that such actions would be forthcoming. As a consequence, disruptions to the real side of the economy today may translate somewhat less into changes in the inflation rate.[20]

In recent years, some researchers have estimated a gap effect of very close to zero, or have estimated the effect to be insignificantly different from zero. It may be tempting to conclude from these results that now inflation is purely expectations-determined; the central bank need only focus on where expectations are, and need no longer concern itself with the influence of gaps on inflation. But as suggested above, this shortcut leads nowhere. One needs to articulate where expectations come from, and all of the leading models of inflation take one back to activity gaps or the equivalent. It is conceivable that one could write a model in which expectations center on a different driver for inflation (money growth?), but the theory and empirics in support of such a view are slim at present.[21]

c. Persistence

In a line of work that extends back a quarter-century or more, a number of authors have debated whether inflation exhibits some persistence (sluggishness in response to shocks) that is not well-accounted for by a slow-moving gap variable and rational expectations.[22] The early work found fairly compelling evidence of such persistence, but more recent data suggest a decline in the importance of this determinant of inflation. Nonetheless, the Board staff's current judgmental model still includes a modest role for persistence defined in this way. In current circumstances, the presence of such persistence is important because, other things equal, it would slow the progress towards the Fed's inflation goal even as the gap variables approach or overshoot zero. However, the estimated coefficient on lagged inflation, which may proxy for such persistence, is currently fairly small—perhaps 0.1 to 0.4 depending on the sample chosen—which implies that persistence *per se* would play a smaller role in slowing inflation's return toward the Fed's inflation goal.

---

[19] If we plot observations on inflation against observations on (say) the unemployment gap, a smaller effect of the gap on inflation will be manifested as a flatter or more shallowly-sloped scatter plot—a smaller change in inflation will correspond to a given change in the gap.

[20] In contract-based inflation models, persistently low inflation could lead to the endogenous lengthening of contracts, so that more prices would be bound by contracts, and thus fewer prices would be able to adjust to economic conditions at any point in time. This would imply a shallower Phillips curve.

[21] One such model is a neo-Fisherian model that centers on the so-called Fisher equation. A linearized approximation to the Fisher equation that abstracts from expectations is $r_t = \rho_t + \pi_t$, where r denotes the nominal interest rate, ρ the real rate of interest, and π the inflation rate. The model suggests that, at some horizon, the appropriate way to raise inflation is to raise nominal interest rates, a proposition that has not been tested in the US in recent decades.

[22] See for example Fuhrer and Moore (1995), Rudd and Whelan (2007), Gali and Gertler (1999), Christiano, Eichenbaum and Evans (2005). Even earlier work by Barsky (1987) considers historical periods during which inflation may have contained a unit root that accounted for most of inflation's variation, i.e. inflation was extremely persistent. The chapter by Fuhrer (2011) summarizes more recent evidence on persistence.

d. The role of wages in the determination of inflation

As suggested above, many theoretically-motivated models imply that inflation is driven by (real) unit labor costs, the deflated difference between wages and productivity. In simple models, one can show that real unit labor costs are proportional to output or unemployment gaps, but the assumptions in these models may not hold. Public commentary often refers to wages as an important indicator of inflation and inflation pressures. But the link between wages and inflation is not terribly strong, as discussed in section 4 below. How we should think about the role of wages, given the slippage between wage and inflation fluctuations—that is, given significant variation in the mark-up from labor costs to final goods prices—is a question of active research interest.[23]

e. The exchange value of the dollar

In many empirical models of inflation, including equation (1.1) above, monetary policy can influence inflation through its effect on the dollar. Consider the effect of a decline in the policy rate, which other things equal is expected to lead to a depreciation of the dollar against other currencies. This influence comprises two channels: (1) the direct effect of the dollar on inflation from increases in the prices of imported goods, to the extent that prices in dollars pass through changes in the exchange rate;[24] and (2) the positive effect of the depreciated dollar on the activity gap, which also tends to increase inflation. The first, which relies on the change in the dollar's value, is generally expected to have temporary effects on inflation, as long as such changes do not become embedded in inflation expectations. The second effect, which spurs the level of output and the activity gap if the level of the dollar remains low, might have a more sustained effect on inflation. In equation (1.1), we explicitly capture only the first, direct effect, which acts primarily as a temporary shifter in the simple Phillips curve. Once the exchange rate stabilizes, the effect on the change in imported prices and thus on the inflation rate ultimately ceases. The second effect is captured indirectly, as even though we do not explicitly model the effect of exchange rates on output, any exchange rate movements that have persistent effects on the unemployment gap will be captured by the unemployment gap term.

3. Empirics

a. The slope of the Phillips curve

Virtually every empirical paper on the Phillips curve of the past ten years has come to the conclusion that the slope of the Phillips curve is smaller in magnitude than it used to be. Activity gaps and marginal cost influences on inflation appear to have diminished. A graphical illustration of this empirical feature appears below in figure 3. The figure plots the deviation of inflation from long-run expectations against the unemployment gap, defined as the difference between the civilian unemployment rate and the CBO's estimate of the natural rate.

---

[23] For a recent analysis of the pass-through of wages to prices, see Peneva and Rudd (2017), which finds little evidence of an independent effect of changes in labor costs on inflation. For a discussion of the value of wages as an indicator of labor market slack, see Lindner, Roberts and Wascher (2015).

[24] See Burstein and Gopinath (2014) for work on the extent of pass-through, and the extent to which foreign suppliers price their exports to the U.S. in dollars, bypassing the exchange rate effect.

Figure 3

The shift in slopes is evident in the graph, although there is no significant difference between the slopes post-1990 and the most recent sample. That said, the slope is not estimated to be zero, and the standard error of the estimate in the most recent period is 0.046, so that zero is about two standard errors away from the estimated value. Of course, this regression is extremely simplistic, as it does not control for short-run expectations, oil prices or import prices, which have acted to shift the Phillips curve in ways that, other things equal, might obscure the relationship between inflation and unemployment.[25]

A host of other researchers have attempted to estimate the slope of the Phillips curve, and some additional evidence is presented below. A reasonable (and I hope objective) way to summarize the evidence is that the coefficient is of modest size, but not zero. For our purposes, a working estimate of about -0.1 is about right. Again, it is critical to emphasize that a slope that is literally zero strains credulity, as it implies that inflation evolves in a way that is unrelated to the overall state of the macroeconomy, whether that is proxied by an output gap, an unemployment gap, or a marginal cost proxy. As a practical matter, identifying the effect of such variables on inflation in aggregate data is a difficult econometric task. But concluding from such difficulties that the economy is characterized by a strict dichotomy between the real side of the economy—employment and output—and the nominal side—money and inflation—is a big leap, and seems unwarranted.

b. How low can the natural rate go?

[25] A more careful assessment of the Phillips correlation as it relates to out-of-sample inflation forecasting can be found in Clark (2017).

As many have documented, estimates of the natural rate of unemployment are subject to considerable uncertainty.[26] Submissions to the SEP have consistently reduced their estimates of the long-run unemployment rate, which is likely closely related to the natural rate, over the past six years, as shown in figure 4.[27] As of March 2012, the midpoint of the central tendency was 5.6%; as of late 2017, the median estimate had dropped a full percentage point.

Figure 4

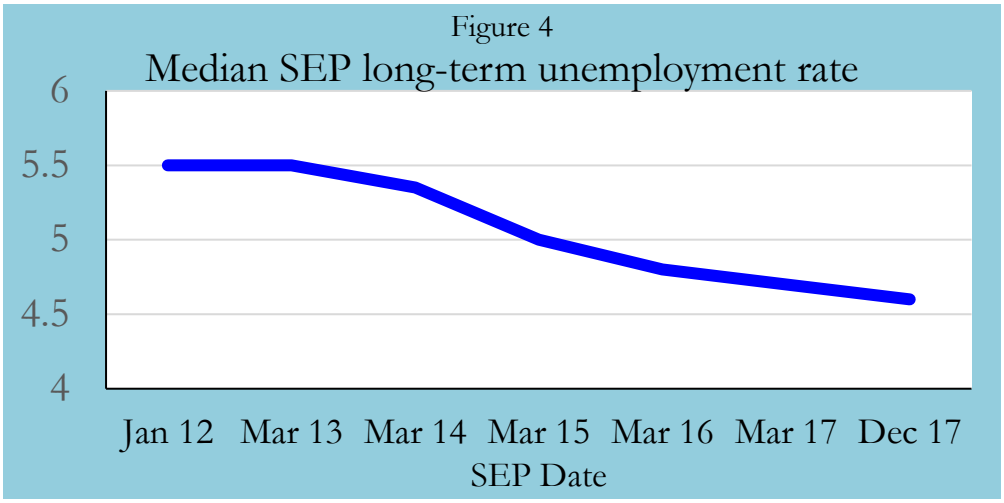

Is there a case for an even lower natural rate (or NAIRU)? Given the historical uncertainty surrounding this concept, the answer is perhaps. But do we have compelling evidence that we have yet to breech the natural rate, or that we have much further to go before we do? Figure 5, which displays wage growth using a variety of measures, suggests not. Whether measured using wage-rate concepts like the ECI compensation excluding incentive-paid occupations, or the AHE (private compensation), or a wage-bill concept like the Employer Costs for Employee Compensation (ECEC), wages have accelerated over the past four years.[28] To be sure, the increase in wage inflation has not been dramatic, but it is roughly consistent with an economy with somewhat subdued productivity growth (below one percent over the past 3-5 years), sub-two percent inflation and near-full employment. While this simple figure cannot prove that our current estimate of the natural rate is correct, it suggests that we are not at present well above the natural rate, and may also be consistent with a natural rate that is somewhat higher than the current unemployment rate.

[26] A representative paper in this literature is Stock, Staiger and Watson (1997), which finds enormous confidence intervals around the estimated NAIRU. This work likely overstates the actual amount of uncertainty that we face today.

[27] The chart depicts the middle of the central tendency or the median, as available in the SEPs. The points charted are for the first SEP of each calendar year, typically March, along with the most recent submission for December 2017.

[28] One might be concerned that, using a compensation measure like the ECEC, one could find increases in overall compensation per hour that are driven by firms' hiring of higher-productivity workers. In principle, this could be the case, but in practice, we have not seen an acceleration in productivity that would be consistent with this story.

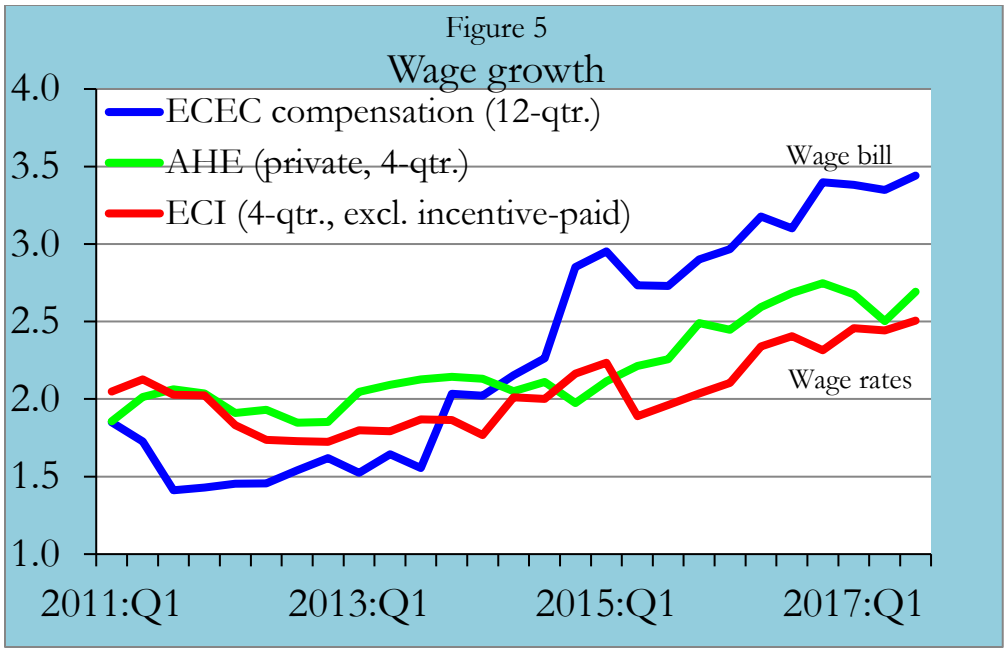

c. Nonlinearity in the Phillips curve near full employment?

In an interesting paper, Barnes and Olivei (2003) examine the possibility of a piecewise-linear Phillips curve, in which the effect of the activity gap on inflation when the economy is closer to full employment differs from the average response. If the effect of the gap is lower when we are close to full employment, this might suggest an additional rationale for inflation's failure to rise toward target in recent years.

Here we alter Barnes and Olivei's model to allow for long-run (or trend) inflation, and estimate a piecewise linear Phillips curve with unknown threshold T. In addition, the threshold here is defined in terms of the gap, rather than in terms of the unemployment rate. Figure 6 depicts a possible version of this specification; the exact specification is displayed in equation (1.2) listed below. In the figure, the slope of the Phillips curve (the response of inflation to the activity gap) declines monotonically, but discontinuously as the gap increases from -1.5 to 4 percentage points. [29]

[29] Note that the Phillips curve shown in figure 6 is similar to that in the original Phillips (1958) work, although that diagram plotted wage inflation against unemployment. The slopes in figure 6 correspond to those estimated for equation (1.2). See also Nalewaik (2016) for a regime-switching implementation of a similar idea, which develops similar results; Murphy (2016) and Barnichon and Jorda (2017) for an examination of the convexity of the Phillips curve, which also reach similar conclusions.

Figure 6
A piecewise linear Phillips Curve

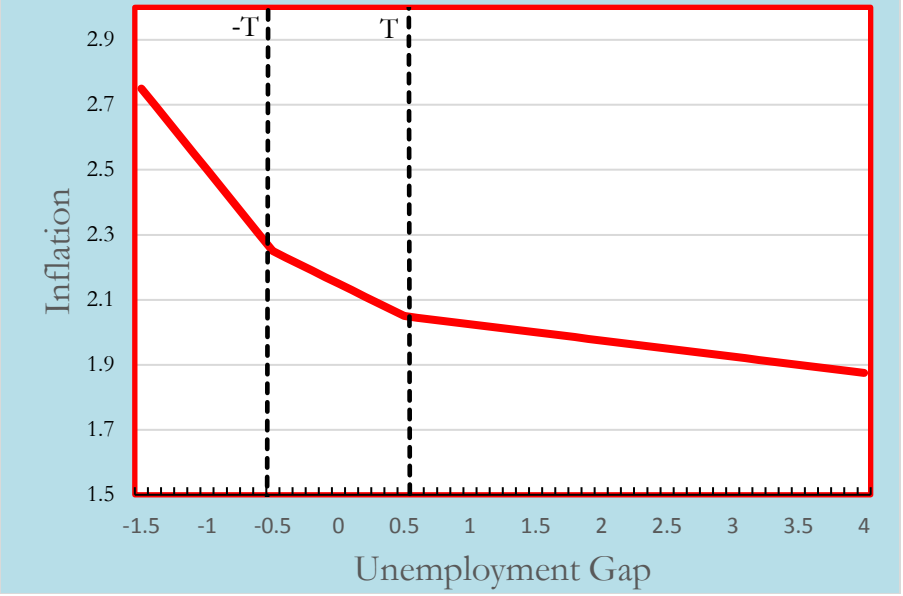

$$\pi_t = a\pi_{t-1} + (1-a)\pi_t^{LR} + \sum_{i=1}^{2} b_i^L \left[(U_{t-i} - U_{t-i}^*) if \left\{ \left(U_{t-i} - U_{t-i}^* \right) < -T \right\} \right] +$$
$$\sum_{i=1}^{2} b_i^H \left[(U_{t-i} - U_{t-i}^*) if \left\{ \left(U_{t-i} - U_{t-i}^* \right) \geq T \right\} \right] + \qquad (1.2)$$
$$\sum_{i=1}^{2} b_i \left[(U_{t-i} - U_{t-i}^*) if \left\{ -T \leq \left(U_{t-i} - U_{t-i}^* \right) < T \right\} \right] + \sum_{i=0}^{1} c_i \Delta RPI_{t-i}$$

Due to the findings that the overall slope of the Phillips curve may have declined in recent years, we estimate this model from 1990 to the present. The estimation chooses the value of *T* that minimizes the sum of squared residuals.

The threshold parameter *T* that minimizes the sum of squared residuals is 0.5, suggesting that the slope of the Phillips curve may vary as indicated in Figure 8.[30] The sum of the b_i, the slope in the near-full-employment region, is -0.22. The two lags of unemployment jointly enter significantly in the equation, but their sum is not significantly different from zero, which suggests that changes in unemployment may matter most in a small neighborhood around full employment.[31] The point estimate for the sum of the b_i^H is smaller at -0.12, although the estimated sum here is significantly different from zero (*t*-statistic of -2.6), implying a significant level effect of the gap. The largest coefficients apply to the unemployment gaps that lie below the near-full-employment region, at

[30] A plot of the sum of squared residuals against the estimated values for *T* reveals a very well-defined minimum at 0.5.

[31] The *p*-value for the restriction that excludes all the lags of this variable is 0.030.

-0.39. While this sum differs from zero at the five percent level, the overall contribution of this variable to the regression is modest at best.[32] Importantly, the overall improvement to the Phillips curve that is afforded by this piecewise linearity is small—the R^2 improves from 0.70 to 0.72.[33]

This simple model develops support for the notion that we may see more of an effect of the activity gap on inflation as we spend more time further away from our estimate of the natural rate—especially as the unemployment rate drops significantly below the natural rate. It might also explain some of the recent behavior in inflation: As the decline in the unemployment rate has slowed, these small changes imply very modest upward pressure on inflation.

d. The roles of short- and long-run expectations

As discussed above, the canonical models for inflation posit roles for both short- and long-run expectations. In many structural models, the roles are somewhat different: The short-run expectations capture the dependence of inflation on a sequence of current and future activity gaps, while the long-run expectation represents the "trend" for inflation, usually taken to be the central bank's inflation goal, to which inflation will ultimately return (given sufficient attention to this goal through the actions of the central bank).

As an empirical matter, recent data might suggest a primary role for long-run expectations, and a subordinate or non-existent role for short-run expectations. To illustrate the empirical success of an inflation forecasting model that uses only long-run expectations, I estimate equation(1.1), which approximates a specification used by the Board staff.[34] I use the FRB/US "PTR" 10-year inflation expectation variable to proxy for π_t^{LR}, long-run inflation expectations. U_t and U_t^* are the civilian unemployment rate and the CBO estimate of the natural rate, respectively, and ΔRPI_t is the change in the relative price of imported goods prices. The estimation sample is from 2007:Q4 to the present. The estimated equation (with *p*-values for individual coefficients or coefficient sums in parentheses) is:

$$\underset{}{\pi_t} = \underset{(0.11)}{0.27}\ \pi_{t-1} + \underset{(0.000)}{0.73}\ \pi_t^{LR} - \underset{(0.003)}{0.13}\ \sum_{i=1}^{2}(U_{t-i} - U_{t-i}^*) + \underset{(0.000)}{0.085}\ \sum_{i=1}^{2}\Delta RPI_{t-i}$$

The next figure displays the actual core PCE inflation rate along with fitted values from this regression over the estimation sample, shown in the dashed red line.

---

[32] The *p*-value for the restriction that excludes this variable is 0.091.

[33] The *p*-value for the *F* statistic that constrains the piecewise linear coefficients to be equal is 0.02 for this value of *T*. Alternative specifications that allow for a continuous nonlinear response of inflation to unemployment develop similar results, finding a small but sometimes significant role for nonlinearity. For example, an exponential model that includes the term $\alpha \exp(\gamma U_t)$ yields estimates of α and γ that produce a relationship that is nearly linear ($\hat{\alpha} = -.27, \hat{\gamma} = 0.06$).

[34] As above, the differences between this specification and the Board's are that the estimation sample differs, the Board does not use the CBO natural rate estimate, and the import price term is not weighted by the share of imports in GDP.

Figure 7

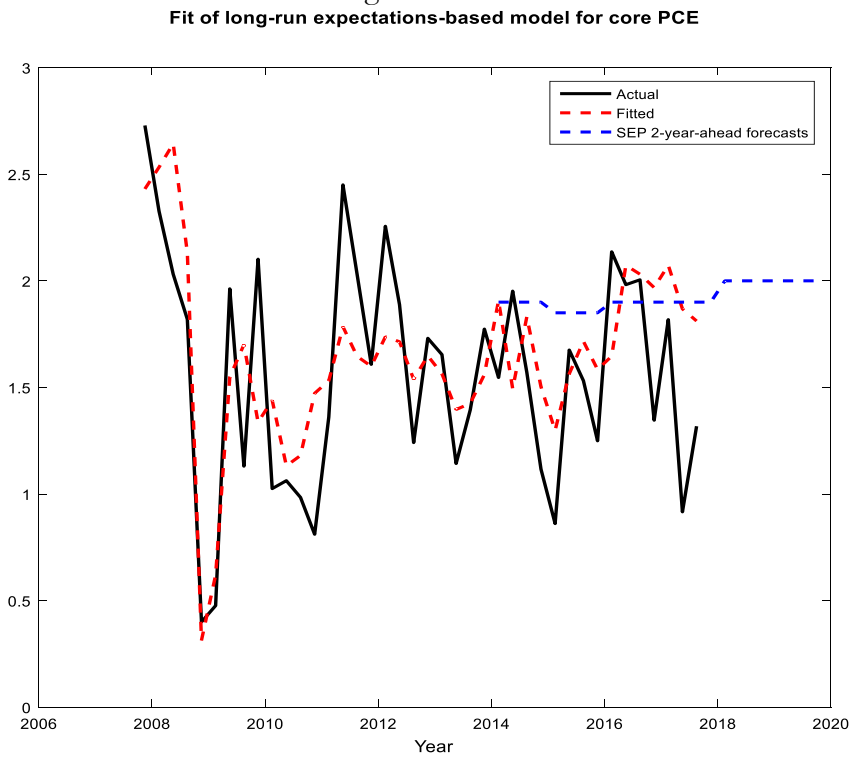

The model fits recent data reasonably well, although it has over-forecasted inflation in recent quarters, but only for the past four quarters or so. Of course, this is an in-sample fit. One can see that, apart from fluctuations due to import prices (and their mildly persistent effect through the lagged inflation term), the model expects a gradual upward tilt to inflation as the unemployment gap closes.

For comparison the chart also provides the SEP forecasts of core PCE inflation for the Q4/Q4 period two years out from the forecast. The forecasts are made in September of the years 2012-2017, and the forecasts are plotted two years ahead from the forecast origin (2014 for the Sep. 2012 forecast, and so on).[35] It is clear that SEP participants have expected inflation to return to two percent (or very close) for several years now, perhaps based on a notion of anchored expectations, but inflation has yet to do so.[36] Again, the magnitude of the miss is modest by historical standards, but of course the longer it persists, the more concerning it should be.

Would the inclusion of short-run expectations improve fit or forecasts? Perhaps a bit, but I don't believe that this solves the current puzzle, despite earlier research that suggests short-run

[35] The chart uses the median forecast for years in which that is available, and the midpoint of the central tendency for years before the median was reported.

[36] It is likely that participants also did not expect the dollar appreciation that occurred during this period, which may have contributed to their over-forecast of inflation.

expectations may well play an important role in inflation dynamics. The next figure displays the in-sample fit of this model.[37] The difference between this and the previous model is scarcely noticeable.

Figure 8

Fit of model including SR expectations, core PCE

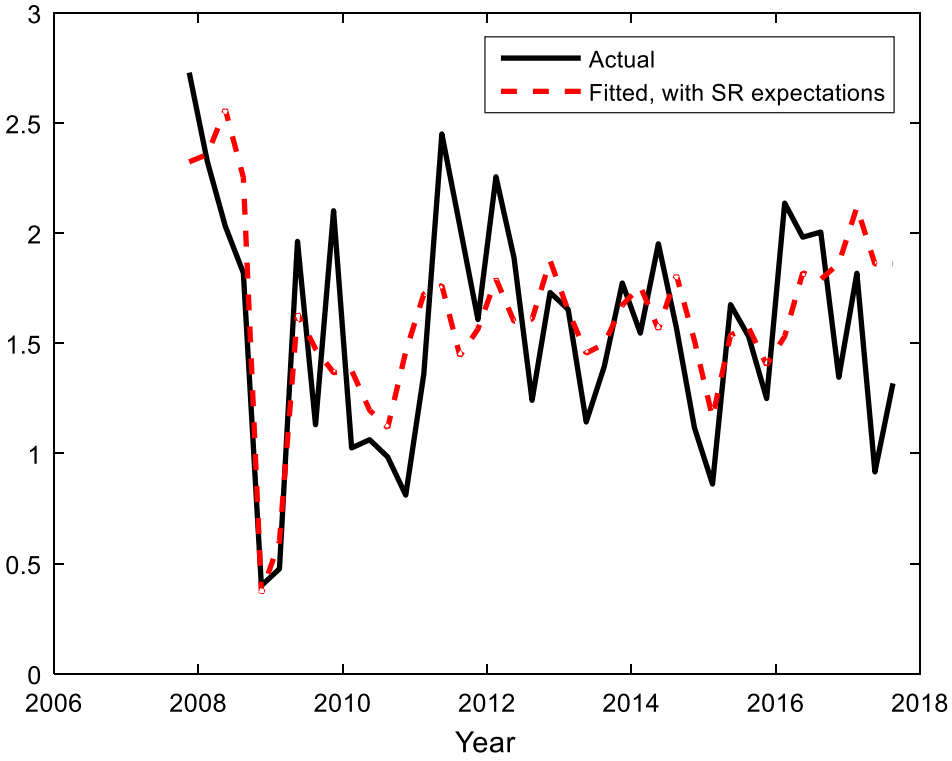

The similarity likely arises because in recent years, short-run expectations have reverted fairly quickly to long-run expectations—consider the episodes in 2004-5 and 2010-11 displayed in figure 9—so that the implications for models that include both variables will of course be similar.[38]

[37] The measure of short-run expectations is the SPF one-quarter ahead median expectation for core PCE inflation. For a model that incorporates both short- and long-run survey expectations in a structural setting, see Fuhrer (2017).

[38] SPF data for the core PCE became available in 2007.

Figure 9

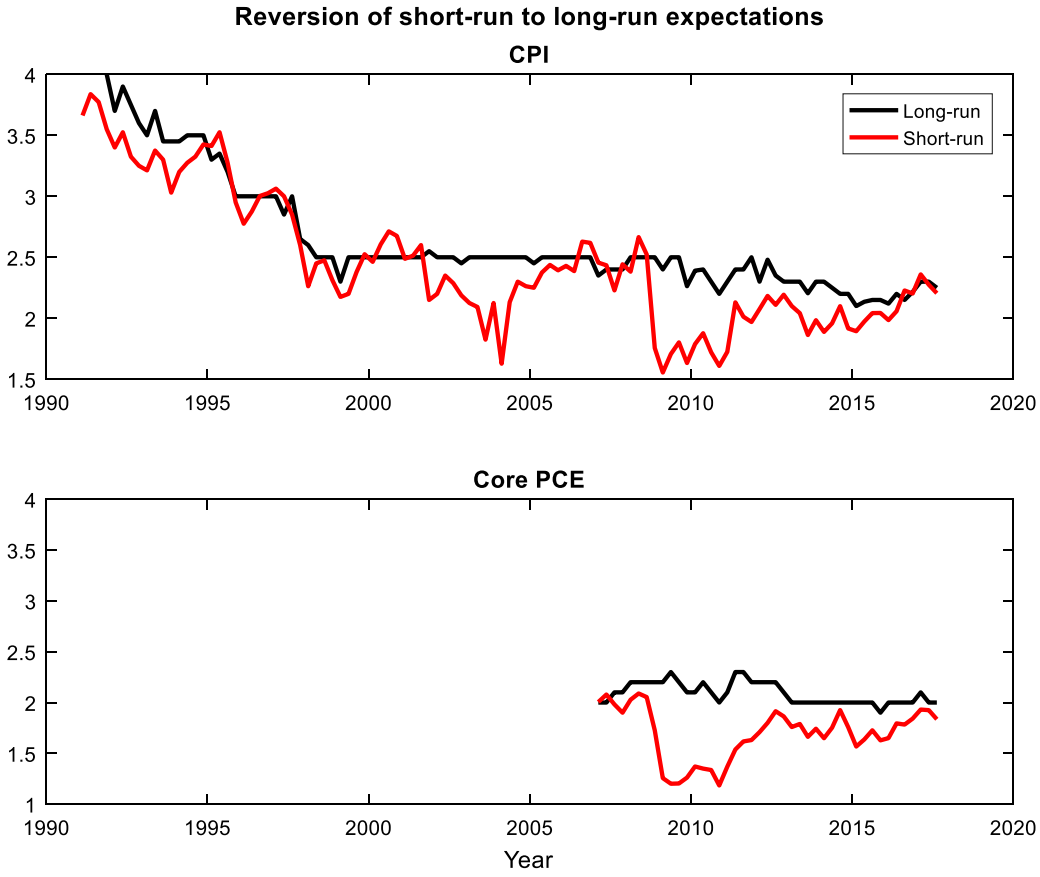

e. Sagging long-run expectations

One explanation for recent low inflation observations is that long-run expectations, which are thought to serve as an anchor for inflation, may themselves have drifted down, perhaps because of the long sequence of below-target inflation realizations. If so, it may take some time for rising inflation to restore long-run expectations to the central bank's target. In the meantime, any slippage in long-run expectations will hold down realized inflation by a modest amount.

Have long-run expectations slipped? It depends on which data you use as a long-run expectations proxy. Figure 10 shows the TIPS 5/5 breakeven inflation measure (without carry or liquidity premium adjustments), the Board's PTR 10-year expectation measure, and the median of the responses to the Michigan question soliciting 5- to 10-year inflation expectations.[39] The decline in the TIPS measure has been larger and more recent than the corresponding trend in the Michigan

[39] As Chan, Clark and Koop (2017) suggest, one can do better than assuming that long-run expectations equal survey or market measures of expectations. Their paper concludes that the trend that they infer—with the aid of such measures—is smoother than many of the proxies, and has exhibited significant stability in the period surrounding the Great Recession.

5- to 10-year expectations (trends are indicated in the dashed lines).[40] No trend is included for the PTR measure, as it has no discernible trend in recent years.[41][42]

Figure 10

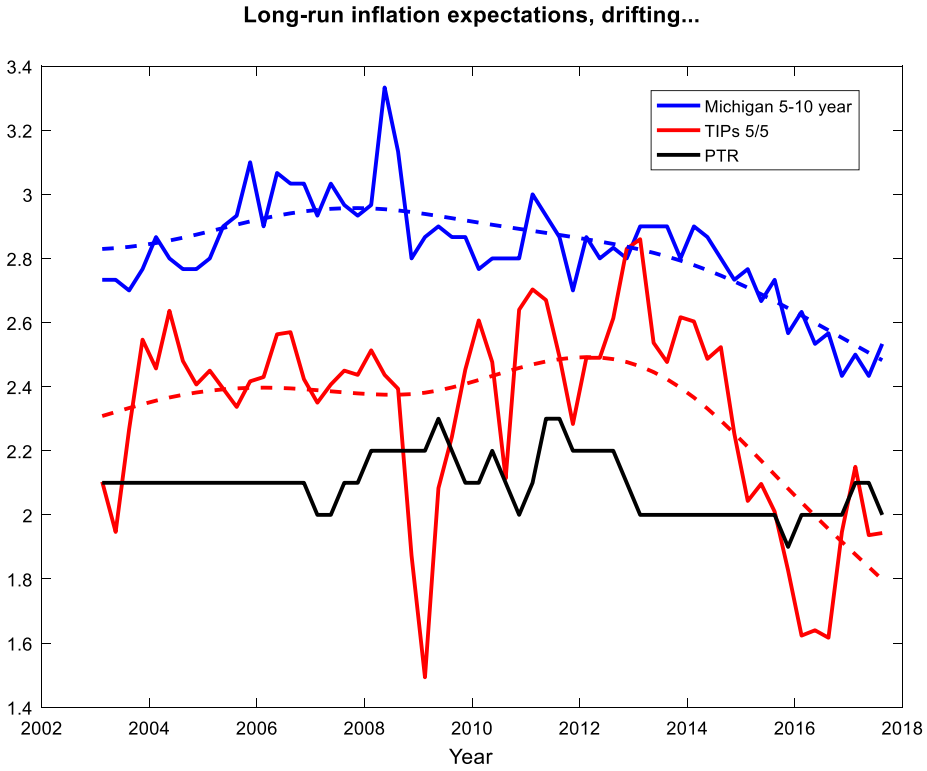

Which of these measures best predicts movements in core PCE inflation? Obviously, this is an empirical question. First, note that in figure 10 above, the PTR measure has differed little from a constant at about 2 or 2.1% since 2003. As a baseline model, we estimate the following simple equation, an unconstrained version equation (1.1), but excluding long-run expectations:

$$\pi_t = k + a\pi_{t-1} + \sum_{i=1}^{2} c_i (U_t - U^*_{t-i}) + d\Delta RPI_t \quad . \tag{1.3}$$

Note that the Michigan and TIPs measures are not based on a core PCE inflation measure, so the intercept k will also allow the regression to adjust for this difference. Because the TIPs data begin in 2003, and because we are most interested in recent developments, we estimate this forecasting equation with different measures for the long-run expectation π_t^{LR} from 2003 or later. The dependent variable is always the one-quarter inflation rate for the core PCE, the unemployment gap

[40] The indicated trends are just HP-filtered versions of the series, in both cases using a smoothing parameter of 1600.

[41] For better comparability, one can compute the five year forward five-year expectation for PCE inflation from the SPF forecast from 2007 forward. Although not displayed in the figure, this measure also shows no sign of slipping down over the past five years.

[42] Note that estimation of a linear trend for the PTR variable yields a very small and perhaps significant negative trend, either since 1998 or since 2003. However, the trend is small enough that it bears essentially no implications for core inflation, and thus we can safely ignore it here.

is the civilian unemployment rate less the CBO estimate of the natural rate, and the change in relative import prices is the one-quarter change in the core import price series constructed by the Board less the lagged four-quarter change in the core PCE inflation rate.

The fit of the baseline model, shown in the black dashed line in figure 11, is not bad, even though it relies only on an intercept, lagged inflation (with a small coefficient of 0.1), the unemployment gap and relative import price inflation.[43] How much does the addition of any of the long-run expectations measures add to this fit?

Figure 11

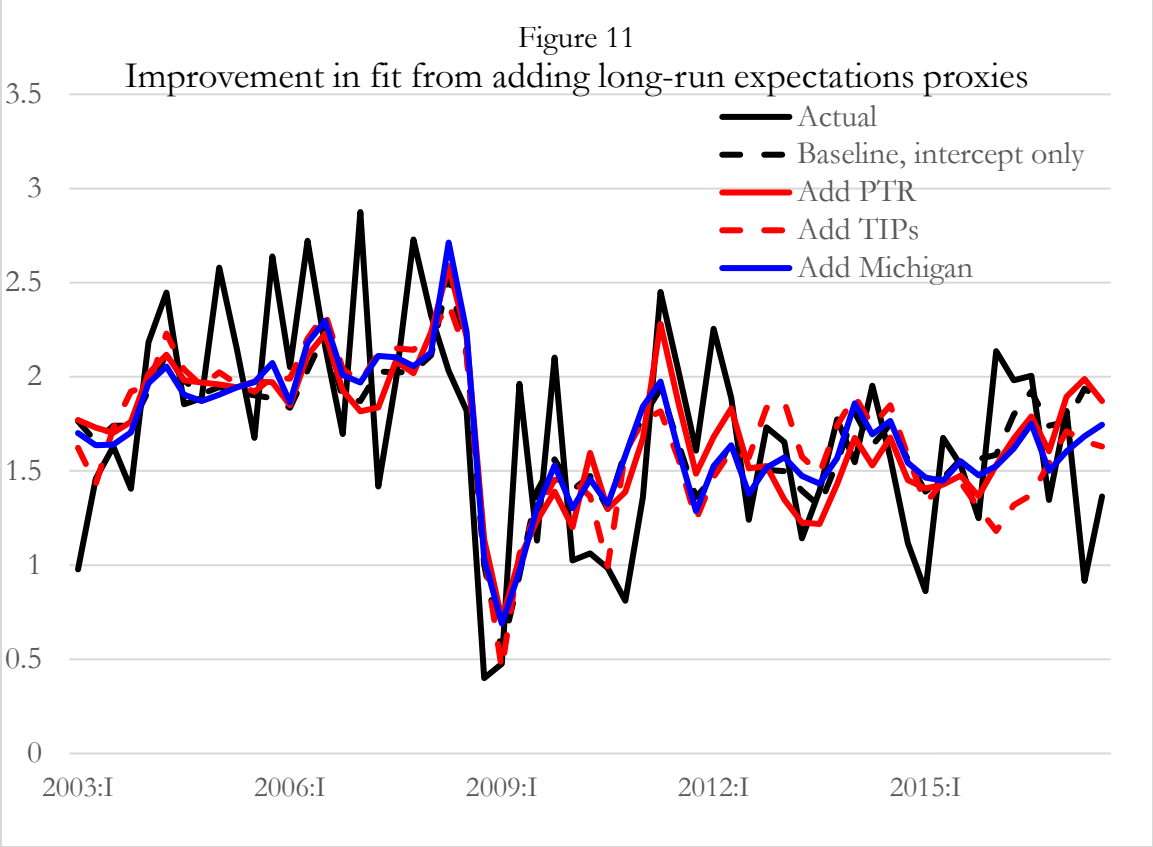

The other lines in figure 11 suggest that the improvement from any of the three measures is rather small. The regressions are very close during the pre-recession period. After the recession, we see a few episodes during which the TIPs regression misses (2011, 2013, 2015), whereas the PTR regression captures some fluctuations a bit better (late 2011). But toward the end of the sample, all three are over-predicting inflation, by nearly the same amount. None of the models with long-run expectations captures the recent decline in inflation.

Table 3 displays the improvement to the regression standard error that is afforded by the addition of long-run expectations measures. As the table indicates, the inclusion of any of these long-run proxies decreases the standard error by a basis point or so, confirming the visual

[43] The estimated intercept is 1.89, with an asymptotic standard error of 0.092, so that one cannot reject an intercept of two with any confidence.

impression gleaned from Figure 11, which is that long-run expectations have explained very little of inflation's fluctuations over the past 15 years.[44]

Table 3
Improvement in standard error for equation (1.3) with different long-run expectations proxies

Baseline—intercept only (SE=0.468)	-
Add PTR	-0.0131
Add TIPs	-0.002
Add Michigan	-0.0135
Add all three	-0.0168

The results can be summarized as follows:

- Long-run expectations proxies sometimes enter significantly in simple forecast regressions such as those presented above.
- However, the improvement from including such measures relative to an intercept—which perhaps proxies for the central bank's inflation target—is negligible since the early 2000's.[45]
- In particular, no long-run expectations measures have explained the recent decline in inflation.
- This suggests somewhat less concern that recent declines in the TIPs or Michigan long-run expectations presage a significant and persistent undershoot of inflation relative to target.

f. Estimated persistence

In earlier years, the overall persistence of inflation was quite high—that is, readings of inflation above target were usually followed by many quarters of readings above target, and vice versa—and thus models of inflation required a significant coefficient on lagged inflation to capture this persistence.[46] In more recent years, this appears to be less true, although the overall quietude of inflation makes it somewhat difficult to know just how persistent inflation might be in the presence of larger or more influential shocks.

In practical terms, this means that the coefficient on lagged inflation in staff inflation models is lower today than it would have been in the 1970s, 80s and 90s. A simple demonstration of that appears in table 4 below, which estimates the long-run expectations model in equation (1.1) above for three different samples. The table also presents the first autocorrelation, a summary measure of the overall persistence of inflation for the sample. As the table suggests, the estimated coefficient on

[44] Including Todd Clark's trend inflation measure, as described in his background memo to the Committee, produces results that are very similar to those for the PTR measure, although the identification of the unemployment gap term is weakened somewhat—the estimated inflation equation yields a smaller and somewhat less significant coefficient on the gap.

[45] Not surprisingly, samples that include the run-up and subsequent decline in inflation, from the 1960s through the 1980s, find a much more significant role for trend inflation measures—because inflation had a trend, both going up and coming down. Inflation has exhibited very little if any trend since the early 2000s.

[46] See Fuhrer(2011) for a full description of the issues surrounding inflation persistence.

lagged inflation indeed has fallen as the overall persistence of inflation has declined. This of course matters because, other things equal, a larger coefficient on lagged inflation implies a slower return of inflation towards target.

Note that the estimated inflation equation above (equation (1.1)) takes into account the decline in inflation persistence, and thus will not mis-forecast inflation (and least not dramatically) for that reason.

Table 4
Declining persistence and the estimated lag coefficient in the inflation equation

Sample	Lagged inflation coefficient	First autocorrelation
1968:1-1989:4	0.7	0.82
1990:1-2007:3	0.6	0.69
2007:4-2017:3	0.27	0.27

The implication of this finding is that when inflation is subject to temporary shocks, inflation is likely to exhibit a less persistent response to such shocks than it did in earlier eras. Similarly, when monetary policy acts so as to return inflation to its goal, inflation will put up less intrinsic resistance to such movement than it did earlier. In current circumstances, lower persistence is generally consistent with the forecasts of the Tealbook and many private forecasters who do not expect the low inflation readings of 2017 to persist in 2018.

4. Wages, prices, productivity and the mark-up over unit labor costs

So far, we have discussed price inflation without reference to wages, productivity and "mark-ups," the margin price-setters add above unit labor cost in determining final goods and services prices.[47] The reason is that, for much of post-war history, the average markup of prices over unit labor costs was fairly constant (for the markup measure displayed below, this is the same as the observation that the labor share of income had been roughly constant). Thus we believed that tight labor and product markets led to wage acceleration and, given productivity growth, would translate into higher price inflation. We just didn't need to worry too much about variations in the markup over unit labor costs, so we adopted a shorthand model of inflation that involved price inflation, expectations and activity gaps.[48] In fact, many DSGE models treat the shocks to the inflation equation as mark-up shocks, which are presumed to be transitory, often lasting for only a quarter or two.

[47] Strictly speaking, firms should mark up prices over marginal, not average cost. However, most researchers use average costs, employing the convenient approximation that when average costs are minimized, marginal cost will equal average cost.

[48] Of course, many DSGE models make inflation a function of "real marginal costs," which are most often proxied by real average unit labor costs, which are the same as labor's share of income. In simpler theoretical models, the average cost of output will be proportional to the deviation of output from equilibrium.

Over the past twenty years, profit margins and the markup of prices over unit labor costs (the inverse of the labor share of income) have risen quite significantly, as figure 12 indicates.[49] One can no longer take the share of income or the markup of prices over unit labor costs as approximately constant over time. This trend in profits and markups stands in contrast to a narrative that suggests that disruption in the retail distribution channel and increased competition more broadly may be responsible for subdued inflation.[50]

Figure 12
Mark-ups and Profit Share, nonfarm business

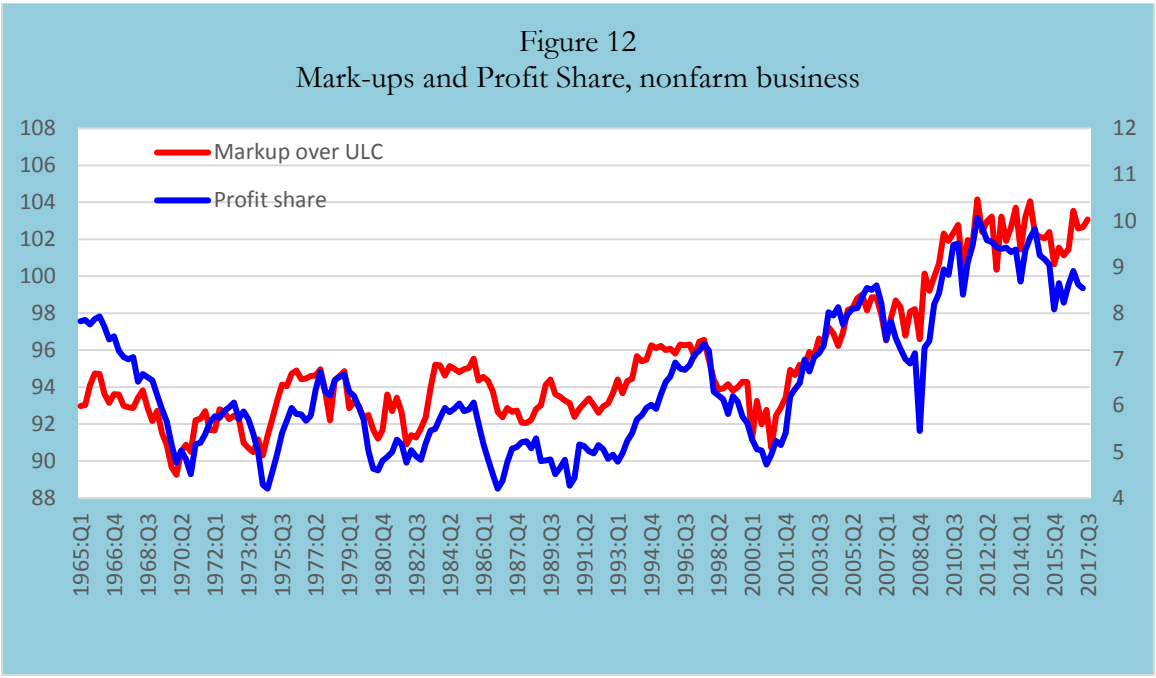

At a mechanical level, the presence of elevated profit margins suggests that the modest increases in wage inflation that we have seen to date have been easier for firms to absorb without passing through to final prices. The most recent flattening or slight downturn in such margins, coupled with continued modest acceleration in wages, might suggest that the cushion provided by profit margins, if any, is eroding, which in turn would imply a higher probability of pass through of labor market pressures into price inflation in coming quarters.[51]

---

[49] Elsby, Hobijn and Şahin (2013) attribute about one-third of the decline in labor's share of income (or correspondingly, the increase in the markup over unit labor costs) to measurement issues, specifically the imputation of labor income for the self-employed. This still leaves a pronounced decline in labor's share of income that represents other, presumably economic factors, some of which are discussed in their paper.

[50] The nonfinancial corporate mark-up has tracked the nonfinancial business mark-up quite closely until recently. Since 2015, the nonfinancial corporate mark-up has fallen more steeply than the nonfinancial business measure.

[51] While not presented here, data from the nonfarm business productivity accounts at the 4- to 6-digit industry level show that the change in current inflation is negatively and significantly correlated with the previous four years' change in the markup, controlling for a number of industry-specific characteristics and for year effects. This correlation is consistent with the notion that industries with sustained growth in markups are better able to cushion final goods inflation from cost shocks. Again, see the work in Peneva and Rudd (2017) for a fuller discussion of passthrough of labor costs to prices.

Of course, this observation is hardly based on deep theory or well-identified empirics. Understanding the source of the rise in mark-ups and margins is essential to understanding their implications for inflation dynamics. One possible explanation for rising margins might lie in an increase in the amount of unmeasured productivity growth in recent years, perhaps related (for example) to the use of advanced technology to improve the quality of the retail goods purchase experience. Under-measured productivity growth will overstate the true increase in unit labor costs; properly measured productivity growth would show a larger buffer between labor costs and prices, allowing firms to hold prices down. Of course, if productivity growth is under-estimated, then more of nominal output should be attributed to increases in real activity, and less to prices, which will further decrease measured price inflation. Suffice it to say that more research is needed in this area.[52]

5. Summary and policy implications

I would summarize our state of knowledge about inflation dynamics as follows:

a. The recent shortfall of inflation relative to target is not large by historical standards, and may not be statistically significant, given the uncertainty surrounding key aspects of our standard inflation framework(s);
b. Relatedly, the magnitude of inflation's shortfall relative to target is small by historical standards. Achieving inflation control to within a few tenths of a percentage point may be beyond our ability;
c. At its heart, variation in inflation around the central bank's target is likely still driven by variation in activity gaps (or real marginal cost);
d. This implies that short-run inflation expectations should also be driven by current and expected gaps;
e. The coefficient linking gaps to inflation—the slope of the Phillips curve—may well be smaller today than it was a quarter-century ago.
f. However, the slope is probably not zero. Recent estimates suggest a reasonably well-estimated coefficient of about -0.1. Holding tightly to an estimate of zero implies that inflation is decoupled from the real economy. Regardless of one's stance on other aspects of dynamics (how expectations are formed, intrinsic persistence, and so on), the assertion that inflation is not fundamentally driven by developments in the overall state of the aggregate economy is very strong.
g. The role of expectations is under active debate.
    i. Conventional models imply that long-run expectations will be well-anchored if the central bank announces a target and acts systematically to move inflation towards it. This suggests that long-run expectations, if centered on the inflation target, should serve as an attractor to inflation over time. Some forecasting models use this reduced-form implication (see, e.g. Faust and Wright (2013)) to predict inflation, often quite successfully. But the attraction of inflation to long-run expectations should be understood in a reduced-form sense—that is, the anchoring of expectations is a reflection of systematic monetary policy that consistently acts to restore inflation to its target. It is

[52] It could also be that the difficulty in measuring productivity in recent years arises from imperfect measurement of hours.

the actions of the central bank, not the expectations *per se*, that are fundamental in returning inflation to its goal.

ii. Conventional models also imply that short-run expectations should figure prominently in inflation dynamics, as they essentially stand in for the role that current and expected activity gaps are assumed to (and estimated to) play in determining inflation.

iii. As a practical matter, a model that abstracts from short-run expectations has done reasonably well over the past decade. While adding short-run expectations has been quite important in earlier periods, it would not help us explain our current predicament. However, one should not assume that ignoring short-run expectations will always be wise, inasmuch as it implies that one can safely ignore the expected course of activity gaps and the like.

h. The long-run expectations measures that have shown the most pronounced shift downward appear not to be clearly linked to the evolution of core inflation in recent years. This observation may reduce somewhat the concern that un-anchored expectations will significantly delay the progress of inflation toward the Fed's inflation goal;

i. It is possible that the activity gap is somewhat more influential when the economy significantly exceeds its potential. Some evidence to that effect is presented above, and in earlier work by Barnes and Olivei (2003), as well as other more recent work cited above. More specifically, the tentative evidence above suggests that inflation responds primarily to changes for small unemployment gaps, and is more responsive for larger negative unemployment gaps—that is, when unemployment falls significantly below the estimated natural rate—than for larger positive gaps.

Policy implications

These rough conclusions suggest some caution in thinking about current policy. While it may be wise to pursue a gradual course of normalization in the face of persistently low inflation, there may also be risks to too much gradualism. If a very gradual strategy is based on the assumption that the Phillips curve slope is zero, so that further declines in the unemployment rate have no cost, the consequences if that assumption is wrong could be significant, a point emphasized in the Board staff memo prepared for the October 31st FOMC meeting.[53] We might be tempted to rely on an inflation forecasting model that simply assumes a gradual reversion towards long-run expectations, without reference to the activity gap, such as the model in Faust and Wright (2013), since that model has forecasted quite well in a number of periods. But that model is silent on the matter of how the Fed actually controls inflation over time, unless we believe the Fed has the power to verbally dictate the level of long-run inflation expectations. Thus such reduced-form models are clearly useful for forecasting, but not for determining appropriate monetary policy.

If the inflation framework I have just sketched out is approximately correct, then we will find fairly soon that inflation is headed back up towards our target. However, we will likely not see as

[53] Erceg, Kiley, Hebden, Lopez-Salido and Tetlow, "Some Implications of Uncertainty and Misperception for Monetary Policy," distributed 24 October 2017.

rapid a return of unemployment towards the natural rate. With some monetary accommodation likely remaining, we risk finding ourselves with inflation near target, unemployment well below the natural rate, and monetary policy calibrated to spur the economy, rather than restrain it appropriately so as to guide both inflation and unemployment to their desired levels.

6. Conclusions

There is much we still need to learn about inflation. The precise role that expectations play is not well understood; as Robert Hall said at a recent conference, the role of expectations in macro models owes largely to assertion, rather than to a strong empirical basis.[54] How expectations are formed by households and firms, and how those expectations affect macroeconomic aggregates remains an area of active study. And the role of activity gaps has become murkier in recent years, as the slope of the Phillips Curve appears to have dwindled.

However, it would be risky to conclude that because activity gap coefficients are small, we really need focus only on expectations. Such a model lacks coherent theoretical foundations. This is also risky because in such a model, there is no direct channel from our actions to expectations. Both in theory and (I believe) in practice, it is through expectations of our actions' effects on (largely) real activity that expectations become well-defined (and perhaps well-anchored). It is also risky because in the absence of a reasonably well-understood channel from monetary policy to inflation, it's not at all clear how we can ever expect to achieve our goal of keeping inflation in the near-neighborhood of our inflation target, as our framework suggests we should.

It is also risky to conclude that, because inflation has declined to an unexpectedly low level over the past couple of quarters, the unemployment rate can continue to fall without adverse consequences. It seems more likely that a protracted period of low unemployment reflects building imbalances for inflation, real activity and (perhaps) financial markets. Whether inflation rises or not, we could find ourselves quite far from any reasonable estimate of the natural rate of unemployment. Our track record in guiding unemployment safely back to but not beyond the natural rate from below is not encouraging.[55]

References

Aaronson, S., De Michelis, A., Doniger, C., Fleischman, C., Gonzalez-Astudillo, M., Tevlin, S. and J. Zickler 2016. "Ramifications of allowing the unemployment rate to undershoot its natural rate," memo to the FOMC, December 1.

[54] Comment at Federal Reserve Bank of Boston annual economic conference, October 2016.

[55] See Aaronson, De Michelis, Doniger, Fleischman, Gonzalez-Astudillo, Tevlin and Zickler (2016) for a somewhat more optimistic take on the history of overshooting full employment.

Barnes, M. and G. Olivei 2003. "Inside and outside bounds: threshold estimates of the Phillips curve," *New England Economic Review*, pp. 3-18.

Barnichon, R. and O. Jorda 2017. "Upside risks to the inflation forecast," Board of Governors memo 10/31-11/1/2017.

Barsky, R. 1987. "The Fisher hypothesis and the forecastability and persistence of inflation," *Journal of Monetary Economics*, vol. 19(1), pages 3-24.

Bewley, Truman 2016. "Discussion of Robert King's presentation," presented at Federal Reserve Bank of Boston's annual economic conference, The Elusive "Great" Recovery: Causes and Implications for Future Business Cycle Dynamics, October 14-15 2016, https://www.bostonfed.org/great-recovery2016/agenda/ .

Blinder, A., Canetti, E., Lebow, D. and J. Rudd 1998. Asking About Prices: A New Approach to Understanding Price Stickiness. Russell Sage Foundation, New York, NY.

Boivin, J., Giannoni, M. and I. Mihov 2009. "Sticky Prices and Monetary Policy: Evidence from Disaggregated US Data," *American Economic Review,* vol. 99 pp. 350-84.

Burstein, A. and G. Gopinath 2014. "International Prices and Exchange Rates." Handbook of International Economics, Elsevier, 4th ed., 4: 391-451.

Chan, J., Clark, T. and G. Coop 2017. "A New Model of Inflation, Trend Inflation, and Long-Run Inflation Expectations, *Journal of Money, Credit and Banking*, forthcoming.

Christiano, L., Eichenbaum, M. and C. Evans 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113(1): 1–45.

Clark, T. 2017. "Inflation, Trends, and Long-Run Expectations: Perspectives from Forecasting Research," memo to the FOMC.

Cogley, T. and A. Sbordone 2008. "Trend inflation, indexation, and inflation persistence in the New Keynesian Philips curve." *American Economic Review,* **98**, No. 5, pp. 2101-2126.

Del Negro, M., Giannoni, M. and F. Schorfheide 2015. "Inflation in the Great Recession and New Keynesian Models," *American Economic Journal: Macroeconomics*, vol. 7, pp. 168-96.

Elsby, M., Hobijn, B. and A. Şahin 2013. "The Decline of the U.S. Labor Share". In: *Brookings Papers on Economic Activity* **44.2**, 1–63.

Erceg, C., Kiley, M., Hebden, J., Lopez-Salido, D. and R. Tetlow 2017. "Some Implications of Uncertainty and Misperception for Monetary Policy," distributed 26 October.

Faust, J. and J. Wright 2013. "Forecasting Inflation", in Handbook of Economic Forecasting (G. Elliott and A. Timmermann (eds.)), Volume 2A, Elsevier.

Fuhrer, J. and G. Moore 1995. "Inflation Persistence." *Quarterly Journal of Economics,* **110(1)**, 127–159.

Fuhrer, J. 2011. "Inflation Persistence," Chapter 9 in the *Handbook of Monetary Economics*, vol. 3A, Benjamin Friedman and Michael Woodford, editors, North-Holland, San Diego CA.

Fuhrer, J. 2017. "Expectations as a source of macroeconomic persistence: Evidence from survey expectations in a dynamic macro model", *Journal of Monetary Economics, vol. 86,* pp. *22-35*.

Galì, J. and M. Gertler 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44: 195–222.

Kydland, F. and E. Prescott (1982). "Time to Build and Aggregate Fluctuations". *Econometrica*. **50** (6): 1345–1370.

Lindner, D., Roberts, J. and W. Wascher 2015. "Compensation and Labor Market Slack," prepared for the October 28, 2015 FOMC meeting.

Long, J., Jr. and C. Plosser (1983). "Real Business Cycles". *Journal of Political Economy*. **91** (1): 39–69.

Maćkowiak, B. and M. Wiederholt, 2009. "Optimal Sticky Prices under Rational Inattention," *American Economic Review* 99(3), 769-803.

Murphy, A. 2016. "Metro Level Evidence on the Convexity of the Phillips Curve," FRB Dallas working paper.

Nalewaik, J. 2016. "Non-Linear Phillips Curves with Inflation Regime-Switching," FEDS Discussion series 2016-078.

Peneva, E. and J. Rudd, 2017. "The Passthrough of Labor Costs to Price Inflation," *Journal of Money, Credit and Banking*, **49(8)**, 1777-1802.

Phillips, A. W. 1958. "The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957," *Economica*, **25**, No. 100 (November), pp. 283-299.

Rudd, J. and K. Whelan 2007. "Modeling Inflation Dynamics: A Critical Review of Recent Research." *Journal of Money, Credit and Banking* 39(Supplement s1): 155–170.

Sims, C. 2003. "Implications of Rational Inattention," *Journal of Monetary Economics* 50(3), 665-690.

Staiger, D., Stock, J. and M. Watson 1997. "How Precise are Estimates of the Natural Rate of Unemployment?" in C. Romer and D. Romer (eds.), *Reducing Inflation: Motivation and Strategy*, University of Chicago Press for the NBER, 195-242.