Finding a Stable Phillips Curve Relationship: A Persistence-Dependent (And Business-Cycle-Phase-Dependent) Regression Model

September 2, 2020

Richard Ashley
Department of Economics, Virginia Tech, Blacksburg, VA 24061-0316
Email: ashleyr@vt.edu

Randal Verbrugge *(Corresponding author)*
Federal Reserve Bank of Cleveland and NBER/CRIW, 1455 E. 6th St., Cleveland, OH 44114
Email: randal.verbrugge@clev.frb.org. 1.301.254.7359.

Abstract

We establish that the Phillips curve is persistence-dependent: inflation responds differently to persistent versus moderately persistent (or versus transient) fluctuations in the unemployment gap. The persistence-dependent relationship we uncover is interpretable as being business-cycle-phase-dependent and is thus consistent with existing theory. Previous work fails to model this dependence, so it finds numerous "inflation puzzles" – such as missing inflation/disinflation – noted in the literature. Our model specification eliminates these puzzles; for example, the Phillips curve has not weakened, and inflation was not "stubbornly low" in 2019. The model's coefficients are stable, and it provides accurate conditional recursive forecasts through the Great Recession.

Keywords: overheating; recession gap; persistence dependence; NAIRU

JEL Classification Codes: E31, E32, C22, C32, E5

There are no potential sources of conflicts of interest.

Acknowledgments: For discussions occurring either as we developed the econometric tools or regarding related papers or this paper specifically, we thank Marianne Baxter, Luca Benati, Larry Christiano, Todd Clark, Oli Coibion, Jim Dolmas, Ayşe Kabukçuoğlu Dur, Rob Engle, Walter Enders, Jonas Fisher, Jordi Gali, Josh Gallin, Barry Jones, Ed Knotek, Markus Kontny, Andy Levin, Kurt Lunsford, Christian Mathes, Sandeep Mazumder, Loretta Mester, Athanasios Orphanides, Katia Peneva, Steve Reed, Robert Rich, Laura Rosner, Jeremy Rudd, Raphael Schoenle, Jim Stock, Christopher Stivers, Murat Tasci, Can Tian, Mark Watson, John C. Williams, and seminar participants at the 2003 North American Summer Meetings of the Econometric Society, the 2004 and 2019 Midwest Macroeconomics Conferences, the 2013 EABCN *Inflation Developments after the Great Recession* conference, the 2020 Econometric Society World Congress, the Board of Governors, the Cleveland Fed, SUNY-Binghamton, and Virginia Tech. All views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Cleveland or the Federal Reserve System.

0. Executive Summary (Not for publication)

Persistence-dependence in the relationship between inflation and the unemployment rate simply means that inflation responds differently to relatively persistent fluctuations in unemployment than it does to relatively transitory ones.

This persistence-dependent relationship sounds exotic, but it is actually just a systematic elaboration – only here applied to the unemployment rate coefficient in the Phillips curve – of the profession's permanent income hypothesis insight in the 1960s, to the effect that the size of the "marginal propensity to consume" coefficient on a fluctuation in disposable income in a simple linear Keynesian consumption function depends on the persistence in this fluctuation. The only difference here is that the unemployment gap is partitioned into three components – with differing persistence levels – that add up to the original unemployment gap data.

Our paper is a detailed investigation into the full nature of the Phillips curve relationship between inflation and unemployment, relaxing two restrictions in the standard linear Phillips curve. This investigation reveals surprisingly deep insights about the nature of the inflation process and its relationship to business cycle dynamics.

Using a one-sided filtering approach – crucial here because of likely feedback in the relationship – we decompose unemployment rate time-series data into three components: a persistent component, composed of all fluctuations that last longer than 48 months; a moderately persistent component, composed of all fluctuations that are completed between 12 and 48 months; and a transient component, which is the remainder. We relate inflation to these three components simultaneously. We further allow for asymmetry (around zero) for each component's coefficient,[1] but permit the data to reject this hypothesis.

Consistent with existing theory, we find that all three of these persistence components have a statistically significant influence on inflation, and we reject symmetry in each of the three

[1] We need to make use of an estimate of the natural rate of unemployment only to partition the persistent component into positive and negative parts. Our results are robust to different estimates of the natural rate.

coefficients. The coefficient pattern is interesting: it is the negative fluctuations in the most persistent unemployment component that enter the model with a significant coefficient, whereas it is the positive fluctuations in the moderately persistent and transient unemployment components that enter with significant coefficients. These three statistically significant coefficients are stable across the sample; in our specification, there is no weakening in the Phillips curve relationship over the Great Recession. Of these coefficients, the one that is most statistically noteworthy is the strong negative coefficient on the positive fluctuations in the moderately persistent unemployment component.

Importantly, these results are economically interpretable: they imply that the observed Phillips curve relationship varies across the phases of the business cycle. In particular, at the onset of a recession, when unemployment is rising rapidly and the moderately persistent and transient components of the unemployment rate both swing positive, we find that there is a very strong downward force on inflation. But as the recovery begins, the variation in these two unemployment components becomes negligible, and our estimated model then indicates that the relationship between labor market slack and inflation essentially vanishes. (As noted above, there is no apparent statistical relationship between inflation and positive fluctuations in the most persistent component.) Only when the economy enters its overheating phase – i.e., when the most persistent component of the unemployment rate drops below the natural rate – does a Phillips curve relationship re-emerge in our model.

Because our model is statistically rich enough to have stable coefficients, it is able to explain why the usual Phillips curve models – which are misspecified in that they ignore the nonlinearities induced by persistence dependence in the relationship – yield estimated Phillips curve relationships that shift about over time and forecast poorly. In contrast, its good inflation forecasting performance validates our model. Using coefficients estimated on the pre-2007 data (and no historical inflation data beyond that point), our model forecasts inflation through the

Great Recession quite well, and its out-of-sample inflation forecasting throughout the sample period of 1985-2016 dominates benchmark models.

Most importantly, however, our persistence-dependent Phillips curve model eliminates the "inflation puzzles" that have so troubled analysts (and monetary policy officials). Thus, for example, our Phillips curve relationship has not deteriorated in recent years, although it does indicate that the relationship between inflation and the unemployment gap is decidedly weaker during an overheating period than it is during a downturn. Therefore, our model does not suffer from "missing inflation;" and it explains why inflation appeared to remain "stubbornly low" from 2010 to 2020.

At the same time, our results are broadly consistent with much extant theoretical work, and the economic significance of the time pattern of a fluctuation in the unemployment gap can potentially point theorists in a productive direction.

Our work further has notable implications for monetary policy, inasmuch as it empirically establishes that the responsiveness of inflation to labor market slack is business-cycle-phase-dependent (or alternatively, that the slack measure relevant for inflation dynamics departs markedly from conventional activity gaps). Our findings indicate that is much more difficult for policy to increase inflation than to decrease inflation. Inflation falls rather sharply when the unemployment rate is rising rapidly (at the onset of a recession). Shortly thereafter, unemployment ceases to influence inflation: it exerts no downward force, and neither does its decline reduce downward force. Inflation, while anchored, only gradually moves towards long-run inflation expectations. Inflation again becomes responsive to the unemployment rate only when the economy overheats – but our coefficient estimates indicate that overheating must be quantitatively large to push inflation up notably.

Finally, our work suggests that inflation will fall well below zero in 2021, and that the U.S. will experience another very prolonged period of low inflation.

1. Introduction

The Phillips curve relationship remains central to macroeconomics and plays an absolutely fundamental role in monetary policy deliberations (see, for example, Brainard, 2019), not least because this relationship lies at the core of the structural models that dominate current monetary policy discussions. But by most accounts, inflation dynamics over the Great Recession seem to have diverged markedly from their previous patterns, posing a number of puzzles to the existing understanding of the Phillips curve relationship. The most prominent puzzle is the missing disinflation: given the large amount of labor market slack that persisted for many years, standard pre-Great-Recession Phillips curve specifications predicted far lower inflation over the Great Recession than actually occurred (see, for example, Ball and Mazumder, 2011; Coibion and Gorodnichenko, 2015). It seems almost universally believed that the Phillips curve relationship has weakened.[2] But the inflation puzzles extend beyond missing disinflation; for instance, the inflation puzzle prior to the COVID collapse was the opposite: shouldn't inflation have been higher, given how long the unemployment rate had been below conventional estimates of its natural rate? Inflation expectations are thought to be currently anchored at target, but this "missing inflation" is believed to threaten this anchoring.

---

[2] Bob Hall famously stated in his AEA presidential address (2011), "It is not news that NAIRU theory is a failure" and that standard New Keynesian Phillips curves "cannot explain the stabilization of inflation at positive rates in the presence of long-lasting slack." Bullard (2017) stated, "The results shown here call into question the idea that unemployment outcomes are a major factor in driving inflation outcomes in the U.S. economy." In July 2019 testimony, Federal Reserve Chair Jerome Powell stated, "The relationship between the slack in the economy or unemployment and inflation was a strong one 50 years ago … and has gone away" (Li, 2019). The Economist (2020) argues that strong jobs reports no longer cause markets to expect rate hikes, owing to the apparently vanished Phillips curve. For similar sentiments, see The Economist (2017), Summers (2017), and Blinder (2018).

These inflation puzzles – as well as a number of other findings and puzzles in the literature, such as the apparent time variation in the relationship,[3] and the odd behavior of the reverse-engineered NAIRU in Coibion and Gorodnichenko (2015) – completely disappear with the persistence-dependent specification of the Phillips curve relationship proposed here.[4] Using coefficients estimated in 2006, recursive forecasts from our specification (conditioned only on the time path of unemployment) well-predict inflation over the *entire* Great Recession: there is no downward-speed puzzle, no missing disinflation, no missing inflation, and neither was inflation in early 2020 "stubbornly low" (FOMC minutes, July 30-31, 2019). Clearly, these facts are highly relevant for policy deliberations today.[5] We emphasize that, under our specification, even the Great Recession did not alter the dynamics of inflation. Moreover, our findings have a natural interpretation in terms of business-cycle-phase dependence and are consistent with extant theory.

Persistence-dependence (or frequency-dependence) in a relationship between two time-series variables *Y* and *X* implies that *Y* responds differently to persistent fluctuations in *X* than it does to transitory fluctuations in *X*. Inflation's persistence-dependent relationship to unemployment sounds exotic, but actually has deep roots in empirical macroeconomics: Friedman (1968) and Phelps (1967, 1968) established that the most persistent unemployment rate

---

[3] See, for example, Clark and McCracken (2006) or Stock and Watson (2007). Similarly, Luengo-Prado et al. (2018) document that the relationship appears to have vanished after 2009. (We agree!)

[4] Persistence-dependence is defined below. We explain Great Recession inflation dynamics without reference to biased inflation expectations (cf. Coibion and Gorodnichenko, 2015) or to the short-term unemployment rate (cf. Ball and Mazumder, 2019). Regarding the latter paper, we agree with those authors that so-called "core" inflation measures are deficient and mask the true Phillips curve relationship; see Appendix A.2.

[5] Indeed, as we note in Section 4, our model at the time of this writing (August 2020) predicts very sharp disinflation over the coming year, bottoming out below -1.5 percent in 2021.

fluctuations, that is, natural rate fluctuations, are unrelated to inflation. Further, persistence-dependence exists in many other macroeconomic relationships. Here are four additional examples: First, the "permanent income hypothesis" implies that consumption should mainly respond to persistent movements in income. Second, RBC modeling was built upon the presumption that business cycle relationships are distinct from low-frequency relationships – and this idea has recently regained traction, as we discuss below.[6] Third, Friedman (1988) and Cochrane (2018) both argued that (transient) measurement error would give rise to what we call persistence-dependence. Finally, data are routinely de-seasonalized under the presumption that relationships at seasonal frequencies differ from those at other frequencies. Until quite recently, however, persistence-dependence was rarely examined, carefully and explicitly.

A nascent body of research, however – building upon earlier work by Comin and Gertler (2006) – is exploring the frequency domain for clues about business cycle drivers and dynamics. Angeletos, Collard, and Dellas (2020) have recently argued that assessing the root causes of business cycles requires delving into the drivers of unemployment, output, consumption, and the like at different frequencies. Beaudry, Galizia, and Portier (2020) likewise highlight medium-frequency cycles and argue that it is crucial to examine the properties of the data at different frequencies in order to discriminate across classes of models. And numerous recent studies have located direct evidence for persistence-dependence in macroeconomic relationships. Cochrane (2018) and Ashley and Verbrugge (2015) find that the velocity of money has a persistence-dependent relationship to interest rates. Blundell, Low, and Preston (2013) find that at the micro

[6] For example, Beaudry et al. (2020) state: "Therefore, in order to evaluate business cycle properties, one needs to find a way to extract properties of the data that are unlikely to be contaminated by the lower-frequency forces that are not of direct interest." Below, we note how standard approaches, such as pre-filtering the data prior to further analysis, will typically lead one astray, and offer a better alternative.

level, consumption responds to persistent movements in after-tax income, and Arellano, Blundell, and Bonhomme (2018) find that persistent earned income shocks are harder to insure against, particularly for young families with low assets. Ashley and Li (2014) find that state-level consumption has a persistence-dependent relationship to movements in both housing wealth and stock wealth. Ciner (2015) finds that stock returns have a persistence-dependent relationship to inflation movements. And Ashley, Tsang, and Verbrugge (2020) find that historical FOMC policy responses to inflation and unemployment are persistence-dependent, with persistent movements in the unemployment rate, and transitory movements in inflation, being ignored by policymakers.[7]

The research question motivating the present research is the following: does the natural rate distinction fully capture persistence-dependence in the Phillips relationship? For instance, the work of Stock and Watson (2010) suggests that the Phillips curve relationship is concentrated at moderate persistence levels and that this relationship is asymmetric – that is, it holds only for *increases* in the unemployment rate. Similarly, transient fluctuations in the unemployment rate seem rather unlikely to influence inflation. And since persistence-dependence implies time-dependence in misspecified regression models, could persistence-dependence explain the recent apparent weakening of the Phillips curve that has occupied so much attention in recent years?

---

[7] Other studies include: (a) Caraiani and Gupta (2018), where the Bank of England responds only to persistent movements in the real exchange rate; (b) Yanfeng (2013), where Japanese industrial activity and inflation have a persistence-dependent relationship to oil prices; (c) Benati (2009), where inflation is mainly related to low-frequency movements in money; (d) Reynard (2007), where the money-inflation relationship depends upon persistence-dependence in velocity; and (e) Cochrane (1989), where interest rates respond negatively to transitory movements in money growth.

With these research antecedents in mind, we here more fully explore the persistence-dependence in this relationship, using recently developed econometric tools that allow the data to speak very transparently as to the nature and form of this dependence, if it exists. Our approach is entirely empirical. We review relevant extant theory (in Section 3.2 and Appendix A.8) and we observe and describe the ways in which our empirical findings are consistent with this body of theory; but we leave the consequent further development of such theoretical modeling efforts for future work. Below we focus instead on the empirical specification of a statistically adequate reduced-form Phillips curve relationship between unemployment fluctuations and suitably-detrended trimmed-mean PCE inflation (whose superiority over "core" PCE inflation is discussed below).

The pattern of persistence-dependence we find has a natural interpretation. We find that there are three distinct empirical Phillips curve relationships, each apparently operative during a different phase of the business cycle. At the beginning of a recession, when unemployment is rising rapidly – so that the moderately persistent component of the unemployment rate becomes positive – we find a big downward impact on inflation. Once the economy has bottomed out and unemployment starts to recover, we find that the Phillips curve relationship essentially vanishes.[8] Then – late in the expansion, when the unemployment rate persistently falls below the natural rate[9] – we find that a Phillips curve relationship reappears, albeit in weaker form. One immediate implication of our findings is that DSGE models with conventional Phillips curves are likely to greatly underestimate the drop in inflation when a recession begins. A second implication is that

[8] See Luengo-Prado, Rao, and Sheremirov (2018). Laxton, Meredith, and Rose (1995) and Debelle and Laxton (1997) also draw attention to the weak relationship between positive slack and inflation. This finding does not fit neatly into a typical empirical specification of a Phillips curve, but aligns nicely with theory; see Appendix A.8.

[9] We discuss the appropriate natural rate concept (and measurement) below.

conventional measures of slack are rather poor proxies for the activity gap that is actually relevant for inflation, particularly during recessions. In particular, we find that the persistence dependence in the data cannot be approximated using a straightforward nonlinear function of the CBO unemployment rate gap.

But does this new reduced-form Phillips curve specification perform well out of sample? We find that it does. As noted above, using data only up through 2006 for coefficient estimation, we find that recursive forecasts from our specification (conditioned only on historical unemployment data) well-predict inflation over the entire Great Recession. As a second test of our specification, we demonstrate that – even though our goal is to understand inflation, not to develop a forecasting model – out-of-sample forecasts generated by a simple variant of our model outperform some standard benchmarks.

2. A Persistence-Dependent Phillips Curve

2.1 Method

Our approach differs fundamentally from the existing Phillips curve literature in that we make use of the "persistence-dependent regression" econometric methodology developed in Ashley and Verbrugge (2009) and Ashley et al. (2020). These methods allow us to use the data to determine the relationship between inflation and movements in the unemployment rate at different persistence levels.

In particular, we take a standard reduced-form Phillips curve specification and then decompose the coefficient on the unemployment rate gap by persistence level (frequency) – allowing for asymmetry in each coefficient,[10] but permitting the data to reject it. We select three

---

[10] Asymmetry has been a hypothesized feature of the Phillips curve from its inception (Phillips 1958), and many previous studies have located evidence for asymmetry or convexity (for example, Debelle and Laxton 1997, Detmeister and Babb, 2017, or Murphy, 2017). Asymmetry is built into the Stock and Watson (2010) recession gap.

persistence components (denoted "persistent," "moderately persistent," and "transient") on largely *a priori* grounds: highly persistent fluctuations encompass all fluctuations whose reversion period exceeds 48 months; moderately persistent fluctuations encompass fluctuations with reversion periods greater than 12 months and less than 48 months; and transient fluctuations encompass all fluctuations with reversion periods of 12 months or less.[11] Our empirical procedure allows the data to inform us as to whether or not this particular partitioning is useful.

The persistence-dependent regression method used here is both straightforward and, to an appropriate degree, flexible, leading to estimation/inference results that are, in the present context, easily interpreted. Essentially, we use one-sided filtering to partition the real-time unemployment rate into components with differing levels of persistence, but which (by construction) sum up exactly to the original unemployment rate series. In particular, we use a moving window to filter the real-time u_t data at each time *t* in a one-sided (backward-looking) manner, partitioning the time *t* observation of the unemployment rate into the three persistence (or, equivalently, frequency) components described above. (We subsequently decompose the persistent component into positive and negative parts using a natural rate estimate, u_t^*.[12]) Since, by construction, these components sum to the original time series – that is, to the real-time $u_t - u_t^*$ data – it is then straightforward to use least squares fitting so as to assess whether inflation movements are related differently to each of these three "persistence components" of the original $u_t - u_t^*$ data.

---

[11] The 48-month cutoff was inspired by the Stock and Watson (2010) recession gap. We present a comparison to that gap in Appendix A.3. Our results do not change appreciably if we use a 60-month cutoff.

[12] Henceforth in the manuscript, for brevity we refer to our procedure as "partitioning $u_t - u_t^*$."

This partitioning, the one-sided filtering, and our restriction of the filtering solely to the $u_t - u_t^*$ data are all essential here. Partitioning is necessary in order that – as noted above – these three components of the unemployment rate gap add up to the original data, so that it is easy to test the null hypothesis that the coefficients with which these three components enter a regression model for the inflation rate are all equal. One-sided filtering is necessary for two reasons. First, one-sided filtering – and only one-sided filtering – sensibly comports with the use of real-time unemployment rate data. And second, two-sided filtering – such as ordinary HP-filtering or ordinary spectral analysis not based on a moving window – inherently mixes up future and past values of the unemployment rate gap in obtaining the persistence components, distorting the causal meaning of inference in the resulting inflation model and limiting its use for practical forecasting and/or policy analysis. These distortions from the use of two-sided filtering are particularly severe when the dependent variable is also filtered and when the key relationship likely (as here) involves feedback from the dependent variable (inflation) to the (filtered) components of $u_t - u_t^*$ being used as explanatory variables. Fundamentally, this is because filtering the dependent variable in a regression model implies that the model error term is similarly filtered.[13] For this reason – and because we want and need a model for inflation itself (rather than for some filtered version of it) – we only decompose the unemployment rate gap into persistence components in our modeling below.

---

[13] For more details, see Sargent (1987) or Ashley and Verbrugge (2009). For the same reason, and because such calculations are incompatible with the use of real-time data, two-sided cross-spectral estimates or filtering with wavelets are similarly ruled out for analyses of the present sort. Even absent feedback, transfer function gain and phase plots are substantially more challenging to interpret than our approach; and in the presence of feedback, Granger describes interpretation of such plots as "difficult or impossible" (Granger, 1969).

The length of the moving window of real-time data on $u_t - u_t^*$ specified for use below in partitioning this unemployment rate gap data into the three persistence components defined above was set here at 48 months. The "persistent" component of $u_t - u_t^*$ is thus aggregating all unemployment rate gap fluctuations with "reversion periods" greater than or equal to 48 months in length, including any trend in $u_t - u_t^*$. Other choices could be made in this regard – for example, a 60-month window – but the inflation modeling results reported below are not particularly sensitive to modest modifications in this window length.[14]

[14]Stata and RATS software to implement the decomposition is available from the authors, and an extensive description of our moving-window-based persistence decomposition algorithm can be found in the Technical Appendix to Ashley et al. (2020). For replicability purposes we note that (per footnote 44 on page 48 of that paper) the standard Christiano-Fitzgerald bandpass filter was iteratively used here. Twelve projected values were used in padding out each window; extension of the data via forecasts to improve filter precision has a long history, starting with Dagum (1975). Here these projections were based on real-time quarterly u_t forecasts from the Survey of Professional Forecasters, converted to monthly frequency (for this purpose) using linear interpolation. As with the window length, the inflation regression model results reported below are not very sensitive to these particular choices. In any case, since the three persistence components are constructed so as to add up to the original u_t - u_t^* series, any nonoptimality in these choices will have only weakened the empirical results we report below.

2.2 Persistence Component Results

Figure 1 displays time plots of four time series. The top portion of this figure plots the real-time unemployment rate (u_t)[15] and estimates of the natural rate of unemployment taken from Tasci (2018).[16] Their difference is the unemployment rate gap ($u_t - u_t^*$), that is here decomposed, as noted above, into the three persistence components:

- the "persistent" gap component, including the sample variation in $u_t - u_t^*$ with reversion periods greater than or equal to 48 months; this component includes any time trend in $u_t - u_t^*$;
- the "moderately persistent" gap component, including the sample variation in u_t with reversion periods less than 48 months but greater than 12 months;
- and the "transient" gap component, comprising the remaining sample variation in u_t, which mean-reverts within 12 or fewer months.

The "transient gap" component is, as one might expect, noisy enough so as to not admit of any clean economic interpretation. Consequently – while it is included in the regression models – this persistence component is, for clarity of exposition, only plotted in Appendix A5 and not plotted in the lower portion of Figure 1, which displays the time variation in the "persistent" and

[15]Sources: Real-time unemployment rates: ALFRED, Federal Reserve Bank of Saint Louis; natural rate series: Tasci (2018).

[16]We prefer this measure of the natural rate of unemployment because it is based on labor market flows, without reference to inflation data. See Occhino (2019) on why it is a poor idea to use a u_t^* estimate based on a Phillips curve in a Phillips curve regression model. We use this measure only to partition the low-frequency component of the unemployment rate into positive and negative parts; however, our results are not sensitive to this choice. The two-sided Tasci (2018) estimate may be used because this natural rate estimate evolves independently of inflation.

"moderately persistent" unemployment rate gap components.

This tripartite persistence decomposition of the unemployment gap results in time series that are readily interpretable in terms of business cycle phases. For instance, turning first to the lower portion of Figure 1, it is evident that the "persistent" gap component recognizably capturing the smooth movements in $u_t - u_t^*$. It traces out the three major recessions during the sample period, and prominently depicts the late-cycle "overheating" periods in these business cycles, wherein this component becomes negative. We demonstrate below that overheating is associated with a notable (but modest) Phillips curve relationship.

The "moderately persistent" component is a central feature in the empirical model developed here. Notice that this time series is fairly smooth, and generally close to zero. But it has a marked tendency to fluctuate upward when the unemployment rate is rising sharply; it then dwindles back to near zero shortly after u_t begins to fall. *In other words, this "moderately persistent" unemployment rate gap component fluctuates upward at the onset of each recession and diminishes shortly after the recovery begins.* In short, this persistence component is a reliable indicator of the onset of a rapid and sustained rise in the unemployment rate. We demonstrate below that, consistent with Stock and Watson (2010), this recession-onset period provides most of the strength in the Phillips curve relationship.

Thus, our persistence/frequency partitioning of $u_t - u_t^*$ can be interpreted *economically* as decomposing the unemployment rate gap into its (highly persistent) overall trend, a moderately persistent "signal" (as it were) for the onset of each recession, and transient noise term. And we demonstrate below that this decomposition results in a more stable and useful Phillips curve model – a model that is business-cycle-phase dependent – with profound implications for understanding inflation dynamics, the limits of monetary policy for fighting low inflation, and business cycle modeling.

2.3 Persistence-Dependent Phillips Curve Regression Model

Our starting point is a relatively standard reduced-form Phillips curve, defined in terms of the unemployment gap. In particular, the baseline model specification is

$$\pi_{t+12}^{12} - \pi_t^* = \alpha + \beta_1 \left(\pi_t^{12} - \pi_{t-12}^* \right) + \beta_2 \left(\pi_{t-12}^{12} - \pi_{t-24}^* \right) + \lambda_1 \left(gap_t \right) + \varepsilon_t \qquad (1)$$

where $\pi_{t+12}^{12} \equiv \ln\left(P_{t+12} / P_t \right)$ denotes the 12-month log-change in the price index, π_t^* is an inflation trend measure relevant for this price index (the "PTR" measure from the Board of Governors, which adjusts and extends forecasts from the Survey of Professional Forecasters), and specification (1) includes a traditional unemployment gap term, gap_t, specified in terms of a natural rate: $gap_t \equiv \left(un_t - un_t^* \right)$, where un_t^* is taken from Tasci (2018).[17] Our inclusion of 24 months of lagged inflation (via our use of the current 12-month inflation rate, and of the 12-month inflation rate from a year ago) is in line with typical practice. Most of our analysis focuses on trimmed-mean PCE inflation, the realized-inflation measure that arguably best removes noise from inflation (Mertens, 2016),[18] but in the Appendix we consider other time series of inflation measures in robustness checks.

---

[17] Regression estimates using the CBO un_t^* estimate are also investigated in the sequel.

[18] As this working paper was undergoing final internal review, we learned of Ball and Mazumder (2019). Like us, these authors eschew the use of core PCE and estimate a Phillips curve using an alternative trend inflation indicator, in their case a weighted median PCE. (We present results for a weighted median PCE in the Appendix.) Similarly, both the Reserve Bank of New Zealand and the Bank of Canada now shun the use of "core" (exclusion-based) inflation measures as measures of trend inflation. See Carroll and Verbrugge (2019) and Verbrugge et al. (2018) for additional evidence regarding the superior forecasting ability of the trimmed mean PCE and median PCE over core PCE in forecasting headline PCE movements. Further, as noted by Boston Fed President Rosengren (2019), when trimmed mean PCE diverges from core PCE, it is the latter that moves to eliminate the gap. (This behavior was verified by Robert Rich and Vicki Consolvo of the Cleveland Fed, and then by the authors.) Stock and Watson (2020) also eschew core inflation and stress the use of an appropriate inflation measure.

Prior studies, such as those of Clark and McCracken (2006), Faust and Wright (2013), Zaman (2013), and Clark and Doh (2014), have shown that inclusion of an accurate trend estimate improves the accuracy of inflation forecasts. In the context of Phillips curve estimation, modeling inflation in terms of deviation from the trend π_t^* amounts to assuming that the Phillips curve relationship is silent about the long-run goals of monetary policy and, instead, pertains to fluctuations in inflation that are more closely related to business cycles. While Phillips curve forecasting models often include other variables such as the relative price of energy or imports, these variables are not found to be helpful for our 12-month projections.[19] Here we estimate our models over the sample period 1985-2016, so as to focus on the recent period during which the Phillips curve is thought to have become weak and/or unstable.

Specification (1) imposes some very strong restrictions on the Phillips curve relationship. One important restriction is that, aside from the movements in the natural rate (un_t^*), the relationship between inflation and movements in the unemployment rate is forced to be the same, regardless of whether those unemployment rate movements are persistent or transient. Whatever the theoretical linkage between the unemployment rate and price setting, it is difficult to believe that highly transitory variation in the unemployment rate influences pricing in the same way that business-cycle variation does.[20] A second important restriction is that this specification imposes linearity in its gap term – that is, positive and negative gaps have the same influence on inflation. This linearity assumption departs from the original Phillips curve (Phillips, 1958) that posited a

---

[19] The PTR series is based upon the median long-term inflation expectations of SPF respondents. This median masks persistent heterogeneity among these respondents; see Binder, Janson, and Verbrugge (2019).

[20] The present paper has incorporated persistence-dependent regression methods into modeling the Phillips curve relationship since its inception in the early 2000s; recent work is now coming around to this view – e.g., see Stock and Watson (2020), which focuses on the business-cycle-frequency relationship of the Phillips curve.

relationship that was steeper at higher levels of economic activity, and numerous papers (e.g., Clark, Laxton and Rose, 2001; Barnes and Olivei, 2003; Fisher and Koenig, 2014; Nalewaik, 2016) have located evidence for this type of nonlinearity in the Phillips curve.

We thus relax these assumptions in a very simple way in our second specification. Define $\left(un_{persist,t}\right)$, $\left(un_{mod.persist,t}\right)$, and $\left(un_{transient,t}\right)$ as the persistent part, the moderately persistent part, and the transient part of un_t, the (real-time) unemployment rate at *t*. Allowing for asymmetry in each term, our specification then extends Equation (1) as follows:

$$\begin{aligned}\pi_{t+12}^{12}-\pi_t^* &= \alpha+\beta_1\left(\pi_t^{12}-\pi_{t-12}^*\right)+\beta_2\left(\pi_{t-12}^{12}-\pi_{t-24}^*\right)+\lambda_1^+\left(un_{persist,t}^+\right)+\lambda_1^-\left(un_{persist,t}^-\right)\\ &\quad+\lambda_2^+\left(un_{mod.persist,t}^+\right)+\lambda_2^-\left(un_{mod.persist,t}^-\right)+\lambda_3^+\left(un_{transient,t}^+\right)+\lambda_3^-\left(un_{transient,t}^-\right)+\varepsilon_t\end{aligned} \tag{2}$$

where $\left(un_{persist,t}^+\right)$ is the positive part of $\left(un_{P,t}-un_t^*\right)$, $\left(un_{mod.persist,t}^+\right)$ is the positive part of $\left(un_{mod.persist,t}\right)$, and $\left(un_{transient,t}^+\right)$ is the positive part of $\left(un_{transient,t}\right)$, and other terms are defined analogously.

Note that asymmetry in the less-persistent components mirrors the asymmetry built into the recession gap term in Stock and Watson (2010). In Section 2.4 below, we examine whether this asymmetry is warranted at each persistence level; i.e., we test the null hypotheses that $\lambda_1^+=\lambda_1^-$, $\lambda_2^+=\lambda_2^-$, or $\lambda_3^+=\lambda_3^-$. Foreshadowing these results, both these tests and information criteria reject a variety of more parsimonious models, including symmetric and/or non-persistence-dependent specifications, such as Equation (1). We also perform a Chow test of coefficient stability to test whether the coefficient estimates change after 2006:12.

2.4 In-Sample Inference Results

2.4.1 Discussion of Coefficient Estimates

Table 1 below displays the OLS parameter estimates for the coefficients in Specification (1) and Specification (2), with estimated t-ratios quoted beneath each coefficient estimate.[21]

Consider first the estimation results for the "standard" Phillips curve specification, Equation (1). Here notice that the evidence for the existence of any Phillips curve relationship between inflation and the unemployment rate gap ($u_t - u_t^*$) is exceedingly weak for this model specification. In fact, with an estimated t-ratio of only -1.54, the estimate of λ, the coefficient on ($u_t - u_t^*$) in Equation (1), has the right sign, but it is not statistically different from zero (on a two-tailed test) at even the 10 percent level of significance. And this coefficient is unstable: a standard Chow test rejects the null hypothesis that the parameters in Equation (1) are stable – when the sample data are partitioned into the period 1985:1 through 2006:12 versus 2007:1 through 2016:12 – with $p = 0.048$.

This first set of results is hardly new or surprising: multiple studies in the literature have found ample evidence for a weak and unstable Phillips curve relationship over the past few decades.

Next consider the analogous estimation and inference results for Equation (2), our persistence-dependent re-formulation of the Phillips curve regression model specification. First, notice that the adjusted R^2 is notably larger for this estimated model than for Equation (1) and that the BIC is notably smaller: evidently the improvement in the fit of Equation (2) to the

[21] These estimated standard errors are (12-month) HAC standard error estimates, as the fitting errors (due to the cumulation in the dependent variable) appear to be serially correlated. Diagnostic checks with regard to possible heteroscedasticity in the model errors are therefore not quoted, as the HAC standard error estimates are consistent in this case anyway. The fitting errors for Equation (2) show no evidence of notable outliers.

sample data more than compensates for its greater complexity. Second, notice that there is ample evidence here for rejecting both the individual null hypotheses of coefficient symmetry on the two most persistent components – that is, $H_0 : \lambda_1^+ = \lambda_1^-$, $H_0 : \lambda_2^+ = \lambda_2^-$ – and even stronger evidence to reject the joint null hypothesis of coefficient symmetry on all three persistence components. (These results are consistent with evidence in Morris, Rich, and Tracy (2019) on asymmetry in the wage-based Phillips curve; see Appendix A.7.)

More importantly, however, note that three of the six coefficients on the unemployment rate gap persistence components are negative and quite statistically significant. In particular, notice that the coefficient on λ_2^+, quantifying the impact of positive fluctuations in the moderately persistent component of ($u_t - u_t^*$), is substantially negative (at -1.67), with an estimated t-ratio of -7.91; so $H_0 : \lambda_2^+ = 0$ can be rejected with $p < 0.00005$. Evidently, the positive fluctuations in the moderately persistent component of ($u_t - u_t^*$), observed (in Figure 1) at the onset of each recession in the sample, match up quite nicely with concomitant drops in the inflation rate at these times. This result is consistent with economic intuition and accords with the findings of Stock and Watson (2010).[22]

The remainder of this section discusses these Equation (2) estimation results in more depth, and interprets the results in terms of business-cycle-phase dependence. The regression

---

[22]We note that qualitatively similar results in this regard also obtain using different measures of inflation, or using the CBO estimates of u_t^* instead of the Tasci (2018) estimates used here; see Appendix A.1. We also obtain qualitatively similar results if we expand the range of the moderately persistent u_t - u_t^* component to include variation with reversion periods up to 60 months in length by using a larger window length, or if we re-specify the model using six-month changes in inflation. While this paper makes no attempt to be a multi-country study, we also note that preliminary analyses using data from Australia yield a similar pattern. These results are available on request.

results in Table 1, as noted above, indicate compelling evidence for asymmetry in the relationship between inflation and all three persistence components. The coefficient estimates differ notably for positive and negative gaps, and the formal hypothesis testing results very clearly reject the null hypotheses that these coefficients are equal, for either the most persistent unemployment component coefficients (λ_1^+ and λ_1^-) and for the moderately persistent unemployment coefficients (λ_2^+ and λ_2^-) – though we can only reject the null hypothesis (λ_3^+ = λ_3^-), corresponding to the transient unemployment components, at the 8 percent level of significance.[23] We argue next that this asymmetry is eminently sensible and aligns well with economic theory.

2.4.2 Business-cycle-phase dependence

These results have a natural interpretation in terms of business-cycle-phase dependence. How so? The three phases of a typical business cycle are the recovery period, the overheating period, and the recession (when the economy moves from a peak to a trough). Our results indicate that the Phillips curve relationship is distinct across these three phases.

[23] The BIC also worsens notably when symmetry is imposed. We do not report these tests for brevity.

Table 1. Phillips Curve Regression Estimation Results[a]

Regressor		Specif. (1)	Specif. (2)
$u_t - u_t^*$	λ	-0.09	
	(t-stat)	(-1.54)	
Persistent component of $u_t - u_t^*$	λ_1^+		0.03
	(t-stat)		(0.64)
	λ_1^-		-0.27
	(t-stat)		(-3.16)
Moderately persistent component of $u_t - u_t^*$	λ_2^+		-1.67
	(t-stat)		(-7.91)
	λ_2^-		0.10
	(t-stat)		(0.14)
Transient component of $u_t - u_t^*$	λ_3^+		-0.51
	(t-stat)		(-2.36)
	λ_3^-		-0.03
	(t-stat)		(-0.23)
Lagged inflation	β_1	0.47	0.28
	(t-stat)	(7.74)	(4.01)
	β_2	0.11	0.34
	(t-stat)	(1.33)	(6.26)
	constant	-0.09	-0.05
	(t-stat)	(-0879)	(-0.52)
Adjusted R-squared		0.56	0.75
BIC		1.04	0.54
Hypothesis Test Rejection P-Values	H_0: $\lambda_1^+ = \lambda_1^-$		<0.005
	H_0: $\lambda_2^+ = \lambda_2^-$		0.02
	H_0: $\lambda_3^+ = \lambda_3^-$		0.08
	H_0: $\lambda_1^+ = \lambda_1^- = \lambda_2^+ = \lambda_2^- = \lambda_3^+ = \lambda_3^-$		<0.005
H_0: (Chow test): {coefficients unchanged before and after 2006:12}		0.048	0.26

a. Figures in parentheses are estimated t-statistics, based on (12-month) HAC standard error estimates.

First focus on $\hat{\lambda}_1^+$ and $\hat{\lambda}_1^-$, coefficient estimates that pertain to the relationship of very persistent (low-frequency) movements in the unemployment rate gap to inflation, those with reversion periods of 48 or more months (see Figure 1). A typical business cycle has a long recovery phase, beginning once the recession has bottomed out and unemployment has begun falling. During this recovery, the unemployment rate is persistently above the natural rate; thus $un_{Persist,t}^+$ is positive and $un_{Persist,t}^-$ is zero. Note that $\hat{\lambda}_1^+$ is both quantitatively negligible and statistically indistinguishable from zero. This result indicates that the estimated impact of the unemployment gap on inflation is essentially *zero* during the recovery phase, despite the fact that the conventional unemployment gap is large. Putting this more starkly, a persistently high unemployment rate *per se* has *no influence* on (more properly, no relationship to) inflation. However, once the unemployment rate moves persistently below its natural rate – i.e., when the economy moves into the overheating phase – a Phillips curve relationship re-emerges. During this phase, and consistent with much previous research, there is a notable upward influence on inflation; in our results, this is reflected in the statistically (and economically) significant coefficient estimate $\hat{\lambda}_1^-$ = -0.27. The estimated size of this coefficient is directly comparable to, and much larger than, the estimate from conventional Phillips curve specification – of -0.09 – given in the first column of Table 1.

Most of the research investigating nonlinearity in the Phillips curve has focused upon a differential force from positive and negative unemployment gaps, but this distinction fails to capture all the subtlety we find in the relationship. We turn to the final phase of a business cycle, the onset of a recession.

Both λ_2^+ and λ_3^+ are negative, statistically significant, and large. Observe that (in the time plot of Figure 1) the moderately persistent component consistently fluctuates upward as a recession begins; the transient component is noisier but nonetheless also fairly consistently fluctuates upward at the same time.[24] Owing to the magnitudes of λ_2^+ and λ_3^+, there is thus a

[24] Admittedly, the positive part of the transient component is less reliably related to a recession, but it nonetheless tends to positively co-move with the moderately persistent component; the correlation is 0.26. See Appendix A.5 for a time plot. In any case, there is less evidence for the importance of $\hat{\lambda}_3^+$; its quantitative magnitude is far smaller than $\hat{\lambda}_2^+$, and it is statistically significant only at the 5% level.

strong downward force on inflation at the *beginning* of a recession – this is the recessionary downward force highlighted by Stock and Watson (2010).[25] The fact that this force vanishes once the recovery begins is consistent with the findings of that paper and explains, for example, the findings in Luengo-Prado, Rao, and Sheremirov (2018): "...we find robust evidence of a structural break in the Phillips curve slope around 2009-2010. The co-movement of sectoral inflation rates and labor market slack has weakened, and it is now almost negligible" (p. 1).

Are these findings exotic, or sensible? Ours is an empirical paper; we provide no new theory here. But our findings are consistent with much previous empirical work (as noted above), with basic business-cycle facts, and with an abundance of extant theory. For instance, why do prices fall so sharply at the onset of a recession? This is consistent with standard industrial organization theory: during collapses, price wars can break out, as firms attempt to steal market share but (once the recovery begins) prices start to edge up again; for macro applications, see, e.g., Gilchrist et al. (2017) and Hong (2019). Furthermore, labor comprises the biggest marginal cost component, but labor share can be a misleading proxy of marginal costs. With short-run labor fixity, standard textbook analysis demonstrates that the marginal cost curve is much steeper. Indeed, overtime labor drops sharply at the onset of a recession, and marginal cost may also drop because workers are underutilized. In this regard, Basu and House (2016) note that "recent empirical studies suggest that the cyclicality of real wages is greater than conventional wisdom would suggest," due to compositional changes and to a wedge between the allocative and remitted wage; they find that allocative wages fall sharply after a monetary tightening. In short, they note, "real marginal cost, properly computed, is strongly procyclical." And why does a negative unemployment gap result in a stronger force than an equally-sized positive

[25] We compare their specification to ours in Appendix A.3.

unemployment gap? This can result from downward price rigidity, from bargaining considerations (see, e.g., Blanchflower and Oswald, 1990 and Moscarini and Postel-Vinay, 2019) or from capacity constraints (see, e.g., Alvarez-Lois 2006 and Kuhn and George 2019). Appendix A.8 provides a broader discussion of extant theory consistent with our findings.

2.4.3 A stable Phillips curve, and comparison to other findings

Note that – in our persistence-dependent model – the coefficients in the Phillips curve relationship are stable over time. In particular, a Chow test of parameter stability for Equation (2) fails to reject the null hypothesis that the model coefficients change after 2006:12: the rejection p-value for this test is 0.26. Evidently, the Great Recession did *not* significantly weaken or alter the Phillips curve relationship. Section 3 below further examines the stability (and the consequential forecasting ability) of the persistence-dependent Phillips curve specification embodied in Equation (2), using out-of-sample validation methods.

In Appendix A.4, we discuss how our results compare to some other prominent findings in the literature. For example, we note there that we can use our results to reinterpret a finding in Coibion and Gorodnichenko (2015) regarding their reverse-engineered NAIRU. Similarly, the findings of Stock and Watson (2020), that attempt to focus solely on business-cycle frequencies, mesh well with the persistence-dependent regression results here. [26] We also discuss there how our results naturally give rise to episodic forecast improvements and to time variation in Phillips curve coefficients for specifications that ignore persistence-dependence. Finally, we argue that

---

[26] Stock and Watson (2020) mainly use a two-sided bandpass filter to remove all but business-cycle frequencies, but also, as a robustness check, use a "year-over-year" filter, which is one-sided. Like the Hamilton (2019) filter, this latter filter does not allow a decomposition of a variable into differing persistence levels, but being one-sided, it is at least not subject to the distortions inherent to two-sided filtering.

our results explain numerous studies that find a convex-concave aspect to the Phillips curve relationship and studies that adduce evidence for regime switching.

3. Out-of-Sample Evidence

The statistical significance of the inference results discussed above strongly support our nonlinear (asymmetric) Phillips curve formulation, disaggregated according to persistence level per Equation (2). These statistical results fundamentally arise from the fact that this specification fits the historical sample data notably better than do the alternatives we considered, even – via consideration of the BIC measure – allowing for the number of additional coefficients estimated in Equation (2).

We find these results persuasive, but not necessarily definitive, in view of the usual concerns as to "data mining." We also, in particular, wondered whether our specification can explain inflation dynamics over the Great Recession. To address these issues, we conducted two additional – "out-of-sample" – exercises. First, we present an analysis of what can be called "partially recursive conditional" forecasts, to examine whether our model can resolve the various inflation puzzles apparently arising during the Great Recession; we also compare our conditional forecasts to those derived from the use of the conventional specification, Equation (1). Further, in a second set of calculations, we examine unconditional forecasts. In particular, we present supporting results based on out-of-sample (OOS) forecasting calculations using the Giacomini-Rossi and Diebold-Mariano testing frameworks. These exercises confirm our in-sample results.

3.1 Conditional Recursive Forecasts

For our forecasting exercises, we drop the Equation (2) regression terms with statistically insignificant estimated coefficients in Equation (2) – for example, λ_1^+, λ_2^-, and λ_3^- – and use the pared-down forecasting model:

$$
\begin{aligned}
\pi_{t+12}^{12} - \pi_t^* &= \alpha + \beta_1 \left(\pi_t^{12} - \pi_{t-12}^* \right) + \beta_2 \left(\pi_{t-12}^{12} - \pi_{t-24}^* \right) \\
&\quad + \lambda_1^- \left(un_{Persist,t}^- \right) + \lambda_2^+ \left(un_{mod.persist,t}^+ \right) + \lambda_3^+ \left(un_{transient,t}^+ \right) + \varepsilon_t
\end{aligned}
\tag{3}
$$

In this section, we use Equation (3) to generate recursive conditional forecasts: these forecasts are conditioned on the historical unemployment values, but they are recursive in that each forecast draws its needed lagged inflation-deviation values from its own recent inflation forecasts. We compare these conditional forecasts to analogous ones obtained from Equation (1), the parallel model that instead conditions on the path of the CBO unemployment gap, and does not allow for the asymmetric business-cycle-phase dependence in Equation (3). In both cases, the model coefficients are fixed at the values estimated using the data prior to 2007:1.

We plot both of these forecasts (along with the actual inflation time series) below in Figure 2. The conditional forecasts generated by Equation (3) do a very respectable job of tracking the broad contours of the evolution of inflation over the Great Recession and the recovery: the sharp dip in inflation, the partial bounceback, and the very slow movement toward long-run expected inflation. [27] Based on our new model specification, the Great Recession apparently did not substantially alter inflation dynamics; thus, in our framework, there are no "inflation puzzles" to worry about.

In sharp contrast, the similarly conditional forecasts generated by the linear model of Equation (1) are quite poor and do generate the well-known set of "puzzles": inflation decelerated far more rapidly than this model predicts,[28] yet overall inflation fell by less than this model predicts, as has often been noted. The divergence between the actual dynamics of inflation

[27] Recall that we specify our models in terms of trimmed mean PCE inflation, detrended by long-run SPF forecasts. Faccini and Melosi (2020) provide a theory that generates low inflationary pressures over most of the recovery.

[28] As Clark (2014) has noted, once one properly accounts for trend inflation, a major disinflation puzzle pertains to why inflation fell *so fast* during the recession; our specification gracefully explains the rapid disinflation.

and this model's predictions is striking. Indeed, tests of this divergence support the visual impressions from Figure 2: changes in the projections using the CBO gap are actually uncorrelated with changes in the inflation path (the estimated correlation is 0.22 +/- 0.19), whereas the correlation between changes in the Equation (3) projections and changes in the inflation path is substantial, at 0.48 +/- 0.10. These out-of-sample forecasting results reinforce a central message of this paper: a failure to properly specify the relationship between the unemployment rate and the inflation rate, allowing for both asymmetry and persistence dependence, yields unstable parameter estimates, strongly counterfactual conditional forecasts, and misleading conclusions about the nature of the inflation process.

At the time this analysis was conducted (June 2019), the unemployment rate had arguably been below the natural rate for some time. Was there, in fact, missing inflation? From the perspective of our model, the answer is again, perhaps surprisingly, "no": from our Equation (3) model, the March 2018 prediction for 12-month trimmed mean PCE inflation in March 2019 was 2.06 percent, only about +0.1 percentage points above its realized value of 1.94 percent. The overheating force in the economy was moderate in early 2018, as implied by our model; one needs a much bigger unemployment rate gap, of the type seen in some previous recessions, to exert a strong upward force on the inflation rate.[29]

In Appendix A.6, we present an extension to Figure 2, which includes projections from both models when the coefficients in each model are estimated a decade later, using data up through 2016:12. This extension figure underscores the stability of our Phillips curve – its recursive forecasts are nearly identical – while the recursive forecasts from the conventional –

[29] For instance, with an unemployment gap of -0.6, the upward force on inflation is (-0.24)(-0.6)=+.014 over 12 months. The June 2019 gap estimate is in the (-0.6,-0.8) range.

Equation (1) – model estimated in 2016 are much flatter, in keeping with the smaller unemployment gap coefficient estimate. This result further emphasizes the lesser fidelity of the conventional model; these circa-2016 conventional-model coefficient estimates heighten the downward speed puzzle while still under-predicting the inflation recovery.

3.2 GR and DM Out-of-Sample Forecast Tests

We conjecture, along the lines of Stock and Watson (2009, 2010), that the forecast improvement generated by our Phillips curve formulation over benchmark models is likely to be episodic, as the unemployment gap terms in our Equation (3) model are only substantially operative during two portions of the business cycle. To examine this issue, we use the Giacomini and Rossi (2010) fluctuation test, in addition to the Diebold-Mariano test of out-of-sample forecasting improvement.

The Giacomini-Rossi (GR) testing framework is well-suited for comparing the historical out-of-sample forecasting performance of competing models when the relative performance of these models may vary over time. However, as the authors point out, it has somewhat low power to detect overall forecasting quality differences; the Diebold-Mariano test is preferable in that context.

The GR "fluctuation: out-of-sample" (F^{OOS}) test statistic is given by

$$F_{t,m}^{OOS} = \hat{\sigma}^{-1/2} m^{-1/2} \left(\sum_{j=t-m}^{t} \hat{\eta}_j^2 - \sum_{j=t-m}^{t} \hat{\varepsilon}_j^2 \right)$$

where $\hat{\sigma}$ is a HAC estimate of the asymptotic variance of the difference; here we set m equal to 48 months. The GR test is two-sided and is based on rolling-window estimates and forecasts. In Figures 3 and 4, we plot the upper and lower 10 percent and 5 percent critical values and the GR F^{OOS} test statistic for two forecast comparisons. When the F^{OOS} statistic rises above the upper critical value, then the forecast performance of the "alternative" (persistence-dependent PC) model is significantly better (over the previous 48 months) compared to the baseline model. Conversely, when the F^{OOS} term falls below the lower critical value, the reverse is true.

There is a large body of research on the performance of inflation forecasts based on economic activity gaps, relative to forecasts based on univariate benchmark models. A classic reference with regard to this point is Atkeson and Ohanian (2001), who famously found that a naïve univariate model generally outperformed the usual Phillips curve (PC) model, although some papers (such as Brayton, Roberts, and Williams (1999) and Stock and Watson (1999)) noted "deterioration" in PC forecasts prior to this study. Stock and Watson (2009) conclude that PC-based forecasts outperform univariate benchmarks sporadically, in particular, during episodes with a large unemployment gap, that is, exceeding positive or negative 1.5 percent. In contrast, previous research investigating the OOS performance of PC-based forecasting models vis-à-vis similar univariate benchmark models over the post-1985 period typically returns negative results, for example, Rossi and Sekhposyan (2010) and Dotsey, Fujita, and Stark (2017). Below we examine the conjecture that the asymmetric and persistence-dependent PC model proposed above – that is, Equation (3) – outperforms conventional models in which those features are omitted.[30]

In particular, in the remainder of this section, we compare forecasts from our Equation (3) against Equation (1) – with the CBO gap – and against an Atkeson-Ohanian-type model. (Analogous comparisons against the Stock-Watson recession-gap model are provided in the Appendix.) Figure 3 depicts the comparison against a standard (non-persistence-dependent) CBO gap model, Equation (1). Short windows are not appropriate here, since our model sharply

[30] Unlike Dotsey, Fujita, and Stark (2017), we study trimmed-mean PCE inflation here, rather than headline inflation; other inflation estimators are examined in the Appendix.

differentiates between different portions of the business cycle. We consequently use a 20-year window and estimate models from 1985:1 onward; thus, our first forecast is for 2005:1, that is, for the 12-month movement in the (detrended) trimmed-mean PCE between 2005:1 and 2006:1. The F^{OOS} statistic looks back m= 48 months, so the GR test itself thus runs from 2009:1 onward.

In Figure 3, the F^{OOS} line is uniformly above zero, indicating that our Equation (3) specification outperforms the baseline CBO specification from 2009:1 to 2016:12. The forecasting improvement gain from the Equation (3) model is statistically significant at the 5 percent level until mid-2009 and from mid-2015 to the end of the sample. The Diebold-Mariano rejection p-value is 0.01, indicating that taking the sample period *as a whole*, the forecast improvement of Equation (3) over the baseline model is convincing.

In Figure 4, we display analogous forecast comparison results comparing the OOS forecast performance of Equation (3) to that of an Atkeson-Ohanian-type model. The latter model "predicts that inflation over the next four quarters is expected to be equal to inflation over the previous four quarters" (Atkeson and Ohanian, 2001, p. 6). Thus, we compare forecasts from Equation (3) against those from the model

$$\pi_{t+12}^{12} - \pi_t^* = \left(\pi_t^{12} - \pi_{t-1}^*\right) + \eta_t \tag{4}$$

Figure 4 displays convincing GR-test evidence for the episodic forecast improvement of our Equation (3) model over the Atkeson-Ohanian model. As in the comparison against the CBO gap model, our Equation (3) forecasts are better on average over the entire comparison period, and these gains are statistically significant at the 5 percent level from late 2010 to late 2012. For this OOS forecast comparison, the Diebold-Mariano rejection p-value is 0.02, indicating that Equation (3) provides better forecasts than the Atkeson-Ohanian-type model (at the 2 percent level) over the forecasting period as a whole.

The test results discussed above show that our Equation (3) re-formulation of the Phillips curve yields statistically significant improvement in out-of-sample forecasting. We take this improved OOS forecasting performance for our re-formulation of the Phillips curve specification to indicate that the statistical inference results quoted in Section 2 reflect a new set of stable statistical regularities in the historical data, rather than merely an improved in-sample fit of a more flexible model specification.[31]

We note that the interpretation of all reduced-form Phillips curve forecasting models is inherently complicated by endogeneity due to extant monetary policy. In particular, as has been known since Kareken and Solow (1963) and Lucas (1976), the empirical (reduced-form) Phillips curve will generally vary with monetary policy.[32] Thus, to the extent that a central bank is successful in controlling inflation, one might expect that the reduced-form Phillips curve relationship will weaken. Further, consider a forecasting model with an unemployment gap term that is seriously misspecified. Least squares estimation of the coefficient on such a gap term might well be biased toward zero in that case, even though the estimated model will still provide unbiased forecasts on average. Consequently, the inflation forecasts generated by such a model will differ substantially from a better-specified model only episodically, for example, when the gap term is quite large, and we find this to be the case in the current exercise. Also, we note that accurate inflation forecasts at horizons longer than 12 months would require accurate forecasts of both un_t and un_t^*.

---

[31] Examining Equation (3) for routine use in forecasting inflation is beyond the scope of the present paper, however.

[32] For recent studies focused on this point, see Fitzgerald and Nicolini (2014) or McLeay and Tenreyro (2018). Occhino (2019) also provides useful intuition. For a recent structural approach in an open economy New Keynesian model that focuses on inflation as a global phenomenon, see Kabukçuoğlu Dur and Martínez-García (2019).

Lastly, we want to again emphasize that our goal in this paper is not to devise an improved forecasting model, but rather to provide insight into inflation dynamics; this will, in turn, be useful for both structural modeling and policy. Still, the conditional forecasting exercise undertaken above is useful in that it buttresses our claim that our persistence-dependent model – Equation (2) – is a better specification than the usual PC specification, as in Equation (1).

4. Conclusions

Being so central a topic to macroeconomics, the Phillips curve is the subject of a vast literature. We have argued above, however, that most of this literature suffers from fairly severe model misspecification in the posited Phillips curve regression equation. This widespread problem has led to erroneous conclusions about the nature of the PC relationship and to the "inflation puzzles" prominently discussed in the recent literature.

We find that our re-specified reduced-form Phillips curve relationship produces stable coefficient estimates across the 1986-2016 sample period, but that this is not a simple linear relationship. Rather, it is what we term "persistence-dependent," with the form of the relationship between inflation and unemployment fluctuations depending significantly – in both the statistical and the economic sense – on the persistence of these unemployment fluctuations.

Reviewing how the empirically stable specification that we obtain better explains inflation variation in the observed macroeconomic historical record, we find that our Phillips curve model is interpretable as a business-cycle-phase-dependent relationship, and one that is consistent with extant theory. When a recession begins – precisely when the unemployment rate is rising rapidly – the moderately persistent and transient components of the unemployment rate become positive. Our coefficient estimates imply that this induces a large reduction in inflation. After the unemployment rate peaks and begins to slowly descend, the aforementioned components effectively return to zero in the historical data, as plotted in Figure 1. The *persistent* component remains, of course, very large and positive during this descent, but our coefficient estimates imply that during this recovery period, the unemployment rate imparts no downward force whatsoever on inflation. In other words, during the recovery – until the unemployment gap actually becomes persistently negative – this gap has little relationship to inflation. Finally, when the recovery turns into the overheating phase late in the expansion – that is, when the persistent component of unemployment becomes negative – our coefficient estimates imply that a Phillips curve relationship re-emerges, with upward force on inflation.

The in-sample fitting and out-of-sample forecasting results described in Sections 2 and 3 above show that our model specification well explains the time-evolution of inflation during the sample period – even over the Great Recession. In particular, under our model specification all of the "inflation puzzles" noted above disappear. Moreover, notably, the full Phillips curve relationship under our model specification has not materially changed over time – nor recently! – although our model does predict that the inflation-unemployment relationship will appear to be essentially nonexistent at certain times. In contrast, estimated versions of the standard Phillips curve specification effectively average the relationship across differing portions of the business cycle; these portions feature differing "persistence profiles" in the unemployment rate gap, and hence differing inflation-unemployment relationships. Owing to the length of the recovery from the Great Recession, these misspecified formulations provide a single estimated Phillips curve model that is currently anomalously weak – so weak that its very existence is called into question. At the time of this writing, the COVID-19 pandemic has induced a rapid rise in the unemployment rate. Hence, the moderately persistent component of the unemployment gap has become positive. Our model suggests that this will coincide with a rapid reduction in inflation, in which case these mis-specified formulations will then apparently "strengthen" and faith in the conventional (Equation (1)) Phillips curve will be (erroneously) restored.[33] When the economy begins to recover, our model predicts a partial rebound of inflation, followed by a very slow movement of inflation back to trend; and as the recovery progresses, we also predict that standard PC formulations will again begin to find "inflation puzzles."

The reduced-form Phillips curve specification developed here is validated by its stable coefficients across the sample and by its historical out-of-sample forecasting effectiveness. It is not, however, presented here primarily as a contribution to the literature on inflation forecasting, although we hope that the present work can and will stimulate progress by others in that direction. Nor, as a reduced-form model, does the model specification formulated here directly contribute to the theoretical literature on inflation, although (as detailed in Appendix A.8) it is consistent with existing theories, both with regard to the asymmetry in its unemployment responsiveness and with regard to the manner in which it varies across the business cycle. And we hope that theorists will see our empirical finding of persistence-dependence in this relationship as a stimulus to investigate why and how this dependence arises. However, we see the main contribution of our work as identifying and documenting an important new statistical regularity – a new "stylized fact" – that any reasonably complete theoretical model for the US macroeconomy "ought to" imply.[34]

---

[33] The current inflation predictions of our Equation (3) model differ sharply from those of standard Phillips curve models. In August 2020, the 2021 Q1 Blue Chip *bottom-10 average* projection was 0.6 percent for core PCE inflation, with a projected rise to 0.8 percent by Q2. Conditional on a currently-reasonable projection of unemployment over 2020 – slowly declining from its April peak of 14.7 percent – the Blue Chip projections are slightly *below* those from Equation (1). Conversely, our Equation (3) model predicts that 12-month trimmed mean PCE inflation will bottom out in mid-2021 below –1.5 percent, and then begin to recover.

[34] Initial results with data on Australia indicate that our results are not unique to the US. Extending this work to a variety of other countries is beyond the scope of the present paper, however.

More broadly, we would like to emphasize the clear implications of the work presented here with regard to current and future monetary policy deliberations. As noted by John Cochrane (quoted in Steelman, Haltom, and Kenney, 2013), "The prevailing theory of inflation these days has nothing to do with money or transactions: the Fed sets interest rates, interest rates affect "demand," and then demand affects inflation through the Phillips curve" (p. 36). The recent experience of year after year of zero nominal interest rates, anchored inflation expectations, and low inflation suggests difficulty in fine-tuning inflation. Even with anchored inflation expectations, evidently the movement of inflation toward its long-run expected level is quite slow, and it takes an appreciable amount of overheating before there is a significant upward force on inflation. Conversely, we find that the Phillips curve mechanism is, in the other direction, very powerful: inflation can be slowed rapidly via a rapid upward movement of the unemployment rate – that is, a recession. The empirical re-formulation of the Phillips relation developed here harmonizes all of that experience in a relatively simple elaboration of the usual – but empirically unstable and unsuccessful – Phillips relation. This re-formulation explains the observed puzzles associated with the usual models, and its empirical implementation is sufficiently stable over the entire sample period as to provide reasonably accurate conditional forecasts of inflation over the Great Recession. These forecasts underscore the notion that inflation in early 2020 was not "stubbornly low" but was – in the re-formulation of the Phillips curve described here – in fact exactly where its pre-2006 dynamics suggest it should have been, given the evolution of the labor market.

References

Alvarez-Lois, Pedro P. (2004). "Capacity constraints, idiosyncratic demand uncertainty and the dynamics of inflation." *Economics Letters*, 83(1): 15-21. doi:10.1016/j.econlet.2003.09.022.

Alvarez-Lois, Pedro P. (2005). "Production inflexibilities and the cost channel of monetary policy." *Economic Inquiry*, 43(1): 170-193. doi:10.1093/ei/cbi012.

Alvarez-Lois, Pedro P. (2006). "Endogenous capacity utilization and macroeconomic persistence." *Journal of Monetary Economics*, 53(8): 2213-2237. doi:10.1016/j.jmoneco.2006.07.001.

Alves, Felipe. (2019) "Job ladder and business cycles." Manuscript, New York University. https://drive.google.com/file/d/16Rzfy_Eu28a1-XMYUTJTe1r4TMjiyYT_/view.

Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme. (2018). "Nonlinear persistence and partial insurance: Income and consumption dynamics in the PSID." *AEA Papers and Proceedings* 108: 281-286. doi: 10.1257/pandp.20181049

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. (2020). "Business Cycle Anatomy." *American Economic Review* (forthcoming).

Ashley, Richard, and Grace Li (2014) "Re-examining the impact of housing wealth and stock wealth on retail sales: Does persistence in wealth changes matter?" *Journal of Housing Economics* 26: 109-118. doi:10.1016/j.jhe.2014.09.003.

Ashley, Richard, Kwok Ping Tsang, and Randal J. Verbrugge. (2020). "A new look at historical monetary policy and the great inflation through the lens of a persistence-dependent policy rule." *Oxford Economic Papers* (forthcoming). (Currently available as Working Paper 18-14r, Federal Reserve Bank of Cleveland. doi:10.26509/frbc-wp-201814r.)

Ashley, Richard, and Randal J. Verbrugge. 2009. “Frequency dependence in regression model coefficients: An alternative approach for modeling nonlinear dynamic relationships in time series.” *Econometric Reviews*, 28(1-3): 4-20. doi:10.1080/07474930802387753.

Ashley, Richard, and Randal J. Verbrugge. (2015). “Persistence dependence in empirical relations: The velocity of money.” Federal Reserve Bank of Cleveland, Working Paper, no. 15-30. doi:10.26509/frbc-wp-201530.

Asplund, Marcus, Rickard Eriksson, and Niklas Strand. (2005). “Prices, Margins and Liquidity Constraints: Swedish Newspapers, 1990–1992.” *Economica* 72 (286): 349–59. doi:10.1111/j.0013-0427.2005.00418.x

Atkeson, Andrew, and Lee E. Ohanian. (2001). “Are Phillips curves useful for forecasting inflation?” *Quarterly Review*, 25(1): 2-11. URL https://ideas.repec.org/a/fip/fedmqr/y2001iwinp2-11nv.25no.1.html.

Bachmann, Rudiger, and Giuseppe Moscarini. (2012). “Business cycles and endogenous uncertainty.” Manuscript, Yale University.

Baghli, Mustapha, Christophe Cahn, and Henri Fraisse. (2007). “Is the inflation-output nexus asymmetric in the Euro area?” *Economics Letters*, 94(1): 1-6. doi:10.1016/j.econlet.2006.04.001.

Ball, Laurence, N. Gregory Mankiw, and David Romer. (1988) “The New Keynesian economics and the output-inflation trade-off.” Brookings Papers on Economic Activity 19(1): 1-65. doi:10.2307/2534424.

Ball, Laurence, and Sandeep Mazumder. (2011). “Inflation dynamics and the Great Recession.” *Brookings Papers on Economic Activity*, 2011(1): 337-381. doi:10.1353/eca.2011.0005.

Ball, Laurence, and Sandeep Mazumder. (2019). "The nonpuzzling behavior of median inflation." Working Paper 25512, National Bureau of Economic Research. doi:10.3386/w25512.

Barnes, Michelle L., and Giovanni P. Olivei. (2003). "Inside and outside bounds: Threshold estimates of the Phillips curve." *New England Economic Review*: 3-18. URL https://ideas.repec.org/a/fip/fedbne/y2003p3-18.html.

Basu, Susanto, and Christopher L. House (2016) "Allocative and Remitted Wages: New Facts and Challenges for Keynesian Models." in John B. Taylor and Harald Uhlig, Eds., *Handbook of Macroeconomics Volume 2*, 297-354. doi:10.1016/bs.hesmac.2016.05.001

Beaudry, Paul, Dana Galizia, and Franck Poitier. (2020). "Putting the cycle back into business cycle analysis." *American Economic Review* 110.1, 1-47. doi:10.1257/aer.20190789.

Benati, Luca. (2009). "Long run evidence on money growth and inflation." Working Paper Series 1027, European Central Bank.

Berger, Tino, Gerdie Everaert, and Hauke Vierke (2016) "Testing for time variation in an unobserved components model for the U.S. economy." *Journal of Economic Dynamics and Control* 69, 179-208. doi:10.1016/j.jedc.2016.05.017.

Bils, Mark, and Peter J. Klenow. (1998). "Using consumer theory to test competing business cycles models." *Journal of Political Economy*, 106(2): 233-261. doi:10.1086/250009.

Binder, Carola. (2017). "Fed speak on main street: Central bank communication and household expectations." *Journal of Macroeconomics*, 52: 238-251. doi:10.1016/j.jmacro.2017.05.003.

Binder, Carola, Wesley Janson, and Randal Verbrugge. (2019). "Thinking outside the box: Do SPF respondents have anchored inflation expectations?" Federal Reserve Bank of Cleveland, Working Paper no. 19-15. doi:10.26509/frbc-wp-201915.

Blanchflower, David G., and Andrew J. Oswald. (1990). "The wage curve." *The Scandinavian Journal of Economics*, 92(2): 215-235. doi:10.2307/3440026.

Blinder, Alan S. (2018). "Is the Phillips curve dead? And other questions for the Fed." *Wall Street Journal*. URL https://www.wsj.com/articles/is-the-phillips-curve-dead-and-other-questions-for-the-fed-1525388237, May 3, 2018.

Blundell, Richard, Hamish Low, and Ian Preston. (2013). "Decomposing changes in income risk using consumption data." *Quantitative Economics* 4: 1-37. doi:10.3982/QE44.

Bils, Mark. (1987). "The cyclical behavior of marginal cost and price." American Economic Review 77(5): 838-855. https://www.jstor.org/stable/1810212.

Bouakez, Hafedh, Emanuela Cardia, and Francisco Ruge-Murcia. (2014). "Sectoral price rigidity and aggregate dynamics." *European Economic Review*, 65: 1-22. doi:10.1016/j.euroecorev.2013.09.009.

Brainard, Lael. (2019). "The disconnect between inflation and unemployment in the new normal." Speech given at the National Tax Association 49th Annual Spring Symposium. URL https://fraser.stlouisfed.org/files/docs/historical/federal%20reserve%20history/bog_members_statements/brainard20190516a.pdf

Brayton, Flint, John M. Roberts, and John C. Williams. (1999). "What's happened to the Phillips curve?" FEDS working paper 99-49, Board of Governors of the Federal Reserve System. doi:10.2139/ssrn.190852.

Briggs, Joseph, and David Mericle. (2020) "Coronavirus and the inflation outlook." Goldman Sachs Economics Research *US Daily*, 3 April 2020.

Bullard, James B. (2017). "Does low unemployment signal a meaningful rise in inflation?" *The Regional Economist*, 25(3), p. 3. URL https://www.stlouisfed.org/publications/regional-economist/third-quarter-2017.

Caraiani, Petre, and Rangan Gupta. (2018). "Is the response of the Bank of England to exchange rate movements frequency-dependent?" Working Papers 201883, University of Pretoria, Department of Economics. URL: https://ideas.repec.org/p/pre/wpaper/201883.html.

Carroll, Daniel, and Randal Verbrugge. (2019). "Behavior of a new median PCE measure: A tale of tails." Federal Reserve Bank of Cleveland Economic Commentary 2019-10. doi:10.26509/frbc-ec-201910.

Cheremukhin, Anton, and Antonella Tutino. (2016). "Information rigidities and asymmetric business cycles." *Journal of Economic Dynamics and Control*, 73: 142-158. doi:10.1016/j.jedc.2016.09.013.

Chevalier, Judith A., and David S. Scharfstein. (1996). "Capital-market imperfections and countercyclical markups: Theory and evidence." *American Economic Review*, 86(4): 703-725. URL http://www.jstor.org/stable/2118301.

Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt (2015) "Understanding the Great Recession." American Economic Journal: Macroeconomics, 7(1): 110-167. doi:10.1257/mac.20140104.

Ciner, Cetin. (2015). "Are equities good inflation hedges? A frequency domain perspective." *Review of Financial Economics* 24.1: 12-17. doi:10.1016/j.rfe.2014.12.001.

Clark, Peter, Douglas Laxton, and David Rose. (2001). "An evaluation of alternative monetary policy rules in a model with capacity constraints." *Journal of Money, Credit and Banking*, 33(1): 42-64. doi:10.2307/2673871.

Clark, Peter B., and Douglas Laxton. (1997). "Phillips curves, Phillips lines and the unemployment costs of overheating." Working Paper 1997/17, International Monetary Fund (IMF). doi:10.5089/9781451843507.001.

Clark, Todd E. (2014) "The importance of trend inflation in the search for missing disinflation." Federal Reserve Bank of Cleveland Economic Commentary 2014-16. doi:10.26509/frbc-ec-201416.

Clark, Todd E., and Taeyoung Doh. (2014). "Evaluating alternative models of trend inflation." *International Journal of Forecasting*, 30(3): 426-448. doi:10.1016/j.ijforecast.2013.11.005.

Clark, Todd E., and Michael W. McCracken. (2006). "The predictive content of the output gap for inflation: Resolving in-sample and out-of-sample evidence." *Journal of Money, Credit and Banking* 38(5): 1127-1148. doi:10.1353/mcb.2006.0068.

Cochrane, John H. (1989). "The return of the liquidity effect: A study of the short-run relation between money growth and interest rates." *Journal of Business and Economic Statistics* 7.1: 75-83. doi:10.1080/07350015.1989.10509715.

Cochrane, John H. (2018). "A brief parable of over-differencing." Manuscript, University of Chicago. URL: https://faculty.chicagobooth.edu/john.cochrane/research/papers/overdifferencing.pdf.

Coibion, Olivier, and Yuriy Gorodnichenko. (2015). "Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation." *American Economic Journal: Macroeconomics*, 7(1): 197-232. doi:10.1257/mac.20130306.

Comin, Diego, and Mark Gertler (2006) "Medium-term business cycles." *American Economic Review* 96.3, 523-551. doi:10.1257/aer.96.3.523.

Dagum, Estela Bee. (1975). “Seasonal factor forecasts from ARIMA models.” *Bulletin of the International Statistical Institute*, 46(3): 203-216.

Daly, Mary C., and Bart Hobijn. (2014) “Downward nominal wage rigidities bend the Phillips curve.” *Journal of Money, Credit and Banking* 46(2): 51-93. doi:10.1111/jmcb.12152.

Debelle, Guy, and Douglas Laxton. (1997). “Is the Phillips curve really a curve? Some evidence for Canada, the United Kingdom, and the United States. *IMF Staff Papers* 44(2): 249-282. doi:10.2307/3867544.

Detmeister, Alan K., and Nathan R. Babb. (2017). “Nonlinearities in the Phillips curve for the United States: Evidence using metropolitan data.” FEDS working paper 2017-070, Board of Governors of the Federal Reserve System. doi:10.17016/feds.2017.070.

Donayre, Luiggi, and Irina Panovska. (2016). “Nonlinearities in the U.S. wage Phillips curve.” *Journal of Macroeconomics*, 48: 19-43. doi:10.1016/j.jmacro.2016.01.004.

Dotsey, Michael, Shigeru Fujita, and Tom Stark. (2017). “Do Phillips curves conditionally help to forecast inflation?” Working Paper 2017-26, Federal Reserve Bank of Philadelphia. doi:10.21799/frbp.wp.2017.26.

Dupraz, Stéphane, Emi Nakamura, and Jón Steinsson. (2019) “A plucking model of business cycles.” NBER Working Paper 26351. doi:10.3386/w26351.

Evans, George. (1985). “Bottlenecks and the Phillips curve: A disaggregated Keynesian model of inflation, output and unemployment.” *The Economic Journal*, 95(378): 345-357. doi:10.2307/2233214.

Faccini, Renato, and Leonardo Melosi. (2020). “Bad jobs and low inflation.” *CEPR Discussion Papers* 13628. URL https://sites.google.com/site/lemelosi/JL20190305.pdf?attredirects=0&d=1.

Faust, Jon, and Jonathan H. Wright. (2013). "Forecasting inflation." In *Handbook of Economic Forecasting*: 2-56. Elsevier. doi:10.1016/b978-0-444-53683-9.00001-3.

Federal Reserve Bank of Cleveland. "Median Consumer Price Index [MEDCPIM158SFRBCLE]." Retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/MEDCPIM158SFRBCLE.

Federal Reserve Bank of Dallas. "Trimmed Mean PCE Inflation Rate [PCETRIM12M159SFRBDAL]." Retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/PCETRIM12M159SFRBDAL.

Federal Reserve Bank of Philadelphia (a). "Real-Time Data Set for Macroeconomists: Unemployment rate" [Data file]. Retrieved from: https://www.philadelphiafed.org/research-and-data/real-time-center/realtime-data/data-files/ruc.

Federal Reserve Bank of Philadelphia (b). "Survey of Professional Forecasters" [Data file]. Retrieved from: https://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters.

Fernández, Cristina, Aitor Lacuesta, Jose Manuel Montero, and Alberto Urtasun. 2015. "Heterogeneity of markups at the firm level and changes during the Great Recession: The case of Spain." Banco de España Working Paper No. 1536. doi:10.2139/ssrn.2699023.

Filardo, Andrew J. 1998. "New evidence on the output cost of fighting inflation." Federal Reserve Bank of Kansas City *Economic Review*, 83(Q III): 33-61. URL https://ideas.repec.org/a/fip/fedker/y1998iqiiinv.83no.3.html.

Fitzgerald, Terry J., and Juan Pablo Nicolini. (2014). "Is there a stable relationship between unemployment and future inflation? Evidence from U.S. cities." Working Papers 713, Federal Reserve Bank of Minneapolis. URL https://ideas.repec.org/p/fip/fedmwp/713.html.

Friedman, Milton. (1968). "The role of monetary policy." *American Economic Review* 58: 1-17. https://assets.aeaweb.org/asset-server/journals/aer/top20/58.1.1-17.pdf.

Friedman, Milton. (1988). "Money and the stock market." *Journal of Political Economy* 96: 221-245 https://www.jstor.org/stable/1833107.

Fuhrer, Jeffrey C., and Giovanni P. Olivei. (2010). "The role of expectations and output in the inflation process: An empirical assessment." Public Policy Brief 10-2, Federal Reserve Bank of Boston. doi:10.2139/ssrn.1633926.

Gasteiger, Emanuel, and Alex Grimaud (2020) "Price Setting Frequency and the Phillips Curve." Manuscript, Catholic University of Milan.

Giacomini, Raffaella, and Barbara Rossi. (2010). "Forecast comparisons in unstable environments." *Journal of Applied Econometrics*, 25(4): 595-620. doi:10.1002/jae.1177.

Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajsek. (2017) "Inflation dynamics during the financial crisis." American Economic Review 107(3): 785-823. doi:10.1257/aer.20150248.

Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajsek. (2018). "Financial heterogeneity and monetary union." FEDS Working Paper 2018-043, Board of Governors of the Federal Reserve System. doi:10.17016/feds.2018.043.

Gottfries, Nils. (1991). "Customer markets, credit market imperfections and real price rigidity." *Economica*, 58(231): 317. doi:10.2307/2554819.

Granger, Clive W. J. (1969). "Investigating causal relations by econometric models and cross-spectral methods." *Econometrica*, 37(3): 424. doi:10.2307/1912791.

Greenspan, Alan. (1994a). "Testimony before the Joint Economic Committee." Board of Governors of the Federal Reserve System. January 31, 1994.

Greenspan, Alan. (1994b). "Testimony before the Subcommittee on Economic Growth and Credit Formation of the Committee on Banking, Finance and Urban Affairs, U.S. House of Representatives." Board of Governors of the Federal Reserve System. July 22, 1994.

Greenspan, Alan. 1995. "Statement by Alan Greenspan, Chairman, Board of Governors of the Federal Reserve System, before the Committee on Finance, US Senate, January 25, 1995." Board of Governors of the Federal Reserve System.

Hamilton, James D. (2018) "Why You Should Never Use the Hodrick-Prescott Filter." *The Review of Economics and Statistics*, 100(5), 832-843. doi:10.1162/rest_a_00706.

Hansen, Gary D., and Edward C. Prescott. (2005). "Capacity constraints, asymmetries, and the business cycle." *Review of Economic Dynamics*, 8(4): 850-865. doi:10.1016/j.red.2005.08.001.

Hong, Sungki. (2019). "Customer capital, markup cyclicality and amplification." 2019 Meeting Papers 959, Society for Economic Dynamics. https://ideas.repec.org/s/red/sed019.html.

Hopenhayn, Hugo A. (1992) "Entry, exit, and firm dynamics in long run equilibrium." *Econometrica* 60(5): 1127-1150. doi:10.2307/2951541.

Huh, Hyeon-seung, and Inwon Jang. (2007). "Nonlinear Phillips curve, sacrifice ratio, and the natural rate of unemployment." *Economic Modelling*, 24(5): 797-813. doi:10.1016/j.econmod.2007.02.003.

Huh, Hyeon-seung, Hyun Hoon Lee, and Namkyung Lee. (2009). "Nonlinear Phillips curve, NAIRU and monetary policy rules." *Empirical Economics*, 37(1): 131-151. doi:10.1007/s00181-008-0226-x.

Kabukçuoğlu Dur, Ayşe, and Enrique Martínez-García (2019). "Mind the Gap! – A Monetarist View of the Open-Economy Phillips Curve." Manuscript, Federal Reserve Bank of Dallas.

Kareken, John, and Robert M. Solow (1963). "Lags in Monetary Policy." In *Stabilization Policies*, Commission on Money and Credit, Prentice Hall: Englewood Cliffs, New Jersey, 14-96.

King, Robert G., and Mark W. Watson. (2012). "Inflation and unit labor cost." *Journal of Money, Credit and Banking*, 44: 111-149. doi:10.1111/j.1538-4616.2012.00555.x.

Klemperer, P. (1995). "Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade." *Review of Economic Studies*, 62(4): 515-539. doi:10.2307/2298075.

Köberl, Eva M., and Sarah M. Lein. (2011). "The NIRCU and the Phillips curve: An approach based on micro data." *Canadian Journal of Economics/Revue Canadienne d'Economique*, 44(2): 673-694. doi:10.1111/j.1540-5982.2011.01649.x.

Kuhn, Florian, and Chacko George. (2019). "Business cycle implications of capacity constraints under demand shocks." *Review of Economic Dynamics* 32: 94-121. doi:10.1016/j.red.2019.01.001.

Kumar, Anil, and Pia M. Orrenius. (2016). "A closer look at the Phillips curve using state-level data." *Journal of Macroeconomics*, 47: 84-102. doi:10.1016/j.jmacro.2015.08.003.

Laxton, Douglas, Guy Meredith, and David Rose. (1995). "Asymmetric effects of economic activity on inflation." *IMF Staff Papers* 42(2): 344-374. doi:10.2307/3867576.

Laxton, Douglas, David Rose, and Demosthenes Tambakis. (1999). "The U.S. Phillips curve: The case for asymmetry," *Journal of Economic Dynamics and Control* 23: 1459-1485.

Layard, Richard, Stephen Nickell, and Richard Jackman. (1991). *Unemployment: Macroeconomic Performance and the Labour Market*. Oxford University Press. URL https://EconPapers.repec.org/RePEc:oxp:obooks:9780198284345.

Lein, Sarah M., and Eva M. Köberl (2009). "Capacity utilisation, constraints and price adjustments under the microscope." KOF Working papers 09-239, KOF Swiss Economic Institute, ETH Zurich. URL https://ideas.repec.org/p/ kof/wpskof/09-239.html.

Li, Yun. (2019). "The Fed chairman says the relationship between inflation and unemployment is gone." CNBC.com, Thursday, July 11, 2019. https://www.cnbc.com/2019/07/11/the-fed-chairman-says-the-relationship-between-inflation-and-unemployment-is-gone.html.

Lindé, Jesper, and Mathias Trabandt (2019) "Resolving the Missing Deflation Puzzle." CEPR Discussion Papers 13690.

Lipsey, Richard G. (1960). "The relation between unemployment and the rate of change of money wage rates in the United Kingdom, 1862-1957: A further analysis." *Economica* 27(105), 1-31. doi:10.2307/2551424.

Lucas, Robert E. (1976). "Econometric policy evaluation: A critique." *Carnegie-Rochester Conference Series on Public Policy*, 1: 19-46. doi:10.1016/s0167-2231(76)80003-6.

Luengo-Prado, María José, Nikhil Rao, and Viacheslav Sheremirov. (2018). "Sectoral inflation and the Phillips curve: What has changed since the Great Recession?" *Economics Letters* 172, 63-68. doi:10.1016/j.econlet.2018.08.016.

Lundin, Magnus, Nils Gottfries, Charlotte Bucht, and Tomas Lindstrom. (2009). "Price and Investment Dynamics: Theory and Plant-Level Data." *Journal of Money, Credit, and Banking* 41 (5): 907–34. doi:10.1111/j.1538-4616.2009.00238.x.

Madeira, Joao. (2014). "Overtime labor, employment frictions, and the New Keynesian Phillips curve." *Review of Economics and Statistics*, 96(4): 767-778. doi:10.1162/rest_a_00457.

McLeay, Michael, and Silvana Tenreyro. (2018). "Optimal inflation and the identification of the Phillips curve." NBER Working Paper 25893. doi:10.3386/w25892.

Mertens, Elmar. (2016). "Measuring the level and uncertainty of trend inflation." *Review of Economics and Statistics* 98(5): 950-967. doi:10.1162/rest_a_00549.

Mikosch, Heiner F. (2012). "Sticky prices, competition and the Phillips curve." KOF Working papers 12-294, KOF Swiss Economic Institute, ETH Zurich. URL https://ideas.repec.org/p/kof/wpskof/12-294.html.

Montero, José Manuel, and Alberto Urtasun. (2014). "Price-Cost Mark-ups in the Spanish Economy: A Microeconomic Perspective." Bank of Spain Working Paper 1407. https://repositorio.bde.es/handle/123456789/7093.

Morris, Michael, Robert Rich, and Joseph Tracy. (2019). "Wage gains and labor market slack." Manuscript, Federal Reserve Bank of Cleveland.

Moscarini, Giuseppe, and Fabien Postel-Vinay. (2017). "The relative power of employment-to-employment reallocation and unemployment exits in predicting wage growth." *American Economic Review*, 107(5): 364-368. doi:10.1257/aer.p20171078.

Moscarini, Giuseppe, and Fabien Postel-Vinay. (2019). "The job ladder: Inflation vs. reallocation." Working paper, Yale University. URL https://cpb-us-w2.wpmucdn.com/campuspress.yale.edu/dist/1/1241/files/2019/10/nominal-rigidities_SA.pdf

Murphy, Anthony. (2017). "Is the US Phillips curve convex? Some metro level evidence." *Manuscript*, FRB Dallas.

Nalewaik, Jeremy. (2016). "Non-linear Phillips curves with inflation regime-switching." FEDS Working Paper 2017-078, Board of Governors of the Federal Reserve System. doi:10.17016/feds.2016.078.

Nechio, Fernanda. (2015). "Have long-term inflation expectations declined?" *FRBSF Economic Letter*, 2015(11). URL https://ideas.repec.org/a/fip/fedfel/00051.html.

Occhino, Filippo. (2019). "The flattening of the Phillips curve: The policy implications depend on the cause." Federal Reserve Bank of Cleveland *Economic Commentary* 2019-11. doi:10.26509/frbc-ec-201911.

Peach, Richard, Robert W. Rich, and M. Henry Linder. (2013). "The parts are more than the whole: Separating goods and services to predict core inflation." *Current Issues in Economics and Finance*, 19(7). URL https://www.newyorkfed.org/medialibrary/media/research/current_issues/ci19-7.pdf.

Peach, Richard W., Robert W. Rich, and Anna Cororaton. (2011). "How does slack influence inflation?" *Current Issues in Economics and Finance*, 17(3). doi:10.2139/ssrn.1895526.

Petrella, Ivan, and Emiliano Santoro. (2012). "Inflation dynamics and real marginal costs: New evidence from U.S. manufacturing industries." *Journal of Economic Dynamics and Control*, 36(5): 779-794. doi:10.1016/j.jedc.2012.01.009.

Phelps, Edmund S. (1967). "PCs, expectations of inflation, and optimal unemployment over time." *Economica* 34, August: 254-81. doi:10.2307/2552025.

Phelps, Edmund S. (1968). "Money-wage dynamics and labor-market equilibrium." *Journal of Political Economy* 76, July-August: 678-711. doi:10.1086/259438.

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(100): 283-299. doi:10.2307/2550759.

Reynard, Samuel. (2007) "Maintaining low inflation: Money, interest rates, and policy stance." Working Paper Series 756, European Central Bank. URL:

https://www.ecb.europa.eu/pub/pdf/scpwps/ecbwp756.pdf?e33f6f7026d50bc42361036f3e679563.

Rosengren, Eric S. (2019). “Weighing the risks to the economic outlook.” Speech given at Stonehill College, September 3, 2019. https://www.bostonfed.org/news-and-events/speeches/2019/weighing-the-risks-to-the-economic-outlook.aspx.

Rossi, Barbara, and Tatevik Sekhposyan. (2010). "Have economic models' forecasting performance for US output growth and inflation changed over time, and when?" *International Journal of Forecasting*, 26(4): 808-835. doi:10.1016/j.ijforecast.2009.08.004.

Sargent, Thomas J. (1987). *Macroeconomic Theory*. Academic Press, second edition.

Shapiro, Carl, and Joseph E. Stiglitz. (1984). “Equilibrium unemployment as a worker discipline device.” *American Economic Review*, 74(3): 433-444. URL https://www.jstor.org/stable/1804018.

Steelman, Aaron, Renee Haltom, and Lisa Kenney. (2013). “Interview: John Cochrane.” Federal Reserve Bank of Richmond *Econ Focus*, Third Quarter: 34-38.

Stock, James H., and Mark W. Watson. (1999). “Forecasting inflation.” *Journal of Monetary Economics*, 44(2): 293-335. doi:10.1016/S0304-3932(99)00027-6.

Stock, James H., and Mark W. Watson. (2007). “Why has U.S. inflation become harder to forecast?” *Journal of Money, Credit and Banking*, 39(s1): 3-33. doi:10.1111/j.1538-4616.2007.00014.x.

Stock, James H., and Mark W. Watson. (2009). “Phillips curve inflation forecasts.” In *Understanding Inflation and the Implications for Monetary Policy*: 101-186. The MIT Press. doi:10.7551/mitpress/9780262013635.003.0003.

Stock, James H., and Mark W. Watson. (2010). "Modeling inflation after the crisis." Working Paper 16488, National Bureau of Economic Research. doi:10.3386/w16488.

Stock, James H., and Mark W. Watson. (2020) "Slack and Cyclically Sensitive Inflation." Forthcoming, *Journal of Money, Credit and Banking*.

Summers, Lawrence. (2017). "America needs its unions more than ever." *The Financial Times*. September 3, 2017.

Tasci, Murat. (2018). "The ins and outs of unemployment in the long run: Unemployment flows and the natural rate." Manuscript, Federal Reserve Bank of Cleveland.

Tasci, Murat, and Randal J. Verbrugge. (2014). "How much slack is in the labor market? That depends on what you mean by slack." *Economic Commentary*, (2014-21). doi:10.26509/frbc-ec-201421.

The Economist. (2020). "Knocking off work: Traders are losing interest in America's jobs figures." *The Economist*. February 8, 2020.

The Economist. (2017). "The Phillips curve may be broken for good; central bankers insist that the underlying theory remains valid, Daily Chart." *The Economist*. November 1, 2017.

US Bureau of Economic Analysis (a). "Personal Consumption Expenditures: Chain-type Price Index [PCEPI]." Retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/PCEPI.

US Bureau of Economic Analysis (b). "Personal Consumption Expenditures [PCE]." Retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/PCE.

US Bureau of Labor Statistics (a). "Consumer Price Index for All Urban Consumers: All Items Less Food and Energy [CPILFESL]." Retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/CPILFESL.

US Bureau of Labor Statistics (b). "Consumer Price Index Research Series (CPI-U-RS) All Urban Consumers: All Items." Retrieved from https://www.bls.gov/cpi/research-series/home.htm#CPI-U-RS%20Data.

US Congressional Budget Office. "Natural Rate of Unemployment (Long-Term) [NROU]." Retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/NROU.

Vavra, Joseph. (2014). "Time-varying Phillips curves." NBER Working Paper No. 19790. Doi: 10.3386/w19790

Verbrugge, Randal. (1997). "Investigating cyclical asymmetries." *Studies in Nonlinear Dynamics and Econometrics* 2.1: 15-22. doi:10.2202/1558-3708.1025.

Verbrugge, Randal. (2020). "Is it time to reassess the focal role of core PCE inflation?" Manuscript, Federal Reserve Bank of Cleveland.

Verbrugge, Randal, Saeed Zaman, and Keerthana Nunna. (2018). "Improving Phillips curve inflation forecasts using a robust asymmetry measure." *Manuscript in preparation*, Federal Reserve Bank of Cleveland.

Yanfeng, Wei. (2013). "The dynamic relationships between oil prices and the Japanese economy: A frequency domain analysis." *Review of Economics and Finance* 3: 57-67. https://ideas.repec.org/a/bap/journl/130205.html.

Xu, Qifa, Xufeng Niu, Cuixia Jiang, and Xue Huang. (2015). "The Phillips curve in the US: A nonlinear quantile regression approach." *Economic Modelling*, 49: 186-197. doi:10.1016/j.econmod.2015.04.007.

Zaman, Saeed. (2013). "Improving inflation forecasts in the medium to long term." *Economic Commentary*, (2013-16). doi:10.26509/frbc-ec-201316.

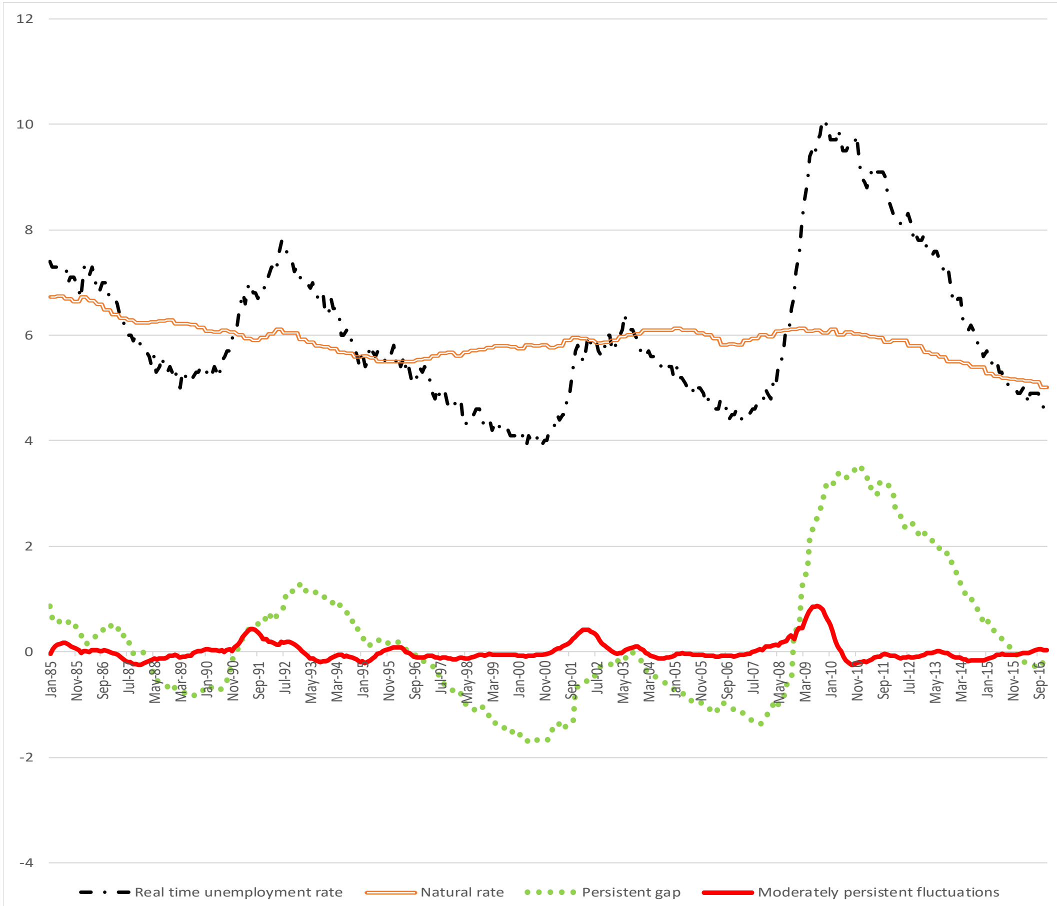

Figure 1: One-Sided Partition of the Unemployment Rate Gap

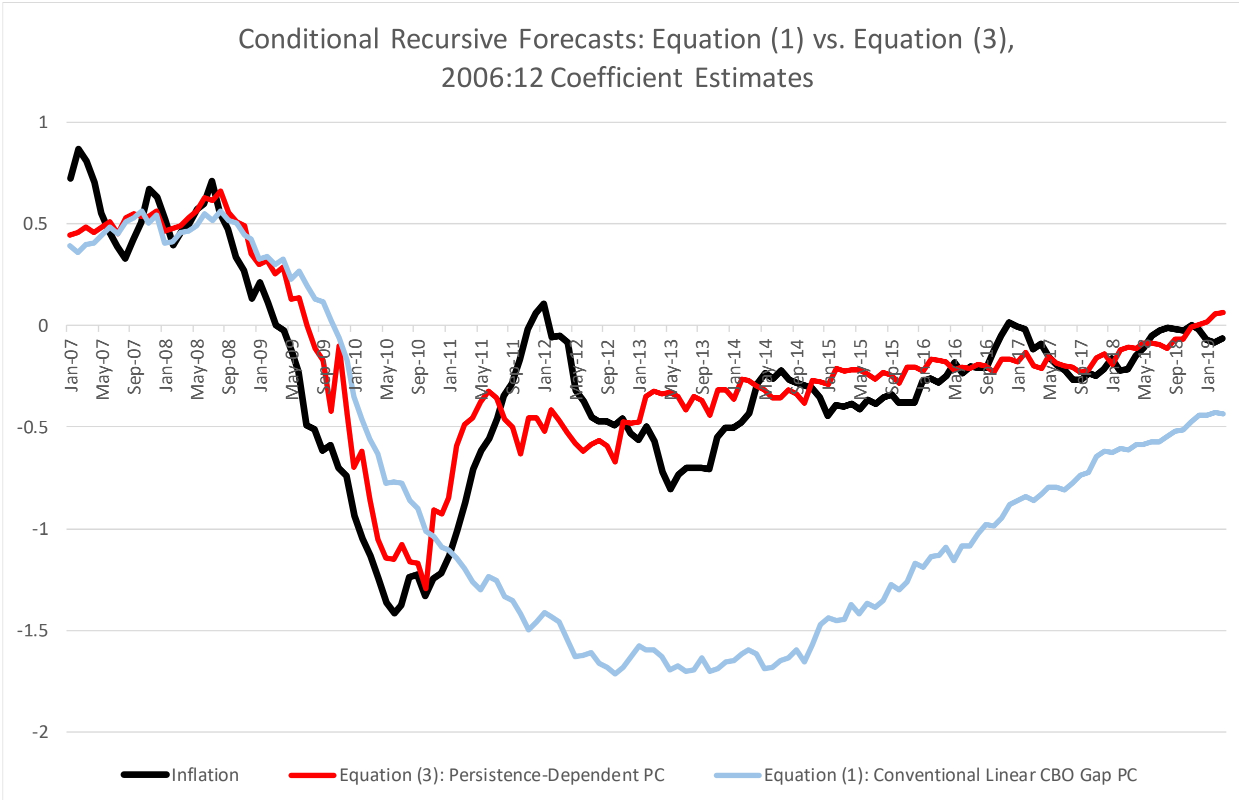

Figure 2: Conditional Recursive Forecasts from Equations (1) and (3)

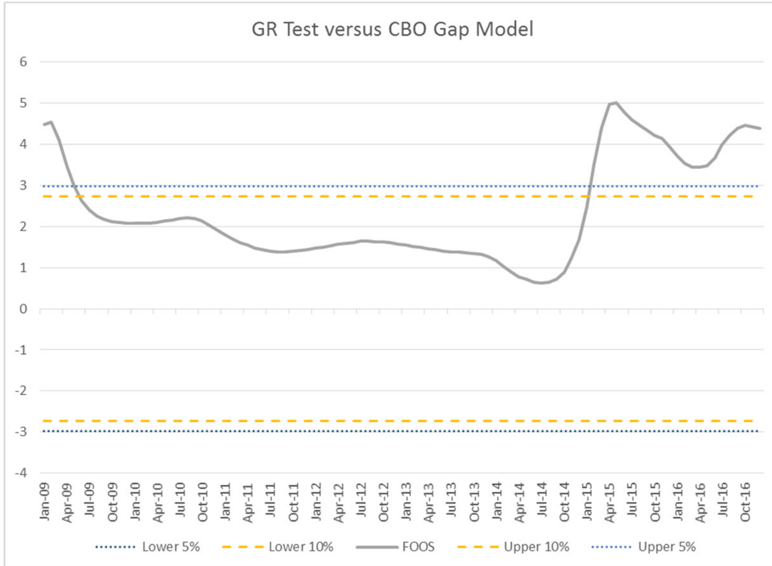

Figure 3. Baseline CBO Model, Equation (1), Versus Asymmetric Persistence-Dependent PC Forecasting Model, Equation (3)

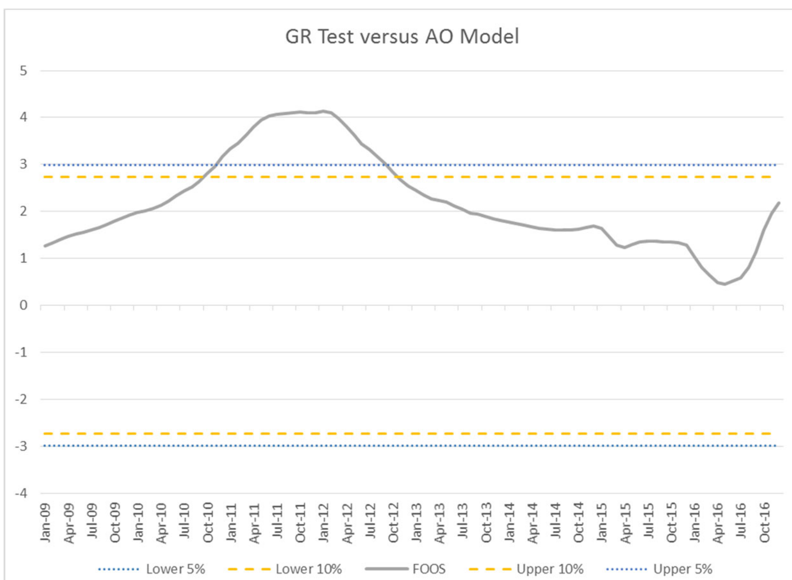

Figure 4: Univariate Atkeson-Ohanian Model, Equation (4), Versus Asymmetric Persistence-Dependent PC Forecasting Model, Equation (3)

Appendix

A.1 Other Inflation Indicators and CBO Gap

In Table A1 we present the results using the CBO u_t^* estimate rather than that of Tasci (2018) (in column 2), and several other inflation indicators. We provide the trimmed PCE results from Table 1 in column 1 for comparison. The final two rows in the table refer to GR and DM forecast comparisons against the baseline CBO model for the same dependent variable.

Table A.1. Other inflation indicators.

	Trimmed PCE	Trimmed PCE with CBO gap	Median PCE	Core PCE	Median CPI	Core CPI
λ_1^+	0.03	0.02	0.05	0.03	0.04	0.01
t-stat	0.67	0.41	2.58	0.62	1.30	0.16
λ_1^-	-0.27	-0.54	-0.35	-0.20	-0.35	-0.24
t-stat	-3.20	-2.72	-11.43	-2.05	-3.14	-2.41
λ_2^+	-1.67	-1.67	-1.91	-0.71	-2.13	-1.55
t-stat	-8.18	-6.98	-7.13	-2.59	-7.79	-5.59
λ_2^-	0.10	-0.08	0.17	-0.91	0.40	0.24
t-stat	0.14	-0.11	0.27	-0.75	0.56	0.25
λ_3^+	-0.51	-0.54	-0.46	-0.48	-0.74	-0.24
t-stat	-2.44	-2.48	-2.10	-1.86	-2.76	-1.06
λ_3^-	-0.03	-0.02	-0.13	-0.05	0.08	-0.09
t-stat	-0.32	-0.32	-0.90	-0.29	0.78	-0.63
H0: $\lambda_1^+ = \lambda_1^-$	0.01	0.01	0.00	0.07	0.00	0.05
H0: $\lambda_2^+ = \lambda_2^-$	0.02	0.04	0.00	0.88	0.00	0.09
H0: $\lambda_3^+ = \lambda_3^-$	0.06	0.05	0.24	0.24	0.02	0.58
H0: $\lambda_1^+ = \lambda_1^- = \lambda_2^+ = \lambda_2^- = \lambda_3^+ = \lambda_3^-$	0.00	0.00	0.00	0.00	0.00	0.00
Adjusted R-squared	0.75	0.74	0.78	0.46	0.74	0.47
GR Test p-value	<0.05	<0.05	<0.05	<0.05	<0.05	<0.10
DM Test p-value	0.01	0.02	0.04	0.06	0.03	0.08

In short, our results do not hinge on using the trimmed-mean PCE as the inflation indicator and, broadly speaking, are robust to using different inflation indicators. The partial exception is core PCE inflation, a topic we turn to next.

A.2 Deficiencies in Core PCE

We acknowledge that the core PCE tests do not reject symmetry in the moderately persistent component. Like Ball and Mazumder (2019), we suggest that this "puzzling" result stems from deficiencies of core PCE inflation as a measure of trend inflation. Theory predicts two major deficiencies of less-food-and-energy ("core") inflation indexes, and both were exhibited in the post-1985 period. First, because the core PCE price index simply excludes items from the basket, core PCE inflation can be subject to bias over prolonged periods. And as Carroll and Verbrugge (2019) indicate, this bias has also been highly unstable over time. For example, between 1995 and 2007, core PCE inflation was downwardly biased by 0.25 percentage points, while it was upwardly biased by 0.3 percentage points between 1980 and 1985. This fact alone raises some doubts about its ability to truly match trend inflation. Second, despite their moniker, core inflation indexes are subject to large idiosyncratic transitory shocks that distort the estimate of trend inflation. (Indeed, the standard deviation of core inflation measures is so large that they are almost always examined in time-averaged form.) Large shocks are not confined to food and energy components. This sensitivity to transitory noise is significant in the present study: transitory shocks can occur at any time, but in the context of analyses that distinguish between phases of the business cycle, these shocks will be especially detrimental if they are correlated with the phase within the sample. One aspect of core PCE inflation is noteworthy: as discussed below, core PCE inflation is sensitive to the movements of prices that are not market-determined, and such movements may well be systematically related to the business cycle. In terms of its ability to reliably reflect trend inflation, as discussed above, when core PCE inflation departs from trimmed-mean PCE inflation, it is core PCE inflation that adjusts to eliminate the gap.

There were only three NBER recessions post-1985. This implies that the moderately persistent component experienced only three nonzero episodes after 1985: starting in 1991, starting in 2001, and starting in mid-2007. During two of these recoveries, core PCE inflation experienced dynamics that were at odds with limited-influence trend inflation indicators such as the trimmed-mean PCE or the median CPI, and even with the other prominent "less food and energy" series, the core CPI. During the aftermath of the 2001 recession, year-over-year core PCE inflation displayed a prominent rebound from early 2002 to early 2003, including one month with 2.4 percent inflation, a reading not seen since the early 1990s.[35] Other limited-influence trend inflation indicators displayed an essentially monotonic decline from 2001Q3 to 2003Q4. During the Great Recession, while other (year-over-year) trend inflation measures displayed an essentially monotonic decline from 2008Q4 through 2010, core PCE inflation again exhibited a strong rebound in the middle of this episode: starting from below 1 percent in September 2009, it rapidly rose to 1.7 percent during the first few months of 2010, then fell gradually back down to end below 1 percent in 2010Q4.

Conversely, during both of these episodes, inflation in the *market-based* core PCE[36] displayed dynamics that were *similar* to other limited-influence trend inflation indicators; see Figure 4. This indicates that core PCE's unusual dynamics during both of these episodes stemmed from the behavior of prices that were not market-determined.[37] In short, core PCE inflation was evidently subject to countervailing idiosyncratic influences during the aftermath of both the 2001 recession and the Great Recession that all but masked trend inflation movements during critical periods. The anomalous behavior of core PCE inflation during these crucial episodes surely calls into serious question its usefulness as a trend inflation estimator.[38]

---

[35] This may have been due to insurance payments related to 9/11 that caused m/m core PCE inflation to run negative in the fall, which showed up in y/y core PCE inflation a year later.

[36] The PCE market-based price index is based primarily on observed market transactions for which there are corresponding price measures. It includes owners' equivalent rent, but excludes most imputed expenditures, such as "financial services furnished without payment," most insurance purchases, gambling, margins on used light motor vehicles, and expenditures by US residents working and traveling abroad.

[37] See also Peach, Rich, and Linder (2013), who display a decomposition into goods and services. The anomalous movements during the Great Recession were almost entirely driven by imputed financial services price movements.

[38] For more details, see Verbrugge (2020).

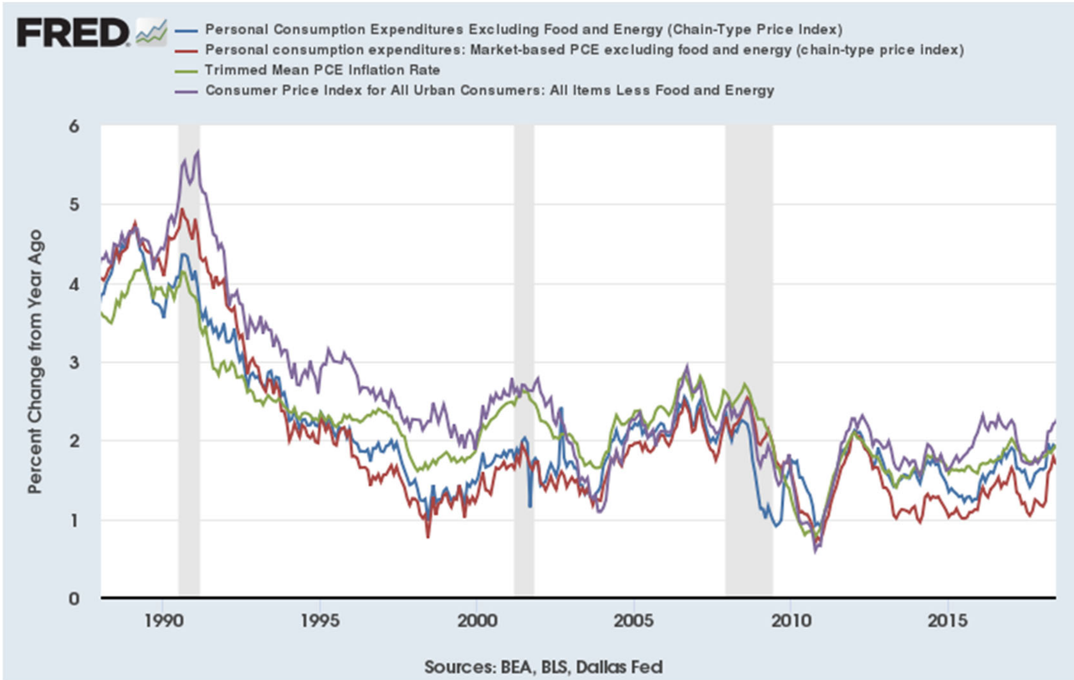

Figure 4: Four Inflation Measures

We plot four trend inflation indicators in Figure 4. Only one of these, "core" PCE, displays a notable inflation rebound in 2010. This came from unusual behavior of nonmarket goods, and in particular, imputed financial services. The susceptibility of core PCE inflation to such movements reduces its usefulness as a trend inflation measure.

A.3 Comparison to Stock-Watson Recession Gap

For the post-1985 period, Figure 5 plots the 12-month trimmed PCE inflation rate (leaded 12 months) along with the (monthly) recession gap term and the positive part of the moderately persistent unemployment rate fluctuations, which we here term the "bust gap."[39]

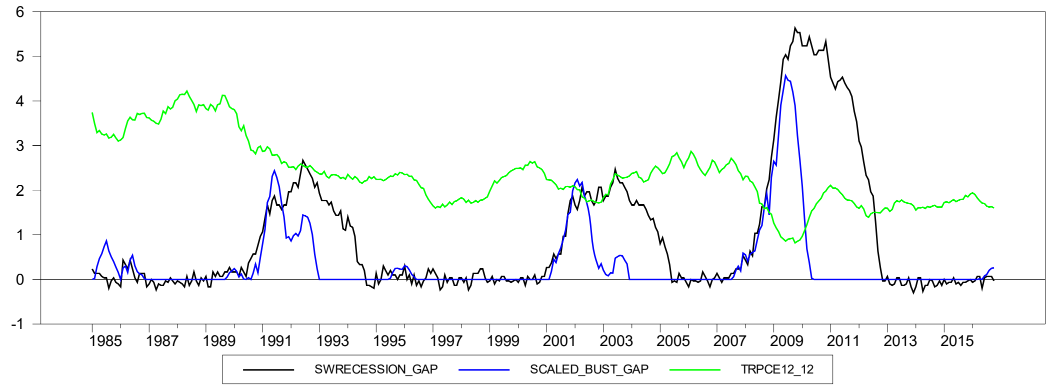

Figure 5: Stock-Watson Recession Gap and Scaled Bust Gap

In this figure, the latter series has been scaled by multiplying it by 5 so as to render its peak magnitude comparable to that of the recession gap during the middle two recession episodes. Regarding ocular econometrics, the bust gap has an edge in timing, in that the peak inflation deceleration is relatively close to the peak of the bust gap (but well prior to the peak of the recession gap) and ends roughly when the bust gap vanishes (while the recession gap stays significantly positive for much longer). However, this is merely suggestive. We now provide out-of-sample forecast evidence that our specification is superior: at least over the post-1985 period, the bust gap better captures the impact of recessions on inflation dynamics.

In Figure 6, we display the Giacomini-Rossi forecast comparison results from our Equation (3) model versus the Stock and Watson recession gap model. While our model

[39] It is worth mentioning that the unemployment gap resulting from the year-over-year filter in Stock and Watson (2020) bears a notable resemblance to the moderately-persistent component in our partitioning. This is not surprising, given the frequency gain of the year-over-year filter.

outperforms the Stock-Watson analogue over the entire period, this is only statistically significant (at the 10 percent level) from September 2011 through June 2012. However, the Diebold-Mariano test, with a p-value of 0.03, indicates that the gain from our model is statistically significant when considering the sample as a whole.

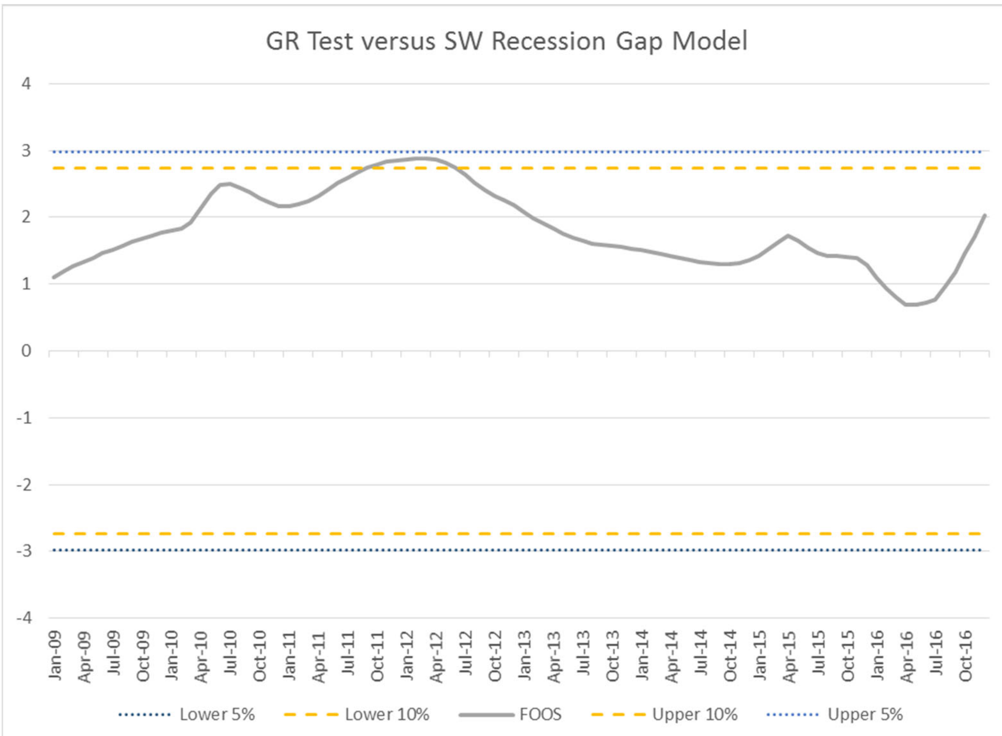

Figure 6: Stock-Watson Gap Model Versus Equation (5) Model

A.4 Relation to Some Other Findings in the Literature

Our results reinterpret a finding in Coibion and Gorodnichenko (2015). These authors constructed the "NAIRU" implied by their estimated model that would be necessary to explain the "missing disinflation" during the Great Recession. While differing in details, the gap implied by the Coibion-Gorodnichenko NAIRU has broad similarities with our bust gap: it starts opening up shortly after the unemployment rate started rising rapidly in 2008 but was virtually back to zero by mid-2009. (We emphasize that inflation data are not used in constructing our gap measures.) These authors concluded that these dynamics were implausible for a NAIRU.

However, our findings indicate that the implausibility of their estimate stemmed not from the possibility that a NAIRU might have dynamics that were at such great odds with conventional estimates, but rather with the notion that a NAIRU is just another way to describe the natural rate of unemployment. As we have noted above, there is no reason that these concepts should coincide.[40] Implicitly, both Stock and Watson (2010) and Coibion and Gorodnichenko (2015) provide evidence supportive of our findings.

A mismeasured gap will likely lead to the conclusion that forecasting performance is episodic (for example, Stock and Watson 2009) or that there is time variation in the inflation process, such as time variation in the coefficient on the activity variable (see evidence in Clark and McCracken 2006, Stock and Watson 2009, Vavra 2013 and Luengo-Prado, Rao, and Sheremirov, 2018). This may explain the forecasting performance of the time-varying unobserved components model of Stock and Watson (2007).

Conversely, the unobserved components model of Berger, Everaert and Vierke (2016) exhibits a stable Phillips curve (note however that the estimated relationship is very weak, and that the point estimates do change over time, but the standard error estimates are large.). These authors implicitly detrend inflation but allow a "persistent transitory component" as well, and relate the inflation gap to an output gap, which in some sense focuses attention on business-cycle frequencies. Stock and Watson (2020) note some disadvantages of such approaches, including the requirement that one specify an explicit model for the trends, a model which could be misspecified. Instead, these authors explicitly eliminate trends (low-frequency variation), via their time-series filters, so as to focus solely on the business-cycle-frequency relationship. As noted above, the findings of Stock and Watson (2020) are consistent with our results: by restriction attention solely to the relationship at business-cycle frequencies, these authors find a strong, and more stable, Phillips curve relationship. The effectiveness of their approach is diminished by its a priori restriction to a consideration of the Phillips curve relation at only this

[40] See further discussion in Tasci and Verbrugge (2014).

particular range of frequencies, however, and their use of a two-sided bandpass filter (for most of their results) is subject to criticisms discussed briefly above, and in more detail in Ashley and Verbrugge (2009) and Ashley, Tan and Verbrugge (2020).

Our findings also reconcile evidence in, for example, Filardo (1998), Barnes and Olivei (2003), Huh and Jang (2007), Baghli, Cahn, and Fraisse (2007), Stock and Watson (2009), Fuhrer and Olivei (2010), Peach, Rich, and Cororaton (2011), and Peach, Rich, and Linder (2013) that the PC is "convex-concave" (see also Xu, Jiang, and Huang (2015)). These studies, among others already noted above, find a steepening of the Phillips curve as slack becomes negative. Similarly, our findings are also consistent with regime-switching studies, such as Huh, Lee, and Lee (2008) or Donayre and Panovska (2016),[41] that find three regimes in the wage Phillips curve. Our viewpoint, though, is that previous studies somewhat mischaracterize the reduced-form Phillips relationships, first because none (aside from Stock and Watson (2010)) can well approximate the positive part of our moderately persistent component, and second because they typically estimate a fixed lower threshold for slack rather than allowing for a time-varying natural rate of unemployment. In sum, the form of nonlinearity we uncover is well-supported in the data and is consistent with economic theory (see Appendix A.8), yet is not cleanly captured by the standard sorts of nonlinearity that most models admit.

[41] See also Nalewaik (2016) for a rich regime-switching approach.

A.5 Relationship between Moderately Persistent and Transient Components

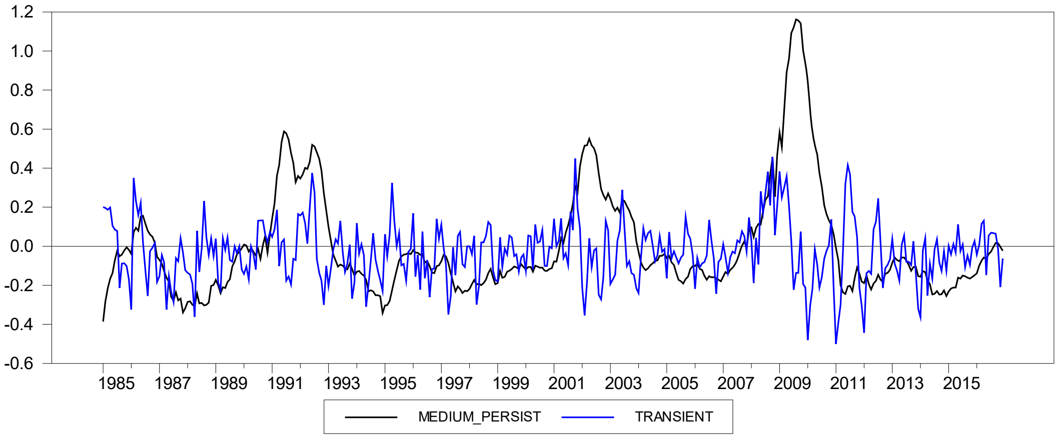

Figure 7: Moderately Persistent and Transient Components of Unemployment

Figure 7 plots un_{modP} (here labelled medium_persist) and $un_{transient}$. These series are positively correlated, especially at the beginning of a recession.

A.6 Conditional Recursive Forecasts Using Both 2006:12 and 2016:12 ("Full Sample")

Coefficient Estimates

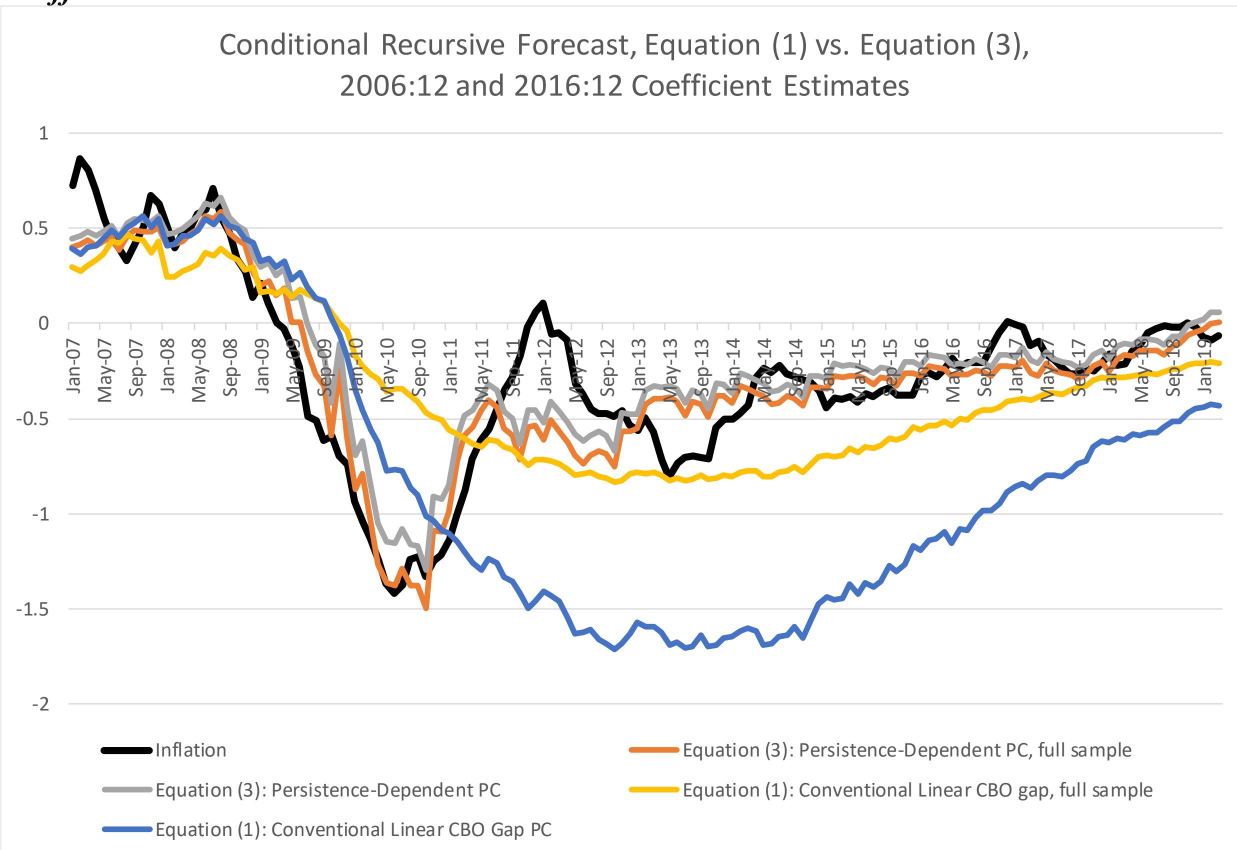

Figure 8: More Conditional Recursive Forecasts

Notice, comparing the gray and orange lines, that the conditional forecasts of our persistence-dependent model are essentially identical, regardless of whether we estimate coefficients in 2006 or 2016. Conversely, conditional forecasts from a standard linear model are notably different if one estimates the Phillips curve in 2016:12 rather than 2006:12. This reflects the purported weakening of the Phillips curve. With coefficients estimated in 2016:12, the fit to the evolution of inflation is still poor, with a far worse fit at the beginning of the Great Recession.

A.7 Comparison to Wage PC of Morris, Rich, and Tracy

In the table below, we compare the wage Phillips curve parameter estimates of Morris, Rich and Tracy (2019) (MRT) to the parameter estimates derived from a similarly specified price Phillips curve. The dependent variable in MRT is a four-quarter growth in average wages term, constructed in MRT on the basis of CPS data, detrended using long-run SPF inflation expectations as in this paper. The specification in MRT is

$$\begin{aligned} GAW_t^{t+4} - \pi_t^* = \alpha_0 + \alpha_3\left(prod_t^{trnd}\right) + \alpha_1^+\left(un_t^Q - un_t^{*,CBO}\right)^+ + \alpha_1^-\left(un_t^Q - un_t^{*,CBO}\right)^- \\ + \alpha_2^+\left[\Delta\left(un_t^Q - un_t^{*,CBO}\right)\right]^+ + \alpha_2^-\left[\Delta\left(un_t^Q - un_t^{*,CBO}\right)\right]^- + \varepsilon_t \end{aligned} \tag{6}$$

where GAW_t^{t+4} refers to growth in average wages (constructed from CPS data, as detailed in MRT), un_t^Q is quarterly unemployment, $un_t^{*,CBO}$ is the CBO estimate of the natural rate of unemployment, and $prod_t^{trnd}$ refers to trend productivity growth at time *t*. For comparison to (6), we re-specify Equation (3) alternatively as

$$\begin{aligned} \pi_t^{12} - \pi_t^* = \alpha + \beta_1\left(\pi_t^{12} - \pi_{t-12}^*\right) + \beta_2\left(\pi_{t-12}^{12} - \pi_{t-24}^*\right) + \lambda_1^+\left(un_t^3 - un_t^*\right)^+ + \lambda_1^-\left(un_t^3 - un_t^*\right)^- \\ + \lambda_2^+\left[\Delta\left(un_t^3 - un_t^*\right)\right]^+ + \lambda_2^-\left[\Delta\left(un_t^3 - un_t^*\right)\right]^- + \varepsilon_t \end{aligned} \tag{7}$$

or in a slightly simplified form,

$$\begin{aligned} \pi_t^{12} - \pi_t^* = \alpha + \beta_1\left(\pi_t^{12} - \pi_{t-12}^*\right) + \beta_2\left(\pi_{t-12}^{12} - \pi_{t-24}^*\right) + \lambda_1^+\left(un_t^3 - un_t^*\right)^+ + \lambda_1^-\left(un_t^3 - un_t^*\right)^- \\ + \lambda_2^+\left[\Delta\left(un_t^3\right)\right]^+ + \lambda_2^-\left[\Delta\left(un_t^3\right)\right]^- + \varepsilon_t \end{aligned} \tag{8}$$

where un_t^3 is a three-month moving average of the real-time unemployment rate.

It is interesting to note that estimates of this wage Phillips curve yield qualitatively similar results, though they differ in details. For instance, there is a modest, but nonzero, downward force on wage growth when labor force slack is high. But the strongest pressures occur on wage growth from upward movements in the unemployment rate (though this force is not as strong as the downward force on prices during these periods), and from overheating (and this force is much stronger than the upward force on prices during these periods).

Table A.2. Comparison to Wage PC of Morris, Rich and Tracy (2019)

Dependent Variable	$(GAW^{t,t+4}-\pi^{e})$ MRT (2019)	$(\pi^{t,t+12}-\pi^{e})$ Eq. (6)	$(\pi^{t,t+12}-\pi^{e})$ Eq. (7)
Variable			
Trend Productivity Growth	0.31		
$(un\text{-}un^{*})^{+}$ [1]	-0.3**	0.0	0.0
$(un\text{-}un^{*})^{-}$ [1]	-1.28**	-0.3**	-0.3**
$[\Delta(un\text{-}un^{*})]^{+}$ [1]	-1.4**	-2.8**	
$[\Delta(un\text{-}un^{*})]^{-}$ [1]	0.9	0.3	
$[\Delta(un)]^{+}$			-2.8**
$[\Delta(un)]^{-}$			0.16
$(\pi^{t-12,t}-\pi^{e,t-13})$		0.45**	0.45**
$(\pi^{t-24,t-12}-\pi^{e,t-25})$		0.20**	0.20**

1. In Morris, Rich, and Tracy (2019) (MRT), this is the CBO gap; in the second and third columns , this is the gap as defined in this paper.
The MRT estimation period is 1983:Q1-2018Q4.
** denotes statistical significance at the 1% level, using Newey-West (1987) standard errors.
Time t subscripts are suppressed in the table.

We again see clear evidence of nonlinearity in the Phillips curve. However, while appealing in their greater simplicity, specifications (7) and (8) do not fit the data quite as well as (3), nor do they yield quite as accurate a conditional forecast of inflation over the post-2006 period.
One might have guessed this, given the estimated differential between λ_2^{+} and λ_3^{+} in conjunction with Appendix A.5. It would be hard for a single term in Δun to capture what appears to be two different relationships.

A.8 Theory
A.8.1 Overheating and Inflation

From its inception in Phillips (1958), it was generally believed (see also Lipsey (1960)) that the general shape of the Phillips curve is convex, so that a negative unemployment gap (an overheating economy) has a bigger price impact than the same percentage positive unemployment gap (slack). Many theories naturally give rise to a convex wage Phillips curve. Layard, Nickell, and Jackman (1991) demonstrate that the shirking model of Shapiro and Stiglitz (1984) implies a nonlinear wage Phillips curve. The "bottlenecks" model of Evans (1985) and the bargaining model of Blanchflower and Oswald (1990) also imply a nonlinear wage Phillips curve. We would expect such convexity to spill over into convexity in the price Phillips curve.

Lindé and Trabandt (2019) demonstrate that the mere inclusion of real rigidities (a la Kimball aggregator) induces a convex shape to the Phillips curve, but this will be erroneously missed unless one eschews linearization and instead uses a nonlinear solution method; see also Gasteiger and Grimaud (2020).

A convex shape to the price Phillips curve is suggested by models in which prices are downwardly rigid, such as Ball, Mankiw, and Romer (1988).[42] In this model, which features menu costs of price adjustment in the presence of generally positive inflation, prices are more sticky downward because the relative price declines can "automatically" occur via inflation. Thus, even if a firm desires a relative price decline, it will optimally choose inaction and wait for inflation to deliver that decline in the near future.

In the standard New Keynesian model, the output gap maps directly into inflationary pressure. In the standard DMP model, the value of unemployment determines the worker's outside option. Moscarini and Postel-Vinay (2017) draw attention to the fact that individual wage growth co-varies more strongly with the aggregate job-to-job transition rate than with the aggregate unemployment rate. Moscarini and Postel-Vinay (2019) provide a New Keynesian job-ladder model that is consistent with this fact and that explains how an overheating labor market can translate into price pressures. In this model, workers' bargaining power derives from the ability to receive outside offers, not from the unemployment outside option. After a downturn, many employed workers are mismatched and easily poachable, and numerous unemployed workers are profitably hired. But late in the cycle, the stiff competition for employed productive workers leads to many outside wage offers being matched by current employers, and these wage increases effectively become cost-push shocks.

Another class of models that naturally deliver a Phillips curve relationship of this sort – that is, strong upward price pressure when the economy is overheating – is capacity-constraints

---

[42] Downward nominal wage rigidity is a classic explanation for a convex wage Phillips curve (see Phillips (1958) and Daly and Hobijn (2014), and see Dupraz, Nakamura and Steinsson (2019) for a recent model delivering asymmetric unemployment fluctuations.

models.[43] Bils and Klenow (1998) find procyclical relative price and TFP movements in highly procyclical consumption good sectors and argue that this suggests the existence of varying capacity utilization with occasionally binding capacity constraints. Capacity constraints naturally induce business cycle asymmetries (Hansen and Prescott, 2005). In the New Keynesian model of Alvarez-Lois (2004), the Phillips curve becomes

$$\pi_t = \beta E\pi_{t+1} + \mu\left(\hat{\theta}_t + \widehat{mc}_t\right)$$

where θ_t is the share of firms in the economy that are operating at full capacity. (See also Alvarez-Lois (2005, 2006) for related models, and Mikosch (2012) and Kuhn and George (2019) for alternative New Keynesian models with capacity constraints.) There is supportive evidence. The paper by Lein and Köberl (2009) is a micro study of Swiss manufacturing firms. These authors find evidence of a strong relationship between price increases and being capacity constrained (either due to labor or due to technical capacity).[44]

[43] The class of models expounded in Clark and Laxton (1997) or Clark, Laxton and Rose (2001) also feature capacity constraints. Alan Greenspan seems to have believed in a convex Phillips curve arising from capacity constraints. For example, in his testimony to the Subcommittee on Economic Growth and Credit Formation (Greenspan 1994b, p.12), he stated: "If the economy were nearing capacity, we would expect to see certain patterns in the statistical and anecdotal information ... To attract additional workers, employers would presumably step up their use of want-ads and might begin to use nonstandard techniques...All of these steps in themselves could add to costs and suggest developing inflationary imbalances." In his testimony before the Joint Economic Committee in January 1994, he noted: "History suggests, however, that higher price inflation tends to surface rather late in the business cycle..." (Greenspan 1994a, p.6). In his testimony before the Committee on Finance in January 1995, he stated: "Knowing in advance our true growth potential obviously would be useful in setting policy because history tells us that economies that strain labor force and capital stock limits tend to engender inflationary instabilities."

[44] Using these same data, Köberl and Lein (2011) find that an aggregated capacity constraint measure is useful in a Phillips curve. Similarly, at the micro level, Mikosch (2012) finds that the slope of the micro Phillips curve is increasing as capacity constraints become tighter, although this effect disappears for firms facing intense competition.

A.8.2 Busts and Inflation

It has been thought puzzling that large labor market slack does not weigh on inflation, leading to the famous inflation puzzle of the Great Recession. Not only is this suggested by a conventional Phillips curve, it is ostensibly an implication of standard New Keynesian theory (see, for example, King and Watson 2012). That paper demonstrates, though, that the low-frequency movements in inflation should line up with low-frequency movements in real unit labor costs. Most of the empirical work in the New Keynesian paradigm has used a variant of labor's share as the proxy for real marginal costs, but Bils (1987), Petrella and Santoro (2012), and Madeira (2014) demonstrate that this can be a misleading proxy. Petrella and Santoro (2012) use the income share of intermediate goods (and stress the importance of disaggregated data; see also Bouakez, Cardia, and Ruge-Murcia 2014); Madeira (2014) constructs a proxy using overtime costs. Both alternatives improve the fit of New Keynesian Phillips curves.

Standard industrial organization theory predicts that, at the onset of a recession, we might see an *initial* drop in inflation, but not *continued* downward pressure – even though slack (as conventionally measured) remains high. In particular, the received wisdom in the industrial organization literature is that demand shortfalls tend to provoke price wars. But this behavior is forward-looking, and price declines are front-loaded. After a time, the price war effect ceases, and prices then start to drift slowly upward again. More generally, as is well known, countercyclical markups will mitigate aggregate price drops during recessions. Fernández et al. (2015) demonstrate that, in Spain, average markups rose in half of the sectors after 2008.[45] Gilchrist et al. (2017) develop a New Keynesian model, extended in Gilchrist et al. (2018), that builds upon these insights, and provides supportive empirical evidence. These authors draw attention to the standard IO theory, but further note a nuance to this basic relationship. In customer markets, pricing decisions are investment decisions, and factors that influence investment will influence pricing. Thus, in the theory of Gottfries (1991) and Chevalier and

[45] In the price experimentation model of Bachmann and Moscarini (2012), a recession might trigger some firms to optimally increase prices, as the costly acquisition of information might allow them to "gamble for resurrection."

Scharfstein (1996), under financial frictions, constrained firms in customer markets facing a fall in demand may find it optimal to maintain, *or even increase*, their prices to boost cash flow and avoid costly external financing. Financially unconstrained firms have the opportunity to reduce prices and invest in market share. In the model of Gilchrist et al. (2017), financial frictions imply that markups remain elevated after the initial adverse demand (or financial) shock.[46] Evidence in both Gilchrist et al. (2017) and Gilchrist et al. (2018) is supportive; for instance, financially constrained firms in the US, on average, raised prices at the onset of the Great Recession, while other firms dropped prices aggressively and increased their market share.[47] Prices remained flat for about a year, then began to rise again. The resulting changes in market share were persistent. See also Hong (2019), who finds that markups are countercyclical (with cyclicality varying systematically across firms) and who develops a customer-capital variant of a Hopenhayn (1992) model consistent with his findings. Finally, Alves (2019) demonstrates that a reduction in job-to-job flows during the recovery can worsen labor productivity, providing an upward force on inflation that is absent in standard models.

[46] This is not the same mechanism as in Christiano, Eichenbaum, and Trabandt (2015), in which a jump in credit spreads increases the cost of working capital, increasing marginal costs. Klemperer (1995) also draws attention to the notion of market share as an investment good, with the concomitant influence of the interest rate on prices. For a model featuring countercyclical markups driven by exit, in the absence of financial frictions, see Cheremukhin and Tutino (2016).

[47] Asplund, Ericksson, and Strand (2005), Lundin et al. (2009), and Montero and Urtasun (2014) find similar evidence. Gilchrist et al. (2018) find a similar dichotomy between firms in financially weak versus financially strong countries in Europe. They further find that the deviations of price trajectories from the predictions of a *standard* Phillips curve can be related to financial constraints.