Variation in the Phillips Curve Relation across Three Phases of the Business Cycle

Richard Ashley and Randal J. Verbrugge

Variation in the Phillips Curve Relation
across Three Phases of the Business Cycle
Richard Ashley and Randal J. Verbrugge

We use recently developed econometric tools to demonstrate that the Phillips curve unemployment rate–inflation rate relationship varies in an economically meaningful way across three phases of the business cycle. The first ("bust phase") relationship is the one highlighted by Stock and Watson (2010): A sharp reduction in inflation occurs as the unemployment rate is rising rapidly. The second ("recovery phase") relationship occurs as the unemployment rate subsequently begins to fall; during this phase, inflation is unrelated to any conventional unemployment gap. The final ("overheating phase") relationship begins once the unemployment rate drops below its natural rate. We validate our findings in a forecasting exercise and find statistically significant episodic forecast improvement. Our analysis allows us to provide a unified explanation of many prominent findings in the literature.

Richard Ashley is at Virginia Tech (ashleyr@vt.edu). Randal J. Verbrugge is at the Federal Reserve Bank of Cleveland and NBER/CRIW (randal.verbrugge@clev.frb.org).

1. Introduction

Few macroeconomic relationships have received as much attention as the Phillips curve, which postulates an inverse relationship between inflation and "slack," the gap between the unemployment rate and its natural rate. Since its discovery in 1958, this relationship has played a key role in macroeconomic thinking. A (structural) Phillips curve is at the heart of the New Keynesian models that dominate current monetary policy discussions. Inflation dynamics and their relationship to labor market slack continue to play a fundamental role in monetary policy deliberations (see, e.g., Brainard, 2017).

Yet despite its continued importance to macroeconomics, some have questioned whether the Phillips curve is a useful characterization of inflation dynamics at all. Given the depth of the Great Recession and the historically large degree of labor market slack, conventional empirical estimates of this relationship predicted sharp and prolonged disinflation (see, e.g., Williams 2010, Ball and Mazumder 2011). Yet core inflation measures did not dip into negative territory even once during this episode. By most conventional measures, slack remained very high even throughout 2013, yet inflation continued to remain positive. Saint Louis Fed President James Bullard has 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" (Bullard, 2017), and such views have clear policy implications. Bob Hall famously stated in his AEA presidential address (2011), "It is not news that NAIRU theory is a failure."[1]

Is a reduced-form Phillips curve, specified in terms of the (overall) unemployment gap, a useful way to think about inflation? We say yes – provided that the specification used allows for its dependence on the business cycle phase. Using a richer model specification, this paper explains the true nature of the Phillips relationship, including its apparent weakening. In addition to rationalizing the missing disinflation puzzle, we provide a unified explanation of several other prominent recent findings, such as the odd behavior of the reverse-engineered NAIRU in Coibion and Gorodnichenko (2015), the fact that the Phillips relationship appears to be only episodically helpful in forecasting (Stock and Watson, 2009, 2010), that it appears to have a convex-concave relationship (e.g., Barnes and Olivei 2003), that it is time-varying (e.g., Clark and McCracken 2006 or Stock and Watson 2007), and that it appears to have vanished after 2009 (Luengo-Prado, Rao, and Sheremirov, 2017). We do this without reference to the short-term

---

[1] For similar sentiments, see The Economist (2017), Summers (2017), and Blinder (2018).

unemployment rate (cf. Ball and Mazumder, 2019) – although we agree with those authors that so-called "core" inflation measures are deficient and mask the true Phillips curve relationship (see Appendix). In Section 6, we further point to a body of extant theory that is consistent with these findings.

We use the recently developed tools of Ashley and Verbrugge (2009) to demonstrate that, from a reduced-form perspective, there are actually three distinct Phillips curve relationships. Each operates during a different phase of the business cycle. In contrast, almost all previous research has conceptualized the Phillips unemployment-inflation association as a single relationship between the unemployment gap and the inflation rate (usually, though not always, assumed to be linear); this model misspecification has resulted in a misleading characterization of the true relationship.[2] We validate our findings via an out-of-sample forecasting exercise. We find – completely in keeping with our in-sample evidence – strong statistical evidence for forecast improvement from our specification compared to that provided by some standard benchmarks.

The first Phillips relationship operates near the beginning of a recession. This is the relationship highlighted by Stock and Watson (2010),[3] who proposed a "recession gap" measure for empirical inflation modeling. Our "bust" gap term, defined below, can be thought of as a refinement of their measure. As those authors found, a significant downward force operates on inflation near the beginning of a recession: Inflation decelerates sharply during the time when unemployment is climbing rapidly. However, shortly after the unemployment rate peaks, the *correlation* between inflation and the unemployment gap vanishes.

This cessation of a Phillips relationship forms the second relationship, which is perhaps more properly termed a *non*-relationship. It is operative during the recovery phase of the business cycle. Irrespective of the amount of labor market slack, such slack is entirely unrelated to inflation during this phase.[4] This fact does not fit neatly into a typical empirical specification of a Phillips curve, but aligns nicely with standard industrial organization theory, as discussed below in Section 6.2.

The third relationship operates when the recovery is well along – and, more precisely, when the unemployment rate falls below its natural rate. (Here, the appropriate natural rate

---

[2] We discuss other nonlinear specifications in the literature below.
[3] See also Meir (2010).
[4] Again, see Luengo-Prado, Rao, and Sheremirov, 2017. Laxton, Meredith and Rose (1995) and Debelle and Laxton (1997) also draw attention to the weak relationship between positive slack and inflation.

concept [and measurement] relates entirely to labor market flows [see Tasci 2018] and is distinct from a NAIRU or NIIRU, as we explain below. See also Tasci and Verbrugge [2014].) From this point onward, our "boom" or "overheating" gap term becomes operative, with labor market tightness being strongly related to subsequent inflation.

Our approach is entirely empirical. We limit the scope of this paper by eschewing any structural modeling and the specification of any new theory. The goal here is to provide empirical insight into inflation dynamics, which will, in turn, inform structural modeling and policy analysis. But we end the paper with a discussion of extant theory that is consistent with our findings. Foreshadowing this discussion, here are three prominent such examples within the New Keynesian framework. First, the overheating relationship can be rationalized by models with capacity constraints or heterogeneous worker ability, as in the New Keynesian models of Alvarez-Lois (2004), Kuhn and George (2019) or Moscarini and Postel-Vinay (2019). Second, Madeira (2014) demonstrates that firms' marginal costs, which are the central driver of inflation in New Keynesian Phillips curves, drop markedly at the onset of a recession, as firms cease using overtime labor. And third, as noted above, standard industrial organization theory predicts the inflation behavior seen at the onset of a recession; see Gilchrist et al. (2017) for an application within a New Keynesian model.

2. Our Empirical Approach

2.1 Description

Our approach to unemployment gap estimation differs fundamentally from the existing Phillips curve literature in that we make use of the "persistence dependent regression" methodology developed in Ashley and Verbrugge (2009) and Ashley, Tsang and Verbrugge (2018). These methods allow us to use the data to determine the relationship between inflation and movements in the unemployment rate at different persistence levels while sidestepping many of the pitfalls often involved in specifying the time behavior of a "natural rate" of unemployment.

Phillips curve forecasting models are usually specified in terms of gaps: for unemployment-based gap models, $gap_t \equiv \left(u_t - u_t^* \right)$, where u_t^* is a natural rate of unemployment,

a "long-term forecast," or a "NAIRU" estimate.[5] Though there is no reason to expect that u_t^* is a fixed constant, much previous research has made this assumption. The vast majority of studies that relax this constancy assumption rely upon the CBO estimate of the natural rate. Less commonly, some studies implicitly assume that u_t^* is an I(1) process and estimate the relationship in differences. And a few studies attempt to model the time evolution of an I(1) u_t^* using a Kalman filter approach (e.g., Brayton, Roberts, and Williams, 1999; Salemi, 1999). Some studies extract an estimate of the time evolution of u_t^* using splines, the HP filter, or bandpass filters, as in Staiger, Stock, and Watson (1997) and Ball and Mankiw (2001).

We are skeptical of the validity of this plethora of approaches and therefore see reliance on any of them as likely to distort the estimation of the relationship between inflation and unemployment. Given the prominence of the CBO estimate in studies of this sort, we discuss its drawbacks for such purposes in some detail. First, the CBO estimation method is not itself constant over time. Historically, the CBO used a Phillips curve to deduce a natural rate (and, indeed, the CBO's natural rate used to be called a NAIRU), but this method has been abandoned. In the 2018 (current vintage) CBO natural rate series, the pre-2004 estimates are still based on Phillips curve estimation methods; in contrast, estimates from 2005 onward are based on a different method (see Shackleton 2018). Surely, as Occhino (2018) notes, a Phillips-curve-based u_t^* estimate of this nature should not generally be used when the object of study is the Phillips curve itself.

Furthermore, there is no compelling theoretical argument in favor of either CBO estimation method. As Tasci and Verbrugge (2014) discuss, there are multiple distinct u_t^* concepts. The u_t^* that is relevant for a Phillips curve inflation relationship need not coincide with a theoretical long-run frictionless rest point of the labor market. Thus, whether or not inflation is closely related to a CBO-based gap measure is an empirical question. We find below that the CBO gap is a poor approximation to the gap measures that do turn out to be relevant for inflation dynamics.

Moving on from the CBO measures, previous filtering approaches – e.g., Kalman filtering approaches – make quite specific assumptions about the way that u_t^* evolves over time, assumptions the validity of which are hardly compelling. Finally, any two-sided filtering –

---

[5] We note in passing that such concepts are distinct but often conflated; see Tasci and Verbrugge (2014) for a discussion and evidence. Clark and Laxton (1997), among others, also distinguish between the NAIRU and the natural rate.

whether bandpass filtering, splines, the CBO approach, or Kalman filtering followed by smoothing – will induce inconsistency in parameter estimation if there is any feedback from inflation to the unemployment rate gap. We discuss this issue briefly below, but for more details, see Ashley and Verbrugge (2009) or Ashley, Tsang and Verbrugge (2018).

We outline our model specification discovery process in Section 2.2 immediately below. Our approach is to allow the data to speak to the relationship between the inflation rate and the unemployment rate, using as few auxiliary assumptions as possible. In keeping with the notion that the Phillips curve is typically estimated in gap form, we do conjecture that there is a long-run "rest point" in the labor market – a "natural rate" – whose movements are basically unrelated to inflation. In this regard, we make use of the natural rate concept and estimate of Tasci (2018); crucially, this estimate is based on trends in labor market flows alone and is thus constructed without reference to inflation; we verify below that this u_t^* estimate is empirically unrelated to inflation. Next, we conjecture that the relationship between inflation and very persistent movements in the unemployment rate may be different from the relationship between inflation and highly transient movements in the unemployment rate. In particular, we find that the Phillips curve relationship is strongest for movements that are on the less persistent part of "business-cycle" frequencies. Our results are consistent with the evidence in Stock and Watson (2010), but – as explained below – our specification goes substantially beyond theirs.

Our methodology is, in fact, related to a traditional frequency-filtering approach and, in that respect, has some resemblance to the older frequency domain methods involving gain and phase. But we note that these more traditional approaches will yield results that, at best, lack any clear interpretation when feedback is present in the relationship.[6] Fundamentally, the two-sided filtering inherent in traditional frequency-domain approaches interacts with any feedback in the relationship to induce correlations between the filtered series and the relevant regression error terms, thus producing inconsistent parameter estimates. (For more details, see Sargent 1987 or Ashley and Verbrugge 2009.)

Instead, we here use the persistence-dependent regression method first introduced in Ashley and Verbrugge (2009). Essentially, we use a one-sided filtering technique to partition the real-time unemployment rate into components with differing levels of persistence, but which (by

---

[6] Transfer function gain and phase plots are substantially more challenging to interpret than the Ashley-Verbrugge (2009) approach, especially where (as here) bi-directional causality is likely. For example, Granger (1969) notes, "In many realistic economic situations, however, one suspects that feedback is occurring. In these situations the coherence and phase diagrams become difficult or impossible to interpret, particularly the phase diagram."

construction) sum up exactly to the original unemployment rate series. It is then straightforward to determine whether inflation movements are related to the persistent part in the same manner in which they are related to less persistent parts. The technique is briefly described as follows; more details are given in Ashley and Verbrugge (2009) and Ashley, Tsang and Verbrugge (2018).

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 various persistence (or frequency) components. One-sided filtering is necessary for two reasons: First, two-sided filtering cannot be conducted using real-time data; second – as noted above – two-sided filtering results in distorted coefficient estimates and generally destroys the ability to make causal statements.[7]

It is well-known that bandpass filters have poor properties near the beginning and the end of the sample, making one-sided filtering challenging. The precision of the partitioning of the unemployment rate into components with different persistence levels is substantially enhanced by extending the data in each window with forecasts or "projections" (see also Dagum (1978),[8] Mise, Kim and Newbold (2005) or Clark and Kozicki (2005)). Thus, for example, to partition the unemployment rate at time t, we apply the Christiano-Fitzgerald filter to a rolling window consisting of 36 monthly observations ending in time period t, post-pended with 12 forecasts, starting at time t+1. (We use the real-time quarterly unemployment rate forecasts from the Survey of Professional Forecasters – converted to monthly frequency by linear interpolation – for this purpose.)[9] The filter is then applied to this resulting 48-month window, and the time-t decomposition value is saved.

2.2 Model Specification Discovery Process

Our final (preferred) specification splits the conventional unemployment gap term into six parts and involves asymmetry. *A priori*, this may sound a bit overly complex. What led us to this specification? Normally, researchers do not report this specification discovery process, but we

---

[7] Two-sided filters (such as the HP filter) applied to both the explanatory and dependent variables, as in classical RBC studies or in recent New Keynesian DSGE modeling, distort relationships among variables that are contemporaneously cross-correlated or in a feedback relationship, because two-sided filtering inherently mixes up future and past values of the time series; see Ashley and Verbrugge (2009) for a more detailed exposition of this point. 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.

[8] Extension of the data via forecasts to improve filter precision has a long history (starting with Dagum 1975), and forms a key part of typical seasonal adjustment procedures. Similarly, the RATS implementation of the Baxter-King filter automatically pads the data with forecasts.

[9] While forecasting recessions is challenging, SPF forecasts appear to be adequate for our analysis.

feel that this explanation is warranted here so as to forestall an erroneous assumption on the reader's part that we conducted an extensive data-mining exercise. (To further reassure the reader, in Section 5 below we provide out-of-sample forecasting evidence in support of our model specification.)

At the outset of this research, our motivating hypothesis was that the relationship between inflation and unemployment rate fluctuations was likely to be frequency-dependent. In particular, we conjectured that highly persistent fluctuations in the unemployment rate – including natural rate fluctuations, but perhaps including somewhat less persistent fluctuations as well – would be unrelated to inflation. Similarly, we conjectured that very high-frequency fluctuations would be unrelated to inflation. In short, we expected to find a Phillips curve relationship restricted to fluctuations of "moderate" persistence levels. Our research plan was to verify this conjecture and let the data themselves speak as to defining "moderate."

Hence, we started with a simple baseline Phillips forecasting model, specified in gap form using an unemployment rate gap; this is detailed in Section 3. We then replaced this unemployment gap by a tripartite partition. In particular, we first formed an unemployment gap by subtracting the natural rate estimate of Tasci (2018) from the real-time unemployment rate,[10] and then partitioned this gap into three parts using our one-sided persistence-decomposition method. Crucially, these three persistence components add up to equal the original gap data; that is why the word "partition" is appropriate here.

Figure 1 below displays four series: a depiction of the two raw series forming the gap, and then two of the three persistence components that partition the gap. (To keep the figure uncluttered, we do not graph the highly transient component; its time series behavior is very noisy. See Appendix A.5.) At the top of the figure are the real-time unemployment rate and the Tasci (2018) natural rate estimate. The difference between these two series is the gap that we partition into three gap components of differing persistence levels. At the bottom of the figure are the two more-persistent components of the gap. The dotted green "Persistent gap" component consists of fluctuations that take longer than 48 months to reverse. The red "Moderately persistent fluctuations" component consists of fluctuations that take between 14 and 48 months to reverse.

---

[10] As noted above, we first tested for feedback in the relationship between the Tasci (2018) natural rate estimate and inflation. It is absent, implying that this two-sided-smoothed series can safely be used in the present study.

Inspecting this figure, it is evident that the persistent gap is basically a smooth version of $\left(u_t - u_t^*\right)$, the overall unemployment gap, *except at the onset of a recession*. Conversely, the moderately persistent component rises rapidly at the beginning of a recession *but then falls almost as rapidly*, so that it is essentially back to zero by the time the unemployment rate peaks. Evidently it coincides with, or signals, the *onset* of a recession. We will return to this observation below; the joint behavior of these two gap terms will play a role in explaining an otherwise puzzling finding in Coibion and Gorodnichenko (2015).

Figure 1: One-Sided Partition of the Unemployment Rate Gap
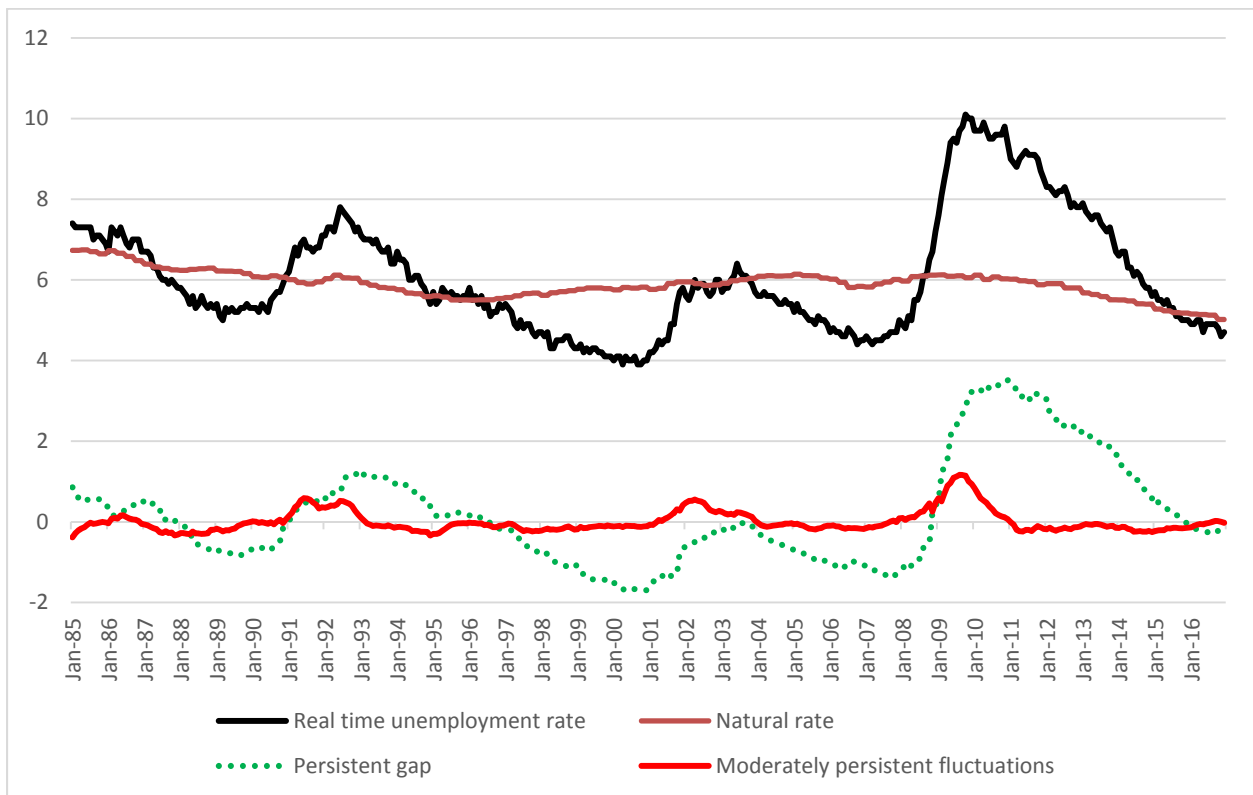
Were our conjectures correct? Yes ... after a fashion. Estimated regression models based on these conjectures did indicate that the persistent gap appeared to be unrelated to inflation, and further that the relationship between transient noise and inflation seemed different and much weaker, with marginal statistical significance. Subsequently, however, we made two crucial observations that led us to refine our specification. First, casual inspection of the relationship between inflation and the medium-persistent part suggested that the comovement between these series was much more evident at the beginning of recessions, when the unemployment rate was climbing rapidly – i.e., during those periods when the moderately persistent component of the

unemployment gap was positive. Moreover, this moderately persistent component is clearly asymmetric (with sharper peaks and long, shallow valleys), indicating a nonlinear data-generating mechanism; this itself suggests that its relationship with inflation might also be asymmetric. Second, after examining a plot of the regression residuals from this initial model, we noticed that the fitting errors were larger near the ends of expansions.[11] Not coincidentally, these precisely corresponded to periods during which the most persistent component of the gap was negative; this component is evidently asymmetric as well.

Motivated by these findings, we investigated whether the relationship between inflation and each of these three gap terms was asymmetric, and we found that all three were. Furthermore, information criteria (AIC and BIC) do not favor a simpler specification that reduces the number of components to four.

This asymmetry might still feel *ad hoc* to some readers. We have three responses. First, the persistent gap asymmetry reflects the nonlinearity in the PC relationship that has been posited from its very first exposition,[12] and that is often historically and presently (2018) described in terms of "overheating."

Second, the asymmetry in the other two gap components mirrors the asymmetry that is built into the recession gap term found in the related work of Stock and Watson (2010). This paper drew attention to the compelling empirical regularity that recessions coincide with a deceleration of inflation. Stock and Watson introduced a "recession gap" measure, which in their study using quarterly data is defined as

$$u_t^{SWgap} \equiv u_t - \min\left[u_t, ... u_{t-11}\right]$$

The point of this measure, of course, is to be able to econometrically address the deceleration in inflation that historically occurs early in recessions, and then evidently ceases. In Appendix A3, we provide a graphical comparison and an out-of-sample forecast comparison to demonstrate that our measure gives rise to forecasts that are superior to those provided by a specification in terms of the Stock-Watson recession gap.

---

[11] We thank a colleague, Ed Knotek, for pointing this out to us.

[12] "When the demand for labor is high and there are very few unemployed we should expect employers to bid wage rates up quite rapidly, each firm and each industry being continually tempted to offer a little above the prevailing rates to attract the most suitable labor from other firms and industries. On the other hand it appears that workers are reluctant to offer their services at less than the prevailing rates when the demand for labor is low and unemployment is high so that wage rates fall only very slowly. The relationship is therefore likely to be highly nonlinear." [William Phillips 1958, p. 283]

Third, as we noted above, upon the imposition of this asymmetry, the usual information criteria improve. And further, in Section 5, we provide out-of-sample forecasting evidence that supports its existence. Thus, at least we can say that the better sample fit and improved forecasting ability of our model specification amply compensates for its greater complexity.

3. Econometric Model

Thus, we posit a Phillips curve forecasting model in which candidate unemployment gap estimates are used in conjunction with lagged inflation to predict the growth rate in the price level between today and 12 months from today, relative to (an estimate of) trend inflation. 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).[13] We find that including 24 months of lagged inflation (by our use of the current 12-month inflation rate, and of the 12-month inflation rate from a year ago) results in a satisfactory specification. Most of our analysis focuses on trimmed-mean PCE inflation, the realized-inflation measure that arguably best removes noise from inflation (Mertens, 2016), but we consider other series for the sake of robustness.

Prior studies, such as those of Clark and McCracken (2008), Kozicki and Tinsley (2012), 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

---

[13] Regression estimates using the CBO un_t^* estimate are investigated below also. As this working paper was going for final internal review, we learned of Ball and Mazumder (2018). 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" inflation measures as measures of trend inflation. See Carroll and Verbrugge (2019) and Verbrugge, Zaman and Nunna (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.

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.

Here we focus on the post-1985 period, during which the Phillips curve is thought to have become far weaker or nearly vanished.

We make two remarks. First, while standard, Equation (1) is linear in the gap. This linearity assumption departs from the original Phillips curve relationship specified in Phillips (1958) (and, indeed, in many undergraduate textbooks) that posited a relationship that was steeper at higher levels of economic activity. Whether or not (1) is approximately linear in the gap term is an empirical question that has been widely addressed.

In particular, numerous papers have located evidence of some form of nonlinearity in the slope of the Phillips curve. For instance, many papers suggest that the US Phillips curve has a convex shape in terms of unemployment, meaning that as unemployment falls below its sustainable level, the upward pressure on inflation rises increasingly, on the margin (see, e.g., Turner, 1995; Clark et al., 1996; Debelle and Laxton, 1997.) Nalewaik (2016) embeds a nonlinear Phillips curve into regime-switching processes for wage and price inflation. The nonlinearity is "a sharp steepening of the Phillips curve after labor-market slack becomes sufficiently negative," a finding consistent with the work of Fisher and Koenig (2014) and Kumar and Orrenius (2016). Neither Murphy (2017) nor Detmeister and Babb (2017) find evidence of this form of nonlinearity at the national level, but they do find some in metropolitan-level data; in particular, the slope of the Phillips curve is steeper at low levels of the unemployment rate.[14] (However, they find that this does not seem to influence forecasts very much, a point we return to below.) Among other authors cited below, Barnes and Olivei (2003) find that the Phillips curve is steeper both at very low *and* very high levels of unemployment, a "convex-concave" shape. Our preferred specification (Equation 4 below) nests this specification, but demonstrates its deficiency as well. In particular, we find that the concave shape pertains only to the recession itself, but *not* to the (prolonged) recovery period.

---

[14] Mercile (2018) also points to city-level evidence of an overheating/inflation link.

Second, Equation (1) imposes that, aside from the movements in un_t^*, the relationship between inflation and movements in the unemployment rate is the same, 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 high-frequency (i.e., low-persistence) variation in the unemployment rate influences pricing in the same way that business-cycle variation does. (For example, the work of Stock and Watson (2010), discussed below, suggests that the relationship is strongest for a particular type of business-cycle fluctuation.) Again, the assumption that there is no persistence-dependence in the Phillips curve relationship is an assumption that can be investigated empirically, and we do so here.

In particular, the specifications considered in this paper extend Equation (1) to include more than one unemployment gap term. First, as an initial extension, we replaced gap_t in (1) with the three gap terms, partitioned into three persistence levels, as discussed above:

$$\pi_t^{12} - \pi_t^* = \alpha + \beta_1\left(\pi_t^{12} - \pi_{t-12}^*\right) + \beta_2\left(\pi_{t-12}^{12} - \pi_{t-24}^*\right)$$
$$+ \lambda_1\left(gap_{P,t}\right) + \lambda_2\left(gap_{modP,t}\right) + \lambda_3\left(gap_{transient,t}\right) + \varepsilon_t \tag{2}$$

In Equation (2) and subsequently, $\left(gap_{P,t}\right)$ refers to the most persistent ("P") part of the unemployment gap (with a reversion period of 48 months or longer), and $\left(gap_{modP,t}\right)$ refers to the moderately persistent part of the unemployment gap (with a reversion period between 12 and 48 months).

As noted above, our data exploration suggested asymmetry in the impact of the first two of these gap terms. But a priori, we expect that a similar asymmetry might be present in the third term as well; we will allow the data to speak. In particular, we allowed the $\left(gap_{P,t}\right)$, $\left(gap_{modP,t}\right)$ and $\left(gap_{transient,t}\right)$ terms to have an asymmetric relationship with inflation:

$$\pi_t^{12} - \pi_t^* = \alpha + \beta_1\left(\pi_t^{12} - \pi_{t-12}^*\right) + \beta_2\left(\pi_{t-12}^{12} - \pi_{t-24}^*\right) + \lambda_1^+\left(gap_{P,t}^+\right) + \lambda_1^-\left(gap_{P,t}^-\right)$$
$$+ \lambda_2^+\left(gap_{modP,t}^+\right) + \lambda_2^-\left(gap_{modP,t}^-\right) + \lambda_3^+\left(gap_{transient,t}^+\right) + \lambda_3^-\left(gap_{transient,t}^-\right) + \varepsilon_t \tag{3}$$

where $\left(gap_{P,t}^+\right)$ equals $\left(gap_{P,t}\right)$ for months such that $\left(gap_{P,t}\right) > 0$, and zero otherwise, and other terms are defined analogously. Both AIC and BIC clearly prefer this specification over (2). Indeed, Equation (3) is our preferred model, as it has a fairly clear interpretation and is consistent with economic theory

13

For completeness, we also tested whether the more parsimonious model obtained from first merging $\left(gap_{modP,t} \right)$ and $\left(gap_{transient,t} \right)$, and then splitting this gap term into its positive and negative parts (Equation (4)), was superior. Both information criteria reject the more parsimonious model.

$$\pi_t^{12} - \pi_t^* = \alpha + \beta_1 \left(\pi_t^{12} - \pi_{t-12}^* \right) + \beta_2 \left(\pi_{t-12}^{12} - \pi_{t-24}^* \right) + \lambda_1^+ \left(gap_{P,t}^+ \right) + \lambda_1^- \left(gap_{P,t}^- \right)$$
$$+ \lambda_2^+ \left(gap_{<48,t}^+ \right) + \lambda_2^- \left(gap_{<48,t}^- \right) + \varepsilon_t \tag{4}$$

To enhance intuition, via more descriptive notation, we will henceforth refer to $gap_{P,t}^-$ as the "boom" gap – since it becomes nonzero in an overheating or boom phase of the business cycle – and $gap_{modP,t}^+$ and $gap_{transient,t}^+$ as "bust" gap terms – since they chiefly become nonzero during the bust phase of the business cycle, when the economy is slipping into a recession.

4. Empirical Results

Our model formulation includes two non-overlapping lags in the dependent variable; there was no apparent need for other lags. HAC estimates (to lag 12) for these (and all other) coefficient estimate standard errors are used here throughout because reference to the sample correlogram of the fitting errors for these models all nevertheless still shows signs of serial correlation, as one might expect, since the dependent variable is in the form of a 12-month cumulative change, and simple model re-specifications did not gracefully eliminate this serial correlation.[15]

The regression results in Table 1 indicate compelling evidence for asymmetry: The coefficient estimates differ notably for positive and negative gaps. And the formal hypothesis testing results for Equations (3) and (4) very clearly reject the null hypotheses that these coefficients are equal, for either the most-persistent unemployment gap component coefficients (λ_1^+ and λ_1^-) and for the moderately-persistent unemployment gap coefficients (λ_2^+ and λ_2^-) – though we can only reject the null hypothesis ($\lambda_3^+ = \lambda_3^-$) at the 6% level of significance. We argue below that this asymmetry is eminently sensible and aligns well with economic theory.

---

[15] Also, it is preferable in the present context to not include in the model specification additional dynamics (in the form of lag structures on the other explanatory variables) that might reduce or eliminate this serial correlation in the errors, as this would complicate the comparison of our results to standard Phillips curve formulations. The HAC standard error estimates are already consistent in the face of any heteroscedasticity in the model errors, so diagnostic checking in this regard is not quoted. Examination of a time plot of the Equation (4) fitting errors shows no outliers.

		Equation (1)	Equation (2)	Equation (3)	Equation (4)
	λ	-0.09			
	(t-stat)	(-1.54)			
	$\lambda_1 \ or \ \lambda_1^+$		-0.06	0.03	0.01
	(t-stat)		(-1.44)	(0.61)	(0.30)
("boom" gap)	λ_1^-			-0.27	-0.26
	(t-stat)			(-3.06)	(-3.07)
("bust" gap)	$\lambda_2 \ or \ \lambda_2^+$		-1.33	-1.68	-1.13
	(t-stat)		(-5.36)	(-7.88)	(-8.60)
	λ_2^-			0.14	-0.00
	(t-stat)			(0.18)	(-0.01)
("transient")	$\lambda_3 \ or \ \lambda_3^+$		-0.30	-0.52	
	(t-stat)		(-1.94)	(-2.38)	
	λ_3^-			-0.03	
	(t-stat)			(0.26)	
	β_1	0.47	0.32	0.28	0.35
	(t-stat)	(7.74)	(4.70)	(4.01)	(6.29)
	β_2	0.11	0.36	0.34	0.27
	(t-stat)	(1.33)	(5.85)	(6.30)	(7.23)
	constant	-0.09	-0.03	-0.05	-0.08
	(t-stat)	(-1.54)	(-0.49)	(-0.45)	(-0.71)
F-test: $\lambda_1^+ = \lambda_1^-$				0.01	0.01
F-test: $\lambda_2^+ = \lambda_2^-$				0.02	0.00
F-test: $\lambda_3^+ = \lambda_3^-$				0.06	
F-test: $\lambda_2^+ = \lambda_3^+$				0.00	
F-test: $\lambda_1^+ = \lambda_1^- = \lambda_2^+ = \lambda_2^- \left(= \lambda_3^+ = \lambda_3^-\right)$				<0.005	<0.005
Adjusted R-squared		0.56	0.69	0.75	0.73
AIC		0.986	0.626	0.438	0.498
BIC		1.038	0.698	0.541	0.580

Table 1. Phillips Curve Regressions.

In Table 1, we quote parameter estimates with t-statistics underneath. The AIC and BIC both suggest that Equation (3) is the best-fitting model, even adjusted for its additional complexity over Equations (2) and (4). In the Appendix, we display parallel results from several other robust inflation estimates. Results are broadly similar. Furthermore, we find that the Equation (1) coefficient estimates are essentially unchanged if we use a "CBO gap" – i.e., $\left(u_t - u_t^{*,CBO}\right)$ – instead of our preferred gap measure, which is based on the Tasci (2018) u_t^* estimates.

Because the null hypothesis $H_0 : \lambda_1^+ = \lambda_1^- = \lambda_2^+ = \lambda_2^- = \lambda_3^+ = \lambda_3^-$ can be rejected with

$p < 0.005$ in our preferred model specification – Equation (4) – these results imply that the

Phillips curve relationship depends on the phase of the business cycle.

There are three statistically significant coefficient estimates: that pertaining to the boom

gap ($gap_{P,t}^-$), that pertaining to the bust gap ($gap_{\mathrm{mod}\,P,t}^+$), and that pertaining to $\left(gap_{transient}^+\right)$. We

note that $\left(gap_{transient}^+\right)$ is negligible except during periods when the unemployment rate is

increasing rapidly.[16] Conversely, the coefficient estimates pertaining to $gap_{P,t}^+$, to $gap_{\mathrm{mod}\,P,t}^-$, and

to $gap_{transient}^-$ are not statistically different from zero. These results have a clear and clean

interpretation: There are effectively three relevant gap terms in the Phillips curve, and this, in

turn, implies that the Phillips curve relationship is phase-dependent; that is, dependent upon the

phase of the business-cycle.

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, with reversion

periods of 48 or more months. Consider the long recovery phase of a typical business cycle.

During much of this recovery, the unemployment rate is above the natural rate, and the

unemployment gap is positive. As long as this relatively persistent (smooth) component of the

unemployment gap is positive, $gap_{P,t}^-$ is zero, while the boom gap, $gap_{P,t}^+$, is positive. Thus, in

contrast to the predictions from alternative Phillips curve specifications, the prediction of this

model is that the estimated impact of the unemployment gap on inflation is essentially *zero*

during the recovery phase. (This follows because $\hat{\lambda}_1^+ \approx 0$ and because the persistent part of the

unemployment gap is the only component which is non-negligible during this period.) Hence,

once the recession has bottomed out and unemployment has peaked, then subsequently (during

the beginning stages of the recovery and for a long while afterward), the unemployment rate is

simply unrelated to inflation. Putting this differently, a persistently high unemployment rate *per

se* has no influence on (more properly, no relationship to) inflation. However, once the

unemployment rate has persistently dropped below the natural rate – which happens late in an

expansion, during the overheating or boom phase – there is evidently a strong upward influence

of a negative (and smooth and persistent) unemployment gap on inflation. The estimated size of

---

[16] Indeed, unreported diagnostic regressions indicate that the importance of this term is driven by *monthly* increases
in the unemployment rate; all remaining portions of the transient component are unimportant.

this coefficient is directly comparable to, and much larger than, the estimate from conventional linear Phillips curves (given in the first column, for Equation (1)). Evidently, the true "upward force" on inflation from this boom gap is obscured by using the traditional gap, which averages the negligible effects of the positive gaps during most of the recovery with the substantial effects of (negative) boom gaps, which enter with a statistically significant ($p = 0.01$) coefficient of -0.27. This differential is the form of nonlinearity in the Phillips curve that much recent research has considered.

Next we focus on $\hat{\lambda}_2^+$ and $\hat{\lambda}_2^-$, coefficient estimates that pertain to moderately persistent fluctuations in the unemployment gap. Inspection of Figure 1 reveals that, by and large, this moderately persistent component only departs from zero when the unemployment rate is rising fairly rapidly. Our Equation (3) and (4) regression estimates bear this out, with $\hat{\lambda}_2^-$ being essentially zero, whereas $\hat{\lambda}_2^+$ is highly significant, at -1.7. This downward force on inflation is effectively reinforced by the transient term, which tends to positively comove with the moderately persistent component (the correlation is 0.26; see Appendix A.5 for a plot of these components alongside one another). We thus conclude that there is a second relationship between unemployment rate fluctuations and price inflation, pertaining to periods of rapidly rising unemployment – e.g., to the "bust" phase of the business cycle. In keeping with the findings of Stock and Watson (2010), at the onset of a recession, there is a pronounced deceleration in inflation. This explains, e.g., the findings in Luengo-Prado, Rao, and Sheremirov (2017): "...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).

In Appendix A4, we discuss how our results compare to some other prominent findings in the literature. For instance, we note that we can reinterpret a finding in Coibion and Gorodnichenko (2015) regarding their reverse-engineered NAIRU. We discuss how our results naturally give rise to episodic forecast improvements and to time-variation in Phillips curve coefficients. We argue that our results explain numerous studies that find a convex-concave aspect to the Phillips curve relationship, and to those that adduce evidence for regime switching.

5. Data Mining? Out-of-Sample Forecasting Results

The statistical significance of the inference results discussed above – which strongly support our nonlinear (asymmetric) Phillips curve formulation, disaggregated according to the phase of the business cycle per Equation (4) – 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 AIC and BIC measures – allowing for the additional number of coefficients estimated in Equation (4).

We find these results persuasive, but not necessarily definitive, in view of the usual concerns as to "data mining." We address this concern in two ways. First, we present supporting results based on out-of-sample (OOS) forecasting calculations, using the Giacomini-Rossi (GR) testing framework. Second, we present what can be called "partially recursive conditional" forecasts, using a variation of our model, and compare these to those that derive from the use of a more conventional specification.

5.1 GR out-of-sample forecast tests

The GR test results discussed below show that our Equation (4) re-formulation of the Phillips curve yields superior out-of-sample forecasting as the GR window moves through the sample data. 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 4 reflect a new set of stable statistical regularities in the historical data rather than the result of torturing a data set into submission. We note, however, that our focus on this forecasting improvement provided by our new specification does not represent an assertion on our part that we have here obtained an inflation forecasting model that can or should supplant all existing ones – that would be the topic of a paper separate from this one.

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, and this conjecture is supported by the GR test results discussed below. This result makes intuitive sense, as the nature of our Equation (4) model is such that it is most likely to make better predictions only during a few, relatively brief portions of the business cycle.

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 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 (exceeding 1.5 percent), either positive or negative. 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 (see, e.g., Rossi and Sekhposyan (2010) and Dotsey, Fujita, and Stark (2017)). In this Section, we examine the conjecture that the better-specified PC model proposed above – specified to include both the (negative) boom gap and the (positive) bust gap – will outperform models in which those terms are omitted.[17]

For our forecasting exercise, we remove all the Equation (4) regression terms with statistically-insignificant estimated coefficients in Equation (4) – e.g., λ_1^+ , λ_2^- , and λ_3^- – and as our forecasting model:

$$\pi_t^{12} - \pi_t^* = \alpha + \beta_1\left(\pi_t^{12} - \pi_{t-12}^*\right) + \beta_2\left(\pi_{t-12}^{12} - \pi_{t-24}^*\right) + \lambda_1^-\left(gap_{P,t}^-\right) \\ + \lambda_2^+\left(gap_{modP,t}^+\right) + \lambda_3\left(gap_{transient,t}^+\right) + \varepsilon_t \tag{5}$$

Because we expect Equation (5) to principally provide episodic forecast accuracy improvements, 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 ideal for comparing the historical out-of-sample forecasting performance of competing models when the relative performance of these models may vary over time. 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.

In the figures below, we plot the upper and lower 10 percent and 5 percent critical values and the GR F^{OOS} test statistic for various forecast comparisons. The test is two-sided, and is based on rolling-window estimates and forecasts. When the F^{OOS} statistic rises above the upper

---

[17] It is worth pointing out that unlike Dotsey, Fujita, and Stark (2017), we study trimmed-mean PCE inflation rather than headline inflation. Other inflation estimators are examined in the Appendix.

critical values, then the forecast performance of the "alternative" (nonlinear PC) model is significantly better (over the previous 48 months) compared to the model to which it is being compared. Conversely, when the F^{OOS} term falls below the lower critical values, the reverse is true.

We compare forecasts from our Equation (5) 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 2 depicts the comparison against a CBO gap model, Equation (1). Short samples are insensible for models like ours, inasmuch as it sharply differentiates between busts, recoveries, and booms. We thus use a 20-year window, and estimate models from 1985:1 onward, so that our first forecast is for 2005:1, i.e., 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 2, the F^{OOS} line remains above zero, indicating that this specification outperforms the baseline CBO specification from 2009:1 to 2016:12. The forecasting improvement gain from the Equation (5) model is statistically significant at the 5% 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 (5) over the baseline model is convincing.

Figure 2: Baseline CBO Model (1) Versus Nonlinear PC Forecasting Model, Equation (5)
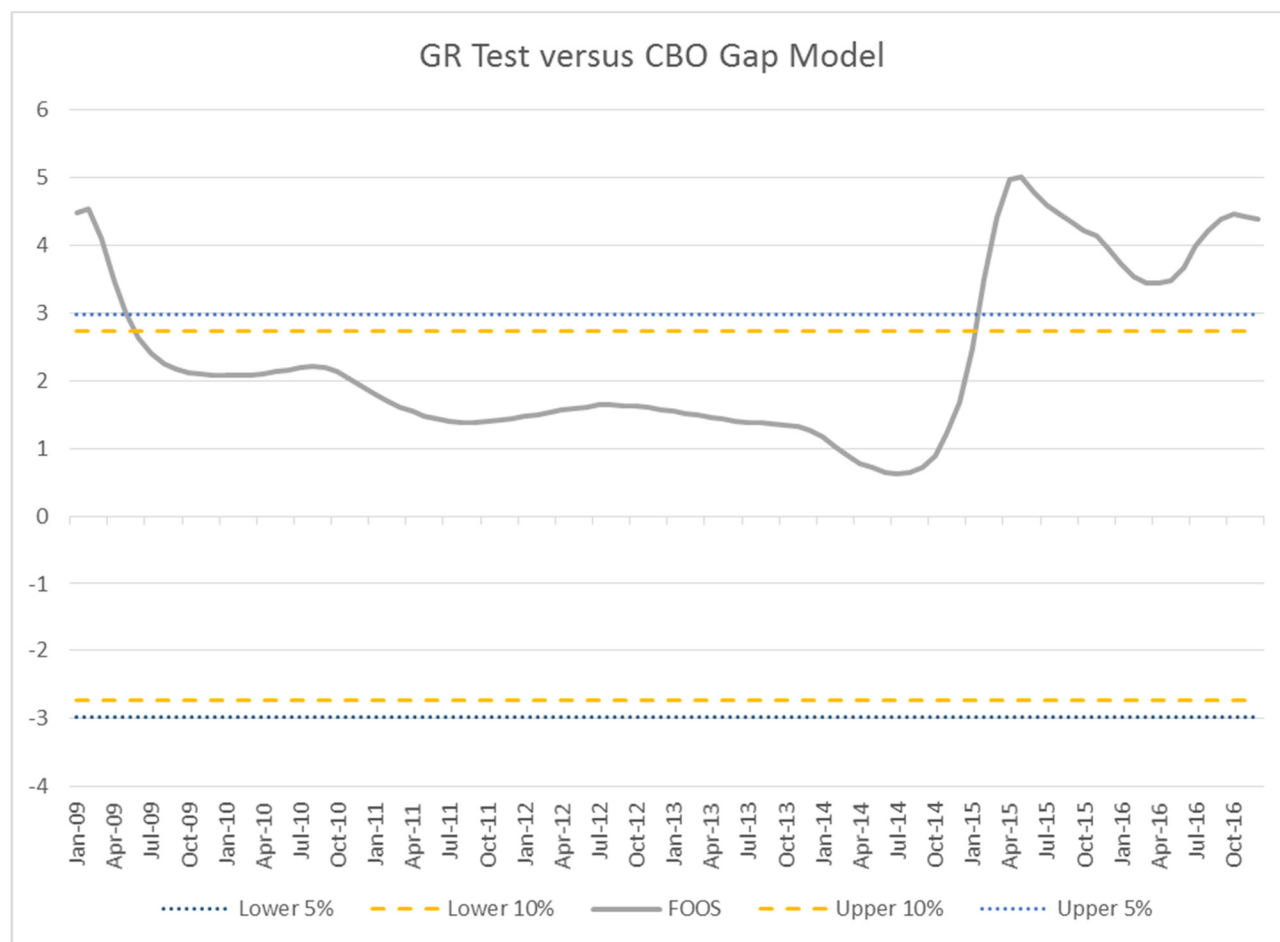

In Figure 3, we display analogous forecast comparison results comparing the OOS forecast performance of Equation (5) to that of an Atkeson-Ohanian-type model. In the latter model, "inflation forecasts over the next four quarters is expected to be equal to inflation over the previous four quarters" (Atkeson and Ohanian, 2001, p.6). Thus, we compare forecasts from Equation (5) against those from the model

$$\pi_t^{12} - \pi_t^* = \left(\pi_{t-12}^{12} - \pi_{t-13}^*\right) + \eta_t \tag{6}$$

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

Figure 3: Univariate Atkeson-Ohanian Model Versus Two-Gap Forecasting Model, Equation (5)
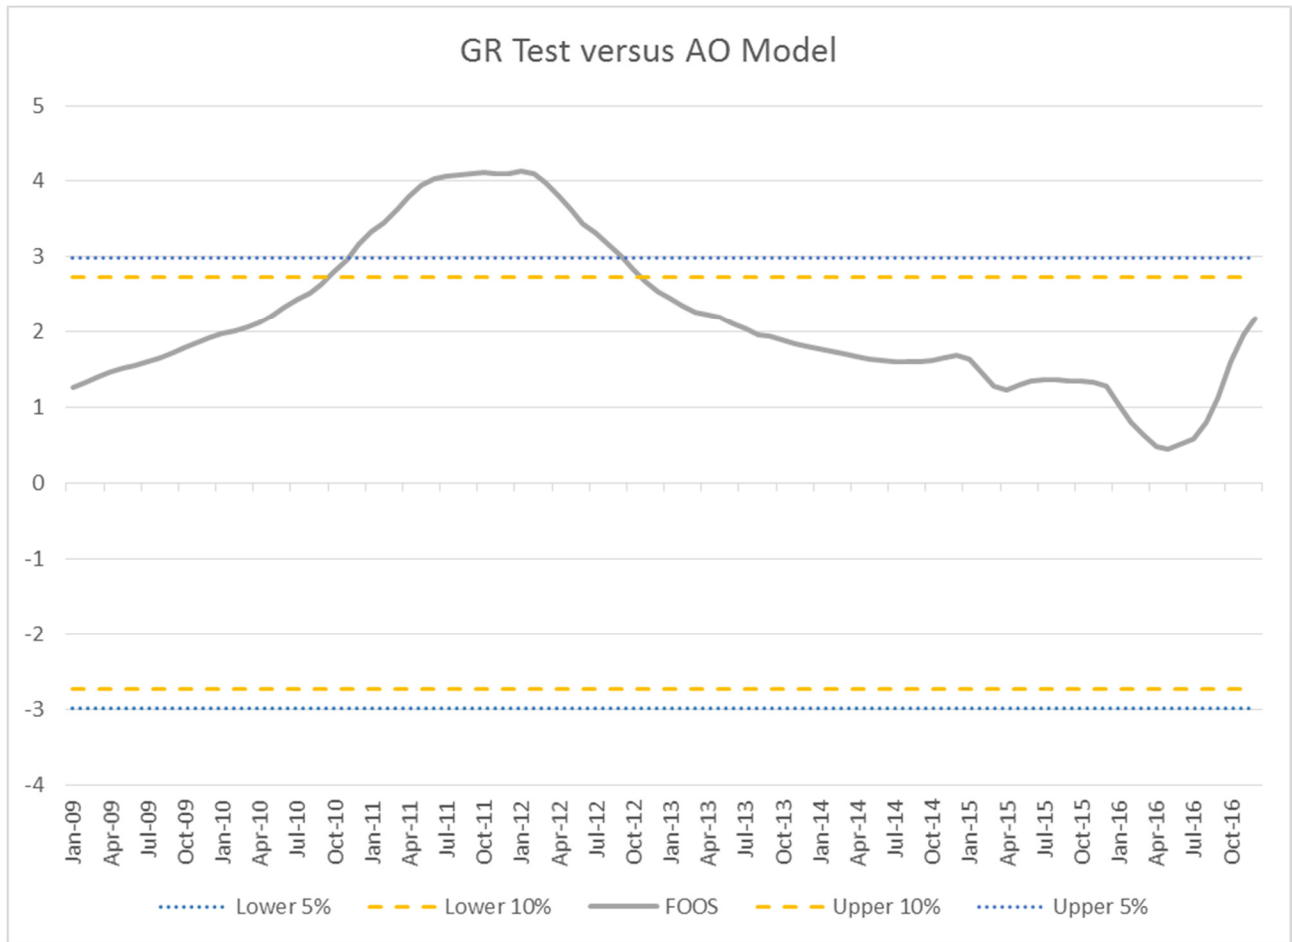
5.2 Partially-recursive conditional forecasts

In this section, we use equation (5) to generate what could be called "partially recursive conditional forecasts;" these are conditioned on the historical unemployment values, but partially recursive in that each forecast draws its needed lagged inflation-deviation values from its own recent inflation forecasts. We compare these conditional forecasts to those from Equation (1), the parallel model that instead conditions on the path of the CBO unemployment gap, not allowing for the asymmetric business-cycle-phase dependence in Equation (5). For reasons that will be apparent below, when we construct these conditional forecasts, we use two distinct sets of parameter estimates. First, we obtain parameter estimates for Equations (1) and (5) using the data through 2006:12 (after which we fix their values); and second, we obtain estimates for both

models derived from the full sample. This yields 4 models that we use to generate conditional forecasts through 2016.[18]

We plot these below in Figure 4. The conditional forecasts generated by Equation (5) do a good job tracking the broad contours of the evolution of inflation over the Great Recession and the recovery. Importantly, this finding holds even if we fix the Equation (5) model parameter values to those estimated at the end of 2006. This result is indicative of stability in this nonlinear relationship: the Phillips curve relationship has *not* weakened.

Conversely, conditional forecasts generated by Equation (1) are quite poor, particularly if coefficient estimates are fixed at 2006:12 levels. The fact that the conditional forecasts from the linear-CBO-gap specification (Equation (1)) are so different indicates that this specification features notable instability; this is the oft-noted purported weakening in the Phillips curve relationship. Furthermore, as Clark (2014) has noted, once one properly accounts for trend inflation, the actual disinflation puzzle pertains to why inflation fell *so fast* during the recession, something our specification gracefully explains. Our specification also gracefully explains the sharp (partial) *recovery* of inflation. Tests of the pronounced divergence between the CBO-gap projections and the actual inflation path support ones impressions: changes in the CBO projections are uncorrelated with changes in the inflation path (the estimated correlation is 0.22 +/- 0.19), whereas the correlation between changes in the Equation (6) projections and changes in the inflation path is substantial (at 0.48 +/- 0.10).[19] This analysis reinforces the message of this paper: a failure to properly specify the relationship between the unemployment rate and the inflation rate yields unstable parameter estimates and strongly counterfactual conditional forecasts.

---

[18] The recursive "forecasts" generated from the "full-sample-estimates" are not truly recursive forecasts, since they incorporate knowledge of the actual inflation path during the recession and recovery. Still, it is interesting to note the extent to which this information would have made a difference.

[19] These both refer to projections from the models which fix coefficient estimates at their 2006:12 level. Correlations with the full-sample CBO-based model are even lower, at 0.14 +/– 0.30; standard error estimates were computing using HAC estimators with 3 lags. Meanwhile correlations from the full-sample Equation (6) model are essentially unchanged, at 0.43 +/– 0.08.

Figure 4: Partially Recursive Conditional Forecasts from Equations (1) and (5)

Conditional Forecasts
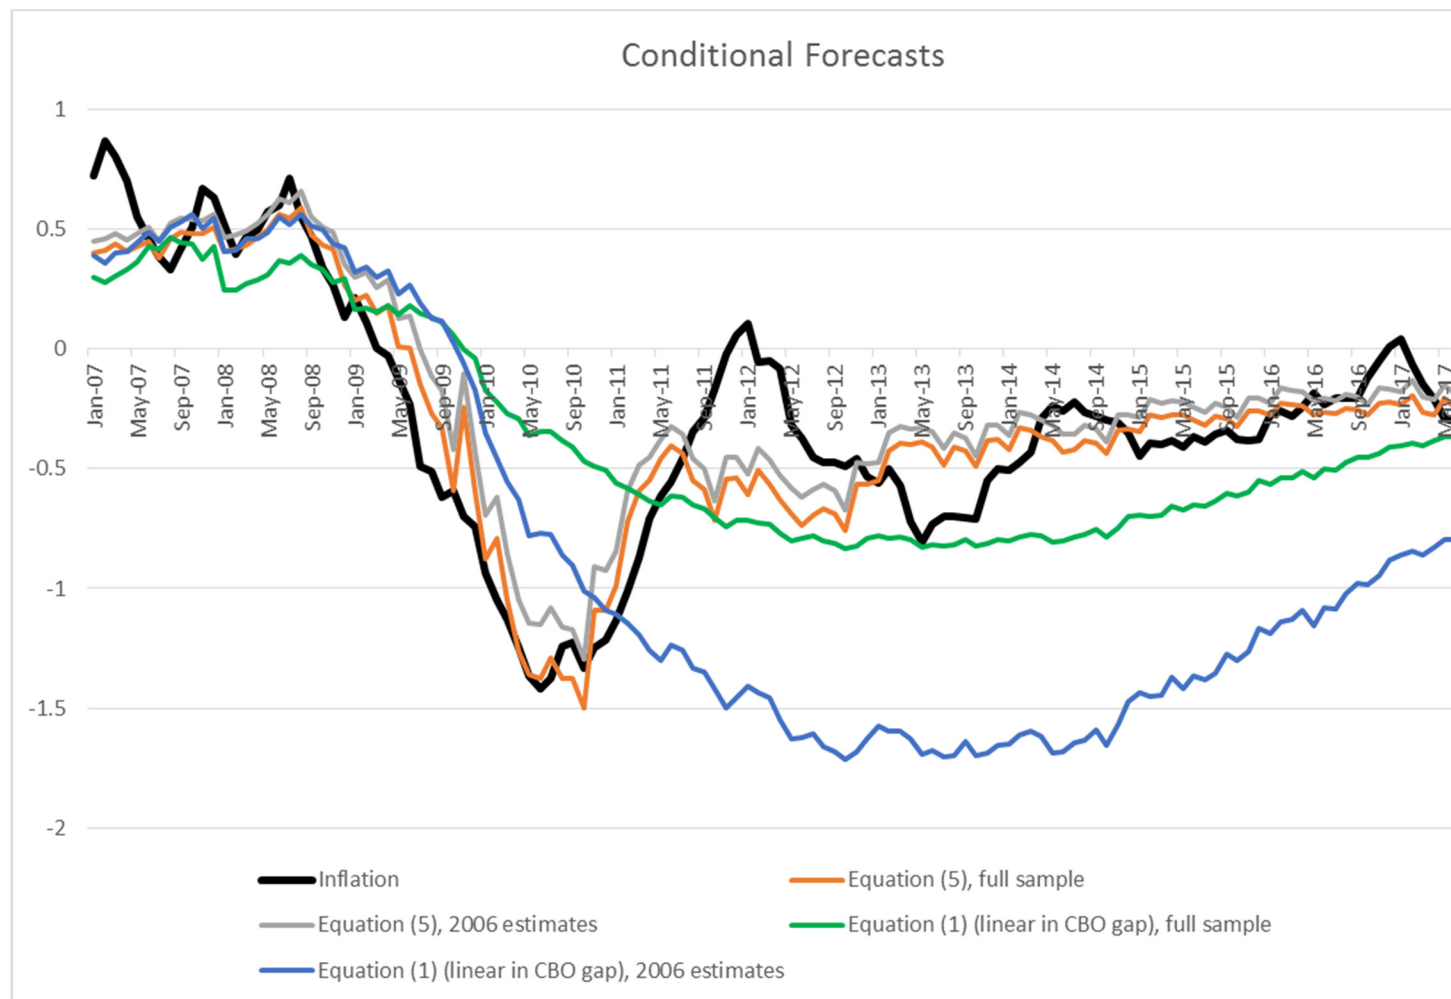
Finally, the reader should note the following several points related to inflation forecasting. First, all reduced-form Phillips curve forecasting models are inherently complicated by endogeneity due to extant monetary policy. In particular, as has been known since Lucas (1976), the empirical (reduced-form) Phillips curve will generally vary with monetary policy.[20] Thus, to the extent that a central bank is successful in controlling inflation, the reduced-form Phillips curve relationship will weaken. Second, consider a forecasting model with a misspecified gap term. Parameter estimation will thus tend to down-weight the gap, but likely still provide unbiased forecasts on average. Consequently, the inflation forecasts generated may not differ substantially from a better-specified gap model only when the gap term is quite large. (We find this to be the case in the current exercise; the projected additional inflation is less than 0.5 percentage points, although this additional forecasting error might still be enough of a

---

[20] For recent studies focused on this point, see Fitzgerald and Nicolini (2014) or McLeay and Tenreyro (2018). Occhino (2018) also provides useful intuition.

difference to influence monetary policy discussions.) Lastly, we again note that our goal in this paper is not to devise an improved forecasting model, but rather to provide insight into inflation dynamics, which will, in turn, be useful for structural modeling and policy analysis.

6. Theory

6.1 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.

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). 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 covaries more strongly with the aggregate job-to-job transition rate than to 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 – i.e., strong upward price pressure when the economy is overheating – is capacity-constraints models.[21] 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

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

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. 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).[22]

At the time of this writing, the unemployment rate has arguably been below the natural rate for well over a year. If our findings are correct, then why are we only now seeing substantial upward pressure in wage and price inflation?

This issue is not unique to this paper; as Faccini and Velosi (2019) note, "Understanding why such tight labor market conditions coupled with low productivity growth have not sparked inflation yet, proves to be very challenging for standard macroeconomic models." There are several alternative explanations. We cannot observe the counterfactual, so there may be

---

[21] 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."

[22] 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.

significant countervailing pressures. It may be that, given the sluggish recovery and prolonged period of low inflation, inflation expectations dropped somewhat and are only slowly recovering.[23] It may be that wage inflation is already "high," given the anemic productivity growth at present. But it may also be that the current recovery is abnormal in some way. Wage inflation is thought by some to depend negatively on both the level and the change in the unemployment rate, as both capture important dimensions regarding the degree of labor market tightness (Blanchard and Gali, 2007, 2010), and this has been a very gradual recovery. Faccini and Melosi (2019) evaluate a model similar to that of Moscarini and Postel-Vinay, and this analysis draws attention to a "malfunctioning" of the job ladder during the last recovery; their analysis suggests that the anomalous behavior of the employment-to-employment flow rate has mitigated the upward inflation force from a tight labor market. From the productive capacity point of view, an economy growing very *slowly* is less likely to hit a "technical" capacity constraint (it can more easily build capacity at a measured pace, and/or is more able to strategically use stockouts and inventory management) and firms may be able to fill vacancies or to make worker-poaching decisions more slowly, which would allow them to use smaller wage increases to attract scarce labor..

6.2 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, e.g., 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.

Further, 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

---

[23] Binder (2017) studies the evolution of US consumer inflation expectations and the FOMC's historical ability to influence them.

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. Gilchrist et al. (2017) develop a New Keynesian model, extended in Gilchrist et al. (2018), that builds upon these insights, and provide 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 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.[24] 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.[25] 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.

7. Conclusion

Being so central a topic to macroeconomics, the Phillips curve is the subject of a vast literature.

---

[24] 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).

[25] 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.

We argue that this literature, to a greater or lesser extent, misspecifies the functional form of the Phillips curve regression equation. This widespread problem has led to erroneous conclusions about the nature of the PC relationship.

We find that the reduced-form PC relationship is "alive and well" – but that this relationship involves two gaps, not one. The first gap operates at the beginning of a recession, when unemployment is rising rapidly. At that time, there is a large reduction in inflation, something Stock and Watson (2010) noted and modeled. However, shortly thereafter, the relationship between unemployment levels and inflation vanishes; high unemployment, even very high unemployment, does not translate into a downward influence on inflation dynamics. Late in expansions, things change again. Shortly after the unemployment rate falls below the natural rate, as measured by Tasci (2018), the economy begins to "overheat" and there is a significant increase in inflation.

We emphasize that for much of the business cycle – for most of the recovery, or the middle phase in our analysis – we find that unemployment (slack) and inflation are apparently not closely related. Thus, studies that effectively average over the entire business cycle will often conclude that the relationship is quite weak, when it emphatically is not weak at all.

The evidence in this paper derives from a reduced-form forecasting model. We are not specifying a structural model, nor providing, or testing, a new theory. On the other hand, we find compelling evidence – including out-of-sample forecasting evidence – in favor of a type of a novel empirical relationship between inflation and the unemployment rate. This relationship depends – in the asymmetric manner specified in our Equation (4) formulation described and estimated in Section 4 above – on the phase of the business cycle. We believe that this empirical result calls out for a theoretical explanation, and we hope that our findings will motivate and guide DSGE modeling efforts, perhaps building upon the theory reviewed in Section 6.

As noted by John Cochrane (Federal Reserve Bank of Richmond, 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." The recent experience of year after year of zero nominal interest rates (supposedly) anchored inflation expectations, and low inflation would seem to illustrate the difficulty of fine-tuning the control of inflation: Apparently inflation expectations do not exert all that strong a

force.[26] In this regard we note that our findings suggest that slack can only increase inflation when the economy is actually overheating, and they suggest that if the Phillips curve is the chief means by which monetary policy influences inflation, then slowing inflation via the Phillips curve mechanism requires a recession.

---

[26] As Cochrane points out, anchored inflation expectations should make the Phillips curve work better.

References

Alvarez-Lois, Pedro P. (2004). "Capacity constraints, idiosyncratic demand uncertainty and the dynamics of inflation." *Economics Letters*, 83(1), pp. 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), pp. 170-193. doi:10.1093/ei/cbi012.

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

Ashley, Richard, Kwok Ping Tsang, and Randal J. Verbrugge (2018). "All fluctuations are not created equal: The differential roles of transitory versus persistent changes in driving historical monetary policy." Working Paper 18-14, Federal Reserve Bank of Cleveland. doi:10.26509/frbc-wp-201814.

Ashley, Richard and Randal J. Verbrugge (2003). "Mis-specification in Phillips curve regressions: Quantifying frequency dependence in this relationship while allowing for feedback." Manuscript, Virginia Tech. (No longer downloadable in this version).

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.

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

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

Ball, Laurence and Sandeep Mazumder (2011). "Inflation dynamics and the Great Recession." *Brookings Papers on Economic Activity*, 2011(1), pp. 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*, pp. 3-18. URL https://ideas.repec.org/a/fip/fedbne/y2003p3-18.html.

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

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

Blanchard, Olivier and Jordi Gali (2007). "Real wage rigidities and the New Keynesian model." *Journal of Money, Credit and Banking*, 39, pp. 35-65. doi:10.1111/j.1538-4616.2007.00015.x.

Blanchard, Olivier and Jordi Gali (2010). "Labor markets and monetary policy: A New Keynesian model with unemployment." *American Economic Journal: Macroeconomics*, 2(2), pp. 1-30. doi:10.1257/mac.2.2.1.

Blanchflower, David G. and Andrew J. Oswald (1990). "The wage curve." *The Scandinavian Journal of Economics*, 92(2), p. 215. 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.

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

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.

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.

Carroll, Daniel and Randal Verbrugge (Forthcoming). "Behavior of a new median PCE measure: a tale of tails." Economic Commentary.

Cheremukhin, Anton and Antonella Tutino (2016). "Information rigidities and asymmetric business cycles." *Journal of Economic Dynamics and Control*, 73, pp. 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." *The American Economic Review*, 86(4), pp. 703-725. URL http://www.jstor.org/stable/2118301.

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), p. 42. doi:10.2307/2673871.

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

Clark, Todd E. and Taeyoung Doh (2014). "Evaluating alternative models of trend inflation." *International Journal of Forecasting*, 30(3), pp. 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.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), pp. 197-232. doi:10.1257/mac.20130306.

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

Dagum, Estela Bee (1978). "Modelling, forecasting and seasonally ad.usting economic time series with the X-11 ARIMA method." *Journal of the Royal Statistical Society, Series D (The Statistician)*, 27(3/4), p. 203. doi:10.2307/2988184.

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, pp. 19-43. doi:10.1016/j.jmacro.2016.01.004.

Dotsey, Michael, Shigeru Fu.ita, 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.

Evans, George (1985). "Bottlenecks and the Phillips curve: A Disaggregated Keynesian model of inflation, output and unemployment." *The Economic Journal*, 95(378), p. 345. doi:10.2307/2233214.

Faccini, Renato and Leonardo Melosi (2019). "Bad jobs and low inflation." *CEPR Discussion Papers* 13628. URL https://sites.google.com/site/lemelosi/JL20190305.pdf.

Faust, Jon and Jonathan H. Wright (2013). "Forecasting inflation." In *Handbook of Economic Forecasting*, pp. 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), pp. 33-61. URL https://ideas.repec.org/a/fip/fedker/y1998iqiiinv.83no.3.html.

Fisher, Richard W. and Evan F. Koenig (2014). "Are we there yet? Assessing progress toward full employment and price stability." *Economic Letter*, 9(13), pp. 1-4. URL https://ideas.repec.org/a/fip/feddel/00015.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.

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.

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

Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakra.sek (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), p. 317. doi:10.2307/2554819.

Granger, Clive W. J. (1969). "Investigating causal relations by econometric models and cross-spectral methods." *Econometrica*, 37(3), p. 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.

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

Hong, Sungki (2019). "Customer Capital, Markup Cyclicality and Amplification." *Manuscript*, Federal Reserve Bank of St. Louis.

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

Huh, Hyeon-seung, Hyun Hoon Lee, and Namkyung Lee (2008). "Nonlinear Phillips curve, NAIRU and monetary policy rules." *Empirical Economics*, 37(1), pp. 131-151. doi:10.1007/s00181-008-0226-x.

King, Robert G. and Mark W. Watson (2012). "Inflation and unit labor cost." *Journal of Money, Credit and Banking*, 44, pp. 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." *The Review of Economic Studies*, 62(4), pp. 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), pp. 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, pp. 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.

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. Kberl (2009). "Capacity utilisation, constraints and price ad.ustments 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.

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

Luengo-Prado, Maria Jose, Nikhil Rao, Viacheslav Sheremirov, et al. (2017). "Sectoral inflation and the Phillips curve: What has changed since the Great Recession?" Current Policy Perspectives 17-5, Federal Reserve Bank of Boston. URL https://EconPapers.repec.org/RePEc:fip:fedbcq:2017_005

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

McLeay, Michael and Silvana Tenreyro (2018). "Optimal inflation and the identification of the Phillips curve." LSE Research Online Documents on Economics 90373, London School of Economics and Political Science. LSE Library. URL https://ideas.repec.org/p/ehl/lserod/90373.html.

Mericle, David (2018). "What makes inflation pop?" Goldman Sachs *US Daily*. 17 August 2018.

Mertens, Elmer (2016) "Measuring the Level and Uncertainty of Trend Inflation." *The Review of Economics and Statistics* 98(5),pp. 950-967. doi:10.1162/rest_a_00549.

Mikosch, Heiner F. (2009). "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.

Moscarini, Giuseppe and Fabien Postel-Vinay (2017a). "The job ladder: Inflation vs. reallocation." Working paper, Yale University. URL https://campuspress.yale.edu/moscarini/files/2017/11/ nominal-rigidities-1p426m8.pdf.

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

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. 2018. "The Flattening of the Phillips Curve: The Policy Implications Depend on the Cause." Manuscript in preparation, Federal Reserve Bank of Cleveland.

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), pp. 779-794. doi:10.1016/j.jedc.2012.01.009.

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), pp. 283-299. doi:10.2307/2550759.

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), pp. 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." *The American Economic Review*, 74(3), pp. 433-444. URL https://www.jstor.org/stable/1804018.

Stock, James H. and Mark W. Watson (1999). "Forecasting inflation." *Journal of Monetary Economics*, 44(2), pp. 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), pp. 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*, pp. 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.

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

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). URL https://ideas.repec.org/a/fip/fedcec/00026.html.

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.

U.S. 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.

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

U.S. 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.

U.S. 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

U.S. 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.

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.

Xu, Qifa, Xufeng Niu, Cuixia Jiang, and Xue Huang (2015). "The Phillips curve in the US: A nonlinear quantile regression approach." *Economic Modelling*, 49, pp. 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). URL https://ideas.repec.org/a/fip/fedcec/y2013inov15n2013-16.html.

Appendix

A1. Other Inflation Indicators

In Table A1 we present the results using several other inflation indicators. We provide the trimmed PCE results from Table 1 for comparison. The final two lines refer to GR and DM forecast comparisons against the baseline CBO model for the same dependent variable.

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

Table A1

Regression results are presented in Table A1. 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.

A2. Whither the Core PCE Results?

We acknowledge that the core PCE tests do not reject symmetry in the bust gap. Like Ball and Mazumder (2018), 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 CPE 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. As we note 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.

There were only three NBER recessions post-1985. This implies that the bust gap only experienced 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.[27] 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[28] 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.[29] 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.

---

[27] 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.

[28] 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.

[29] See also Peach, Rich, and Linder (2013), who display a decomposition into goods and services.

Figure 4: Four Inflation Measures

Sources: BEA, BLS, Dallas Fed
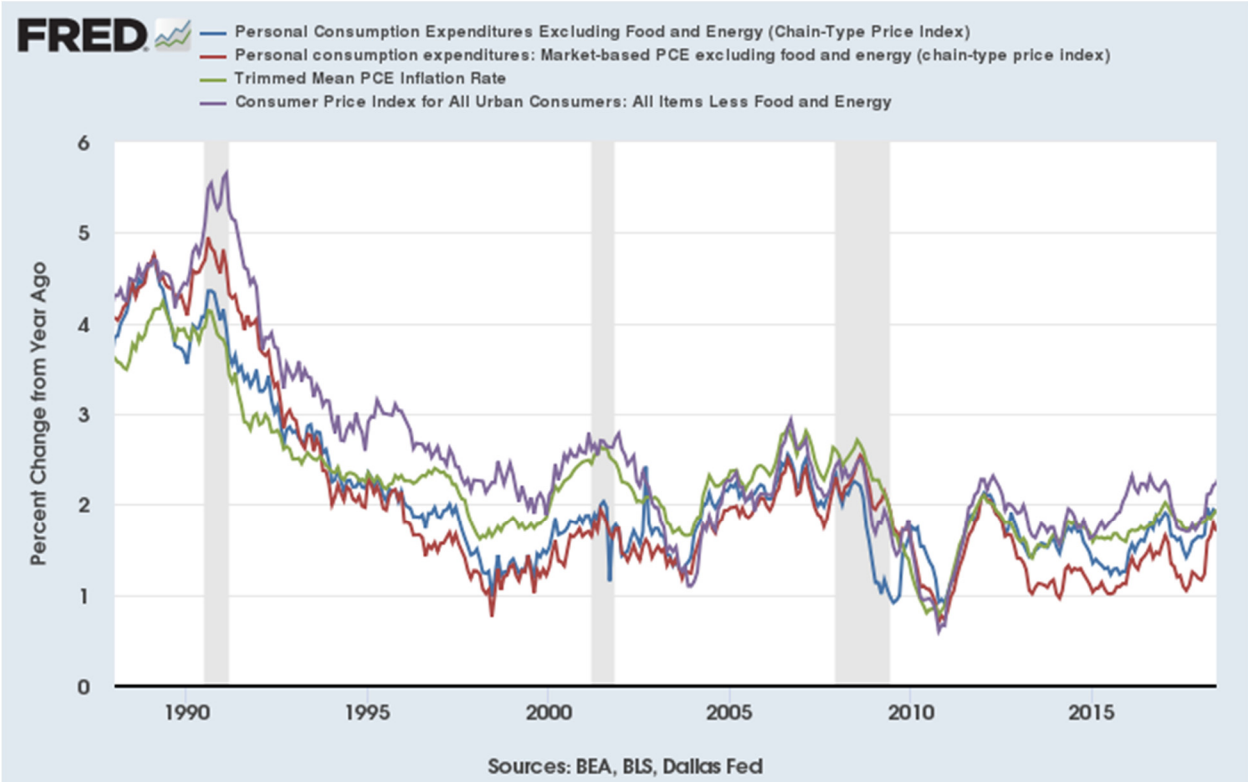
A3. Comparison to Stock-Watson Recession Gap

For the post-1985 period, Figure 5 plots the 12-month trim PCE inflation rate (leaded 12 months) along with the (monthly) recession gap term and our bust gap.

Figure 5: Stock-Watson Recession Gap and Scaled Bust Gap
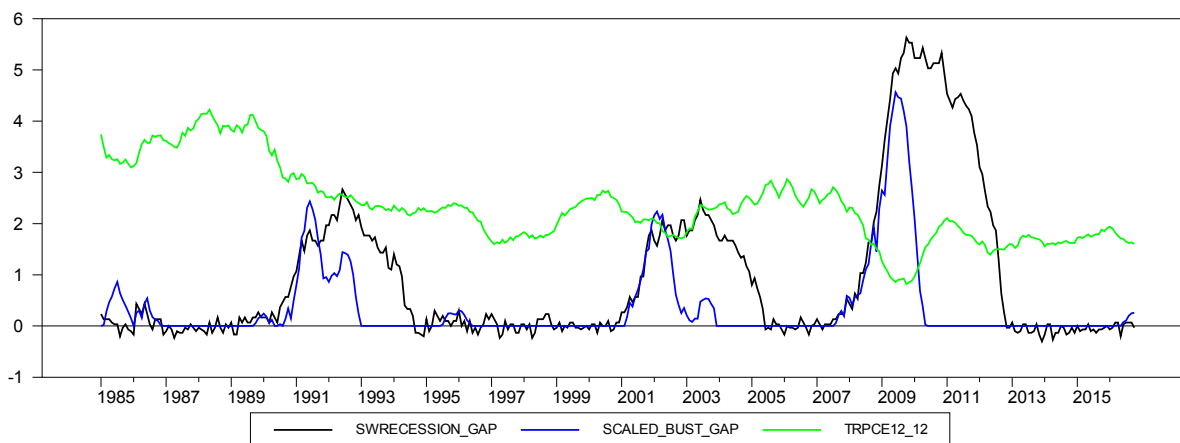
In this figure, the latter series has been scaled by multiplying it by 5 so as to render its peak magnitude comparable to that of the recession gap during the middle two recession episodes. Regarding ocular econometrics, the bust gap has an edge in timing, in that the peak inflation deceleration is relatively 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 (5) model versus the Stock and Watson recession gap model. While our model outperforms the Stock-Watson analogue over the entire period, this is only statistically-significant (at the 10% level) from September 2011 through June 2012. 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
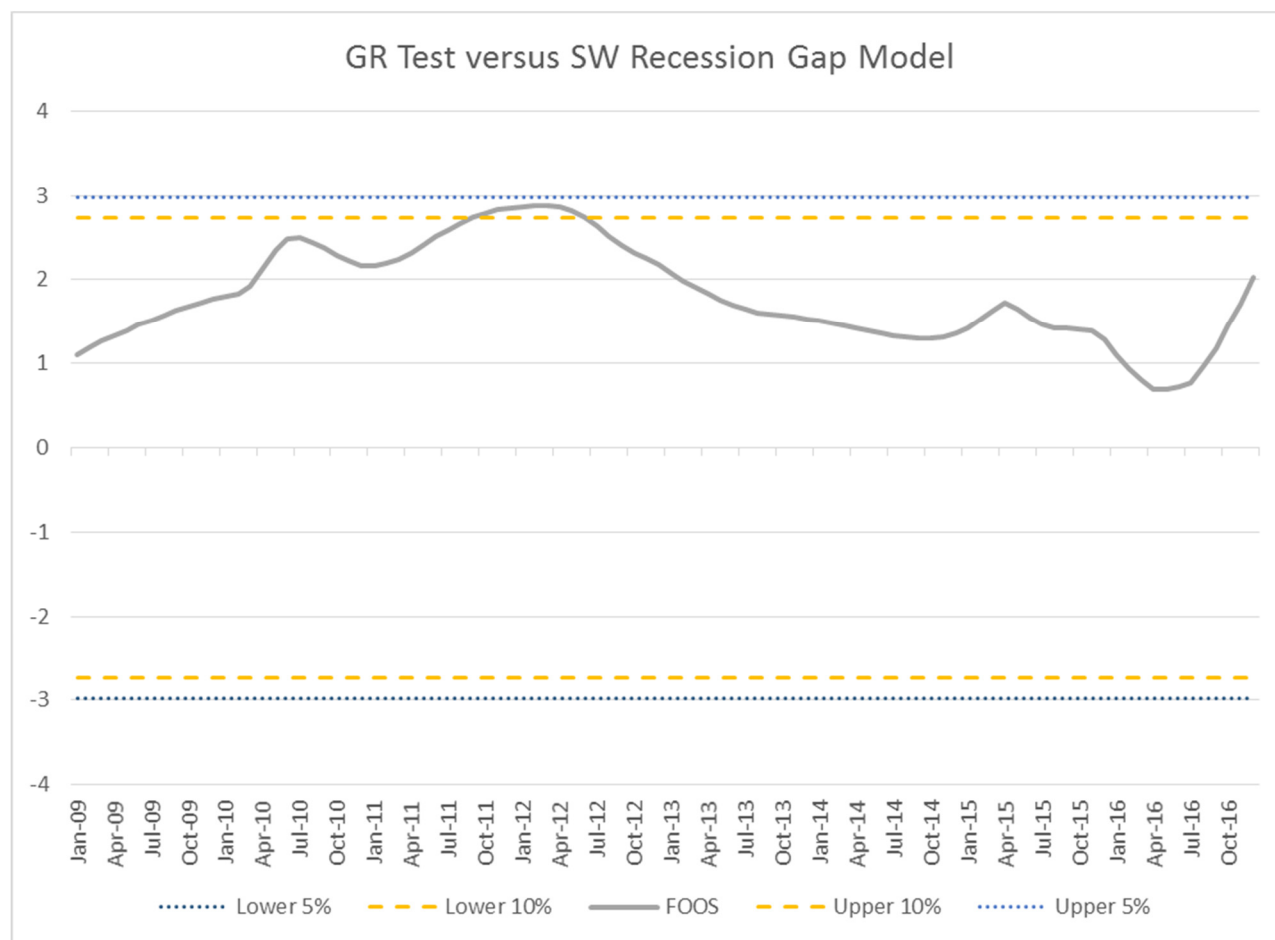
A4. Relation to Some Other Findings in the Literature

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. But as we have noted above, there is no reason that these concepts should

coincide. 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 (e.g., 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 2014 and Luengo-Prado, Rao and Sheremirov, 2017). This may explain the forecasting performance of the time-varying unobserved components model of Stock and Watson (2007).

Our findings also reconcile evidence in, e.g., 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 (2009) or Donayre and Panovska (2016),[30] 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 our bust gap, 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 (as noted above), yet is not cleanly captured by the standard sorts of nonlinearity that most models admit.

---

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

A.5 Relationship between moderately-persistent gap and transient gap