Federal Reserve Bank of Cleveland Working Paper Series

The Intermittent Phillips Curve: Finding a Stable (But Persistence-Dependent) Phillips Curve Model Specification

Richard Ashley and Randal J. Verbrugge

Working Paper No. 19-09R2

February 2023

Suggested citation: Ashley, Richard and Randal J. Verbrugge. 2023. "The Intermittent Phillips Curve: Finding a Stable (But Persistence-Dependent) Phillips Curve Model Specification." Working Paper No. 19-09R2. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-201909r2.

Federal Reserve Bank of Cleveland Working Paper Series
ISSN: 2573-7953

The Intermittent Phillips Curve: Finding a Stable (But Persistence-Dependent) Phillips Curve Model Specification

February 13, 2023

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

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

Abstract

We establish that the Phillips curve is persistence-dependent: inflation responds differently to persistent versus moderately persistent (or versus transient) fluctuations in the unemployment rate gap. This persistence-dependent relationship appears to align with business-cycle stages and is thus consistent with existing theory. Previous work fails to model this dependence, thereby finding numerous "inflation puzzles" – e.g., missing inflation/disinflation – noted in the literature. Our specification eliminates these puzzles; for example, the Phillips curve has not weakened, nor was inflation "stubbornly low" in 2019. The model's coefficients are stable, and it provides accurate conditional recursive forecasts through the Great Recession. There are important monetary policy implications.

Keywords: overheating; recession gap; persistence dependence; NAIRU; Phillips curve.

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

There are no potential sources of conflicts of interest.

Acknowledgments: For discussions occurring either as we developed the econometric tools or regarding related papers or this paper specifically, we thank Guido Ascari, Ruedi Bachmann, Christiane Baumeister, Marianne Baxter, Luca Benati, David Berger, Larry Christiano, Tim Cogley, Todd Clark, Oli Coibion, Jim Dolmas, Ayşe Kabukçuoğlu Dur, Rob Engle, Walter Enders, Jonas Fisher, Jordi Gali, Josh Gallin, Ana Galvao, Luca Gambetti, Pablo Guerron, Barry Jones, Ed Knotek, Markus Kontny, Andy Levin, Kurt Lunsford, Christian Mathes, Sandeep Mazumder, Loretta Mester, Athanasios Orphanides, Katia Peneva, Steve Reed, Ricardo Reis, Robert Rich, Laura Rosner, Juan Rubio-Ramirez, Jeremy Rudd, Ayşegūl Şahin, Raphael Schoenle, Jim Stock, Christopher Stivers, Murat Tasci, Can Tian, Mark Watson, John C. Williams, and seminar participants at the 2003 North American Summer Meetings of the Econometric Society, the 2004 and 2019 Midwest Macroeconomics Conferences, the 2013 EABCN *Inflation Developments after the Great Recession* conference, the 2020 Econometric Society World Congress, the Board of Governors, the Cleveland Fed, SUNY-Binghamton, and Virginia Tech. All views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Cleveland or the Federal Reserve System. There were no outside sources of funding for this research.

1. INTRODUCTION

The Phillips curve relationship remains central to macroeconomics and plays an absolutely fundamental role in monetary policy deliberations (see, for example, Brainard, 2019), not least because this relationship lies at the core of the structural models that dominate current monetary policy discussions.

By most accounts, however, inflation dynamics over the Great Recession seem to have diverged markedly from their previous patterns, posing a number of puzzles to the existing understanding of the Phillips curve relationship. The most prominent such puzzle is "missing disinflation" (e.g., Ball and Mazumder, 2011; Coibion and Gorodnichenko, 2015), prompting most analysts to conclude that the Phillips curve had weakened (e.g., Hall 2011; Bullard 2017). But over the 2016-2019 period, many wondered why inflation had not yet reached the inflation target (e.g., Heise et al., 2022). This "missing inflation" was believed to threaten the anchoring of long-run inflation expectations, and partly motivated the shift to average inflation targeting.

These inflation puzzles – as well as other findings and puzzles in the literature, such as the apparent time variation in the relationship,[1] and the odd-looking reverse-engineered NAIRU in Coibion and Gorodnichenko (2015) – completely disappear with the persistence-dependent specification of the Phillips curve relationship proposed here.[2] Based upon coefficients estimated using only data through 2006, recursive forecasts from our model specification (conditioned only on the time path of unemployment) well-predict inflation over the *entire* Great Recession and recovery: there is no downward-speed puzzle (Clark, 2014), no missing disinflation, and neither

[1] See, for example, Clark and McCracken (2006), Stock and Watson (2007), and Luengo-Prado et al. (2018).
[2] Moreover, we need not refer to biased inflation expectations (cf. Coibion and Gorodnichenko, 2015) or to the short-term unemployment rate (cf. Ball and Mazumder, 2019).

was inflation in early 2019 "stubbornly low" (FOMC minutes, July 30-31, 2019). Our model coefficients are stable: the Great Recession did not alter the dynamics of inflation; the Phillips curve did not flatten. Additionally, our findings have a natural interpretation in terms of stages of the business cycle, and are consistent with extant theory. And as we will explain shortly, our findings are highly relevant for monetary policy.

What is persistence-dependence (or, equivalently, frequency-dependence) in the relationship between, say, $Y(t)$ and an explanatory variable $X(t)$? It means that $Y(t)$ responds differently to persistent versus transitory fluctuations in $X(t)$. The notion of a persistence-dependent relationship of inflation to unemployment might sound exotic, but the idea actually dates back to the 1960s: Friedman (1968) and Phelps (1967) both noted that, while many studies had found an inverse "overall" relationship between inflation unemployment, (highly-persistent) natural rate fluctuations are unrelated to inflation. Indeed, macroeconomics is rife with persistence-dependent relationships. Perhaps the most famous is the "permanent income hypothesis," where consumption responds mainly to persistent movements in income. Persistence-dependence also motivates addressing ongoing (transient) measurement errors, which do not impact $Y(t)$ (see Hannan (1963) and Cochrane (2018)), and undertaking seasonal adjustment, which presumes that seasonal relationships are distinct from non-seasonal ones.

Furthermore, RBC modeling was built upon the presumption that business-cycle relationships are distinct from low-frequency relationships. This idea has recently regained traction. Angeletos et al. (2020) and Beaudry et al. (2020), building upon Comin and Gertler (2006), argue that examining the drivers and dynamics of macroeconomic variables by frequency allows one to better assess the causes of business cycles and discriminate across models; and Williams (2017) argues that the next generation of DSGE models must feature shocks with different frequency profiles. Furthermore, numerous recent studies, including Blundell et al.

(2013), Arellano et al. (2018), and Ashley et al. (2020), uncover evidence for persistence-dependence in important macroeconomic relationships (see Appendix D, which surveys numerous such studies).

Typical approaches to addressing persistence-dependence have involved two-sided filtering of many or all of the variables – dependent and explanatory – in the models, frequently using the Hodrick-Prescott filter. In Ashley and Verbrugge (2022a), we review the literature critiquing this practice and demonstrate the inferential distortions almost necessarily created by these approaches (see also Ashley and Verbrugge 2009, our summary in Appendix I, and related work by Doppelt 2021). We note that the persistence-dependent econometric methodology used here to re-specify the Phillips curve model is also highly applicable to estimation of the "next-generation" DSGE models noted above. That work is beyond the scope of the current paper but is presently under way.

Motivated by these research antecedents – as well as by the work of Stock and Watson (2010), hereafter SW, whose work suggests that the Phillips curve relationship is mainly a business-cycle relationship – in the present work we carefully and comprehensively explore the persistence-dependence in this relationship, using recently developed econometric tools that allow the data to speak very transparently as to the nature and form of this dependence, if it exists.[3] Our approach is entirely empirical; our results do not hinge on accepting strong – indeed, any – structural assumptions. We review relevant extant theory (in Section 3.2 and Appendix H) and we observe and describe the ways in which our empirical findings are consistent with it; but we leave the consequent further development of such theoretical modeling efforts for future work.

[3] King and Watson (1994) were perhaps the first to suggest that the Phillips curve varies with frequency; see also Lee (1995) and Pakko (2000). However, aside from SW, all previous work used two-sided filtering techniques (or cospectra) that yield unreliable inferences (see Ashley and Verbrugge 2009, 2022a,b and Doppelt 2021).

We find that the effect of the unemployment gap on inflation is both asymmetric and persistence-dependent; that is, it depends on both the sign and the persistence of the gap. Furthermore, the stable pattern of persistence-dependence we uncover has a natural (albeit somewhat informal) interpretation in terms of business-cycle stages. Moderately persistent movements in the gap apparently exert a very strong influence on inflation – but only when they are positive; this coincides with a recession and for a few months afterwards. But as the recovery continues, the Phillips curve vanishes. Very persistent movements in the unemployment gap exert a notable influence on inflation – but only when they are *negative*, i.e., when the economy is "overheating." Thus, the Phillips curve relationship is "intermittent."

Our findings are highly relevant for monetary policy. DSGE models that imply conventional (linear) Phillips curve specifications are likely to severely underestimate both the inflationary force when the economy is overheating, and also the deflationary force of a recession. Furthermore, conventional slack measures – even nonlinear functions of same – will poorly approximate the true Phillips curve relationship. Despite anchored inflation expectations, inflation moves sluggishly toward the policy target. If inflation is too low, this sluggish pace can be increased only by notable overheating. If inflation is too high, inflation won't recover to the target rapidly, absent a recession. But contra conventional estimates, the recession need not be excessively prolonged or severe, as Lawrence Summers (Wolf, 2022) and Ball et al. (2022) have argued.

As in any empirical project, overfitting is a danger; hence, we validate our in-sample findings with out-of-sample forecasts. We find that our specification performs quite well indeed. As noted above, a conditional recursive forecast accurately forecasts inflation *over the entire Great Recession and recovery*; unconditional out-of-sample forecasts also perform well. We take these out-of-sample results as very strong confirmation that our re-specification of the Phillips curve relationship is a notable and valuable improvement on the standard formulation.

2. A PERSISTENCE-DEPENDENT PHILLIPS CURVE

2.1 Frequency/Persistence Decomposition Method

The "persistence-dependent regression" methodology used below was developed in Ashley and Verbrugge (2009) and Ashley et al. (2020), and is briefly reviewed in Appendix I.

Herein we use this methodology to analyze a standard reduced-form Phillips curve specification relating the inflation rate to the unemployment rate gap, $u_t - u_t^*$.[4] We disaggregate the gap regression coefficient into three distinct persistence coefficients by partitioning the gap times series into three persistence components:[5]

$$u_t - u_t^* = gap_{hi-persist,t} + gap_{mod-persist,t} + gap_{transient,t}$$

where $gap_{hi-persist,t}$ subsumes the highly persistent fluctuations in $u_t - u_t^*$, $gap_{mod-persist,t}$ subsumes the moderately persistent fluctuations, and $gap_{transient,t}$ subsumes all of the remaining ("transient") fluctuations. Retaining any lag structure, these components are directly subsituted for the original $u_t - u_t^*$ explanatory variable in the original Phillips curve specification; we estimate a separate coefficient for each.

To ensure that OLS regression estimation and inference remain consistent and valid, this decomposition is obtained as a backward-looking (*one-sided*) partitioning. (Ashley and Verbrugge (2022a) show that any two-sided filtration would mix future and past values of $u_t - u_t^*$ together, inducing endogeneity in these persistence components.)

[4] We use the u_t^* measure from Tasci (2018). Tasci provided us with an update through 2019. This natural-rate measure is suitable for our analysis since (a) it is constructed without reference to inflation data (Occhino (2019) provides the rationale), and (b) it is found to be uncorrelated with inflation. We also use the CBO measure below, and in Appendix A. We use the real-time unemployment rate, obtained from ALFRED, Federal Reserve Bank of St. Louis.
[5] In our application, three components are both economically interpretable and statistically manageable.

First, we obtain the most persistent component of u_t by using a one-sided bandpass filter to extract all of the fluctuations in u_t with time variations that mean-revert on a timescale greater than $\tau_{persist}$ = 48 months. These highly persistent variations in u_t, $u_t^{persist}$ – which (as explained in Appendix I) include a nonlinear trend – are hived off first, since this improves the performance of the subsequent bandpass filtrations of $\left(u_t - u_t^{persist}\right)$, i.e., the remaining less-persistent components. We selected $\tau_{persist}$ = 48 months on macroeconomic grounds, thereby obtaining interesting and interpretable results in this setting.[6] The natural rate of unemployment u_t^* is subtracted (only) from this most persistent component, to form $gap_{hi-persist,t}$.

We next partition $\left(u_t - u_t^{persist}\right)$ into two components, by extracting the moderately persistent fluctuations that mean-revert on a time scale greater than $\tau_{transient}$ = 12 months; the resultant term we label $gap_{mod-persist,t}$. The residual – composed of the fluctuations that mean-revert on a time scale less than or equal to $\tau_{transient}$ – we label $gap_{transient,t}$. Again, our parametric choice of $\tau_{transient}$ was made on economic grounds: we take it as economically meaningful to consider unemployment rate fluctuations that mean-revert within a year as "transient" – as contrasted to our choice in taking fluctuations that only mean-revert on a timescale of four years or more as "highly persistent." Note that by construction, our three persistence "gap" components add up to the original unemployment rate gap time series, $u_t - u_t^*$.

Computationally, we use a standard two-sided Christiano-Fitzgerald (2003) bandpass

[6] The value of $\tau_{persist}$ is irrelevant absent any persistence-dependence in the relationship, but could be critical when the form of the persistence-dependence varies sharply. The robustness results in Appendix J indicate that the form of the persistence-dependence in the present setting is sufficiently smooth that our results are not highly sensitive to minor variations in this choice.

filter, within a sequence of moving windows.[7] In this way we obtain a completely backward-looking (i.e., one-sided) filtration of the data.

To conserve space, the remaining technical details of the decomposition of u_t into its persistence components are mostly deferred to Appendix I. We note here, however, that in windowed filtering applications, it is generally desirable to pad estimation windows with η periods of projections (or forecasts), to avoid well-known "end effects." Accordingly, we use η = 12 months of projections. As to the overall estimation window length, we use $\kappa = 48$ months below, but we note that our econometric results are not materially different if $\kappa = 60$ or 72 months were chosen instead.[8]

Following numerous suggestions in the literature to the effect that the impact of the unemployment gap on inflation is likely to be asymmetric – including Verbrugge (1997), Stock and Watson (2010), and Dupraz et al. (2019) – our preferred specification allows for sign-asymmetry in each component. Testing for such asymmetry and/or for persistence-dependence amounts to straightforward parameter restriction inferences.

As demonstrated in Appendix J, the Phillips curve results discussed below are not particularly sensitive to the details of how the persistence components are calculated, so long as the persistence components are obtained (via moving windows) using one-sided filtering. What ***does*** matter, in this and other macroeconomic settings we have considered, is whether or not one (appropriately) allows for persistence-dependence at all.

---

[7] An alternative, used in our earliest work, is the "AV" filter from Ashley and Verbrugge; those results are typically similar (see Appendix J and Ashley and Verbrugge (2022b)). Because it is so basic, the AV filter is easier to explain in depth and easier to program *de novo.* The CF filter, on the other hand, requires neither explanation nor programming; it is implemented in Stata and most other econometric packages.

[8] Thus in the main body of the paper, we filter a 48-month window that contains 36 months of data and 12 months of projections. Projections were obtained by interpolating quarterly unemployment rate projections from *Blue Chip Economic Indicators*, a resource of Wolters Kluwer Legal and Regulatory Solutions U.S., available monthly. In Appendix J, we demonstrate that our results are relatively insensitive to a wide range of choices for parameters such as κ and η, and even variations in $\tau_{persist}$.

2.2 Persistence Component Results

Figure 1 displays time plots of five time series. The top portion of this figure plots the real-time unemployment rate (u_t) and the natural rate (u_t^*). The bottom portion plots the three components of the unemployment gap.

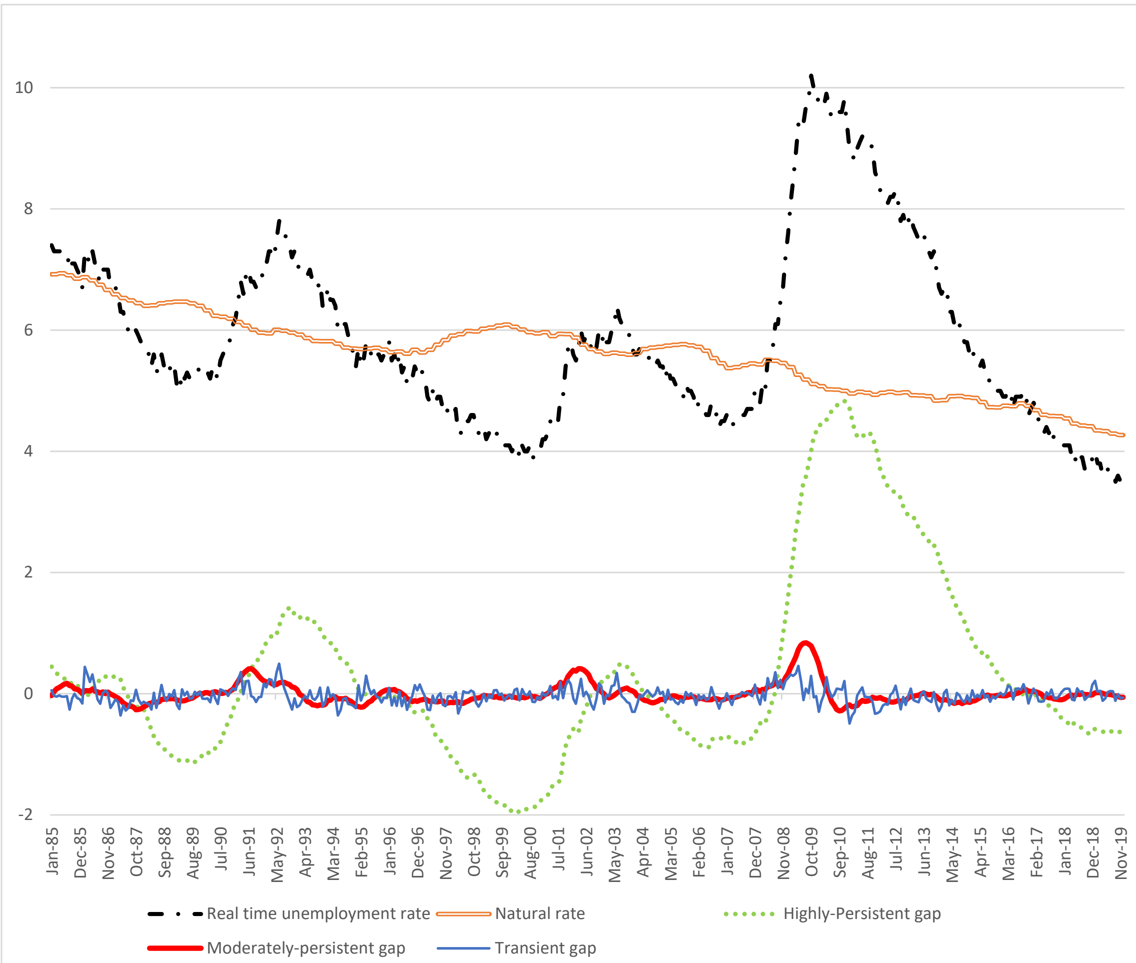

Figure 1: One-Sided Partition of the Unemployment Rate Gap

In our view, the three component time series have a relatively clear economic interpretation in terms of business-cycle stages. The highly persistent gap component is recognizably capturing the smooth movements in $u_t - u_t^*$. It traces out the three major recessions during the sample period.[9] Its negative dips comprise the late-cycle "overheating" periods of these business cycles, demonstrated below to be important for the overall Phillips curve relationship. The moderately persistent component is generally close to zero, but has a marked tendency to fluctuate upward at the onset of recessions, peak a few months after NBER-dated recessions end, then descend. These positive surges are also important for the Phillips curve relationship. (The transient component is mostly acyclical, but weakly correlated with the moderately persistent component.[10])

2.3 Persistence-Dependent Phillips Curve Regression Model

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

$$\left(\pi_{t+12}^{12}-\pi_{t}^{*}\right)=\alpha+\beta_{1}\left(\pi_{t}^{12}-\pi_{t-12}^{*}\right)+\beta_{2}\left(\pi_{t-12}^{12}-\pi_{t-24}^{*}\right)+\lambda_{1}\left(gap_{t}\right)+\varepsilon_{t} \qquad (1)$$

In this "Equation (1)" specification, gap_t is a traditional unemployment gap term, specified in terms of a natural rate: $gap_t \equiv \left(un_t - un_t^*\right)$;[11] $\pi_{t+12}^{12} \equiv \ln\left(P_{t+12}/P_t\right)$ denotes the 12-month log-change in a price index; and π_t^* is an inflation trend measure, discussed below. We focus on 12-month inflation both to remove noise, and because this is the chief focus for most policymakers; indeed, in the US, the inflation target is thus specified (although our results are similar if we use six-month inflation). Our inclusion of 24 months of lagged inflation (via our use of the current 12-month inflation rate, and of the 12-month inflation rate from a year ago) is in line with typical practice. Most of our analysis focuses on trimmed-mean PCE inflation, the realized-inflation

[9]Note that NBER-dated recessions over this period end 4-20 months prior to the date of peak unemployment. It is also worth noting that we are herein – for expository clarity – informally attaching interpretations to movements in the most persistent and moderately persistent components of the unemployment rate that are more intuitively accessible to the reader because we verbally identify them as linked to phases of the business cycle.

[10] The correlation is 0.34. Frequency components can have a modest non-zero correlation with each other if they are estimated using one-sided filters. Appendix E provides a useful plot.

[11] We use the unemployment rate because it is a very good indicator of business cycles, and unemployment rate gaps are far better measured than output gaps.

measure that arguably best removes noise from inflation (Mertens, 2016),[12] but in Appendix A we consider other time series of inflation measures in robustness checks.

The inflation trend estimate is the "PTR" measure from the FRB/US model of the Board of Governors.[13] Prior studies (e.g., Clark and McCracken (2006), Faust and Wright (2013), and Clark and Doh (2014)) have shown that inclusion of an accurate inflation trend estimate improves forecast accuracy. In the context of Phillips curve estimation, modeling inflation in terms of deviation from the trend π_t^* amounts to abstracting from the long-run goals of monetary policy and focusing on fluctuations in inflation that are more closely related to business cycles.[14] While Phillips curve forecasting models sometimes include other variables such as the relative price of energy or imports, over the post-1985 period these variables are not found to be helpful for our 12-month projections.

The model estimation sample is set here at 1985-2019. We start our analysis in 1985 because it is well known that inflation dynamics experienced a break sometime near 1985. Furthermore, by most accounts, the post-1985 period coincides with weakness or instability in the Phillips curve, a time when univariate models (famously, the Atkeson and Ohanian (2001) model) began to outperform Phillips curve models. We end the sample estimation period in 2019 because the data in 2020 pose serious challenges for estimation and inference, with extreme realizations in many macroeconomic series for reasons that are well known.[15] However, related work (Verbrugge and Zaman, 2023) extends our reduced-form specification and embeds it in a

[12] "…the trimmed-mean rate of PCE inflation stands out as a particularly strong signal of trend inflation" (Mertens, 2016, p. 966) and, when core PCE and trimmed-mean PCE diverge, the former moves to eliminate the gap. Core PCE is dominated along many dimensions (see Appendix B and Verbrugge 2022). As this paper was undergoing final internal review, we learned of Ball and Mazumder (2019), who also eschew core PCE and use median PCE. See also Verbrugge and Zaman (2022).

[13] The PTR series adjusts and extends median long-term forecasts from the Survey of Professional Forecasters. By modeling inflation in gap form against this series, we effectively impose anchored long-run inflation expectations. However, individual SPF respondents may not be anchored; see Binder et al. (2022).

[14] Putting this differently, the verticality (or not) of the Phillips curve is important for monetary policy but somewhat tangential to our study.

[15] This is true along many dimensions, such as the extreme plunge in payroll employment, the extreme spike in temporary layoffs (which experienced a 20-standard-deviation shock), the very large fiscal stimulus, the extreme supply-chain disruptions, etc. In time-series and DSGE modeling, it is now widely accepted that the pandemic period necessitates special treatment, but the form that this special treatment should take remains unsettled (see, e.g., Schorfheide and Song 2020, Lenza and Primiceri 2020, and Carriero et al. 2021).

nonlinear structural model. Among other things, this model enriches the Phillips curve equation via inclusion of a supply-side price pressures variable, necessary both for appropriately analyzing post-2019 inflation dynamics, and for assessing prospects going forward.

The "baseline" Equation (1) specification imposes some very strong testable restrictions on the Phillips curve relationship. It imposes the restriction that, aside from the distinction between natural rate fluctuations and other fluctuations, all fluctuations – whether persistent or transient – have the same relationship to inflation. However, previous work cited above has suggested that the relationship at business-cycle frequencies is notably stronger.[16] A second important restriction is that this baseline specification imposes linearity and symmetry on its gap term – that is, positive and negative gaps have the same influence on inflation. This linearity assumption departs from the original Phillips curve (Phillips, 1958) that posited a relationship that was steeper at higher levels of economic activity; further, numerous papers (see Appendix D) have located evidence for this type of nonlinearity in the Phillips curve. To preview our results, both restrictions are strongly rejected below.

The first restriction is relaxed in our second specification, Equation (2), where we partition the unemployment gap by persistence level:

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

where $\left(gap_{hi-persist,t}\right)$, $\left(gap_{mod-persist,t}\right)$ and $\left(gap_{transient,t}\right)$ are defined in Section 2.2.

We relax the second restriction by allowing for sign-asymmetry (about zero)[17] in each term:

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

[16] The present paper has incorporated persistence-dependent regression methods into modeling the Phillips curve relationship since its inception in the early 2000s; recent work is now coming around to this view – e.g., see Stock and Watson (2020), which focuses on the business-cycle-frequency relationship of the Phillips curve.
[17] We did no threshold search.

where $\left(gap^{+}_{hi-persist,t}\right)$ is the positive part of $\left(gap_{hi-persist,t}\right)$, $\left(gap^{-}_{hi-persist,t}\right)$ is the negative part of $\left(gap_{hi-persist,t}\right)$, and other terms are defined analogously.

Equation (3) is our preferred model, as (in our view) it has a fairly clear interpretation and is broadly consistent with economic theory. For forecasting purposes, we propose a specification that drops terms with statistically insignificant coefficient estimates:

$$\begin{aligned}\left(\pi^{12}_{t+12}-\pi^{*}_{t}\right) &= \alpha+\beta_1\left(\pi^{12}_{t}-\pi^{*}_{t-12}\right)+\beta_2\left(\pi^{12}_{t-12}-\pi^{*}_{t-24}\right)+\lambda^{-}_{1}\left(gap^{-}_{hi-persist,t}\right) \\ &+\lambda^{+}_{2}\left(gap^{+}_{mod-persist,t}\right)+\lambda^{+}_{3}\left(gap^{+}_{transient,t}\right)+\varepsilon_t\end{aligned} \tag{4}$$

Note that sign-asymmetry in the less-persistent components mirrors the asymmetry built into the SW recession gap term. As seen below, symmetry is strongly rejected, and Equation (3) is unambiguously preferred by all of our formal tests – even by the BIC, which strongly penalizes larger models. We also perform a Chow test of coefficient stability to test whether the coefficient estimates change after the Great Moderation ended, i.e., after 2006:12.

2.4 In-Sample Inference Results

2.4.1 Discussion of coefficient estimates

Table 1 below displays the OLS parameter estimates for the coefficients in Equations (1), (2), (3), and (4), with estimated t-ratios quoted beneath each coefficient estimate. It also reports measures-of-fit statistics such as the BIC, and p-values for various tests.

Consider first the estimation results for the "standard" Phillips curve specification, Equation (1). For this model specification, one might question the existence of the Phillips curve relationship. With an estimated t-ratio of only -1.54, the Phillips curve coefficient estimate λ is not statistically different from zero (on a two-tailed test) at even the 10 percent level of significance. Further, it is unstable: a standard Chow test rejects the null hypothesis that the parameters in Equation (1) are stable – when the sample data are partitioned into the period

1985:1 through 2006:12 versus 2007:1 through 2019:12 – with $p = 0.03$. This first set of results is not surprising: multiple studies in the literature have found ample evidence for a weak and unstable Phillips curve relationship over the past few decades.

Next consider the analogous estimation and inference results for Equation (2), which admits persistence-dependence. The data convincingly reject the persistence-independence restriction in Equation (1). A test of the equality of coefficients across persistence components is rejected with p-value < 0.005. Moreover, the BIC is notably smaller – so that the improvement in fit of Equation (2) over (1) more than compensates for its greater complexity. Our Equation (2) regression results by themselves would suggest that the Phillips curve is more or less confined to fluctuations of moderate persistence in the unemployment rate. This is in keeping with previous research findings, e.g., Stock and Watson (2020).

Finally, Equation (3) admits sign-asymmetry in each component. Symmetry is clearly rejected: the data clearly reject both $H_0 : \lambda_1^+ = \lambda_1^-$ and $H_0 : \lambda_2^+ = \lambda_2^-$ (though $H_0 : \lambda_3^+ = \lambda_3^-$ is only rejected at the 7 percent level), and the data reject the joint null hypothesis of coefficient symmetry on all three persistence components. The BIC also improves notably, from 0.67 to 0.48, so the increased complexity of this model is more than made up for by the improvement in the in-sample fit. (Out-of-sample evidence below compellingly reinforces this point.) Equation (3) gives rise to results that have a clear interpretation – one that is markedly different – and that align well with extant theory, as we discuss below. Notice that in moving from Equation (2) to (3), the highly persistent gap is shown to matter after all.

Dropping insignificant terms, Equation (4) leads to a modest improvement in the BIC, but does not change the interpretation. Hence we confine our discussion here to our baseline model, (3).

Table 1. Phillips Curve Regression Estimation Results

Regressor		Specif. (1)	Specif. (2)	Specif. (3)	Specif. (4)
$u_t - u_t^*$	λ	-0.08			
	(t-stat)	(-1.54)			
Persistent component	λ_1^+			0.03	
of $u_t - u_t^*$	(t-stat)		-0.07	(0.87)	
	λ_1^-		(-1.54)	-0.29	-0.27
	(t-stat)			(-3.57)	(-3.50)
Moderately persistent component	λ_2^+			-1.81	-1.72
of $u_t - u_t^*$	(t-stat)		-1.27	(-9.16)	(-9.68)
	λ_2^-		(-4.48)	0.23	
	(t-stat)			(0.31)	
Transient component	λ_3^+			-0.53	-0.51
of $u_t - u_t^*$	(t-stat)		-0.27	(-2.00)	(-2.10)
	λ_3^-		(-1.64)	0.07	
	(t-stat)			(0.68)	
Lagged inflation	β_1	0.48	0.31	0.26	0.24
	(t-stat)	(9.65)	(3.88)	(3.83)	(3.51)
	β_2	0.09	0.29	0.32	0.34
	(t-stat)	(1.11)	(4.08)	(6.86)	(9.28)
	constant	-0.08	0.33	-0.07	-0.06
	(t-stat)	(-0.88)	(1.32)	(-0.67)	(-0.73)
Adjusted R-squared		0.55	0.68	0.74	0.74
BIC		0.97	0.67	0.47	0.45
Hypothesis Test	H_0: $\lambda_1^+ = \lambda_1^-$			<0.005	
Rejection P-Values	H_0: $\lambda_2^+ = \lambda_2^-$			0.01	
	H_0: $\lambda_3^+ = \lambda_3^-$			0.07	
H_0: $\lambda_1^+ = \lambda_1^- = \lambda_2^+ = \lambda_2^- = \lambda_3^+ = \lambda_3^-$			<0.005	<0.005	<0.005
H_0: (Chow test): {λ coefficients unchanged before and after 2006:12}		0.03	0.50	0.36	0.74

NOTE: Figures in parentheses are estimated t-statistics, based on (13-month) HAC standard error estimates. Given this choice, diagnostic checks regarding heteroscedasticity are not quoted. Fitting errors for Equations (2), (3), and (4) display no evidence of notable outliers.

Notice in Equation (3) that three of the six gap component coefficients are negative and statistically significant; the other three are quantitatively negligible and not statistically significant. The λ_1^+ coefficient estimate, quantifying the impact of positive fluctuations in the highly persistent gap, is zero. Contra conventional wisdom, large persistent positive gaps in the unemployment rate evidently have no impact on inflation. The λ_1^- coefficient estimate, quantifying the impact of negative fluctuations in the highly persistent gap, is negative (at -0.29), with an estimated t-ratio of -3.57. The λ_2^+ coefficient estimate, quantifying the impact of positive fluctuations in the moderately persistent gap, is substantially negative (at -1.81), with an estimated t-ratio of -9.16. [18] Quantitatively, λ_1^- and λ_2^+ are the predominant Phillips curve relationships; see Appendix E for a graphical illustration.

2.4.2 Association with business-cycle stages

In our view, these Equation (3) results have a natural interpretation in terms of stages of the business cycle: the recession, the recovery period, and the overheating period. Most of the research investigating nonlinearity in the Phillips curve has focused on a differential force from positive and negative unemployment gaps, but this distinction fails to capture all of the subtlety we find in the relationship.

As discussed above in Section 2.2, the moderately persistent gap becomes non-negative during a recession (and for several more months). Its movement is associated with a very strong subsequent downward force on inflation (see Appendix E, and SW): the intermittent Phillips curve is powerfully evident. [19] This result is consistent with economic intuition, and accords well

---

[18]Qualitatively similar results in this regard also obtain using different measures of inflation, or using the CBO estimates of u_t^* instead of the Tasci (2018) estimates used here; see Appendix A. We also obtain qualitatively similar results if we: base our within-window forecasts on univariate models; use the final-vintage unemployment rate instead of the real-time rate; filter the unemployment gap itself, rather than the real-time unemployment rate; expand the range of the moderately persistent u_t - u_t^* component to include variation with reversion periods up to 60 or 72 months in length; impose symmetry in the transient component; or if we use the jobless unemployment rate (see Hall and Kudlyak 2020). While this paper makes no attempt to be a multi-country study, we also note that preliminary analyses using data from Australia yield a similar pattern. These results are available on request.

[19] Our estimate of λ_3^+ is also statistically significant, albeit only at the 5 percent level. The transient unemployment-gap component is noisier, but tends to comove with the moderately persistent component, and thus plays a reinforcing role in a recession; see Appendix E.

with both the SW "recession gap" findings (see Appendix C) and with evidence in Morris et al. (2019) on asymmetry in the wage-based Phillips curve (see Appendix G).

The highly persistent gap becomes positive shortly after a recession begins and remains positive during the recession and the recovery. But since $\hat{\lambda}_1^+$ is both quantitatively negligible and statistically indistinguishable from zero, this highly persistent part of the unemployment gap has no impact on inflation. Hence, the intermittent Phillips curve *vanishes* some months after the recession ends, irrespective of the size of the gap. Putting this more starkly, a persistently high unemployment rate *per se* does not reduce inflation.[20]

The Phillips curve remains dormant until the highly persistent gap becomes negative, i.e., once the economy "overheats." During this stage, consistent with much previous research, there is a notable upward influence on inflation: $\hat{\lambda}_1^- = -0.29$. The estimated size of this coefficient is directly comparable to, and much larger than, the estimate from a conventional Phillips curve specification (namely -0.08), given in the first column of Table 1. All of these important insights would have been missed had we not allowed for both persistence-dependence and asymmetry in our specification.

Are these findings exotic, or sensible? Ours is an empirical paper; we provide no new theory here. But our findings are consistent with much previous empirical work (as noted above), with basic business-cycle facts, and with an abundance of extant theory. Appendix H provides a broader discussion of this theory; here are a few highlights. First, a rapid decline in prices at the onset of a recession is consistent with standard industrial organization theory: during collapses, price wars can break out, as firms attempt to steal market share, but – shortly after the recovery

[20] Recall the findings of Luengo-Prado et al. (2018), who locate "robust evidence" for a structural break around 2009-2010, rendering the Phillips curve "negligible."

begins – prices start to edge up again; for macro applications, see, e.g., Gilchrist et al. (2017) and Hong (2019). Second, labor is the biggest marginal cost component. The marginal cost curve is quite steep, if labor is fixed in the short run. Overtime labor drops sharply at the onset of a recession, and marginal costs drop further due to underutilization. Basu and House (2016) provide a revealing discussion, summing it up by stating: "real marginal cost, properly computed, is strongly procyclical." Third, why does a persistently negative unemployment gap result in a stronger force than an equally sized persistently positive unemployment gap? This can result from downward price rigidity, from bargaining considerations (see, e.g., Moscarini and Postel-Vinay 2019) or from capacity constraints (see, e.g., Alvarez-Lois 2006 and Kuhn and George 2019).

2.4.3 A stable Phillips curve

In our persistence-dependent model, the coefficients in the Phillips curve relationship are stable over time. In particular, a Chow test of parameter stability for Equation (3) fails to reject the null hypothesis that the model coefficients change after 2006:12: the rejection p-value for this test is 0.36.[21] Evidently, the Great Recession did *not* significantly weaken or alter the Phillips curve relationship. Further evidence of coefficient stability is given by out-of-sample forecasts; see Section 3.

2.4.4 A comparison to other findings

In Appendix D, we discuss how our results compare to some other prominent findings in the literature. For example, our results reinterpret a finding in Coibion and Gorodnichenko (2015) regarding their reverse-engineered NAIRU. Similarly, the findings of Stock and Watson (2020),

[21] We also performed CUSUM and CUSUMSQ tests. These tests suggest parameter instability but are not indicative as to the source of said instability. Results from a Hansen (1992) stability test were inconclusive. We also split the sample into four sub-samples, based upon NBER recession peaks, and tested – and failed to reject – parameter stability in the two key Phillips curve coefficients.

who attempt to focus solely on business-cycle frequencies, mesh well with the persistence-dependent regression results here.[22] Finally, we argue that our results explain numerous studies that find a convex-concave aspect to the Phillips curve relationship and studies that adduce evidence for regime switching.

Our results naturally give rise to episodic forecast improvements and to time variation in Phillips curve coefficients for common specifications. An advantage of our approach is that we need not make specific (semi-structural) assumptions about models for trends, as these models could be misspecified (see Stock and Watson 2020). That said, our results are conditional on our persistence decomposition.

There is a large literature using semi-structural unobserved-component (UC) models to study the Phillips curve, trend inflation, and the natural rate of unemployment, and a more recent literature examines time-varying parameter models. We review some of this literature in Appendix D. Our contribution to this debate is that our particular respecification of the unemployment gap provides an intuitively appealing and relatively simple tweak to the model that resolves a number of empirical puzzles in the literature and appears to be relatively stable across the sample period, compared to the usual alternatives. This observation suggests that out-of-sample validation should be unusually critical in evaluating our claim that this new model specification is an improvement on those previously considered. Section 3 below takes up this topic.

[22] Stock and Watson (2020) mainly use a two-sided bandpass filter to remove all but business-cycle frequencies, but also, as a robustness check, use a "year-over-year" filter, which is one-sided. Like the Hamilton (2018) filter, this latter filter does not allow a decomposition of a variable into differing persistence levels, but being one-sided, it is at least not subject to the distortions inherent to two-sided filtering.

3. OUT-OF-SAMPLE SUPPORTING EVIDENCE

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

We find these results persuasive, but not necessarily definitive, in view of the usual concerns as to “data mining.” Out-of-sample (OOS) tests almost always fail to confirm any benefits of nonlinear models (Stock and Watson, 2009). We were also curious about whether our specification can explain inflation dynamics over the Great Recession and recovery. To address these issues, we conducted two additional OOS exercises. First, we present a recursive conditional forecast, to examine whether our model can resolve the various inflation puzzles apparently arising during the Great Recession; we also compare these to “conventional” (Equation (1)) conditional forecasts. Second, we examine unconditional forecasts. In particular, we present supporting results based on OOS forecasting calculations using the Giacomini-Rossi and Diebold-Mariano testing frameworks. Stock and Watson (2009) state, “…at least since 1985, Phillips curve forecasts do not outperform univariate benchmarks on average.” Forecasts based upon our specification outperform those based upon conventional Phillips curves *and* those based upon the celebrated Atkeson-Ohanian univariate benchmark. While it is not our goal to produce a forecasting model, our OOS tests taken together demonstrate that our in-sample results do not reflect overfitting, and underscore our conviction that we have uncovered something fundamental about the inflation process.

3.1 Conditional Recursive Forecasts

We use Equation (3) to generate recursive conditional forecasts out to 12 years, from 2007:1 through 2019:12. These forecasts are conditioned on the historical unemployment values during this forecast period, but they are recursive in that each forecast draws its needed lagged inflation-deviation values from its own recent inflation forecasts. These conditional forecasts are compared to analogous ones obtained from Equation (1); these forecasts condition on the path of the CBO unemployment gap, and impose symmetry and persistence-independence.[23]

In both cases, the model coefficients are estimated over 1985-2006 (the period of the Great Moderation), and then fixed – that is, not updated.

We also constructed a monthly Stock-Watson (2007) unobserved component-stochastic volatility (UCSV) model of 12-month trimmed-mean PCE inflation; the trend estimate from this model at time *t* (here, 2006:12) serves as its prediction for 12-month inflation going forward.[24] Its poor performance underscores the relevance of the unemployment gap for inflation dynamics.

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

---

[23] If we use our preferred U* measure in Equation (1), the forecast improvement from (3) is even more stark. Conditional recursive forecasts from Equation (4) improve upon those from Equation (3), but we omit these to avoid cluttering the figure.

[24] We thank Saeed Zaman for constructing this model and providing these estimates.

In sharp contrast, the similarly conditional forecasts generated by the linear model of Equation (1) are quite poor, and generate the well-known set of "puzzles." Inflation decelerated more rapidly than this model predicts.[25] Then, overall inflation recovered by about a full percentage point, even as this linear model was calling for further decline (through mid-2012). Hence, from August 2010 through December 2013, this model was underpredicting inflation by more than a full percentage point; this is the "missing disinflation" that has often been noted. The divergence between the actual dynamics of inflation and this model's predictions is striking. These out-of-sample forecasting results reinforce a central message of this paper: a failure to properly specify the relationship between the unemployment rate and the inflation rate, allowing for both asymmetry and persistence-dependence, yields unstable parameter estimates, strongly counterfactual conditional forecasts, and misleading conclusions about the nature of the inflation process.

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

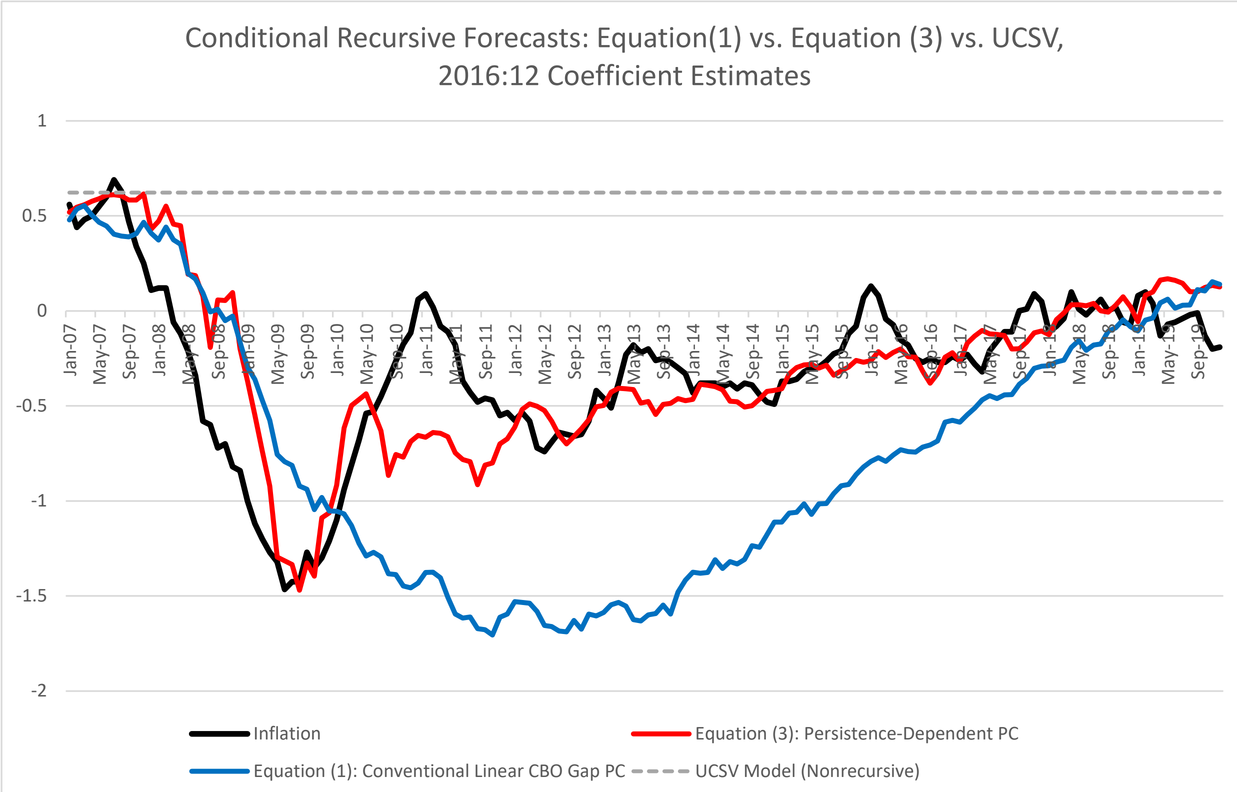

Figure 2: Conditional Recursive Forecasts from Equations (1) and (3), and from a UCSV model

By the end of 2019, the unemployment rate had arguably been below the natural rate since early 2017. Was there, in fact, missing inflation? From the perspective of our model's recursive forecast, the answer is "not really." The entire inflation trajectory over the recovery – the very sluggish movement toward long-run inflation expectations – was completely in line with inflation dynamics prior to 2006. At the end of 2019, perhaps one could say that there were 0.3 ppts of "missing inflation," but this deviation from our prediction could well have proven transitory, as previous deviations had been.

In Appendix F, we present an extension to Figure 2, which includes projections from both models using coefficients estimated over the full sample (but which are still recursive, based upon inflation data from 2006:12). This exercise underscores the stability of Equation (3) – its recursive forecasts are *extremely similar* – and emphasizes the lesser fidelity of the conventional model, Equation (1) – recursive forecasts using 2019:12 coefficients are *notably flatter*, reflecting the purported weakening in the Phillips curve. These flatter forecasts heighten the downward speed puzzle, and still under-predict the inflation recovery.

3.2 GR and DM Out-of-Sample Forecast Tests

We conjecture, along the lines of Stock and Watson (2009, 2010), that forecast improvements generated by our Phillips curve formulation over benchmark models are likely to be episodic for two reasons. First, the unemployment gap terms in our Equation (3) model are only substantially operative during two stages of the business cycle. Second, models with mis-specified unemployment gap terms will still yield unbiased forecasts, with notable deficiencies only when the gap is large. To examine this issue, we use the Giacomini and Rossi (2010) fluctuation test, in addition to the conventional Diebold-Mariano forecasting test.

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

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

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

where $\hat{\eta}^2$ and $\hat{\varepsilon}^2$ are the respective forecast errors, and $\hat{\sigma}$ is a HAC estimate of the asymptotic variance of the difference; here we set *m* equal to 60 months. The GR test is two-sided and is based on rolling-window estimates and forecasts. In Figure 3, we plot the upper and lower 10 percent and 5 percent critical values and the GR F^{OOS} test statistic for the two forecast comparisons described below. When the F^{OOS} statistic rises above the upper critical value, then the forecast performance of the "alternative" (persistence-dependent PC) model is significantly better (over the previous 60 months) compared to the baseline model. Conversely, when the F^{OOS} term falls below the lower critical value, the reverse is true.

A large literature compares inflation forecasts based on economic activity gaps, relative to univariate forecasts. A classic reference is Atkeson and Ohanian (2001), who famously found that a naïve random walk model generally outperformed the usual Phillips curve (PC) model. Almost all previous research investigating the OOS performance of PC-based forecasting models vis-à-vis similar univariate benchmark models over the post-1985 period returns negative results (e.g., Rossi and Sekhposyan 2010 and Dotsey et al. 2017). Below we examine the conjecture that Equation (3) outperforms conventional benchmark models.

In particular, we compare forecasts from our Equation (3) against Equation (1) – a standard Phillips curve, with the CBO gap – and against the (monthly) Atkeson-Ohanian (AO) model, i.e., 12-month inflation. Figure 3A depicts the comparison against Equation (1). Short estimation windows are not appropriate here, since our model sharply differentiates between different portions of the business cycle. We consequently use a 20-year window and estimate models from 1985:1 onward; thus, our first forecast is for 2005:1, that is, for the 12-month movement in the (detrended) trimmed-mean PCE between 2005:1 and 2006:1. The F^{OOS} statistic averages over the previous $m = 60$ months, so the GR test itself thus runs from 2010:1 onward.

In Figure 3A, the F^{OOS} line is almost always above zero, indicating that our Equation (3) specification outperforms the baseline CBO. The difference is statistically significant at the 5 percent level (over a five-year window) in early 2010 and late 2019. The Diebold-Mariano rejection p-value is less than 0.02, indicating that taking the sample period *as a whole*, the forecast improvement of Equation (3) over the baseline model is compelling.

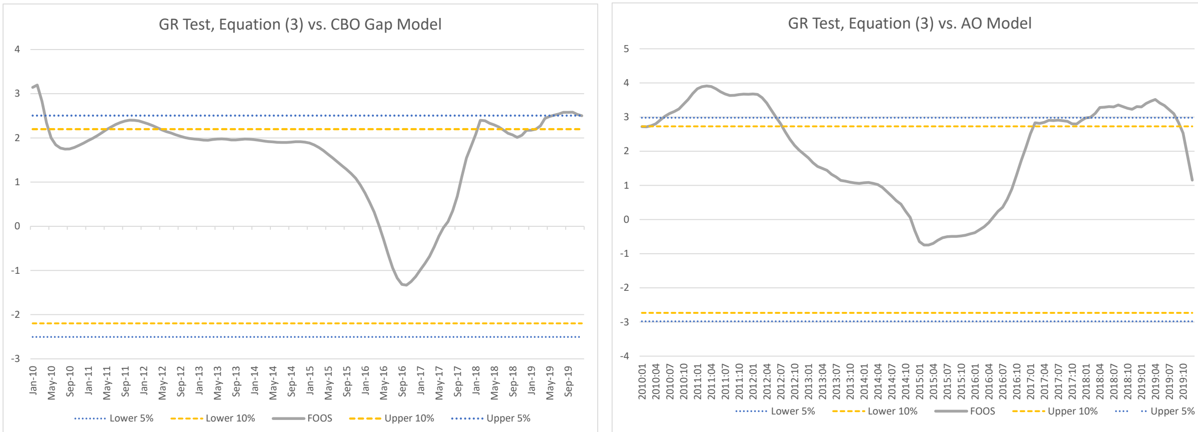

(A) Versus Equation (1) (B) Versus AO Model

Figure 3. GR Fluctuation Test, Equation (3) versus: (A), Baseline CBO Model, Equation (1); and (B), Atkeson-Ohanian Model

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

$$\left(\pi_{t+12}^{12}-\pi_{t}^{*}\right)=\left(\pi_{t}^{12}-\pi_{t-1}^{*}\right)+\eta_{t}$$

Figure 3B indicates quite compelling evidence for episodic forecast improvement of our Equation (3) model over the Atkeson-Ohanian model. The F^{OOS} line is almost always above zero, and the forecast gains are statistically significant at the 5 percent level from mid-2010 to mid-2012 and from early-2018 to mid-2019. For this OOS forecast comparison, the Diebold-Mariano rejection p-value is less than 0.005. Analogous comparisons against the UCSV model (another common benchmark in the forecasting literature) and against the SW recession-gap model are provided in Appendix C. Against those models as well, the Equation (3) model provides notably better forecasts.

The test results discussed above show that our Equation (3) re-formulation of the Phillips curve yields statistically significant improvement in out-of-sample forecasting. Our goal in this paper is not to devise an improved forecasting model, but rather to provide insight into inflation dynamics. These forecasting exercises are useful because they buttress our claim that our persistence-dependent model is an improvement over the usual PC specification. Further, they indicate that the statistical inference results quoted in Section 2 reflect a new set of stable statistical regularities in the historical data, rather than merely an improved in-sample fit of a more flexible model specification.

4. CONCLUSIONS

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

We find that our re-specified reduced-form Phillips curve relationship produces stable coefficient estimates across the 1985-2019 sample period, but that this is not a simple linear relationship. Rather, it is what we term "persistence-dependent," with the form of the relationship between inflation and unemployment fluctuations depending significantly – in both the statistical and the economic sense – on the persistence (and sign) of these unemployment fluctuations.

The empirically stable specification that we obtain better explains inflation variation in the observed macroeconomic historical record. We find that the Phillips curve is "intermittent" and, in our view, both naturally interpretable as one that varies across the stages of the business cycle, and one that is gracefully consistent with extant theory. When a recession begins – precisely when the unemployment rate is rising rapidly – the moderately-persistent and transient components of the unemployment rate become positive. Coefficient estimates imply that this induces a large reduction in inflation. After the unemployment rate peaks and begins to slowly descend, the aforementioned components effectively return to zero in the historical data (see Figure 1). The *highly persistent* component remains very large and positive during this descent, but coefficient estimates imply that as long as this highly persistent component remains positive, it imparts no downward force on inflation. Thus, the Phillips curve vanishes. Finally, when the recovery turns into the overheating stage, late in the expansion – that is, when the highly persistent component becomes negative – coefficient estimates imply that a Phillips curve relationship re-emerges, with notable upward force on inflation.

The in-sample fitting and out-of-sample forecasting results described in Sections 2 and 3 above show that our model specification well explains the time evolution of inflation during the sample period – even over the Great Recession and recovery. In particular, under our model specification all of the “inflation puzzles” noted above disappear. Moreover, notably, the relationship remained essentially unchanged over the Great Recession and recovery.

Estimated versions of the usual (standard) Phillips curve specification effectively average the three distinct relationships across differing portions of the business cycle: two strong relationships, and one non-relationship. Hence, they are bound to suggest that the Phillips curve is weak. Moreover, because the recovery period is associated with a negligible Phillips curve, the very prolonged recovery from the Great Recession caused conventional Phillips curve coefficient estimates to fall substantially. Conversely, our coefficient estimates are essentially the same when estimated using 2006:12 data or 2019:12 data.

The reduced-form Phillips curve specification developed here is validated by its stable coefficients across the sample and by its historical out-of-sample forecasting effectiveness. It is *not*, however, presented here primarily as a contribution to the literature on inflation forecasting, although we hope that the present work can and will stimulate progress by others in that direction. Nor, as a reduced-form model, does the model specification formulated here directly contribute to the theoretical literature on inflation, although (as detailed in Appendix H) it is consistent with existing theories, both with regard to the asymmetry in its unemployment responsiveness and with regard to the manner in which it varies across the business cycle. We do hope that theorists will see our empirical finding of persistence-dependence in this relationship as a stimulus to investigate why and how this type of dependence arises. However, we see the main contribution of our work as identifying and documenting an important new statistical regularity – a new "stylized fact," as it were – that any reasonably complete theoretical model for the US macroeconomy "ought to" imply.[26]

More broadly, we would like to emphasize the clear implications of the work presented here with regard to current and future monetary policy deliberations. As noted by John Cochrane (quoted in Steelman et al., 2013, p.36), "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." Even prior to the current challenges presented by the COVID collapse and recovery, the recent experience – of year after year of zero nominal interest rates, anchored inflation expectations, and low inflation – suggests difficulty in fine-tuning inflation. Even with anchored inflation expectations, the movement of inflation toward its long-run expected level is evidently quite slow. What can monetary policy do to speed

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

up this journey? If inflation is too low, it takes an appreciable amount of overheating before there is a significant upward force on inflation. If inflation is too high, it can be slowed rapidly – but only via a rapid upward movement in the unemployment rate, i.e., a recession.

The empirical re-formulation of the Phillips relationship developed here harmonizes much of the post-1985 experience in a relatively simple elaboration of the usual – but empirically unstable and unsuccessful – Phillips relationship. This re-formulation explains the observed puzzles associated with the usual models, and its empirical implementation is sufficiently stable as to provide reasonably accurate conditional forecasts of inflation over the Great Recession and accompanying recovery. These forecasts underscore the notion that inflation in 2018 and 2019 was not "stubbornly low" but was – in the re-formulation of the Phillips curve described here – in fact just where its pre-2006 dynamics suggest it should have been, given the evolution of the labor market. The related work of Verbrugge and Zaman (2023), which builds upon our model, underscores our conclusions and demonstrates the difficulties facing monetary policymakers at present.

Funding: We received no outside funding for this work.

REFERENCES

Álvarez-Lois, Pedro P. 2006. "Endogenous Capacity Utilization and Macroeconomic Persistence." *Journal of Monetary Economics* 53 (8): 2213–37. https://doi.org/10.1016/j.jmoneco.2006.07.001.

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2020. "Business-Cycle Anatomy." *American Economic Review* 110 (10): 3030–70. https://doi.org/10.1257/aer.20181174.

Arellano, Manuel, Richard Blundell, and Stephane Bonhomme. 2018. "Nonlinear Persistence and Partial Insurance: Income and Consumption Dynamics in the PSID." *AEA Papers and Proceedings* 108 (May): 281–86. https://doi.org/10.1257/pandp.20181049.

Aruoba, S. Borağan, Francis X. Diebold, and Chiara Scotti. 2009. "Real-Time Measurement of Business Conditions." *Journal of Business & Economic Statistics* 27 (4): 417–27. https://doi.org/10.1198/jbes.2009.07205.

Ashley, Richard, Kwok Ping Tsang, and Randal Verbrugge. 2020. "A New Look at Historical Monetary Policy (and the Great Inflation) through the Lens of a Persistence-Dependent Policy Rule." *Oxford Economic Papers* 72 (3): 672–91. https://doi.org/10.1093/oep/gpaa006.

Ashley, Richard, and Randal J. Verbrugge. 2008. "Frequency Dependence in Regression Model Coefficients: An Alternative Approach for Modeling Nonlinear Dynamic Relationships in Time Series." *Econometric Reviews* 28 (1–3): 4–20. https://doi.org/10.1080/07474930802387753.

———. 2022a. "Death to Regression Modelling on Bandpass-Filtered Data: New Results on the General Undesirability of Filtering Dynamic Regressions, with a Proposed Solution." Manuscript, Virginia Tech.

———. 2022b. "Time Series Filtering for Studying Frequency Dependence: What Works, What Doesn't." Manuscript, Federal Reserve Bank of Cleveland.

Atkeson, Andrew, and Lee E. Ohanian. 2001. "Are Phillips Curves Useful for Forecasting Inflation?" *Quarterly Review* 25 (1): 2–11. https://doi.org/10.21034/qr.2511.

Ball, Laurence M., Daniel Leigh, and Prachi Mishra. 2022. "Understanding US Inflation during the COVID Era." In *BPEA Conference Drafts, September 8-9, 2022*. https://www.brookings.edu/bpea-articles/understanding-u-s-inflation-during-the-covid-era/.

Ball, Laurence M., and Sandeep Mazumder. 2011. "Inflation Dynamics and the Great Recession." *Brookings Papers on Economic Activity* 42 (1): 337–81. https://doi.org/10.1353/eca.2011.0005.

Basu, Susanto, and Christopher L. House. 2016. "Allocative and Remitted Wages: New Facts and Challenges for Keynesian Models." In *Handbook of Macroeconomics*, 2:297–354. Elsevier. https://doi.org/10.1016/bs.hesmac.2016.05.001.

Beaudry, Paul, Dana Galizia, and Franck Portier. 2020. "Putting the Cycle Back into Business Cycle Analysis." *American Economic Review* 110 (1): 1–47. https://doi.org/10.1257/aer.20190789.

Binder, Carola Conces, Wesley Janson, and Randal Verbrugge. forthcoming. "Out of Bounds: Do SPF Respondents Have Anchored Inflation Expectations?" *Journal of Money, Credit and Banking*. https://doi.org/10.1111/jmcb.12968.

Blundell, Richard, Hamish Low, and Ian Preston. 2013. "Decomposing Changes in Income Risk Using Consumption Data." *Quantitative Economics* 4 (1): 1–37. https://doi.org/10.3982/QE44.

Brainard, Lael. 2019. "The Disconnect Between Inflation and Employment in the New Normal." In *"Certain Uncertainty: Tax Policy in Unsettled Times" National Tax Association 49th Annual Spring Symposium*. Washington, DC. https://www.federalreserve.gov/newsevents/speech/brainard20190516a.htm.

Bullard, James B. 2017. "Does Low Unemployment Signal a Meaningful Rise in Inflation?" *The Regional Economist* 25 (3). https://ideas.repec.org//a/fip/fedlre/00152.html.

Carriero, Andrea, Todd E. Clark, Massimiliano Marcellino, and Elmar Mertens. 2021. "Addressing COVID-19 Outliers in BVARs with Stochastic Volatility." Working paper 21-02R. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202102r.

Christiano, Lawrence J., and Terry J. Fitzgerald. 2003. "The Band Pass Filter." *International Economic Review* 44 (2): 435–65. https://doi.org/10.1111/1468-2354.t01-1-00076.

Clark, Todd E. 2014. "The Importance of Trend Inflation in the Search for Missing Disinflation." *Economic Commentary*, no. 2014-16 (August). https://doi.org/10.26509/frbc-ec-201416.

Clark, Todd E., and Taeyoung Doh. 2014. "Evaluating Alternative Models of Trend Inflation." *International Journal of Forecasting* 30 (3): 426–48. https://doi.org/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): 1127–48. https://doi.org/10.1353/mcb.2006.0068.

Cochrane, John H. 2012. "A Brief Parable of Over-Differencing." Manuscript, University of Chicago, Booth School of Business. https://www.johnhcochrane.com/s/overdifferencing-889r.pdf.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015. "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation." *American Economic Journal: Macroeconomics* 7 (1): 197–232. https://doi.org/10.1257/mac.20130306.

Comin, Diego, and Mark Gertler. 2006. “Medium-Term Business Cycles.” *American Economic Review* 96 (3): 523–51. https://doi.org/10.1257/aer.96.3.523.

Doppelt, Ross. 2021. “Should Macroeconomists Use Seasonally Adjusted Time Series? Structural Identification and Bayesian Estimation in Seasonal Vector Autoregressions.” Manuscript, Michigan State University. https://www2.gwu.edu/~forcpgm/Doppelt_Seasonality_2021.pdf.

Dotsey, Michael, Shigeru Fujita, and Tom Stark. 2018. “Do Phillips Curves Conditionally Help to Forecast Inflation?” *International Journal of Central Banking* 14 (4): 43–92. https://www.ijcb.org/journal/ijcb18q3a2.htm.

Dupraz, Stéphane, Emi Nakamura, and Jón Steinsson. 2019. “A Plucking Model of Business Cycles.” Working paper 26351. National Bureau of Economic Research. https://doi.org/10.3386/w26351.

Faust, Jon, and Jonathan H. Wright. 2013. “Chapter 1 - Forecasting Inflation.” In *Handbook of Economic Forecasting*, 2, Part A:2–56. Elsevier. https://doi.org/10.1016/B978-0-444-53683-9.00001-3.

Friedman, Milton. 1968. “The Role of Monetary Policy.” *American Economic Review* 58 (1): 1–17. https://www.jstor.org/stable/1831652.

Giacomini, Raffaella, and Barbara Rossi. 2010. “Forecast Comparisons in Unstable Environments.” *Journal of Applied Econometrics* 25 (4): 595–620. https://doi.org/10.1002/jae.1177.

Gilchrist, Simon, Raphael S. Schoenle, Jae Sim, and Egon Zakrajšek. 2017. “Inflation Dynamics during the Financial Crisis.” *American Economic Review* 107 (3): 785–823. https://doi.org/10.1257/aer.20150248.

Hall, Robert E. 2011. “The Long Slump.” *American Economic Review* 101 (2): 431–69. https://doi.org/10.1257/aer.101.2.431.

Hall, Robert E., and Marianna Kudlyak. 2022. “The Unemployed with Jobs and without Jobs.” *Labour Economics* 79 (December): 102244. https://doi.org/10.1016/j.labeco.2022.102244.

Hamilton, James D. 2018. “Why You Should Never Use the Hodrick-Prescott Filter.” *The Review of Economics and Statistics* 100 (5): 831–43. https://doi.org/10.1162/rest_a_00706.

Hansen, Bruce E. 1992. "Testing for Parameter Instability in Linear Models." *Journal of Policy Modeling* 14 (4): 517–33. https://doi.org/10.1016/0161-8938(92)90019-9.

Heise, Sebastian, Fatih Karahan, and Ayşegül Şahin. 2022. "The Missing Inflation Puzzle: The Role of the Wage-Price Pass-Through." *Journal of Money, Credit and Banking* 54 (S1): 7–51. https://doi.org/10.1111/jmcb.12896.

Hong, Sungki. 2019. "Customer Capital, Markup Cyclicality, and Amplification." In *2019 Meeting Papers*. Society for Economic Dynamics. https://ideas.repec.org/p/red/sed019/959.html.

King, Robert G., and Mark W. Watson. 1994. "The Post-War US Phillips Curve: A Revisionist Econometric History." *Carnegie-Rochester Conference Series on Public Policy* 41 (December): 157–219. https://doi.org/10.1016/0167-2231(94)00018-2.

Kuhn, Florian, and Chacko George. 2019. "Business Cycle Implications of Capacity Constraints under Demand Shocks." *Review of Economic Dynamics* 32 (April): 94–121. https://doi.org/10.1016/j.red.2019.01.001.

Lee, Jim. 1995. "The Phillips Curve Behavior over Different Horizons." *Journal of Economics and Finance* 19 (3): 51–69. https://doi.org/10.1007/BF02920614.

Lenza, Michele, and Giorgio Primiceri. 2020. "How to Estimate a VAR after March 2020." Working paper 27771. National Bureau of Economic Research. https://doi.org/10.3386/w27771.

Luengo-Prado, María José, Nikhil Rao, and Viacheslav Sheremirov. 2018. "Sectoral Inflation and the Phillips Curve: What Has Changed since the Great Recession?" *Economics Letters* 172 (November): 63–68. https://doi.org/10.1016/j.econlet.2018.08.016.

Mertens, Elmar. 2016. "Measuring the Level and Uncertainty of Trend Inflation." *Review of Economics and Statistics* 98 (5): 950–67. https://doi.org/10.1162/REST_a_00549.

Moscarini, Giuseppe, and Fabien Postel-Vinay. 2019. "The Job Ladder: Inflation vs. Reallocation." Manuscript. Yale University. https://campuspress.yale.edu/moscarini/files/2019/02/nominal-rigidities_SA-29sq2lo.pdf.

Occhino, Filippo. 2019. "The Flattening of the Phillips Curve: Policy Implications Depend on the Cause." *Economic Commentary*, no. 2019-11 (July). https://doi.org/10.26509/frbc-ec-201911.

Pakko, Michael R. 2000. “The Cyclical Relationship between Output and Prices: An Analysis in the Frequency Domain.” *Journal of Money, Credit and Banking* 32 (3, pt.1): 382–99. https://doi.org/10.2307/2601171.

Phelps, Edmund S. 1967. “Phillips Curves, Expectations of Inflation and Optimal Unemployment over Time.” *Economica* 34 (135): 254–81. https://doi.org/10.2307/2552025.

Phillips, A. W. 1958. “The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957.” *Economica* 25 (100): 283–99. https://doi.org/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): 808–35. https://doi.org/10.1016/j.ijforecast.2009.08.004.

Schorfheide, Frank, and Dongho Song. 2020. “Real-Time Forecasting with a (Standard) Mixed-Frequency VAR during a Pandemic.” Working paper 20-26. Federal Reserve Bank of Philadelphia. https://doi.org/10.21799/frbp.wp.2020.26.

Steelman, A., Haltom, R., and Kenney, L. (2013), “Interview, John Cochrane.” Federal Reserve Bank of Richmond *Econ Focus*, no. Q3: 34-38. https://www.richmondfed.org/publications/research/econ_focus/2013/q3/full_interview.

Stock, James H., and Mark W. Watson. 2007. “Why Has U.S. Inflation Become Harder to Forecast?” *Journal of Money, Credit and Banking* 39 (s1): 3–33. https://doi.org/10.1111/j.1538-4616.2007.00014.x.

———. 2009. “Chapter 3 - Phillips Curve Inflation Forecasts.” In *Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective*, edited by Jeff Fuhrer, Yolanda K. Kodrzycki, Jane Sneddon Little, and Giovanni P. Olivei, 99–202. Cambridge, MA: MIT Press. https://doi.org/10.7551/mitpress/8380.003.0005.

———. 2010. “Modeling Inflation After the Crisis.” Working paper 16488. National Bureau of Economic Research. https://doi.org/10.3386/w16488.

———. 2020. “Slack and Cyclically Sensitive Inflation.” *Journal of Money, Credit and Banking* 52 (S2): 393–428. https://doi.org/10.1111/jmcb.12757.

Tasci, Murat. 2018. "The Ins and Outs of Unemployment in the Long Run: Unemployment Flows and the Natural Rate." Working paper 12-24 (updated). Federal Reserve Bank of Cleveland.

Verbrugge, Randal J. 1997. "Investigating Cyclical Asymmetries." *Studies in Nonlinear Dynamics & Econometrics* 2 (1). https://doi.org/10.2202/1558-3708.1025.

———. (2022) "Is it Time to Reassess the Focal Role of Core PCE Inflation in Assessing the Trend in PCE Inflation?" *Economia*, 45(89), 73-101. https://doi.org/10.18800/economia.202201.004

Verbrugge, Randal J., and Saeed Zaman. 2022. "Improving Inflation Forecasts Using Robust Measures." Working paper 22-23. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202223.

———. 2023. "The Hard Road to a Soft Landing: Evidence from a Nonlinear Structural Model." Federal Reserve Bank of Cleveland, Working Paper No. 23-03. https://doi.org/10.26509/frbc-wp-202303.

Williams, John C. 2017. "DSGE Models: A Cup Half Full." In *DSGE Models in the Conduct of Policy: Use as Intended*, edited by Refet S. Gürkaynak and Cédric Tille, 16–22. VoxEU.org, CEPR. http://www.thethao247web.net/content/dsge-models-conduct-policy-use-intended.

Wolf, Martin. 2022. "Larry Summers: 'The Destabilisation Wrought by British Errors Will Not Be Confined to Britain'." *Financial Times*, October 6, 2022, sec. Global Economy. https://www.ft.com/content/20117143-2084-4ac1-98a5-5c48fae7fc23.