Inflation Expectations and Nonlinearities in the Phillips Curve*

Alexander Doser† Ricardo Nunes‡ Nikhil Rao§ Viacheslav Sheremirov¶

August 2022

Abstract

This paper examines the presence of nonlinearities in the Phillips curve. We allow for a flexible form of nonlinearity and estimate a threshold regression model with the number and location of thresholds determined directly from the data. Because the literature emphasizes inflation expectations in the linear case, we study in detail how different measures of inflation expectations affect the role of nonlinearities. Over the estimation period starting in the late 1960s, we document a weak case for nonlinearity. Depending on the inflation measure and identification strategy, the preferred model is either linear or mildly nonlinear. In the nonlinear as well as linear specifications, consumer inflation expectations play an important role, while professional forecasters' expectations appear less important. Moreover, not controlling for consumer expectations may lead the econometrician to overestimate the degree of nonlinearity and the significance of professional forecasters' expectations. While on the whole the case for nonlinearities is not strong, we identify some specifications and historical episodes during which nonlinearities play a more important role, such as the missing disinflation.

Keywords: Consumer expectations; Missing disinflation; Nonlinearities; Volcker disinflation
JEL classification: D84, E24, E31, E32

*We would like to thank Marco Del Negro, Martin Eichenbaum, Jeff Fuhrer, Yuriy Gorodnichenko, Giovanni Olivei, Søren Hove Ravn, Andrea Tambalotti, Jenny Tang, Jon Willis, Randy Wright, and anonymous referees as well as seminar participants at the Federal Reserve Bank of Kansas City, Federal Reserve Bank of Boston, University of Surrey, Midwest Macro at the University of Wisconsin–Madison, Midwest Macro at the Michigan State University, the EEA-ESEM 2018 meeting in Cologne, the Money Macro and Finance 2019 at the LSE, and the Royal Economic Society 2021 for comments. The views expressed herein are those of the authors and are not necessarily those of the Federal Reserve Bank of Boston or the Federal Reserve System.

†Northwestern University: AlexanderDoser2023@u.northwestern.edu.
‡University of Surrey: Ricardo.Nunes@surrey.ac.uk.
§University of Michigan: NikhRao@umich.edu.
¶Federal Reserve Bank of Boston: Viacheslav.Sheremirov@bos.frb.org.

1. Introduction

Inflation dynamics have long been a subject of prolific economic research. In the 1970s and '80s, much research in this area was dedicated to understanding the causes and costs of high inflation and how to disinflate effectively. Afterward, the focus shifted to understanding the determinants of inflation and the role of expectations in the context of low or moderate inflation. Despite important new developments in the study of the Phillips curve, its exact functional form depends on specific modeling assumptions and remains heavily debated empirically (e.g., Clark and McCracken, 2006; Nason and Smith, 2008; Mavroeidis, Plagborg-Møller, and Stock, 2014; Ball and Mazumder, 2011, 2019, 2020; Stock and Watson, 2020).

In this paper, we explore whether nonlinearity can reconcile the Phillips curve with evidence from several prominent episodes of extreme fluctuations in inflation and/or economic activity: the Great Inflation of the 1970s, the Volcker disinflation, the missing disinflation following the Great Recession, and the missing inflation of the late 2010s. It turns out that the answer to this question depends crucially on the measurement of inflation expectations. In our preferred model, which, following Coibion and Gorodnichenko (2015), emphasizes the survey measures of households' expectations, nonlinearity plays a limited role. It helps to improve the fit of the model during the missing disinflation but not during the other episodes. In the full sample, the linear model cannot be rejected in favor of a nonlinear model. But if the expectations process is based on professional forecasters only, the estimated nonlinearity appears more important, because the linear model with mismeasured expectations has less bearing in the data. In contrast, we document that households' expectations play an important role in inflation dynamics during all episodes, even when the model allows for nonlinearity.

The novelty of our paper is threefold. First, we estimate a Phillips curve without assuming a particular form of nonlinearity, by approximating an arbitrary function with a piecewise-linear function for which the number of segments is estimated directly from the data. The previous empirical literature on nonlinear Phillips curves typically imposes a particular functional form and therefore provides limited evidence on nonlinearity in general. Second, by revisiting nonlinearity of the Phillips curve with the 2010s data, we significantly extend the high-unemployment sample and shed more light on inflation dynamics during such episodes. The post–World War II data prior to the Great Recession, on which much of the existing results are based, contain limited observations of high unemployment.[1] Moreover, understanding inflation dynamics in the high-unemployment environment is especially relevant for pandemic

[1] If the Phillips curve becomes flatter during recessions, the high levels of unemployment would not lead to a sharp decrease in inflation. In fact, the empirical relationship between wage inflation and unemployment originally documented by Phillips (1958) is represented by a convex curve, not a linear relationship. Hooper, Mishkin, and Sufi (2020) also show interesting evidence along these lines.

recessions, which are characterized by particularly high levels of unemployment. Third, our paper considers possibly spurious nonlinearity of the Phillips curve that may result from measurement error in inflation expectations. Such measurement issues have been brought to the forefront by the recent literature but had not been considered in the context of nonlinear models prior to this paper. Hence, we provide a framework to merge two strands of the literature on the inability of Phillips curves to explain inflation dynamics during high-volatility episodes.

Studying whether the empirical Phillips curve is nonlinear can also shed light on the relevant mechanism in workhorse aggregate models, especially those that emphasize state dependence. For instance, downward nominal wage rigidity could break the link between economic activity and prices during deep recessions, making the Phillips curve flatter. If, instead, prices change more frequently during a deep recession (e.g., due to state dependence or rational inattention, whereby firms easily observe a deep recession but not a mild one), the Phillips curve should become steeper.[2] Since different mechanisms can push the slope of the Phillips curve in opposite directions—and can also create several regions with different slopes—this paper takes an empirical approach that has the advantage of detecting any type of nonlinearity stemming from any structural parameter.

We estimate the Phillips curve using single-equation methods. To properly identify this structural relationship and examine the robustness of the results, we consider several different identification strategies. First, we control for an extensive list of supply-side shifters considered in the literature. Second, similar to Barnichon and Mesters (2020), we isolate exogenous variation in demand, using monetary policy shocks identified in the literature (Romer and Romer, 2004) to instrument for the unemployment gap. Finally, we exploit regional variation in inflation and unemployment, thereby removing the effects of aggregate shocks (Hazell et al., 2022). Because controlling for supply-side shifters allows us to explore the longest time period, we use this method as a baseline.[3]

To estimate nonlinearity without imposing any particular functional form, we combine these identification strategies with threshold regression methods (Hansen, 2000, 2017). The threshold regression allows us to determine the number and location of thresholds directly from the data, and therefore can be used to approximate any arbitrary form of nonlinearity. A one-threshold model (i.e., two regions with different slopes) could detect changes in the slope

[2]We discuss these effects in more detail in Online Appendix A. In the workhorse New Keynesian models, the slope of the Phillips curve depends on the Calvo probability of resetting prices, the intertemporal elasticity of substitution, the Frisch elasticity of labor supply, and the elasticity of substitution across product varieties. In addition, nonlinearities in the microfoundations (e.g., a kink in individual demand curves) may result in nonlinear aggregate relationships.

[3]The Romer and Romer shocks are obtained with a five-year lag, while regional data are not available during the early period. Hence, we would not be able to study jointly the Great Inflation of the 1970s, the Volcker disinflation of the '80s, and the recent missing-inflation episode with either method.

due to downward nominal wage rigidity and kinked demand, among other reasons. A two-threshold model could provide an even better fit to a nonlinear relationship, and could also detect an inaction (flat) region surrounded by two steep regions, as one would expect from a state-dependent pricing perspective.

In line with the previous literature, we emphasize survey measures of inflation expectations. While the importance of such surveys has been documented for linear models, we extend this result to the nonlinear case. In our empirical specification, we include survey measures of inflation expectations by consumers and professional forecasters as well as backward-looking adaptive expectations. We note that in the New Keynesian model, the expectations relevant for the Phillips curve are those by the price-setters (i.e., firms). However, recent surveys of U.S. firms do not extend back far enough to cover our sample period. Conventionally, consumer expectations are thought to proxy for expectations by small firms, such as mom-and-pop stores, whereas professional forecasters expectations' may reflect the expectations of large firms. Still, backward-looking expectations could proxy for agents with myopia or rational inattention. Hence, distinguishing between these components is interesting not only from the measurement perspective but the economic one.

We estimate our model for a range of measures of inflation and labor market slack. Our preferred inflation measure is headline CPI inflation, because consumer inflation expectations correspond most closely to this measure. We discuss other headline measures such as PCE and GDP deflator inflation as well as core inflation. Because the methodology that the BLS uses to construct the CPI has significantly evolved during our sample period, we also consider a current-method CPI series. Our preferred measure of labor market slack is the output gap, which accounts for changes in the natural rate of unemployment. The results based on the labor share are similar.

We document several findings. First, over the estimation period, we find at best weak evidence for nonlinearity. Depending on the inflation measure and identification strategy, we either cannot reject a linear model in favor of a nonlinear model or, when we do, the preferred nonlinear model has a high unemployment threshold, with a "flat" regime that is rarely operational. Moreover, if consumer expectations are included in the model and the inflation measure corresponds to those expectations, estimated nonlinearities are not statistically significant. Second, whether the preferred model is linear or nonlinear, consumer inflation expectations play a pivotal role, whereas professional forecasters' expectations appear less important. Third, while nonlinearities play a limited role over the full sample period, they help improve the model during some historical episodes. In particular, nonlinearities improve the fit and out-of-sample forecasts during the early 2010s, when following the Great Recession the unemployment rate was particularly high. In contrast, survey consumer expectations help the model during all the historical episodes considered. Therefore, macroeconomic models emphasizing beliefs should be more successful in matching the data than the models focusing on nonlinearities.[4]

Our ability to detect nonlinearities is highly influenced by the Great Recession and its aftermath, an episode when the unemployment gap was persistently high. However, we see it as a contribution rather than a limitation of our paper. One possible reason why the previous literature could not establish a convincing case for nonlinearities is that nonlinearities matter only at very high rates of unemployment. Our baseline estimates of the threshold confirm this conjecture. Nonetheless, despite the significant role of the Great Recession in our analysis, we learn from other episodes as well: about 40% of observations for the high-unemployment regime come from three different decades in the previous century.

This paper contributes to a broad literature examining nonlinearities in macroeconomic relationships in general (e.g., Kilian and Vigfusson, 2011; Auerbach and Gorodnichenko, 2012; Santoro et al., 2014; Tenreyro and Thwaites, 2016; Berge, De Ridder, and Pfajfar, 2021; Barnichon, Debortoli, and Matthes, 2022) and specifically in the Phillips curve (e.g., Debelle and Laxton, 1997; Barnes and Olivei, 2003; Kumar and Orrenius, 2016; Hooper, Mishkin, and Sufi, 2020; Ascari, Bonomolo, and Haque, 2022). Our paper is related to the recent studies examining whether nonlinearities can explain the missing disinflation (e.g., Beaudry and Portier, 2018; Kumhof and Wang, 2021; Harding, Lindé, and Trabandt, 2022).[5] We contribute to this literature by adopting a flexible estimation approach and by examining jointly nonlinearities and consumer expectations. In agreement with these concurrent papers, we find that nonlinearities can contribute to explaining the missing disinflation. We also find, however, that once consumer expectations are incorporated, nonlinearities become muted, whereas the significance of consumer expectations is robust to nonlinearity. Our flexible estimation approach, which can detect any form of nonlinearity, ensures that this result is not driven by ex ante limiting nonlinearity to a particular functional form.

Our paper also contributes to the literature emphasizing survey measures of inflation expectations (e.g., Leduc, Sill, and Stark, 2007; Adam and Padula, 2011; Chan, Clark, and Koop, 2018; Coibion, Gorodnichenko, and Kamdar, 2018; Pfajfar and Roberts, 2022). Since many such papers emphasize robust deviations from rationality in surveys (e.g., Andrade and Le Bihan, 2013; Pfajfar and Santoro, 2013; Ehrmann, Pfajfar, and Santoro, 2017; Fuhrer, 2018;

[4]See, for instance, Mertens (2016) and Gerko (2018) for recent work focusing on beliefs in macro models.

[5]Nonlinearities can arise due to interest-rate bounds (Beaudry and Portier, 2018; Brunnermeier and Koby, 2018; Kumhof and Wang, 2021), a kinked demand curve (Kimball, 1995; Harding, Lindé, and Trabandt, 2022), downward nominal wage rigidity (Daly and Hobijn, 2014), and state-dependence in structural parameters such as the Calvo rate (Gagnon, 2009; Alvarez et al., 2018; Petrella, Santoro, and de la Porte Simonsen, 2018). Asymmetries in the transmission of monetary policy and their consequences are examined in Schaling (2004) and Santoro et al. (2014).

Bordalo et al., 2020), and there is no widely accepted approach to incorporate nonrational and heterogeneous expectations to multiequation DSGE models, we focus directly on the Phillips curve equation. Finally, our paper contributes to the literature analyzing the economic environments of the 1970s and '80s, characterized by an inflation runup followed by a drastic disinflation.[6]

The paper proceeds as follows. Section 2 describes the empirical methodology, including identification strategies. Section 3 briefly summarizes our data and sources. In Section 4, we present our main results. We discuss our baselines estimates for measures of labor market slack, inflation, inflation expectations, and a number of estimation approaches. Section 5 uses our baseline estimates to better understand inflation dynamics during the historical episodes. We examine in-sample model fit during each episode and discuss some out-of-sample forecasting performance of the nonlinear models during the recent period. Section 6 examines robustness of the results. Section 7 concludes.

2. Methodology

We estimate the structural Phillips curve using single-equation methods. We consider multiple approaches to identify this structural relationship. Our baseline approach follows a long list of papers that estimated the Phillips curve by controlling for an extensive list of supply-side shifters. In addition, we consider two other approaches. First, in order to isolate demand-driven variation, we instrument the slack variable with the monetary policy shocks separately identified in the literature. Following Barnichon and Mesters (2020), we focus on Romer and Romer (2004) shocks. Second, we focus on regional variation, thereby removing aggregate effects such as oil shocks and endogenous monetary policy responses using time fixed effects. Despite some natural quantitative discrepancy between these approaches, all of them qualitatively paint a similar picture. While this estimation approach focuses on a single equation, our empirical specification can be derived from standard New Keynesian models. Hence, by isolating variation specific to demand shifters, our estimates have structural interpretation relevant for the New Keynesian Phillips curve.

We begin with a linear version of the expectations-augmented Phillips curve:

$$\pi_t = \mu + \mathbb{E}_t\, \pi_{t+1} + \kappa\, u_t + \boldsymbol{\varphi}\, \mathbf{z}_t + \varepsilon_t, \tag{1}$$

[6] Popular explanations of the macroeconomic dynamics of that time include, among many others, a policy-regime change (Blanchard, 1984; Clarida, Galí, and Gertler, 2000), changes in NAIRU and real-time measurement issues (Ball, 1997; Orphanides, 2001), supply-side shifters (Barsky and Kilian, 2002; Blinder and Rudd, 2013), rising disagreement in inflation forecasts (Mankiw, Reis, and Wolfers, 2004), and the evolution of policymakers' beliefs (Primiceri, 2006). Erceg and Levin (2003), Goodfriend and King (2005), Bordo et al. (2007), and Nunes (2009) also examine these issues.

where π_t is the rate of inflation, $\mathbb{E}_t\, \pi_{t+1}$ is expected inflation, u_t is a measure of real economic activity (e.g., the unemployment gap), $\mathbf{z}_t$ is a vector of controls observed in period t, ε_t is the error term, and μ, κ, and $\boldsymbol{\varphi}$ are estimated parameters.[7] The coefficient κ measures the slope of the linear Phillips curve. When κ is large in absolute value, inflation is sensitive to changes in economic activity. Before we proceed to nonlinearities, we discuss a few modeling choices in the context of the linear model: the choice of the slack variable u_t, the control variables $\mathbf{z}_t$, and the treatment of expectations $\mathbb{E}_t\, \pi_{t+1}$.

As a baseline measure of slack, we use the unemployment gap. By using the gap—rather than the unemployment *rate*—we focus on nonlinearities that occur over the business cycle and separate them from possible secular trends in the labor-market dynamics. In the New Keynesian tradition, one can use a microfounded model (e.g., Rotemberg and Woodford, 1997) to derive a relationship between inflation and marginal cost. Therefore, an empirical analog of this model calls for a direct measure of marginal cost such as the labor share of income (Galí and Gertler, 1999). However, the downward trend in the labor share observed since the early 2000s makes this measure problematic. For this reason, we use the unemployment gap as a benchmark measure, but we also report estimates of the model using the unemployment rate, the labor share, and an adjusted measure that accounts for the downward trend in the raw labor share by setting the labor fraction of proprietors' income to its historical average (Armenter, 2015).[8]

In a very interesting analysis, McLeay and Tenreyro (2019) discuss the pitfalls of estimating the Phillips curve due to optimal monetary policy and the presence of cost-push shocks. One of the solutions they propose is to include a set of variables that control for cost-push shocks. We follow that approach and include in our set of controls, $\mathbf{z}_t$, two lags of the growth of the relative price of food and energy, two lags of the change in the nominal exchange rate, and the Gordon (1982) price- and wage-control variable. McLeay and Tenreyro (2019) also mention the importance to control for inflation expectations—as we do in our baseline specification—as well as financial frictions, which we analyze more carefully in a later section. As discussed before, we also present results when marginal costs are the driving force of inflation because, as McLeay and Tenreyro discuss, the econometric identification is improved when the central bank's stabilization variable differs from the one driving the inflation process.

[7]We do not impose the constraint that $\mu = 0$, because our focus is on the empirical relationship between inflation and unemployment. If the models that give rise to this constraint are correct, the econometric cost of including the intercept is negligible. But if they are wrong, omitting the intercept would bias our estimates of the slope. Moreover, the constant controls for the mismatch between the levels of expectations and actual inflation in surveys. For instance, consumers in the Michigan survey expect on average higher inflation than observed in the data.

[8]We extend the original variable through 2019:Q4.

We estimate our model using single-equation methods.[9] The advantage of our approach is that its econometric implementation is straightforward and can be easily combined with available techniques to estimate nonlinearity. Despite the well-known identification challenges of this approach, we make significant inroads by considering a variety of different strategies and conducting extensive sensitivity analyses. Adopting single-equation estimation also facilitates comparison with extensive previous literature.

2.1. Expectations Process

Following the literature, we model inflation expectations as a combination of backward-looking and forward-looking terms (e.g., Fuhrer, 2010; Nunes, 2010). We employ the University of Michigan's Surveys of Consumers, which asks consumers about their expectation of inflation over the following year,[10] as well as the Survey of Professional Forecasters (SPF), which collects forecasts from expert forecasters for various inflation variables, including at a one-quarter horizon.[11] The SPF may best capture how large firms set prices, while the University of Michigan's Surveys of Consumers (UMSC) reflects consumers' expectations and may best capture expectations of small businesses.

In addition, a large body of literature emphasizes inflation persistence (e.g., Fuhrer and Moore, 1995; Fuhrer, 2006), which can be reconciled with backward-looking expectations wherein a firm's forecast of future inflation is a weighted average of past inflation. Hence, we also include lags of actual inflation. To cover a full year of observations and to minimize the Akaike information criterion for both the labor-share and unemployment-gap specifications, we choose five lags, but our results are not sensitive to this choice.[12]

To summarize, the expectations process is modeled as follows:

$$\mathbb{E}_t\, \pi_{t+1} = \sum_{i=1}^{5} \delta_i\, \pi_{t-i} + \alpha_1\, \mathbb{E}_t^{\mathrm{SPF}}\, \pi_{t+1} + \alpha_2\, \mathbb{E}_t^{\mathrm{UMSC}}\, \pi_{t+1}, \tag{2}$$

[9] Alternatively, the Phillips curve could be estimated using full information from a DSGE model, as in Smets and Wouters (2007), among many others. This method relies on cross-equation restrictions and priors (e.g., Del Negro and Schorfheide, 2008). Our approach complements that strand of literature.

[10] While the New Keynesian Phillips curve mandates the use of one-period-ahead expectations, the Michigan survey does not ask about one-quarter-ahead expectations. The strong performance of this measure, however, suggests that this measurement error is likely immaterial for our analysis.

[11] While we focus on short-term inflation expectations, as motivated by workhorse models, practitioners have recently turned their attention to long-term expectations (e.g., Yellen, 2015). We note that long-term expectations are persistent and thus have less potential to explain the puzzling inflation dynamics that we study in this paper. We verify that our estimates are not affected in a material way by controlling for the 10-year inflation expectations used in the Federal Reserve Board's FRB/US model.

[12] By controlling for inflation lags, we also focus on inflation innovations over the business cycle and remove the lower-frequency fluctuations in inflation.

where $\mathbb{E}_t^{\text{SPF}}\,\pi_{t+1}$ is one-quarter-ahead expected inflation in the SPF, $\mathbb{E}_t^{\text{UMSC}}$ is the expected inflation from the UMSC, and δ_i, α_1, α_2 are estimated parameters. Denoting $\sum_{i=1}^{5}\delta_i$ by α_0, we constrain the coefficients on the expectation terms to sum to one (i.e., $\alpha_0+\alpha_1+\alpha_2=1$); thus, $\hat{\alpha}_0$, $\hat{\alpha}_1$, and $\hat{\alpha}_2$ represent the relative weights of the backward-looking and forward-looking components in the expectations process.[13]

Because we focus directly on the Phillips curve equation, we can allow the expectations process to deviate from the full information rational expectations benchmark. That is, instead of treating $\mathbb{E}_t\,\pi_{t+1}$ as the mathematical expectation of next-period inflation, as would be the case in FIRE DSGE models, we can think of it as a separate—though possibly endogenous—process, which may depend both on the underlying inflationary trend and subjective beliefs. This approach is consistent with recent evidence on nonrationality of expectations obtained, in particular, from the survey data that we use (e.g., Andrade and Le Bihan, 2013; Coibion, Gorodnichenko, and Kamdar, 2018; Fuhrer, 2018; Bordalo et al., 2020). One can loosely think of aggregate expectations in Equation (2) as a weighted average of expectations of backward-looking agents; large firms, approximated by professional forecasters; and small firms, approximated by households, with the corresponding weights α_i.

2.2. Nonlinear Specification

To estimate nonlinearities, we employ the Hansen (1996, 2000) threshold regressions. This method approximates a nonlinear curvature by a piecewise-linear function in which the number of kinks (thresholds) is determined endogenously. Relative to other nonlinear estimation methods, the threshold regression has several advantages. Each linear segment can be estimated by ordinary least squares, and therefore estimation and inference are straightforward. There is no need to assume any specific form of nonlinearity: the data decide how much nonlinearity (i.e., how many thresholds) there is. Further, this method has been used before and therefore allows for comparison with previous studies (e.g., Barnes and Olivei, 2003). Its major shortcoming is the tendency to produce wide confidence intervals for thresholds. That is, even though we can improve the fit of the model and test explicitly for nonlinearities, we may not be able to determine the thresholds' location with certainty.

We estimate the model with a continuity constraint.[14] Without it, even in the absence of shocks, infinitesimal changes in unemployment would lead to jumps in the inflation rate. Discontinuity could also result in a lack of equilibrium, which would be difficult to reconcile

[13]In a standard New Keynesian model, $\alpha_0+\alpha_1+\alpha_2=\beta$, the discount factor. At a quarterly frequency, $\beta\approx 1$. The unconstrained regressions support this restriction in the vast majority of cases. See Online Appendix A for additional details.

[14]Note that continuity imposes a constraint on the intercept across the regimes. With one threshold (two regimes), we have one free intercept parameter, while the other one is pinned down by the constraint.

with the U.S. time-series and most standard economic models. Hansen (2017) describes in detail the econometric apparatus for the linear constraint case that we analyze here.

A piecewise-linear Phillips curve with a vector $\boldsymbol{\gamma} = (\gamma_1, \ldots, \gamma_m)$ containing m thresholds can be written as follows:

$$\begin{aligned} \pi_t =& \underline{\mu}(\boldsymbol{\gamma}) + \mathbb{E}_t \pi_{t+1} + \underline{\kappa}(\boldsymbol{\gamma}) u_t + \boldsymbol{\varphi} \mathbf{z}_t + \varepsilon_t, \\ \underline{\mu}(\boldsymbol{\gamma}) =& \sum_{j=1}^{m+1} \mu_j \, \mathbb{I}\big(\gamma_{j-1} \le u_t < \gamma_j\big), \\ \underline{\kappa}(\boldsymbol{\gamma}) =& \sum_{j=1}^{m+1} \kappa_j \, \mathbb{I}\big(\gamma_{j-1} \le u_t < \gamma_j\big), \end{aligned} \tag{3}$$

where $\mathbb{I}(\gamma_{j-1} \le u_t < \gamma_j)$ is an indicator function of a condition $\gamma_{j-1} \le u_t < \gamma_j$, assuming $\gamma_0 = -\infty$ and $\gamma_{m+1} = +\infty$. In a one-threshold case, this definition results in two regimes: $u_t < \gamma$ (regime "L") and $u_t \ge \gamma$ (regime "R"). This threshold allows for shifts in the Phillips curve over the range of u_t. To compute the optimal thresholds $\boldsymbol{\gamma}^*$, an OLS regression is estimated sequentially for all possible values of $\boldsymbol{\gamma}$. We then choose $\boldsymbol{\gamma}^*$ that minimizes the residual sum of squares.[15]

We then test the null hypothesis of the linear model against the alternative of a one-threshold model. Let S_0 and S_1 be the residual sums of squares under the null hypothesis and under the alternative, respectively. For a sample of n observations, the F-statistic of this test is of the form:

$$F = n \frac{S_0 - S_1}{S_1},$$

with a distribution that can be approximated through a bootstrap procedure.[16] Since the critical values of this test depend on parameters of the model, it is more informative to report its p-values. This test can be extended for the null of an arbitrary number of thresholds $\ell \ge 0$ against the alternative of $\ell + k$ thresholds, $k \ge 1$. The optimal number of thresholds can be determined by running the test sequentially, starting from $\ell = 0$ and $k = 1$ and then increasing ℓ by one if the null is rejected. The optimal number of thresholds ℓ^* is the lowest ℓ for which the null is not rejected.

[15]To maintain statistical power of the test, we constrain the grid for $\boldsymbol{\gamma}$ to ensure that each regime contains no less than 10% of the sample size.

[16]Confidence intervals for the threshold are computed using a similar bootstrap procedure; see Hansen (2017) for more details.

3. Data

We use aggregate data, obtained from public sources, from 1968 through 2019 at a quarterly frequency. The start date is motivated by availability of data on consumer expectations. Our baseline inflation measure is the seasonally adjusted, annualized quarterly growth rate of the CPI. We focus on the headline CPI, because surveys of consumer expectations ask their respondents about the cost of living, which corresponds to headline consumer prices, as food and energy expenses take a significant share in U.S. households' budgets. Because our paper emphasizes expectations as a key factor of inflation dynamics, it is important to avoid a mismatch between the measure of inflation and the corresponding expectations. We also study in detail the specifications based on the growth rates of the PCE price index and the GDP deflator as well as the core measures of consumer prices (i.e., net of food and energy prices).

We measure unemployment with the total civilian population (over 16 years old) unemployment rate. To construct the gap, we subtract from this measure the natural unemployment rate, produced by the CBO. As an alternative slack variable, we consider the seasonally adjusted labor share for the nonfarm business sector. We also use the Armenter (2015) procedure to account for the secular downward trend in labor's share.

To measure expectations, we rely on two major sources: the Survey of Professional Forecasters (SPF) and the University of Michigan's Surveys of Consumers (UMSC). Both surveys have been widely used in the literature. From the SPF, we take the one-quarter-ahead median expectations of the GDP deflator, because the corresponding series for the CPI is available starting only in the 1980s. From the UMSC, we take the mean expectations of the cost-of-living increase over the following 12 months period.[17] For controls and robustness checks, we use a plethora of other variables summarized in Table B.1 of Online Appendix B.

4. Empirical Results

4.1. Baseline Estimates for a Range of Slack and Regime Measures

Table 1 presents estimates of the Phillips curve for a range of slack and regime measures. In Panel A, we show estimates of a standard linear model. Our baseline estimates are based on the unemployment gap as a slack and regime measure (column 1). In the linear case, the slope of the Phillips curve is -0.30, which is consistent with the previous literature. The slope is also negative and significant when we use the unemployment rate as a slack and regime

[17] It is well known that the SPF mean and median inflation expectations are similar, but the UMSC mean and median expectations are not. We use the mean for the UMSC because it is available in the earlier period. The main results are similar if we use the median UMSC expectations for a shorter period with available data.

Table 1: Linear versus Nonlinear Phillips Curve: Various Slack/Regime Measures

	Unemployment		Labor Share	
	Gap	Rate	Raw	Adjusted
	(1)	(2)	(3)	(4)
Panel A: Linear Model				
Slope, $\hat{\kappa}$	-0.30^{***}	-0.14^{*}	0.20^{***}	0.32^{***}
	(0.08)	(0.07)	(0.04)	(0.05)
Expected inflation				
UMSC, $\hat{\alpha}_2$	0.79^{***}	0.77^{***}	1.06^{***}	0.93^{***}
	(0.17)	(0.16)	(0.18)	(0.16)
SPF, $\hat{\alpha}_1$	−0.12	−0.14	−0.14	0.02
	(0.27)	(0.27)	(0.22)	(0.23)
Sum of lags, $\hat{\alpha}_0$	0.33^{***}	0.38^{***}	0.08	0.05
	(0.11)	(0.12)	(0.12)	(0.12)
Panel B: Threshold Model				
Slopes				
left, $\hat{\kappa}_L$	-0.50^{***}	−0.04	0.08	0.62^{***}
	(0.18)	(0.21)	(0.07)	(0.16)
right, $\hat{\kappa}_R$	0.08	-0.19^{**}	0.33^{***}	0.26^{***}
	(0.19)	(0.09)	(0.10)	(0.07)
Expected inflation				
UMSC, $\hat{\alpha}_2$	0.73^{***}	0.77^{***}	1.04^{***}	0.99^{***}
	(0.15)	(0.16)	(0.18)	(0.16)
SPF, $\hat{\alpha}_1$	−0.07	−0.14	−0.05	−0.07
	(0.27)	(0.27)	(0.24)	(0.24)
Sum of lags, $\hat{\alpha}_0$	0.34^{**}	0.38^{**}	0.02	0.08
	(0.17)	(0.16)	(0.19)	(0.17)
Threshold, $\hat{\gamma}$				
point estimate	1.95	5.87	−1.95	−3.11
95% confidence interval	[−0.82, 2.93]	[4.23, 8.47]	[−7.63, 3.21]	[−4.27, 2.54]
No. of thresholds, p-value				
0 vs. 1, H_0: 0	0.20	0.95	0.18	0.41
1 vs. 2, H_0: 1	0.92	0.38	0.41	0.34
R^2	0.77	0.75	0.79	0.79
N	205	205	205	205

Notes: The estimation sample is 1968:Q4 through 2019:Q4. The dependent variable is the seasonally adjusted, annualized CPI inflation rate. Alternative threshold variables are in columns (1)–(4). Inflation expectations include consumer expectations from the University of Michigan's Surveys of Consumers, SPF forecasts of one-quarter-ahead GDP inflation, and five lags of inflation. (The SPF forecasts of CPI inflation are not available for the early sample.) Control variables (estimates not reported) are two lags of the growth rate of the relative price of food and energy, two lags of the change in the nominal exchange rate, and the price- and wage-control measure. The adjusted labor share (column 4) is obtained by setting the labor fraction of proprietors' income to its historical average. The threshold point is estimated using the Hansen regression-kink method. Newey–West standard errors, allowing for autocorrelation of up to five lags, are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

variable (column 2). We estimate a significant slope also when we focus on the labor share measures. In column (3), we use the raw measure because it relates directly to the unobserved output gap, while in column (4) we use the adjusted measure based on the historical average of proprietors' income allocated to labor. Because the labor share moves in the opposite direction from unemployment, the slope in this case is positive, as expected. In all these cases, consumer expectations, measured by the University of Michigan survey, are a dominant component of the inflation expectations process. Across these specifications, the UMSC coefficient is at least 0.77, whereas the sum of coefficients on inflation lags is at most 0.38. The SPF carries almost no weight. Consumer inflation expectations are particularly important in the labor-share specifications, with a corresponding weight close to one.

Next, we compare the linear model with a one-threshold model, whose estimates are shown in Panel B. We estimate the threshold value of the unemployment gap at 1.95%.[18] The 95% confidence interval is rather wide, corresponding to unemployment rates between 4.7% and 8.4%. The curve is relatively steep left of the threshold (with a slope of −0.50) and essentially flat right of the threshold. Note that, even for the left portion of the curve (low unemployment), we cannot rule out estimates of a small response of inflation to unemployment, suggesting that the piecewise-linear Phillips curve may still be overall flat, and it is especially so when unemployment is high. Overall, we cannot reject the linear model in favor of the threshold model, suggesting that nonlinearities played a limited role in the full sample. The p-value of the test for the linear model (null hypothesis) versus the one-threshold model ranges 0.18–0.95, depending on the slack and regime variable.[19] The results are qualitatively similar for the unemployment-rate and labor-share specifications. Similar to the linear case, consumer expectations of inflation play a prominent role. The baseline UMSC coefficient is 0.73, whereas the sum of coefficients on inflation lags is 0.34. The SPF expectations remain irrelevant once consumer expectations are controlled for.

To better understand why nonlinearities play only a limited role in the full sample, Figure 1 highlights the episodes when the unemployment gap was high and the Phillips curve was flatter than usual. The longest episode of the high-unemployment regime corresponds to the Great Recession and its aftermath. In addition, our procedure identifies three other, shorter episodes, which are spread in time and are each associated with an economic downturn. The timing of our unemployment regimes, however, differs from the timing of recessions defined as two consecutive quarters of a GDP decline or by a more comprehensive procedure employed by the NBER. Typically, a high-unemployment regime also contains early recovery periods, when output is on an upward trend but the unemployment gap remains high. Hence, such regimes better capture the state of labor markets than output dynamics. Importantly, our model puts thresholds in different non-contiguous time periods and therefore provides a mechanism dif-

[18]According to Congressional Budget Office data, the natural rate of unemployment averages at 5.47% during our baseline estimation period. Hence, the estimated unemployment gap threshold corresponds to an unemployment rate of 7.42%.

[19]In the table, we also report a test for one versus two thresholds. This test suggests an even smaller likelihood of the second thresholds.

Figure 1: Regimes in the Nonlinear Model

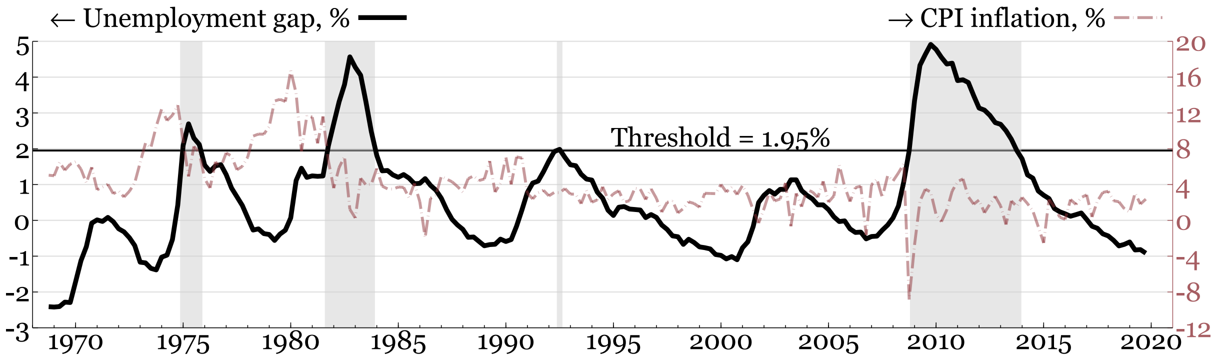

Notes: The shaded areas correspond to the periods when the unemployment gap is above the estimated threshold (a flatter segment of the Phillips curve). The unemployment gap (solid line, left axis) is constructed as the difference between the unemployment rate and the CBO measure of the natural rate. Inflation (broken line, right axis) is measured as the annualized quarterly growth rate of the headline CPI.

ferent from a structural break resulting in a flattening of the Phillips curve (e.g., Roberts, 2006; Simon, Matheson, and Sandri, 2013; Luengo-Prado, Rao, and Sheremirov, 2018).

Overall, our baseline estimates provide a compelling case that, as a whole, nonlinearities do not matter much. The "best" case for nonlinearities favors one threshold that gives rise to two regimes: a dominant regime that prevails in the sample and an infrequent regime driven by a few deep recessions. While this asymmetry makes it difficult to estimate the difference between the regimes precisely, we note that, en route to the best model, our grid-search procedure considers cases where the threshold values give rise to regimes occurring with similar frequencies. However, such symmetric models are rejected by the data in favor of (locally) linear models. Hence, to the extent nonlinearities matter at all, they do so only rarely.

4.2. The Role of Inflation and Inflation Expectations Measures

In this section, we discuss how the measurement of inflation and inflation expectations affects estimated nonlinearity. We show that overall the mismatch between the measures of inflation and inflation expectations puts nonlinearities in a more prominent light. We begin by excluding consumer expectations from the baseline specification based on headline CPI inflation. These results are shown in column (1) of Table 2.[20] In this case, we still identify the threshold at 1.95, and the left regime is steeper than the right regime. However, the linear model is rejected at the 5% significance level and the coefficient on the SPF expectations becomes large and highly significant. Hence, not accounting for consumer expectation leads the econometri-

[20] To save space, in Panel A we report only the slopes of the linear models, as in this section we focus mainly on the threshold model. Similarly to the baseline, the coefficients on inflation expectations are not materially affected by allowing for nonlinearity in the effect of the slack variable.

Table 2: Nonlinearity and Inflation Measurement

	CPI (1)	PCE (2)	GDP Deflator (3)	Core CPI (4)	Core PCE (5)	CPI-U-RS (6)
			Panel A: Linear Model			
Slope, $\hat{\kappa}$	-0.31^{***}	-0.15^{**}	-0.15^{***}	-0.22^{***}	-0.15^{***}	-0.16^{*}
	(0.10)	(0.06)	(0.05)	(0.05)	(0.05)	(0.09)
			Panel B: Threshold Model			
Slopes						
left, $\hat{\kappa}_L$	-0.63^{***}	-0.36^{***}	-0.60^{***}	-0.59^{**}	-0.60^{***}	-0.46^{**}
	(0.19)	(0.13)	(0.16)	(0.24)	(0.16)	(0.20)
right, $\hat{\kappa}_R$	0.31	0.24^{*}	-0.08	-0.19^{***}	-0.08	0.28
	(0.27)	(0.14)	(0.05)	(0.05)	(0.05)	(0.18)
Expected inflation						
UMSC, $\hat{\alpha}_2$		0.57^{***}	0.29^{***}	0.35^{***}	0.29^{***}	0.66^{***}
		(0.08)	(0.07)	(0.13)	(0.07)	(0.19)
SPF, $\hat{\alpha}_1$	0.73^{***}	0.13	0.24^{*}	0.16	0.24^{*}	0.38^{**}
	(0.16)	(0.16)	(0.12)	(0.23)	(0.12)	(0.19)
Sum of lags, $\hat{\alpha}_0$	0.27	0.30^{**}	0.47^{***}	0.49^{***}	0.51^{***}	-0.05
	(0.16)	(0.12)	(0.08)	(0.13)	(0.09)	(0.18)
Threshold, $\hat{\gamma}$						
point estimate	1.95	1.95	-0.32	-0.82	2.47	1.95
95% CI	$[-0.67, 2.93]$	$[-0.82, 2.93]$	$[-0.82, 1.27]$	$[-0.82, 2.93]$	$[-0.80, 2.93]$	$[-0.38, 3.31]$
Test p-value						
0 vs. 1, H_0: 0	0.03	0.01	0.04	0.69	0.04	0.05
1 vs. 2, H_0: 1	0.94	0.72	0.77	0.34	0.36	0.56
R^2	0.74	0.83	0.90	0.87	0.91	0.71
N	205	205	205	205	205	163

Notes: See notes to Table 1. CPI-U-RS refers to the CPI research series that uses current methods.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

cian to overestimate the role of nonlinearities and the importance of professional forecasters' expectations.

While we use headline CPI inflation as a baseline measure, there are several good reasons to consider other inflation measures as well. First, the Federal Reserve's preferred measure changed in 2000 from CPI to PCE inflation; hence, significant media coverage relates to this measure. Second, the GDP deflator has the advantage of having the longest forecast series in the SPF. Third, headline inflation tends to be volatile due to energy and food prices, which are not directly accounted for by DSGE model. While our controls partially address this issue, using core inflation measures is a more straightforward way to remove these components. This of course comes at a cost that survey measures of consumer expectations ask the respondents about prices in general, and food and energy prices can have disproportionate effects on expectations due to their salience. Finally, the methods used to construct the CPI have evolved during the sample period. The BLS publishes a research series bases on current methods (CPI-U-RS). We consider these alternative measures of inflation next.

Column (2) of Table 2 presents our baseline-specification estimates when inflation is measured using the PCE price index. In this specification, the estimated threshold equals the baseline value, and the linear model is rejected at a 1% level. The relative weight of the consumer expectations decreases from 0.73 to 0.57, whereas the weight of the professional forecasters' expectations increases relative to the baseline in Table 1, but overall remains small. Despite the stronger statistical case for nonlinearities, the economic interpretation is similar to the baseline. That is, even if nonlinearities matter, they do so at rather high levels of unemployment, thereby giving rise to a regime that occurs relatively infrequently.

Column (3) of Table 2 focuses on the GDP deflator. This inflation measure has been often preferred in the DSGE literature (e.g., Smets and Wouters 2007). The threshold is again significant, at a 5% level. Importantly, the estimated threshold is -0.32, giving relatively more prominence to the high-slope regime than in the baseline.[21] As before, the curve appears relatively steep when the unemployment gap is high and nearly flat when it is low. The UMSC weight in this model is 0.29, even lower than in the PCE model, while the SPF weight is 0.24. The more prominent role of the SPF expectations can be explained by the fact that both actual inflation and predicted inflation are measured with the GDP deflator.

Our conclusions remain qualitatively the same when we consider the core measures of consumer prices, but there are some quantitative differences. For the core CPI (column 4), the threshold moves to the left. However, the threshold estimate is insignificant (p-value at 0.69), and the slopes on both sides of the threshold are strongly negative and statistically indistinguishable from each other. Simply put, with the core CPI, the threshold disappears completely. This result supports our general conclusion that nonlinearities do not provide a consistent explanation of inflation dynamics in our study. For the core PCE price index (column 5), the high-unemployment regime occurs even less frequently than in the baseline, as the estimated threshold is about 0.5 percentage points larger than our benchmark estimate.[22]

Finally, in column (6) we show estimates obtained using constant methods CPI inflation (CPI-U-RS).[23] Starting in 1978, this series is shorter than the baseline. Nonetheless, the overall conclusions remain intact. The threshold value is the same as in our baseline based on the longer sample. The left portion of the curve is flatter than the right portion. Consumer expectations have the largest weight in the inflation expectations process. As in some other exercises, and unlike in our baseline, the formal test rejects linearity in favor of one threshold. But, as before, the baseline threshold gives rise to asymmetric regimes, with the high-unemployment regime being relatively infrequent.

To summarize, this section establishes two main points. First, despite the mismatch, consumers' inflation expectations are important not only for CPI inflation but for a wide range of inflation measures. Second, while nonlinearities are either marginally significant or insignifi-

[21] This point is visualized by Figure B.1 in Online Appendix B.

[22] In Online Appendix B, Tables B.2–B.7 present a set of results for other measures of slack (the unemployment rate and labor's share). They are consistent with our main conclusions.

[23] For a description and a recent application of this series, see Stewart and Reed (2000) and Bolhuis, Cramer, and Summers (2022).

cant in specifications with CPI inflation and consumer expectations (of CPI inflation), nonlinearities may regain significance when consumer expectations are absent or when the measures of inflation and consumer expectations no longer coincide, most notably for GDP inflation.

4.3. Evidence from Other Identification Strategies

External Shocks as Instruments

In our baseline specification, we trace the Phillips curve by controlling for cost-push shocks. Including these controls addresses various issues with identification of the Phillips curve.[24] A complementary approach is to use instrumental variables. In recent work, Barnichon and Mesters (2020) estimate the Phillips curve using monetary policy shocks as instruments.[25] Combining the Hansen (2017) threshold regression with the Barnichon and Mesters (2020) approach is a very interesting but challenging endeavor, as the econometric tools to do so are currently unavailable.

Instead, we make an inroad as follows. We estimate the model using two-stage least squares with grid search over potential threshold values. And then we assess nonlinearities by testing for equality of the estimated regimes' slopes in the nonlinear case. However, we are unable to formally test the linear model versus nonlinear models or to compute confidence intervals for the threshold. To implement two-stage least squares, we follow the literature and instrument the unemployment gap with separately identified monetary policy shocks. We use the Romer and Romer (2004) shocks updated by Wieland and Yang (2020). These shocks are constructed using the Federal Reserve's confidential forecasts, which are released to the public with a five-year lag. Therefore, our sample period ends in 2015. We consider two version of these shocks: One is based on the federal funds rate, as in the original paper. The other is based on the Wu and Xia (2016) shadow rate, which we use to account for unconventional monetary policies during the period when the federal funds rate was constrained by the zero lower bound.

Results in columns (1) and (2) of Table 3 show that our qualitative conclusions remain mostly unchanged. The slope parameter in the linear regression (−0.30 for the shadow rate specification in Panel A, column 2) is similar to the OLS linear slope, thereby indicating that cost-push shocks can address endogeneity issues. In the nonlinear case (Panel B), the threshold is estimated at 1.13, with a left slope (−0.59) higher in absolute value than the right slope (−0.06). However, we cannot reject the null hypothesis that the slopes in the two threshold regions are equal (p-value is 0.51).[26] In accordance with the main analysis, consumer expecta-

[24] In Section 6.1, as a robustness exercise, we also control for financial frictions, because they could simultaneously affect the left-hand-side and right-hand-side variables.

[25] See also Del Negro et al. (2020) for an estimation approach conditional on demand shocks.

[26] We provide additional results in Online Appendix D. In particular, we consider the case where we fix the

Table 3: Alternative Identification Strategies

Method	Two-stage least squares		MSA panel
Slack / regime	Unemployment gap		Unemployment rate
Romer & Romer shocks based on	Federal funds rate	Wu & Xia shadow rate	—
Fixed effects	—	—	MSA, Time
	(1)	(2)	(3)
	Panel A: Linear Model		
Slope, $\hat{\kappa}$	−0.27	−0.30	−0.41***
	(0.18)	(0.19)	(0.06)
	Panel B: Threshold Model		
Slopes			
left, $\hat{\kappa}_L$	−0.65	−0.59	−0.45***
	(0.64)	(0.54)	(0.07)
right, $\hat{\kappa}_R$	−0.10	−0.06	−0.34***
	(0.32)	(0.31)	(0.08)
Expected inflation			
UMSC, $\hat{\alpha}_2$	0.84***	0.84***	0.00
	(0.20)	(0.19)	(0.11)
SPF, $\hat{\alpha}_1$	−0.27	−0.27	
	(0.36)	(0.35)	
Sum of lags, $\hat{\alpha}_0$	0.43**	0.43**	0.07**
	(0.15)	(0.15)	(0.04)
Threshold, $\hat{\gamma}$	1.08	1.13	6.8
Test p-value (H_0 : linear model)	0.55	0.51	0.94
N	168	168	1,248

Notes: Estimates obtained using two-stage least squares with grid search and aggregate data are shown in columns (1) and (2). The estimation sample is 1969:Q1 through 2015:Q4 at a quarterly frequency. The dependent variable is CPI inflation. The slack and regime variable is the unemployment gap. The Romer and Romer (2004) shocks, extended by Wieland and Yang (2020), are used to instrument the gap. In column (1), the shocks are based on the federal funds rate. In column (2), the shocks are based on the Wu and Xia (2016) shadow rate. The 2SLS regressions include 20 lags of the shocks. The threshold estimates in Panel B are obtained by minimizing the residual sum of squares.

Estimates of the panel regression in Equation (4) applied to metropolitan-level data are shown in column (3). The estimation sample includes 24 MSAs during the period 1991:H1 through 2017:H2 at a semiannual frequency. The dependent variable is the core CPI inflation rate. The slack and regime variable is the unemployment rate.The threshold is estimated following Kremer, Bick, and Nautz (2013). The threshold p-value is computed using a residual wild bootstrap clustered at the MSA level (Cameron, Gelbach, and Miller, 2008). Both MSA and time fixed effects are included. Standard errors are clustered at the MSA level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

tions are highly significant and dominate other measures of inflation expectations. The coefficient on the UMSC measure is 0.84, while the sum of lag coefficients is 0.43. As in the baseline case, the SPF expectations remain irrelevant when consumer expectations are included.

threshold at the level estimated using our baseline method. We find qualitatively similar results. However, we note that the slopes in all these exercises are not statistically different from zero, indicating limited power of this method.

Exploiting Regional Variation

Yet another approach to identify the Phillips curve exploits regional variation. Kumar and Orrenius (2016), Babb and Detmeister (2017), Hooper, Mishkin, and Sufi (2020), and Hazell et al. (2022), among others, provide very interesting analyses along this dimension. A particular advantage of this approach is that, by focusing on deviations of regional inflation and unemployment from their corresponding aggregate rates, one can remove variation due to aggregate supply shocks (e.g., oil price shocks), monetary policy, and other factors determined at the national level.

Building on McLeay and Tenreyro (2019), we introduce a threshold into a Phillips curve estimated at the metropolitan area level. Following their approach, we approximate consumer inflation expectations in a given MSA by the UMSC expectations for the broad geographical region in which the metro area is located.[27] In our preferred specification, we include the unemployment rate as a slack and regime variable, metropolitan area fixed effects (to partially account for regional variation in the natural unemployment rates), and time fixed effects to account for aggregate shocks. The threshold is estimated as in Kremer, Bick, and Nautz (2013). Due to the requirement of a balanced panel as well as general data availability, we estimate the model for 24 MSAs during the period 1991:H1–2017:H2 at a semi-annual frequency.[28]

Specifically, we estimate the following specification:

$$\pi_{i,t} = \mu_i + \zeta_t + \delta\, \pi_{i,t-1} + \alpha\, \mathbb{E}^{\text{UMSC}}_{j(i),\,t}\, \pi_{j(i),t+1} + \kappa_L\, u_{i,t} \times \mathbb{I}(u_{i,t} < \gamma) + \kappa_R\, u_{i,t} \times \mathbb{I}(u_{i,t} \geq \gamma) + \varepsilon_{i,t}, \quad (4)$$

where $\pi_{i,t}$ is the inflation rate at MSA i and semester t, $\mathbb{E}^{\text{UMSC}}_{j(i),\,t}\, \pi_{j(i),t+1}$ is the expected next year inflation in geographical region j containing MSA i, $u_{i,t}$ is the unemployment rate, $\varepsilon_{i,t}$ is the error term, and $\mu_i, \zeta_t, \delta, \alpha, \kappa_L, \kappa_R, \gamma$ are estimated parameters. In this specification, location fixed effects, μ_i, help account for variation in the natural unemployment rates across metropolitan areas, and time fixed effects, ζ_t, for aggregate effects, including endogenous monetary policy responses.

We show estimates of the regional specification in column (3) of Table 3. Overall, they are consistent with those based on aggregate data. The linear slope, at −0.41, is slightly higher than in the aggregate case. In the nonlinear model, we again find that the Phillips curve is steeper in the left region (with a slope of −0.45) than in the right region (−0.34). The difference between the slopes is smaller in magnitude than in the aggregate case. We estimate

[27]The four broad regions for which the UMSC measure is available are North Central, North East, South, and West. We note that metropolitan inflation varies substantially within these regions, and likely so do the corresponding expectations. Alternative measures with more detailed regional variation are available only for a short period.

[28]To enhance comparability with McLeay and Tenreyro (2019), we focus on core CPI inflation.

the threshold at an unemployment rate of 6.8%, with a 95% confidence interval from 5.4% to 7.5%. However, the linear model cannot be rejected in favor of a one-threshold model, with a corresponding p-value of 0.94. The coefficient on consumer inflation expectations is virtually zero, likely due to mismeasured expectations.[29] We note that, due to data limitations, these results call for further research, as more detailed data on regional variation in consumer expectations as well as longer time-series become available.

To sum up, whether we identify the Phillips curve using supply-side control variables, exogenous demand shocks, or regional variation, we find no strong evidence of significant nonlinearities in the Phillips curve. In contrast, we find that consumer expectations of aggregate inflation play an important role in inflation dynamics, which cannot be fully captured by professional forecasters' expectations or inflation inertia. While nonlinearities may be limited on average, in the next section we consider historical episodes during which nonlinearities could play a more prominent role.

5. Analyses of Historical Episodes

We now examine the merits of nonlinearities and consumer expectations during four different periods of large deviations of inflation and/or unemployment from their trends. To distinguish the role played by consumer expectations in these episodes, we isolate the innovation component in consumers' inflation expectations (i.e., the inflation expectations that cannot be forecast by data available in the previous quarter). To this end, we allow consumer expectations to depend on the lags of real-time inflation, the federal funds rate, the SPF forecast, and the change in oil prices.[30] We also add four lags of consumer inflation expectations, since Fuhrer (2018) shows that, at a micro level, UMSC participants tend to revise their inflation forecasts in response to the lagged central tendency of survey inflation expectations. Such a mechanism should render persistence in survey expectations.[31] That is, we estimate the following specification:

$$\mathbb{E}_t^{\text{UMSC}}\,\pi_{t+1} = a + \sum_{i=1}^{4} \rho_i\, \mathbb{E}_{t-i}^{\text{UMSC}}\,\pi_{t-i+1} + \sum_{i=1}^{4} b_i\, \pi_{t-i|t} + c\, r_{t-1} + d\, \mathbb{E}_{t-1}^{\text{SPF}}\,\pi_t + f\, \frac{\Delta P_t^{\text{oil}}}{P_{t-1}^{\text{oil}}} + \varepsilon_t^{\text{UMSC}}, \quad (5)$$

[29]We present additional results in Online Appendix E. When we do not include time fixed effects, which absorb aggregate inflation expectations, the coefficient on regional consumer expectations is positive and significantly different from zero. Thus, consumer inflation expectations are likely still important at the regional level. See Online Appendix Table E.1 for more details. Note also that time fixed effects fully absorb our measure of professional forecasters' expectations, as these data are available only at the aggregate level.

[30]The real-time data go back to 1994:Q3. We use revised data for the period when real-time data are not available.

[31]Binder (2017) finds that many respondents round their forecasts to the nearest zero or five. If inflationary shocks are small, this mechanism can also generate persistence in the measured expectations.

Table 4: Model In-Sample Fit

Inflation	Great Inflation: Peak-to-Trough Change, ppt (1)	Great Inflation: RMSE (2)	Volcker Disinflation: Peak-to-Trough Change, ppt (3)	Volcker Disinflation: RMSE (4)
CPI Inflation	13.1		−16.5	
Linear model				
estimated UMSC expectations	4.4	2.59	−6.9	2.82
actual expectations	6.6	1.90	−9.9	2.70
Threshold model				
estimated	4.6	2.63	−5.8	3.06
actual	6.6	1.98	−9.0	2.85

Notes: The Great Inflation episode is defined as the period 1976:Q2 through 1980:Q1, based on the lowest and highest CPI inflation values around the crisis. The Volcker disinflation episode is 1980:Q1 through 1983:Q1.

where $\pi_{t-1|t}$ is real-time inflation in period $t-1$ as observed in period t, r_t is the nominal federal funds rate, and P_t^{oil} is the oil price. With this inflation expectations process in mind, we estimate the model that combines Equation (2) with either Equation (1) or (3), wherein we use either the actual UMSC variable $\mathbb{E}_t^{\text{UMSC}}\pi_{t+1}$ or its fitted value, $\widehat{\mathbb{E}}_t^{\text{UMSC}}\pi_{t+1}$.

5.1. The Great Inflation and the Volcker Disinflation

We now look at the performance of the model during the Great Inflation of the 1970s and the subsequent Volcker disinflation. We define the former period as 1976:Q2 through 1980:Q1 and the latter period as 1980:Q1 through 1983:Q1. While we would like to focus on out-of-sample forecasts, the sample period preceding these episodes is too short to reliably estimate a threshold model. Therefore, in this section we focus on in-sample performance.[32] We use two criteria: (1) the change in inflation from the beginning to the end of the period, and (2) the models' fit during the entire episode measured by the root mean squared error (RMSE).[33] Table 4 shows the results.

During the Great Inflation, the runup in inflation was 13.1 percentage points (column 1). We show that nonlinearities do not help explaining this runup, while the models with actual UMSC expectations perform better than the models with estimated consumer expectations. For instance, conditional on expectations, switching from the linear model to the nonlinear one can explain only up to additional 0.2 percentage points of the inflation runup, whereas conditional on (non)linearity, switching from predicted to actual inflation expectations explains additional

[32] In the next section, where we study the missing disinflation and the missing inflation of the 2010s, we focus mainly on out-of-sample forecasts, with in-sample results relegated to the appendix.

[33] Figure B.2 in Online Appendix B shows the models' fit during the full estimation period. Figures B.3 and B.4 provide additional evidence on the two episodes.

2.0 to 2.2 percentage points of inflation. About half of the total runup, however, remains unexplained by these models. Similar conclusions follow from the RMSE analysis (column 2). The lowest RMSE is achieved by the linear model with actual inflation expectations (1.90), while both threshold models have higher RMSEs than the corresponding linear models.

The Volcker period is characterized by a disinflation of 16.5 percentage points (column 3). The best model is again the linear model with actual inflation expectations, which can explain 9.9 percentage points of disinflation (column 3) and also has the lowest RMSE at 2.70 (column 4). We conclude that not only nonlinearities do not help explain this early episode, but they can even harm the performance of the Phillips curve.

5.2. The Missing Disinflation and the Missing Inflation of the 2010s

In contrast, during the early 2010s both nonlinearities and consumer expectations can explain the missing disinflation.[34] One key piece of evidence comes from one-step-ahead out-of-sample forecasts of inflation. Panel A of Figure 2 compares the absolute forecast errors obtained from the linear model with backward-looking and SPF expectations (red dashed line) and the corresponding threshold model (green thin dash-dot line). It is well known that inflation forecasts of different models are usually similar. Still, during the missing disinflation period, the threshold model's forecast improves upon the linear model's forecast. For 2009:Q3, the linear model's forecast is improved by about half a percentage point. For the period 2010:Q4 through 2011:Q2, the threshold model produces quarterly forecasts that are again significantly closer to actual inflation. In 2012 and 2013, as the missing disinflation episode is near its end, the two forecasts converge to each other.

Panel A of Figure 2 also plots forecasts from the linear model with consumer expectations (orange dash-dot line) and from the baseline model with one threshold and consumer expectations (blue solid line). Adding consumer expectations to the model clear has a larger quantitative effect than allowing for nonlinearities. For example, in 2009:Q3 and 2010:Q4 through 2011:Q2 consumer expectations reduce absolute forecast error by up to 2.5 percentage points. But while adding nonlinearities to the model with consumer expectations improves the forecast further, the gain is relatively small. Hence, even though nonlinearities, overall, improve the fit of the model during this episode, their role is muted when expectations are properly measured.[35]

In the late 2010s, a *missing-inflation* puzzle emerged, providing yet another evidence on

[34]Figures B.5 and B.6 in the online appendix show that nonlinearities help explain the missing disinflation episode for the case of naive, backward-looking expectations and evaluate the role of the continuity constraint.

[35]In the online appendix, we also show in-sample fit (Tables B.8, B.9; Figure B.7) and out-of-sample dynamic forecasts (Figure B.8), employing the innovations obtained from Equation (5). These results support our conclusions.

Figure 2: Out-of-sample Forecasts

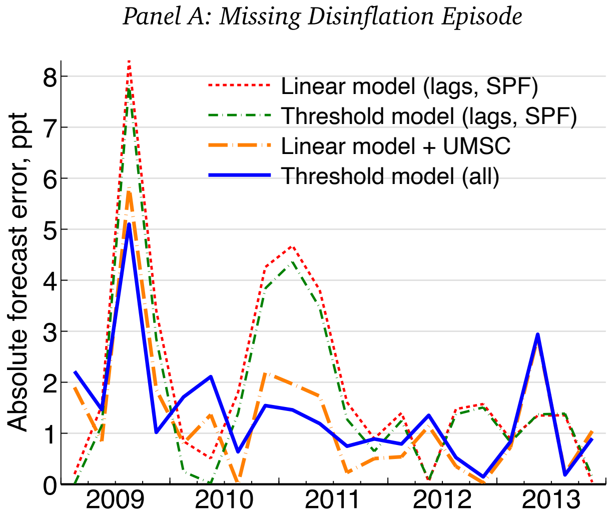

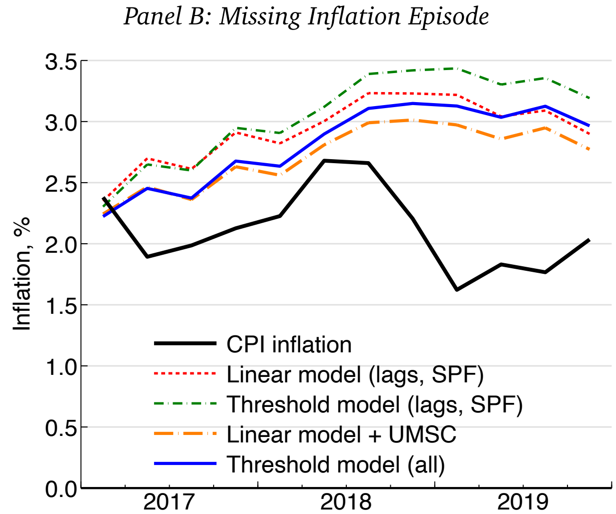

Notes: The left panel focuses on the missing disinflation episode and shows absolute out-of-sample forecast errors for the linear model with inflation lags and SPF (red dashed line); for the corresponding threshold model (green short dash–dot line); for the linear model with the lags, SPF, and UMSC inflation expectations (orange long dash–dot line); and for the threshold model with all expectations based on estimates in column (1) of Table 1 (blue thin line). The right panel shows actual out-of-sample forecasts as well as the actual, seasonally adjusted, annualized CPI inflation rates (thick black line). As quarterly changes in the headline CPI were volatile during the missing inflation period, for better visibility in Panel B we show averages of the annualized quarterly rates over the preceding four-quarter period.

nonlinearities and expectations. This episode featured very low unemployment rates and gaps, indicating possible overheating (CBO, 2019; Crump et al., 2019). In 2017:Q1, for example, the unemployment gap fell below zero and continued to decline until it reached its trough of nearly −1% in the second half of 2019. According to both the linear and nonlinear models, inflation should have picked up during this period. Instead, inflation remained stable and even declined by some measures. At the same time, consumers expected low inflation, which could bridge the gap between the models and the data. We find that consumer expectations improve the fit during this period; however, the effects are not large.

Panel B of Figure 2 shows the forecasts obtained from the four versions of the Phillips curve we considered earlier, as well as actual CPI inflation, during the period 2017–2019. All four models produce inflation dynamics similar to one another: that is, they forecast rising inflation. However, the threshold models diverge from the data more than the corresponding linear models, because the Phillips curve is relatively steeper in the low-unemployment regime. In contrast, the models with consumer expectations are closer to the observed inflation than the corresponding models that allow for only SPF and backward-looking expectations. The best forecasts are obtained from the linear model with consumer expectations (orange thick dash-dot line). Yet, the differences between the models' forecasts are small relative to the corresponding forecast errors. For instance, consumer expectations account for a 0.1 percentage point improvement in forecast accuracy, while the thresholds harm the forecasts by up to 0.2

percentage points. However, even the best forecasting model misses the data by 0.75–1.5 percentage point by 2019. Hence, from a Phillips curve prospective, the missing inflation remains a puzzle even when nonlinearities and various measures of expectations are considered.[36]

Thus, while allowing for a high-unemployment regime helps the Phillips curve match the data in the early 2010s (during the missing disinflation), it does not help in the 1970s (during the Great Inflation), in the 1980s (during the Volcker disinflation), and in the late 2010s (during the missing inflation). We also confirm this evidence formally, by testing nonlinearities in each of these episodes. The F-test rejects the linear model for the missing-disinflation period but favors it for the other episodes (Online Appendix C). Yet, consumers' inflation expectations play a crucial role in all these episodes. For this reason, when allowing for a Phillips curve with both consumer expectations and nonlinearities, the former are highly significant in the full sample, whereas the latter are not.

6. Robustness

6.1. Financial Frictions

The recent global financial crisis brought to the forefront financial frictions as a factor affecting economic fluctuations. Gilchrist, Yankov, and Zakrajšek (2009), Philippon (2009), and others emphasize the predictive content of corporate bond spreads for consumption, output, and inflation. Moreover, in a structural model, Christiano, Eichenbaum, and Trabandt (2015) and Del Negro, Giannoni, and Schorfheide (2015), among others, propose financial frictions as yet another explanation for the missing disinflation.[37] If financial frictions correlate with expectations or the sources of nonlinearities, controlling for them may affect our baseline estimates. Therefore, we test whether the state of financial markets overturns our results on nonlinearities and expectations.

In column (1) of Table 5, we control for the Baa–Aaa corporate bond spread, a popular measure used in the literature.[38] We find that this measure does not change the results qualitatively. In fact, the results strengthen (see Table 1, column 1). The coefficient on consumer expectations increases and remains statistically significant. Thresholds remain insignificant, and now even more so.

[36]For explanations of the missing inflation favoring changing trends in competition, see Heise, Karahan, and Şahin (2022).

[37]The two papers, however, provide alternative explanations. Christiano, Eichenbaum, and Trabandt (2015) emphasize a decrease in TFP growth and an increase in the cost of working capital, whereas Del Negro, Giannoni, and Schorfheide (2015) focus on the economic slack during the Great Recession.

[38]To reflect firms' financial costs, we add the federal funds rate to all spread measures. The qualitative results are unchanged if we employ the spread measures directly. The results for other measures of inflation are also qualitatively similar to the baseline and are reported in Table B.10 in Online Appendix B.

Table 5: Controlling for Credit Spreads

	Baa–Aaa Spread (1)	GZ Credit Spread (2)	Excess Bond Premium (3)	Fixed Sample (4)
	Panel A: Linear Model			
Slope, $\hat{\kappa}$	-0.21^{***}	-0.18^{**}	-0.16^{**}	-0.29^{***}
	(0.08)	(0.08)	(0.07)	(0.09)
	Panel B: Threshold Model			
Slopes				
left, $\hat{\kappa}_L$	-0.34^{**}	-2.01	-1.82	-3.57^{**}
	(0.15)	(1.41)	(1.24)	(1.45)
right, $\hat{\kappa}_R$	0.04	-0.15^{*}	-0.13^{*}	-0.22^{**}
	(0.23)	(0.08)	(0.08)	(0.09)
Expected inflation				
UMSC, $\hat{\alpha}_2$	1.08^{***}	1.09^{***}	1.15^{***}	0.83^{***}
	(0.20)	(0.21)	(0.22)	(0.20)
SPF, $\hat{\alpha}_1$	-0.10	-0.17	-0.20	-0.21
	(0.17)	(0.24)	(0.23)	(0.36)
Sum of lags, $\hat{\alpha}_0$	0.03	0.08	0.05	0.38^{*}
	(0.18)	(0.21)	(0.22)	(0.21)
Spread coefficient	0.22^{***}	0.21^{***}	0.22^{***}	
	(0.05)	(0.06)	(0.06)	
Threshold, $\hat{\gamma}$				
point estimate	1.95	-0.71	-0.71	-0.71
95% confidence interval	[−0.82, 2.93]	[−0.71, 3.14]	[−0.71, 3.14]	[−0.71, 2.30]
No. of thresholds, p-value				
0 vs. 1, H_0: 0	0.38	0.62	0.68	0.21
R^2	0.81	0.81	0.81	0.78
N	205	188	188	188

Notes: The slack variable is the unemployment gap. Inflation is measured with the CPI inflation. The baseline estimation sample (column 1) is 1968:Q4 through 2019:Q4; the sample in columns (2)–(4) is 1973:Q1 through 2019:Q4. For estimation details, see the notes to Table 1. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

In column (2) of Table 5, we control for the Gilchrist and Zakrajšek (2012) credit spread, a measure based on corporate bonds traded in the secondary market and shown to be highly informative about economic activity. In addition to the component measuring countercyclical movements in expected defaults—which is similar to the Baa–Aaa corporate bond spread—this variable captures the excess bond premium, measuring changes in the relationship between default risk and credit spreads. In column (3), we estimate the specification with the excess bond premium.

Similar to the results with the Baa–Aaa spread, these financial variables strengthen our findings. This can be seen by comparing the results in columns (2) and (3) with those in column (4), which presents estimates of our baseline specification for the same sample period.[39] Con-

[39]The credit-spread data are available starting in 1973. Relative to the baseline sample, therefore, we lose more than four years of observations with high inflation. The results in column (1) can be directly compared to the baseline (Table 1, column 1).

Table 6: Robustness to Structural Break: Post-1990 Period

	Unemployment		Labor Share	
	Gap	Rate	Raw	Adjusted
	(1)	(2)	(3)	(4)
Slopes				
left, $\hat{\kappa}_L$	-0.64^{**}	-0.37^{**}	0.17^{**}	0.72^{***}
	(0.25)	(0.17)	(0.07)	(0.17)
right, $\hat{\kappa}_R$	0.24	0.14	0.59^{***}	0.22^{*}
	(0.33)	(0.46)	(0.21)	(0.12)
Expected inflation				
UMSC, $\hat{\alpha}_2$	1.35^{***}	1.24^{***}	1.46^{***}	1.59^{***}
	(0.34)	(0.34)	(0.34)	(0.23)
SPF, $\hat{\alpha}_1$	-0.51	-0.52	0.77^{**}	-0.08
	(0.33)	(0.33)	(0.37)	(0.35)
Sum of lags, $\hat{\alpha}_0$	0.16	0.28	-1.23^{**}	-0.51
	(0.36)	(0.36)	(0.49)	(0.42)
Threshold, $\hat{\gamma}$				
point estimate	2.50	8.03	-0.80	-2.77
95% confidence interval	[−0.59, 3.14]	[4.07, 8.27]	[−7.75, 1.46]	[−4.95, 0.32]
No. of thresholds, p-value.				
0 vs. 1, H_0: 0	0.12	0.65	0.13	0.22
R^2	0.47	0.43	0.51	0.49
N	120	120	120	120

Notes: The estimation sample is 1990:Q1 through 2019:Q4. For estimation details, see the notes to Table 1. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

sumer expectations remain significant and their relative weight increases. The linear model statistically dominates the nonlinear one. The presence of thresholds was statistically insignificant, and now becomes even more so. In addition, the presence of nonlinearities would still have to be confronted with other popular explanations of the missing inflation, such as smaller movements in short-term unemployment (Ball and Mazumder, 2019) or in expected future marginal costs (Del Negro, Giannoni, and Schorfheide, 2015).

6.2. Structural Break

To provide more evidence that the regimes identified in the threshold model are not due to a structural break in the linear model, we estimate our baseline specification when the sample is split about a possible break period considered in the literature. While the exact reasons for the flattening of the Phillips curve are debated, inflation overall was more stable and less responsive to unemployment starting in the early 1990s than in the previous decades. We therefore split our sample about 1990:Q1.[40]

Table 6 shows estimates of the threshold model in the post-1990 period.[41] They are overall

[40] This choice is consistent with procedures based on optimization over possible break points.

[41] We focus on the later subsample because it is more relevant for the contemporaneous Phillips curve. Table B.11

consistent with the baseline estimates. The threshold is estimated at an unemployment gap of 2.5% (column 1), a little higher than the baseline. The threshold is (marginally) insignificant at the 10% level, as in the baseline case. The Phillips curve appears relatively steep in the low-unemployment regime, and flat in the high-unemployment regime. As in the baseline, consumer expectations dominate other expectation components. The results are also similar to the baseline for other measures of slack (columns 2–4). Thus, we conclude that our findings are not driven by a structural change in the Phillips curve slope and they remain materially unchanged when we focus on the recent period.

7. Conclusion

In this paper, we examine the presence of nonlinearities in the Phillips curve. It turns out that it is important to examine jointly the roles of consumer expectations and nonlinearities in inflation dynamics. While each factor is important to some degree, we find that consumer expectations are a robust and salient feature of the data, whereas nonlinearities are muted overall and important only in some specific cases. Since we find a more prominent role for consumer expectations, it is important that our approach does not limit the source and the form of nonlinearities. This is particularly relevant as recessions can have countervailing effects on the slope of the Phillips curve, producing intricate functional forms. Yet, even when we employ a flexible estimation technique, which can capture nonlinearity stemming from any structural parameter, we find a prominent role for consumer expectations and only a muted, episodic role for nonlinearities.

To understand these findings, our paper examines in detail several episodes wherein inflation was especially hard to predict: the Great Inflation of the 1970s, the Volcker disinflation, the missing disinflation following the 2008–09 financial crisis, and the missing inflation of the late 2010s. The missing disinflation can be explained with either nonlinearities or consumer expectations. The stability of consumer expectations in this period provided an anchor for inflation, and a flatter Phillips curve helped stabilize inflation as well.

In contrast, the Volcker disinflation provides key evidence in favor of consumer expectations. This period was characterized by high unemployment and a recession similar to the missing-disinflation episode. During the Volcker disinflation, however, inflation was reduced; therefore, the same type of nonlinearity that helps explain the missing disinflation (through a flatter Phillips curve) makes the Volcker disinflation too slow. In contrast, consumer expectations improve the fit of the model and help match the data during the Volcker disinflation. The Great Inflation in the 1970s is yet another episode that sheds light on this important issue.

in Online Appendix B shows estimates for the pre-1990 period.

This episode brings additional evidence because in the 1970s, unlike in the other two episodes, economic activity was not subdued. In this episode, too, consumer expectations help the model match the data.

A more recent episode of the *missing-inflation*, which emerged around 2017–2019, also favors the expectations channel, since expectations were anchored at the time and inflation did not pick up. However, despite consumer inflation expectations' trending downward during this period, expected inflation did not decline enough to ensure a good fit of the Phillips curve. Hence, from the Phillips curve perspective, the missing inflation episode remains a puzzle, even when nonlinearities as well as consumer expectations are included in the model.

References

Adam, Klaus, and Mario Padula. 2011. "Inflation Dynamics and Subjective Expectations in the United States." *Economic Inquiry* 49 (1): 13–25.

Alvarez, Fernando, Martin Beraja, Martín Gonzalez-Rozada, and Pablo A. Neumeyer. 2018. "From Hyperinflation to Stable Prices: Argentina's Evidence on Menu Cost Models." *Quarterly Journal of Economics* 134 (1): 451–505.

Andrade, Philippe, and Hervé Le Bihan. 2013. "Inattentive Professional Forecasters." *Journal of Monetary Economics* 60 (8): 967–982.

Armenter, Roc. 2015. "A Bit of a Miracle No More: The Decline of the Labor Share." *FRB Philadelphia Business Review* 2015 (Q3): 1–9.

Ascari, Guido, Paolo Bonomolo, and Qazi Haque. 2022. "The LONG Run Phillips Curve is ... a Curve." Manuscript.

Auerbach, Alan J., and Yuriy Gorodnichenko. 2012. "Measuring the Output Responses to Fiscal Policy." *American Economic Journal: Economic Policy* 4 (2): 1–27.

Babb, Nathan R., and Alan K. Detmeister. 2017. "Nonlinearities in the Phillips Curve for the United States: Evidence Using Metropolitan Data." Finance and Economics Discussion Series 2017-070, Board of Governors of the Federal Reserve System.

Ball, Laurence M. 1997. "Disinflation and the NAIRU." In *Reducing Inflation: Motivation and Strategy*, edited by Christina D. Romer and David H. Romer, 167–194. Chicago: University of Chicago Press.

Ball, Laurence M., and Sandeep Mazumder. 2011. "Inflation Dynamics and the Great Recession." *Brookings Papers on Economic Activity* 42 (1): 337–405.

———. 2019. "A Phillips Curve with Anchored Expectations and Short-Term Unemployment." *Journal of Money, Credit and Banking* 51 (1): 111–137.

———. 2020. “The Nonpuzzling Behavior of Median Inflation.” In *Changing Inflation Dynamics, Evolving Monetary Policy*, edited by Gonzalo Castex, Jordi Galí, and Diego Saravia, 49–70. Santiago: Central Bank of Chile.

Barnes, Michelle L., and Giovanni P. Olivei. 2003. “Inside and Outside Bounds: Threshold Estimates of the Phillips Curve.” *New England Economic Review* 2003: 3–18.

Barnichon, Regis, Davide Debortoli, and Christian Matthes. 2022. “Understanding the Size of the Government Spending Multiplier: It’s in the Sign.” *Review of Economic Studies* 89 (1): 87–117.

Barnichon, Regis, and Geert Mesters. 2020. “Identifying Modern Macro Equations with Old Shocks.” *Quarterly Journal of Economics* 135 (4): 2255–2298.

Barsky, Robert B., and Lutz Kilian. 2002. “Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative.” In *NBER Macroeconomics Annual*, vol. 16, edited by Ben S. Bernanke and Kenneth Rogoff, 137–198. Cambridge, Mass.: MIT Press.

Beaudry, Paul, and Franck Portier. 2018. “Real Keynesian Models and Sticky Prices.” Working Paper No. 24223, National Bureau of Economic Research.

Berge, Travis, Maarten De Ridder, and Damjan Pfajfar. 2021. “When Is the Fiscal Multiplier High? A Comparison of Four Business Cycle Phases.” *European Economic Review* 138: 103852.

Binder, Carola C. 2017. “Measuring Uncertainty Based on Rounding: New Method and Application to Inflation Expectations.” *Journal of Monetary Economics* 90: 1–12.

Blanchard, Olivier J. 1984. “The Lucas Critique and the Volcker Deflation.” *American Economic Review* 74 (2): 211–215.

Blinder, Alan S., and Jeremy B. Rudd. 2013. “The Supply-Shock Explanation of the Great Stagflation Revisited.” In *The Great Inflation: The Rebirth of Modern Central Banking*, edited by Michael D. Bordo and Athanasios Orphanides, 119–175. Chicago: University of Chicago Press.

Bolhuis, Marijn A., Judd N. L. Cramer, and Lawrence H. Summers. 2022. “Comparing Past and Present Inflation.” Working Paper No. 30116, National Bureau of Economic Research.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. 2020. “Overreaction in Macroeconomic Expectations.” *American Economic Review* 110 (9): 2748–2482.

Bordo, Michael D., Christopher J. Erceg, Andrew T. Levin, and Ryan Michaels. 2007. “Three Great American Disinflations.” In *Monetary Policy, Transparency, and Credibility*, proceedings of a Federal Reserve Bank of San Francisco conference. Conference draft available at frbsf.org/economic-research/events/2007/march/monetary-policy-transparency-credibility/levin_bordo_etal.pdf.

Brunnermeier, Markus K., and Yann Koby. 2018. "The Reversal Interest Rate." Working Paper No. 25406, National Bureau of Economic Research.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics* 90 (3): 414–427.

Chan, Joshua C.C., Todd E. Clark, and Gary Koop. 2018. "A New Model of Inflation, Trend Inflation, and Long-Run Inflation Expectations." *Journal of Money, Credit and Banking* 50 (1): 5–53.

Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt. 2015. "Understanding the Great Recession." *American Economic Journal: Macroeconomics* 7 (1): 110–167.

Clarida, Richard, Jordi Galí, and Mark Gertler. 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *Quarterly Journal of Economics* 115 (1): 147–180.

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.

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.

Coibion, Olivier, Yuriy Gorodnichenko, and Rupal Kamdar. 2018. "The Formation of Expectations, Inflation, and the Phillips Curve." *Journal of Economic Literature* 56 (4): 1447–1491.

Congressional Budget Office (CBO). 2019. "The Budget and Economic Outlook: 2019 to 2029." Report available at cbo.gov/publication/54918.

Crump, Richard K., Stefano Eusepi, Marc Giannoni, and Ayşegül Şahin. 2019. "A Unified Approach to Measuring u*." *Brookings Papers on Economic Activity* 50 (1): 143–214.

Daly, Mary C., and Bart Hobijn. 2014. "Downward Nominal Wage Rigidities Bend the Phillips Curve." *Journal of Money, Credit and Banking* 46 (S2): 51–93.

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.

Del Negro, Marco, Marc P. Giannoni, and Frank Schorfheide. 2015. "Inflation in the Great Recession and New Keynesian Models." *American Economic Journal: Macroeconomics* 7 (1): 168–196.

Del Negro, Marco, Michele Lenza, Giorgio E. Primiceri, and Andrea Tambalotti. 2020. "What's Up with the Phillips Curve?" *Brookings Papers on Economic Activity* : 301–357.

Del Negro, Marco, and Frank Schorfheide. 2008. "Forming Priors for DSGE Models (and How It Affects the Assessment of Nominal Rigidities)." *Journal of Monetary Economics* 55 (7): 1191–1208.

Ehrmann, Michael, Damjan Pfajfar, and Emiliano Santoro. 2017. "Consumers' Attitudes and Their Inflation Expectations." *International Journal of Central Banking* 13 (1): 225–259.

Erceg, Christopher J., and Andrew T. Levin. 2003. "Imperfect Credibility and Inflation Persistence." *Journal of Monetary Economics* 50 (4): 915–944.

Fuhrer, Jeffrey C. 2006. "Intrinsic and Inherited Inflation Persistence." *International Journal of Central Banking* 2 (3): 49–86.

———. 2010. "Inflation Persistence." In *Handbook of Monetary Economics*, vol. 3, edited by Benjamin M. Friedman and Michael Woodford, 423–486. Amsterdam: Elsevier.

———. 2018. "Intrinsic Expectations Persistence: Evidence from Professional and Household Survey Expectations." Working Paper No. 18-9, Federal Reserve Bank of Boston.

Fuhrer, Jeffrey C., and George R. Moore. 1995. "Inflation Persistence." *Quarterly Journal of Economics* 110 (1): 127–159.

Gagnon, Etienne. 2009. "Price Setting During Low and High Inflation: Evidence from Mexico." *Quarterly Journal of Economics* 124 (3): 1221–1263.

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

Gerko, Elena. 2018. "Expectations and Monetary Policy." Working Paper.

Gilchrist, Simon, Vladimir Yankov, and Egon Zakrajšek. 2009. "Credit Market Shocks and Economic Fluctuations: Evidence from Corporate Bond and Stock Markets." *Journal of Monetary Economics* 56 (4): 471–493.

Gilchrist, Simon, and Egon Zakrajšek. 2012. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review* 102 (4): 1692–1720.

Goodfriend, Marvin, and Robert G. King. 2005. "The Incredible Volcker Disinflation." *Journal of Monetary Economics* 52 (5): 981–1015.

Gordon, Robert J. 1982. "Price Inertia and Policy Ineffectiveness in the United States, 1890–1980." *Journal of Political Economy* 90 (6): 1087–1117.

Hansen, Bruce E. 1996. "Inference When a Nuisance Parameter Is Not Identified Under the Null Hypothesis." *Econometrica* 64 (2): 413–430.

———. 2000. "Sample Splitting and Threshold Estimation." *Econometrica* 68 (3): 575–604.

———. 2017. "Regression Kink with an Unknown Threshold." *Journal of Business & Economic Statistics* 35 (2): 228–240.

Harding, Martín, Jesper Lindé, and Mathias Trabandt. 2022. "Resolving the Missing Deflation Puzzle." *Journal of Monetary Economics* 126 (C): 15–34.

Hazell, Jonathon, Juan Herreño, Emi Nakamura, and Jón Steinsson. 2022. "The Slope of the Phillips Curve: Evidence from U.S. States." *Quarterly Journal of Economics* 137 (3): 1299–1344.

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.

Hooper, Peter, Frederic S. Mishkin, and Amir Sufi. 2020. "Prospects for Inflation in a High Pressure Economy: Is the Phillips Curve Dead or Is It Just Hibernating?" *Research in Economics* 74 (1): 26–62.

Kilian, Lutz, and Robert J. Vigfusson. 2011. "Nonlinearities in the Oil Price–Output Relationship." *Macroeconomic Dynamics* 15 (S3): 337–363.

Kimball, Miles S. 1995. "The Quantitative Analytics of the Basic Neomonetarist Model." *Journal of Money, Credit and Banking* 27 (4): 1241–1277.

Kremer, Stephanie, Alexander Bick, and Dieter Nautz. 2013. "Inflation and Growth: New Evidence from a Dynamic Panel Threshold Analysis." *Empirical Economics* 44 (2): 861–878.

Kumar, Anil, and Pia M. Orrenius. 2016. "A Closer Look at the Phillips Curve Using State-Level Data." *Journal of Macroeconomics* 47 (A): 84–102.

Kumhof, Michael, and Xuan Wang. 2021. "Banks, Money, and the Zero Lower Bound on Deposit Rates." *Journal of Economic Dynamics and Control* 132 (C): 104208.

Leduc, Sylvain, Keith Sill, and Tom Stark. 2007. "Self-fulfilling Expectations and the Inflation of the 1970s: Evidence from the Livingston Survey." *Journal of Monetary Economics* 54 (2): 433–459.

Luengo-Prado, María José, Nikhil Rao, and Viacheslav Sheremirov. 2018. "Sectoral Inflation and the Phillips Curve: What Has Changed since the Great Recession?" *Economics Letters* 172: 63–68.

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers. 2004. "Disagreement About Inflation Expectations." In *NBER Macroeconomics Annual*, vol. 18, edited by Mark Gertler and Kenneth Rogoff, 209–270. Cambridge, Mass.: MIT Press.

Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H. Stock. 2014. "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve." *Journal of Economic Literature* 52 (1): 124–188.

McLeay, Michael, and Silvana Tenreyro. 2019. "Optimal Inflation and the Identification of the Phillips Curve." In *NBER Macroeconomics Annual*, vol. 34, edited by Martin S. Eichenbaum, Erik Hurst, and Jonathan A. Parker, 199–255. Chicago: University of Chicago Press.

Mertens, Elmar. 2016. “Managing Beliefs about Monetary Policy under Discretion.” *Journal of Money, Credit and Banking* 48 (4): 661–698.

Nason, James M., and Gregor W. Smith. 2008. “Identifying the New Keynesian Phillips Curve.” *Journal of Applied Econometrics* 23 (5): 525–551.

Nunes, Ricardo. 2009. “Learning the Inflation Target.” *Macroeconomic Dynamics* 13 (2): 167–188.

———. 2010. “Inflation Dynamics: The Role of Expectations.” *Journal of Money, Credit and Banking* 42 (6): 1161–1172.

Orphanides, Athanasios. 2001. “Monetary Policy Rules Based on Real-Time Data.” *American Economic Review* 91 (4): 964–985.

Petrella, Ivan, Emiliano Santoro, and Lasse de la Porte Simonsen. 2018. “Time-varying Price Flexibility and Inflation Dynamics.” Working Paper. Available at econ.ku.dk/esantoro/images/PSS_Prices_9Nov.pdf.

Pfajfar, Damjan, and John M. Roberts. 2022. “The Role of Expectations in Changed Inflation Dynamics.” *International Journal of Central Banking* 18 (1): 199–238.

Pfajfar, Damjan, and Emiliano Santoro. 2013. “News on Inflation and the Epidemiology of Inflation Expectations.” *Journal of Money, Credit and Banking* 45 (6): 1045–1067.

Philippon, Thomas. 2009. “The Bond Market’s q.” *Quarterly Journal of Economics* 124 (3): 1011–1056.

Phillips, A.W. 1958. “The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957.” *Economica* 25 (100): 283–299.

Primiceri, Giorgio E. 2006. “Why Inflation Rose and Fell: Policy-Makers’ Beliefs and U. S. Post-war Stabilization Policy.” *Quarterly Journal of Economics* 121 (3): 867–901.

Roberts, John M. 2006. “Monetary Policy and Inflation Dynamics.” *International Journal of Central Banking* 2 (3): 193–230.

Romer, Christina D., and David H. Romer. 2004. “A New Measure of Monetary Shocks: Derivation and Implications.” *American Economic Review* 94 (4): 1055–1084.

Rotemberg, Julio J., and Michael Woodford. 1997. “An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy.” In *NBER Macroeconomics Annual*, vol. 12, edited by Ben S. Bernanke and Julio J. Rotemberg, 297–361. Cambridge, Mass.: MIT Press.

Santoro, Emiliano, Ivan Petrella, Damjan Pfajfar, and Edoardo Gaffeo. 2014. “Loss Aversion and the Asymmetric Transmission of Monetary Policy.” *Journal of Monetary Economics* 68: 19–36.

Schaling, Eric. 2004. "The Nonlinear Phillips Curve and Inflation Forecast Targeting: Symmetric Versus Asymmetric Monetary Policy Rules." *Journal of Money, Credit and Banking* 36 (3): 361–386.

Simon, John, Troy Matheson, and Damiano Sandri. 2013. "The Dog That Didn't Bark: Has Inflation Been Muzzled or Was It Just Sleeping?" In *World Economic Outlook: Hopes, Realities, Risks*, 79–96. Washington: International Monetary Fund.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (3): 586–606.

Stewart, Kenneth J., and Stephen B. Reed. 2000. "Data Update Consumer Price Index Research Series Using Current Methods (CPI-U-RS), 1978–1998." *Industrial Relations: A Journal of Economy and Society* 39 (1): 161–164.

Stock, James H., and Mark W. Watson. 2020. "Slack and Cyclically Sensitive Inflation." *Journal of Money, Credit and Banking* 52: 393–428.

Tenreyro, Silvana, and Gregory Thwaites. 2016. "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions." *American Economic Journal: Macroeconomics* 8 (4): 43–74.

Wieland, Johannes F., and Mu-Jeung Yang. 2020. "Financial Dampening." *Journal of Money, Credit and Banking* 52 (1): 79–113.

Wu, Jing Cynthia, and Fan Dora Xia. 2016. "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound." *Journal of Money, Credit and Banking* 48 (2-3): 253–291.

Yellen, Janet L. 2015. "Inflation Dynamics and Monetary Policy." Philip Gamble Memorial Lecture, University of Massachusetts, Amherst.