

# THE QUARTERLY JOURNAL OF ECONOMICS

---

Vol. 137

2022

Issue 3

---

## THE SLOPE OF THE PHILLIPS CURVE: EVIDENCE FROM U.S. STATES\*

JONATHON HAZELL  
JUAN HERREÑO  
EMI NAKAMURA  
JÓN STEINSSON

We estimate the slope of the Phillips curve in the cross section of U.S. states using newly constructed state-level price indices for nontradeable goods back to 1978. Our estimates indicate that the slope of the Phillips curve is small and was small even during the early 1980s. We estimate only a modest decline in the slope of the Phillips curve since the 1980s. We use a multiregion model to infer the slope of the aggregate Phillips curve from our regional estimates. Applying our estimates to recent unemployment dynamics yields essentially no missing disinflation or missing reflation over the past few business cycles. Our results imply that the sharp drop in core inflation in the early 1980s was mostly due to shifting expectations about long-run monetary policy as opposed to a steep Phillips curve, and the greater stability of inflation between 1990 and 2020 is mostly due to long-run inflation expectations becoming more firmly anchored. *JEL Codes:* E30.

\*We thank Tomas Breach, Massimiliano Cologgi, and Victoria de Quadros for excellent research assistance. We thank John Molino, our project coordinator at the BLS, for all his help throughout the project. We thank Hassan Afrouzi, Regis Barnichon, Olivier Blanchard, Chi-Young Choi, Jordi Galí, Yuriy Gorodnichenko, Pierre-Olivier Gourinchas, Greg Kaplan, Chen Lian, Silvana Tenreiro, John Williams, and seminar participants at various institutions for valuable comments and discussions. We thank the National Science Foundation, the Alfred P. Sloan Foundation, the Smith Richardson Foundation, the Institute for New Economic Thinking, and the Julis-Rabinowitz Center for Public Policy and Finance for financial support.

© The Author(s) 2022. Published by Oxford University Press on behalf of the President and Fellows of Harvard College. All rights reserved. For Permissions, please email: [journals.permissions@oup.com](mailto:journals.permissions@oup.com)

*The Quarterly Journal of Economics* (2022), 1299–1344. <https://doi.org/10.1093/qje/qjac010>. Advance Access publication on February 16, 2022.

## I. INTRODUCTION

The Phillips curve is a formal statement of the common intuition that if demand is high in a booming economy, this will provoke workers to seek higher wages and firms to raise prices. A well-known formulation is the New Keynesian Phillips curve:

$$(1) \quad \pi_t = \beta E_t \pi_{t+1} - \kappa (u_t - u_t^n) + v_t.$$

According to this formulation, inflation  $\pi_t$  is determined by three factors: expected inflation  $E_t \pi_{t+1}$ , the output gap—measured here as the difference between unemployment  $u_t$  and the natural rate of unemployment  $u_t^n$ —and cost-push shocks  $v_t$ . The slope of the Phillips curve  $\kappa$  represents the sensitivity of inflation to the output gap (i.e., to an increase in demand).

The episode in U.S. economic history that has perhaps most strongly influenced the profession's thinking regarding the slope of the Phillips curve is the Volcker disinflation. In the early 1980s, Paul Volcker's Federal Reserve sharply tightened monetary policy. Unemployment rose sharply and inflation fell sharply. The conventional interpretation of this episode is that it provides evidence for a relatively steep Phillips curve.

One way to formalize this conventional interpretation is to assume that inflation expectations are adaptive:  $\beta E_t \pi_{t+1} = \pi_{t-1}$  in equation (1). This yields the accelerationist Phillips curve:

$$(2) \quad \Delta \pi_t = -\kappa (u_t - u_t^n) + v_t.$$

[Stock and Watson \(2019\)](#) estimate  $\kappa$  in this equation and refer to it as the “Phillips correlation.” They measure  $\Delta \pi_t$  by the annual change in 12-month core personal consumption expenditures (PCE) inflation, and  $u_t - u_t^n$  by the Congressional Budget Office (CBO) unemployment gap, both at a quarterly frequency. [Figure I](#) reproduces this analysis. It suggests that the slope of the Phillips curve was steep prior to and during the Volcker disinflation (0.67 for the period 1960–1983) but has flattened considerably since then (to only 0.03 for the period 2000–2019q1).<sup>1</sup>

The insensitivity of inflation to changes in unemployment between 1990 and 2020 led many economists to suggest that the Phillips curve had disappeared—or was “hibernating.” During the Great Recession, unemployment rose to levels comparable to those during the Volcker disinflation, yet inflation fell by

1. See also [Ball and Mazumder \(2011\)](#), [Kiley \(2015b\)](#), and [Blanchard \(2016\)](#).

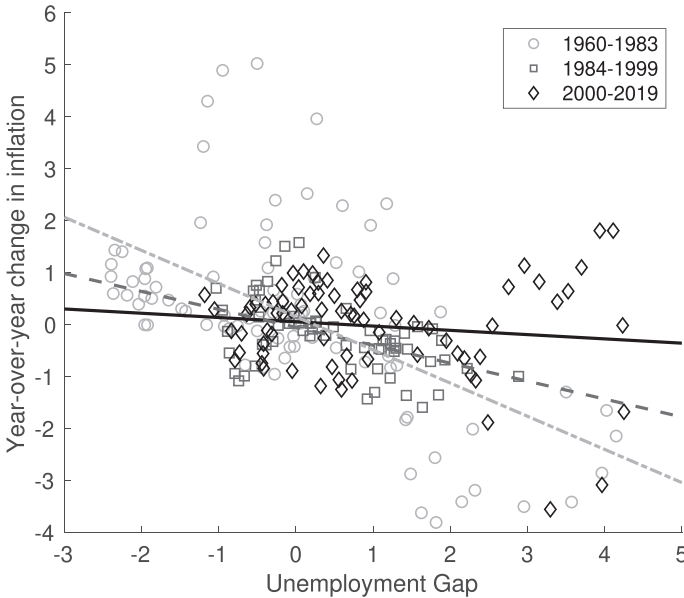


FIGURE I

## Stock and Watson's Changing Phillips Correlation

The black solid line is a regression line for 2000–2019. The dark gray broken line is a regression for 1984–1999. The light gray dash-dot line is a regression line for 1960–1983. The year-over-year change in inflation is the four-quarter change in the (backward-looking) four-quarter moving average of headline PCE inflation. The unemployment gap is the four-quarter (backward-looking) moving average of the gap between the unemployment rate and the natural rate of unemployment. Authors' calculations. The figure replicates figure 1 from [Stock and Watson \(2019\)](#).

much less. The “missing disinflation” during and after the Great Recession then gave way to “missing reflation” in the late 2010s as unemployment fell to levels not seen in 50 years but inflation inched up only slightly. A similar debate raged in the late 1990s, when unemployment was also very low without this leading to much of a rise in inflation. Some have argued that the apparent flattening of the Phillips curve signals an important flaw in the Keynesian model.

There is an alternative interpretation of these facts that emphasizes the anchoring of long-term inflation expectations in the United States ([Bernanke 2007](#); [Mishkin 2007](#)). [Figure II](#) plots long-term inflation expectations from the Survey of Professional Forecasters. During the 1980s, long-term inflation expectations fluctuated a great deal. In particular, they fell rapidly over the

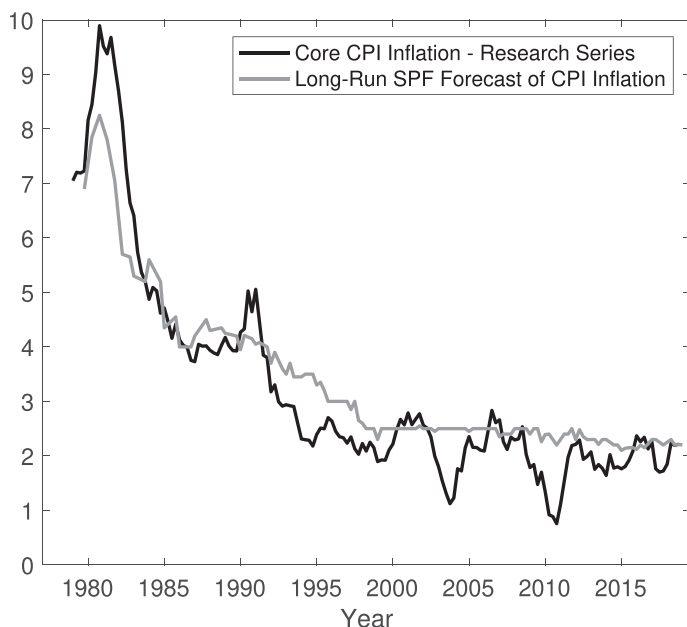


FIGURE II

## PCE Core Inflation and Long-Term Inflation Expectations

The gray line plots 10-year-ahead inflation expectations for the Consumer Price Index (CPI). From 1990 onward, these come from the Survey of Professional Forecasters. For the 1980s, these come from Blue Chip and are available on the Research and Data site of the Federal Reserve Bank of Philadelphia. The black line plots 12-month core CPI inflation using the Bureau of Labor Statistics' research series. This research series uses current methods to calculate inflation back in time.

period of the Volcker disinflation. In sharp contrast, since 1998, long-term inflation expectations have been extremely stable.

An alternative to the standard narrative of the Volcker disinflation is that the decline in inflation was driven not by a steep Phillips curve but by shifts in beliefs about the long-run monetary regime in the United States that caused the rapid fall in long-run inflation expectations we observe in Figure II. To see how this can be the case, it is useful to solve equation (1) forward and assume for simplicity that unemployment follows an AR(1) process. This yields

$$(3) \quad \pi_t = -\psi \tilde{u}_t + E_t \pi_{t+\infty} + \omega_t,$$

where  $\tilde{u}_t$  denotes the deviation of unemployment from its long-run expected value,  $E_t \pi_{t+\infty}$  represents long-term inflation

expectations, and the parameter  $\psi$  is proportional to  $\kappa$  in [equation \(1\)](#). ([Section II](#) presents a more detailed derivation.) This formulation of the Phillips curve makes clear that changes in beliefs about the long-run monetary regime feed strongly into current inflation: the coefficient on  $E_t\pi_{t+\infty}$  in [equation \(3\)](#) is one. Furthermore, in the presence of substantial variation in  $E_t\pi_{t+\infty}$ , the relationship between  $\pi_t$  and  $\tilde{u}_t$  may be essentially uninformative about the slope of the Phillips curve ( $\psi$  and  $\kappa$ ). In particular, if changes in  $E_t\pi_{t+\infty}$  comove negatively with  $\tilde{u}_t$  (as they would during an imperfectly credible shift in the long-run inflation target) the Phillips curve would appear to be steeper than it actually was.

[Sargent \(1982\)](#) emphasizes that hyperinflations tend to end quickly—much too quickly to be explained by even a very large value of  $\kappa$  in the Phillips curve. In these episodes, it is clear that the primary cause of the abrupt fall in inflation is an abrupt fall in  $E_t\pi_{t+\infty}$  associated with an abrupt change in the policy regime. Volcker’s monetary policy constituted a sharp regime shift that was imperfectly credible at the outset but became gradually more credible as time passed ([Erceg and Levin 2003](#); [Goodfriend and King 2005](#); [Bianchi and Ilut 2017](#)). This regime shift led to a large and sustained decline in long-term inflation expectations over the 1980s but also a transitory rise in unemployment. Perhaps it was this large change in inflation expectations that was the primary cause of the rapid fall in inflation over this period rather than high unemployment working through a steep Phillips curve.

This discussion highlights an important identification problem researchers face when they seek to estimate the slope of the Phillips curve: inflation expectations may covary with the output gap. Standard methods for estimating the Phillips curve aim to address this issue by controlling for inflation expectations  $E_t\pi_{t+1}$  when estimating [equation \(1\)](#). A challenge with this approach is that estimates are quite sensitive to details of the specification. [Mavroeidis, Plagborg-Møller, and Stock \(2014\)](#) show that reasonable variation in the choice of data series, the specification, and the time period used yield a wide range of estimates for  $\kappa$  roughly centered on a value of zero (i.e., they are equally likely to have the “right” as the “wrong” sign). [Mavroeidis, Plagborg-Møller, and Stock \(2014, 124\)](#) point to a weak-instruments problem in driving these results: there simply isn’t enough variation in the aggregate data to separately identify the coefficients on unemployment and expected inflation. They conclude: “the literature has reached a limit on how much can be learned about the New Keynesian Phillips curve from aggregate macroeconomic time series. New

identification approaches and new datasets are needed to reach an empirical consensus.”

In addition to the identification problem discussed above, researchers seeking to estimate the slope of the Phillips curve face the classic simultaneity problem of distinguishing demand shocks from supply shocks. Supply shocks ( $u_t^n$  and  $v_t$ ) yield positive comovement of inflation and unemployment (stagflation). If the variation used to identify the slope of the Phillips curve is contaminated by such shocks, the estimated slope will be biased toward zero and may even have the “wrong” sign. [Fitzgerald and Nicolini \(2014\)](#) and [McLeay and Tenreyro \(2019\)](#) point out that a central bank conducting optimal monetary policy will seek to offset aggregate demand shocks. If the central bank is successful, the remaining variation in inflation will be only due to supply shocks, a worstcase scenario for the simultaneity problem.

Can cross-sectional data help overcome these problems? Several recent papers have argued that they can. [Fitzgerald and Nicolini \(2014\)](#) and [McLeay and Tenreyro \(2019\)](#) show that using regional data helps overcome the simultaneity problem of distinguishing demand and supply shocks: central banks cannot offset regional demand shocks using a single national interest rate. These papers as well as [Kiley \(2015a\)](#), [Babb and Detmeister \(2017\)](#), [Hooper, Mishkin, and Sufi \(2019\)](#), and [Fitzgerald et al. \(2020\)](#) make use of city-level inflation data produced by the Bureau of Labor Statistic (BLS) to estimate regional Phillips curves. [Beraja, Hurst, and Ospina \(2019\)](#) use regional wage data to estimate wage Phillips curves.

We contribute to this regional Phillips curve literature in several ways. First, we show formally how estimating the Phillips curve using regional data provides a solution to the problem of shifting values of  $E_t\pi_{t+\infty}$  confounding the estimation of the slope of the Phillips curve. We derive a regional Phillips curve in a simple benchmark multiregion model of a monetary union. The model clarifies the interpretation of the slope of regional Phillips curves relative to that of the aggregate Phillips curve. We also use the model to show that changes in the long-run monetary regime are absorbed by time fixed effects when the regional Phillips curve is estimated using a panel data specification. The intuition is that such long-run regime changes are common to all regions and therefore cancel out across regions in the monetary union.

Using our cross-section specification, we estimate a modest flattening of the Phillips curve when we split our sample in 1990: the Phillips curve in the post-1990 sample is flatter by a

factor of two. This contrasts sharply with empirical specifications that make use of time series variation: a specification without time fixed effects yields a 50–100 times steeper Phillips curve for the pre-1990 sample. We interpret this as evidence that shifting long-run inflation expectations seriously confound estimates of the Phillips curve based on time series variation in the pre-1990 sample.

Our cross-sectional estimates indicate that the slope of the Phillips curve is small and was small even during the 1980s. Combining our estimate of the slope of the Phillips curve with an estimate of the persistence of fluctuations in unemployment, we find that a 1 percentage point increase in unemployment reduces inflation by about 0.34 percentage points, that is,  $\psi = 0.34$  in [equation \(3\)](#). This implies that only a modest fraction of the large changes in inflation in the early 1980s can be accounted for by the direct effect of increasing unemployment working through the slope of the Phillips curve. In contrast, movements in long-run inflation expectations were large over this period, as is evident from [Figure II](#). In particular, long-run inflation expectations fell by about 4 percentage points from 1981 to 1986, accounting for about two-thirds of the fall in core inflation during this period. We conclude that a majority of the rapid decline in core inflation during the Volcker disinflation arose from a rapid decline of long-term inflation expectations, associated with a rapidly changing monetary regime.<sup>2</sup>

Our estimates of the slope of the Phillips curve imply essentially no missing disinflation during the Great Recession or missing re-inflation in the late 2010s or late 1990s. In other words, our cross-sectional estimates are consistent with the magnitude of movements in aggregate inflation after 1990. We conclude that the stability of inflation since 1990 is due to long-run inflation expectations becoming more firmly anchored. These conclusions echo those of [Jorgensen and Lansing \(2019\)](#).

Our analysis uses new state-level consumer price indices for the United States that we have constructed back to the 1970s. Prior to our work, state-level price indices based on BLS micro-price data have not existed. The BLS has published city-level inflation series for a group of relatively large cities.

2. [Carvalho et al. \(2021\)](#) reach a similar conclusion using very different methods. They propose a model for long-run inflation expectations and show how their model generates the result that the Volcker disinflation was driven by shifting long-run inflation expectation and also that long-run inflation expectations become anchored in the 1990s onward.

But it has refrained from reporting inflation indices for smaller metropolitan areas (and for states). Our new state-level price indices use all the available underlying micro data gathered by the BLS. We also construct state-level price indices for nontradeables and tradeables. We focus our analysis on the behavior of the prices of nontradeable goods. This is important. For prices set at the national level—as is more likely for tradeables—the slope of the regional Phillips curve will be zero no matter how large the slope of the aggregate Phillips curve is.

A notable conclusion of the recent regional Phillips curve literature has been that the estimated slope of the regional Phillips curve has tended to be steeper than the slope estimated for the aggregate Phillips curve. The theoretical framework we develop helps explain why this is the case. We show that panel data estimates of the regional Phillips curve by prior researchers are estimates of  $\psi$  in [equation \(3\)](#) as opposed to estimates of  $\kappa$  in [equation \(1\)](#). This means that they are not directly comparable to much of the aggregate literature. We discuss how researchers can convert estimates of  $\psi$  to  $\kappa$  and explain what other statistics this conversion depends on (primarily the degree of persistence of the unemployment variation used to estimate  $\psi$ ). Our analysis highlights the importance of the exact specification used in estimating regional Phillips curves.<sup>3</sup>

The regional setting, along with our new inflation indices, allow us to leverage new forms of variation in estimating the Phillips curve. We develop a new tradeable-demand spillovers instrument building on insights from [Nguyen \(2014\)](#). This instrument is based on the idea that supply shocks in tradeable sectors will differentially affect demand in nontradeable sectors in regions that are differentially exposed to the shocked tradeable sectors: for example, an oil boom will increase demand for restaurant meals in Texas. In carrying out our regional analysis, we are careful to account for the fact that roughly 42% of the expenditure weight in core inflation is on the shelter component of housing services, which are measured by rents.<sup>4</sup> We estimate the slope of the regional Phillips curve for rents and show that it is substantially

3. For example, [Nishizaki and Watanabe \(2000\)](#) find evidence of Phillips curve flattening in their baseline specification with no time fixed effects, but this evidence changes dramatically when time fixed effects are added.

4. Much of the expenditure weight for housing derives from owner-occupied housing. However, rents are used to measure inflation for all shelter, due to the difficulty of backing out the user cost of housing from actual house prices in a theoretically appealing way. The expenditure weight of the CPI less food and



steeper than the regional Phillips curve for nontradeables excluding housing. We use the combination of these estimates to predict the behavior of aggregate core inflation, which includes rents, and show that these predictions match the greater aggregate cyclicity of core inflation than core inflation excluding housing, a fact emphasized by [Stock and Watson \(2019\)](#). We conclude from this that the behavior of rent prices play an important role in determining the slope of the regional and aggregate Phillips curves.

In addition to the papers already mentioned, our work builds on the vast empirical and theoretical literature on the Phillips curve. The literature originates with [Phillips \(1958\)](#) and [Samuelson and Solow \(1960\)](#). [Friedman \(1968\)](#) and [Phelps \(1967\)](#) emphasized the importance of including an inflation expectations term in the Phillips curve. [Gordon \(1982\)](#) emphasized the importance of supply shocks. Important early papers that estimate the New Keynesian Phillips curve include [Roberts \(1995\)](#), [Fuhrer and Moore \(1995\)](#), [Galí and Gertler \(1999\)](#), and [Sbordone \(2002\)](#), but see also works cited in [Mavroeidis, Plagborg-Møller, and Stock \(2014\)](#). Important recent publications estimating the Phillips curve include [Ball and Mazumder \(2011, 2019\)](#), [Coibion and Gorodnichenko \(2015\)](#), [Stock and Watson \(2019\)](#), [Barnichon and Mesters \(2020\)](#), and [Del Negro et al. \(2020\)](#). Our article is also related to a recent literature that assesses the missing disinflation during the Great Recession (see [Christiano, Eichenbaum, and Trabandt 2015](#); [Del Negro, Giannoni, and Schorfheide 2015](#); [Gilchrist et al. 2017](#); [Crump et al. 2019](#)).

The article proceeds as follows. [Section II](#) derives [equation \(3\)](#) and explains the problem of regime change in estimating the Phillips curve. [Section III](#) describes our main framework for interpreting the regional Phillips curve. [Section IV](#) describes our new state-level inflation indices. [Section V](#) presents our empirical results. [Section VI](#) concludes.

## II. THE POWER AND PROBLEM OF LONG-RUN INFLATION EXPECTATIONS

To appreciate the value of using regional variation to estimate the slope of the Phillips curve, it is useful to understand the central

---

energy is 77.7%, and 32.3 percentage points out of this expenditure weight are rents.

role of long-run inflation expectations in determining aggregate inflation. To this end, we solve [equation \(1\)](#) forward to get

$$(4) \quad \pi_t = -\kappa E_t \sum_{j=0}^{\infty} \beta^j u_{t+j} + \omega_t,$$

where  $\omega_t \equiv E_t \sum_{j=0}^{\infty} \beta^j (\kappa u_{t+j}^n + v_{t+j})$ . This equation illustrates how inflation at time  $t$  is determined by the path of unemployment out into the infinite future.<sup>5</sup> We can furthermore decompose the variation in future unemployment  $u_{t+j}$  into a transitory and permanent component. Define the transitory component of variation in unemployment to be  $\tilde{u}_t = u_t - E_t u_{t+\infty}$ , where  $E_t u_{t+\infty}$  is the permanent component of the variation in unemployment. Using these concepts, we can rewrite [equation \(4\)](#) as

$$(5) \quad \pi_t = -\kappa E_t \sum_{j=0}^{\infty} \beta^j \tilde{u}_{t+j} - \frac{\kappa}{1-\beta} E_t u_{t+\infty} + \omega_t.$$

Assuming that shocks to  $u_t^n$  and  $v_t$  are transitory, [equation \(1\)](#) implies that  $E_t \pi_{t+\infty} = -\frac{\kappa}{1-\beta} E_t u_{t+\infty}$ . We can then rewrite [equation \(5\)](#) as

$$(6) \quad \pi_t = -\kappa E_t \sum_{j=0}^{\infty} \beta^j \tilde{u}_{t+j} + E_t \pi_{t+\infty} + \omega_t.$$

5. While the most popular microfoundation of the New Keynesian Phillips curve—and the one we develop in [Section III](#)—is based on the price rigidity assumptions in [Calvo \(1983\)](#), this equation or something very similar arises from several other microfoundations. [Roberts \(1995\)](#) shows that the same Phillips curve arises from [Rotemberg's \(1982\)](#) quadratic costs of price adjustment model and [Taylor's \(1979, 1980\)](#) model of staggered contracts (the timing of the output gap term is slightly different in the Taylor model). Furthermore, [Gertler and Leahy \(2008\)](#) develop the same Phillips curve as a linear approximation of a model with Ss foundations. In the case of the Rotemberg model in continuous time, the derivation does not rely on a linear approximation around a zero-inflation steady state. Models based on information frictions yield Phillips curves that are not forward looking. These models, however, typically assume no price rigidity. Incorporating price rigidity into these models would make their Phillips curves forward looking as well. [Sbordone \(2002\)](#), [Galí, Gertler, and Lopez-Salido \(2005\)](#), and [Rudd and Whelan \(2005\)](#) develop approaches to estimating the Phillips curve on aggregate data using versions of [equation \(4\)](#).

Finally, let's assume for simplicity that  $\tilde{u}_t$  follows an AR(1) process with autocorrelation coefficient equal to  $\rho_u$ . In this case  $E_t \tilde{u}_{t+j} = \rho_u^j \tilde{u}_t$  and we can rewrite equation (6) as

$$(7) \quad \pi_t = -\psi \tilde{u}_t + E_t \pi_{t+\infty} + \omega_t,$$

where  $\psi = \frac{\kappa}{1-\beta\rho_u}$ .

This way of writing the Phillips curve highlights the importance of long-run inflation expectations in determining inflation at the aggregate level. Long-run inflation expectations  $E_t \pi_{t+\infty}$  appear with a coefficient of one in equation (7). In other words, current inflation moves one for one with changes in long-run inflation expectations. These long-run expectations are determined by the private sector's beliefs about the long-run monetary regime being followed by the central bank (the long-run inflation target). Variation in beliefs about the long-run monetary regime therefore have very large effects on current inflation.<sup>6</sup>

Equation (7) implies that inflation can vary dramatically without any variation in  $\tilde{u}_t$  if there is substantial variation in long-run inflation expectations. In this case, the relationship between inflation and  $\tilde{u}_t$  may be entirely uninformative about the slope of the Phillips curve. Worse still, variation in long-run inflation expectations may be correlated with variation in  $\tilde{u}_t$ . For example, it seems very plausible that Volcker's willingness to allow unemployment to rise to very high values in the early 1980s—and the fact that he was not forced to resign—signaled to the public that he was serious about bringing down inflation (and had the backing of the president to do this). Such a correlation will impart an upward bias on estimates of the slope of the Phillips curve unless variation in inflation expectations can be controlled for. But in practice, controlling for inflation expectations is hard because of weak instruments (Mavroeidis, Plagborg-Møller, and Stock 2014) and because direct measures of inflation expectations may be imperfect. So, a rapid drop in inflation expectations may masquerade as a steep Phillips curve.

Why has the Phillips curve appeared to flatten over the past few decades? Figure II shows that since roughly 1998, long-term

6. Equations (6) and (7) remain valid in the case where the coefficient on  $E_t \pi_{t+1}$  in equation (1) is equal to one rather than  $\beta$ . In this case, the long-run Phillips curve is vertical, and long-run unemployment and inflation are unrelated. That is, inflation still satisfies  $\pi_t = -\kappa E_t \sum_{j=0}^{\infty} \tilde{u}_{t+j} + \omega_t + E_t \pi_{t+\infty}$ , but  $E_t \pi_{t+\infty}$  is now independent of  $E_t \pi_{t+\infty}$ . The forward sum is bounded even though  $\beta = 1$ , because  $\tilde{u}_t$  has zero unconditional mean.

inflation expectations have been firmly anchored at close to 2%. This has led to a collapse of the covariance between  $E_t\pi_{t+\infty}$  and unemployment and therefore eliminated any bias associated with poorly proxied variation in inflation expectations. A fall in this bias will appear from the perspective of the (misspecified) accelerationist Phillips curve (such as the one we discuss in the introduction) as a flatter curve.

One piece of corroborating evidence for this view is the close relationship between  $\pi_t$  and  $E_t\pi_{t+1}$  in the data. Recall that the standard formulation of the New Keynesian Phillips curve—[equation \(1\)](#)—implies that it is the gap between  $\pi_t$  and  $\beta E_t\pi_{t+1}$ —let's call this the inflation gap—that must be explained by demand pressure (the  $\kappa u_t$  term) or supply shocks ( $\kappa u_t^n + v_t$ ). [Figure III](#) plots SPF forecasts of inflation over the next year along with four different measures of current inflation. The difference between the two series is approximately equal to the inflation gap  $\pi_t - \beta E_t\pi_{t+1}$ .

The measure of current inflation plotted in the top left panel of [Figure III](#) is the 12-month change in the overall CPI. This conventional way of comparing current inflation and inflation expectations over the next year suggests that these series are closely related, but that there is nevertheless substantial variation in the gap between them (the inflation gap). Moving to the top right panel, we measure current inflation by the 12-month change in core CPI inflation, excluding food and energy. The inflation gap measured this way is quite a bit smaller. Evidently, commodities account for a large part of the inflation gap for the overall CPI. However, a substantial inflation gap remains in the early 1980s.

The measure of current inflation plotted in the bottom left panel of [Figure III](#) is the 12-month change in the core PCE. The advantage of this series is that it makes use of current measurement methods, retroactively applied. In this case, the inflation gap is very small. A similar message emerges in the bottom right panel using the 12-month change in the core CPI research series published by the BLS. This series also uses consistent, modern methods to calculate inflation back in time. A particularly important measurement change for our purposes occurred in 1983, when the BLS switched to using rent inflation as a proxy for overall housing inflation, including for owner-occupied housing (rental equivalence). Before that time, housing services inflation in the CPI was constructed from a weighted average of changes in house prices and mortgage costs (i.e., interest rates). This

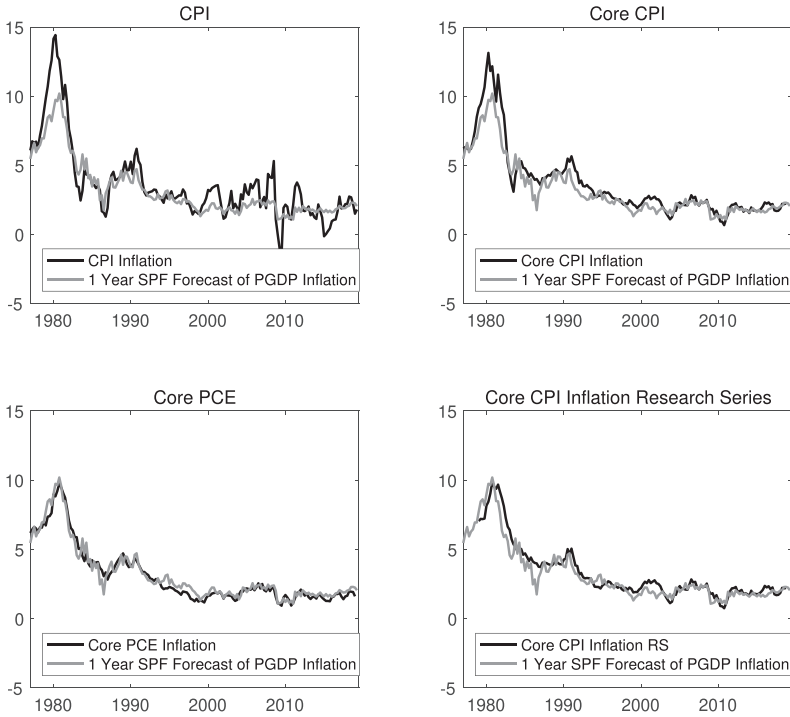


FIGURE III

$$\pi_t \text{ versus } E_t \pi_{t+1}$$

Each panel shows the comparison of the one-year-ahead forecast of the GDP deflator coming from the Survey of Professional of Forecasters and a measure of inflation. The top left panel uses the published headline CPI. The top right panel excludes food and energy by plotting the published measure of the core CPI. The bottom panels correct for changes in the methodology of inflation measurement. The bottom left panel uses PCE inflation, which has maintained a stable methodology, and the bottom right panel uses the Constant Methodology Research Series for core CPI published by the Bureau of Labor Statistics. We use forecasts of the GDP deflator because forecasts for the CPI are not available before 1980.

earlier approach essentially baked in a strong relationship between Volcker's actions to curb the Great Inflation and measured CPI inflation, since interest rates (and house prices) fed directly into the CPI.<sup>7</sup>

7. These choices are consequential because the housing component of the CPI has a weight of roughly one-third in the overall CPI. [Online Appendix B.2](#) presents our attempt to replicate the pre-1983 BLS housing methodology on more modern

The overall message that emerges from [Figure III](#) is that the inflation gap for core inflation measured using modern methods is tiny throughout our sample period. Importantly, this includes the period of the Volcker disinflation. This is suggestive evidence that the slope of the Phillips curve was small throughout our sample period: unemployment varied a great deal in the early 1980s and again in the Great Recession without much variation in the inflation gap. However, the panels in [Figure III](#) illustrate well that this conclusion is sensitive to the details of how inflation is measured.<sup>8</sup> It is also sensitive to whether the expectations data used come from the SPF or from the Michigan Survey of Consumers as [Coibion and Gorodnichenko \(2015\)](#) emphasize, and is also sensitive to the exact timing of the variables.

### III. A MODEL OF THE REGIONAL PHILLIPS CURVE

We develop a two-region, New Keynesian, open-economy model featuring tradeable and nontradeable sectors. We derive a regional Phillips curve in this model and show how it relates to the aggregate Phillips curve. The model demonstrates a chief benefit of regional data: time and state fixed effects “difference out” changes in long-run inflation expectations. The model also illustrates the importance of using nontradeable inflation when estimating the slope of the Phillips curve using regional data.

#### *III.A. Model Setup*

Our model consists of two regions that belong to a monetary and fiscal union. We refer to the regions as Home (H) and Foreign (F). The population of the entire economy is normalized to one. The population of the home region is denoted by  $\zeta$ . Labor is immobile across regions. In each region, there is a single labor market. Household preferences, market structure, and firm behavior take the same form in both regions. Below we describe the economy of the home region. All prices in the economy are denominated in “dollars,” a digital currency issued by the federal government.<sup>9</sup> Throughout, we adopt the following conventions

---

data. The main conclusion from this is that this methodology would have led to much more variable (and cyclical) inflation over the past few decades.

8. We discuss this in more detail in [Online Appendix B.1](#).

9. In other words, we are considering an economy in the cashless limit ([Woodford 1998, 2003](#)).

unless otherwise stated. Lowercase variables are the logs of uppercase variables. Hatted variables denote the percentage deviation of a variable from its steady-state value. Steady-state values are recorded without time subscripts.

1. *Households.* The representative household in the home region seeks to maximize the utility function

$$E_0 \sum_{t=0}^{\infty} \beta^t u(C_{Ht}, N_{Ht}),$$

where  $C_{Ht}$  is per capita consumption of a composite consumption good,  $N_{Ht}$  is per capita employment, and  $\beta$  is the household's subjective discount factor. We follow Greenwood, Hercowitz, and Huffman (1988) in assuming that the function  $u(C_{Ht}, N_{Ht})$  takes the form

$$(8) \quad u(C_{Ht}, N_{Ht}) = \frac{\left( C_{Ht} - \chi \frac{N_{Ht}^{1+\varphi-1}}{1+\varphi-1} \right)^{1-\sigma^{-1}}}{1 - \sigma^{-1}},$$

where  $\varphi$  is the household's Frisch elasticity of labor supply,  $\sigma$  determines the household's elasticity of intertemporal substitution, and  $\chi$  governs the intensity of the household's disutility of labor. We refer to this preference specification as GHH preferences.

The composite consumption good  $C_{Ht}$  is a constant elasticity of substitution (CES) index over tradeables  $C_{Ht}^T$  and nontradeables  $C_{Ht}^N$  given by

$$C_{Ht} = \left[ \phi_N^{\frac{1}{\eta}} C_{Ht}^N{}^{\frac{\eta-1}{\eta}} + \phi_T^{\frac{1}{\eta}} C_{Ht}^T{}^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}},$$

where  $\eta$  is the elasticity of substitution between tradeables and nontradeables and  $\phi_T$  and  $\phi_N$  are the household's steady-state expenditure shares on tradeable and nontradeable goods, respectively.  $C_{Ht}^N$  and  $C_{Ht}^T$  are themselves composite goods described further below. Nontradeable goods are only consumed in the region where they are produced. In contrast, the market for tradeable goods is completely integrated across regions. Hence, home and foreign households may face different prices for nontradeables but face the same prices for tradeable goods. The expenditure share

on tradeable and nontradeable goods must sum to 1, that is,  $\phi_N + \phi_T = 1$ .

The composite nontradeable good  $C_{Ht}^N$  is given by

$$C_{Ht}^N = \left[ \int_0^1 C_{Ht}^N(z)^{\frac{\theta-1}{\theta}} dz \right]^{\frac{\theta}{\theta-1}},$$

where  $C_{Ht}^N(z)$  denotes consumption of variety  $z$  of nontradeable goods in the home region. The home price of this nontradeable variety is  $P_{Ht}^N(z)$ . The parameter  $\theta > 1$  denotes the elasticity of substitution between different nontradeable varieties.

Home tradeable consumption  $C_{Ht}^T$  is a CES aggregate over tradeable goods produced in the home and foreign regions given by

$$(9) \quad C_{Ht}^T = \left[ \tau_{Ht}^H \frac{1}{\eta} C_{Ht}^{TH \frac{\eta-1}{\eta}} + \tau_{Ht}^F \frac{1}{\eta} C_{Ht}^{TF \frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}},$$

where  $C_{Ht}^{TH}$  and  $C_{Ht}^{TF}$  are home consumption of composite tradeable goods produced in the home and foreign regions, respectively. We assume (for simplicity) that the elasticity of substitution between home-produced and foreign-produced tradeables is  $\eta$  (the same as the elasticity of substitution between tradeables and nontradeables). Demand for home-produced and foreign-produced tradeables is subject to shocks denoted by  $\tau_{Ht}^H$  and  $\tau_{Ht}^F$ , respectively. We normalize  $\tau_{Ht}^H + \tau_{Ht}^F = 1$ . For simplicity, we do not allow for home bias in tradeable consumption. Thus, we set  $\tau_{Ht}^H = \tau_{Ht}^F = \zeta$ , that is, the share of spending on goods from the home region in each region is equal to the size of the home region.

The home and foreign composite tradeable goods are CES indices given by

$$C_{Ht}^{TH} = \left[ \int_0^1 C_{Ht}^{TH}(z)^{\frac{\theta-1}{\theta}} dz \right]^{\frac{\theta}{\theta-1}} \quad \text{and} \quad C_{Ht}^{TF} = \left[ \int_0^1 C_{Ht}^{TF}(z)^{\frac{\theta-1}{\theta}} dz \right]^{\frac{\theta}{\theta-1}},$$

where  $C_{Ht}^{TH}(z)$  and  $C_{Ht}^{TF}(z)$  are home consumption of varieties of tradeable goods produced in the home and foreign region, respectively. The prices of these home-produced and foreign-produced tradeable good varieties are  $P_{Ht}^T(z)$  and  $P_{Ft}^T(z)$ , respectively.



Households maximize utility subject to a sequence of budget constraints:

$$C_{Ht}^N P_{Ht}^N + C_{Ht}^{TH} P_{Ht}^T + C_{Ht}^{TF} P_{Ft}^T + E_t [M_{Ht,t+1} B_{H,t+1}] \leq B_{Ht} + W_{Ht} N_{Ht} + \Xi_{Ht}^N + \Xi_{Ht}^T,$$

where  $B_{Ht}$  is a random variable denoting payoffs of the state-contingent portfolio held by households in period  $t$ ;  $M_{Ht,t+1}$  is the one-period-ahead stochastic discount factor of the home representative household;  $P_{Ht}^N$ ,  $P_{Ht}^T$ , and  $P_{Ft}^T$  are price indices that give the minimum cost of purchasing a unit of  $C_{Ht}^N$ ,  $C_{Ht}^{TH}$ , and  $C_{Ht}^{TF}$ , respectively;  $W_{Ht}$  is the nominal wage received by workers in region  $H$ ; and  $\Xi_{Ht}^N$  and  $\Xi_{Ht}^T$  are the profits of nontradeable and tradeable firms in the home region. There is a complete set of financial markets across the two regions. To rule out Ponzi schemes, we assume that household debt cannot exceed the present value of future income in any state.

We present the first-order necessary conditions for household optimization in [Online Appendix A.1](#). As we noted, the problem of the foreign household is analogous. We therefore refrain from describing it in detail here. For simplicity, we do not allow for tradeable-demand shocks to foreign tradeable consumption as we do for home tradeable consumption.

2. *Firms.* There is a continuum of firms in the tradeable and nontradeable sectors. Firms are indexed by  $z$  and firm  $z$  specializes in the production of differentiated good  $z$ . Labor is the only variable factor of production used by firms.

We begin by discussing the nontradeable sector. The output of good  $z$  in the nontradeable sector is denoted  $Y_{Ht}^N(z)$ . The production function of firm  $z$  in this sector is

$$(10) \quad Y_{Ht}^N(z) = Z_{Ht}^N N_{Ht}^N(z),$$

where  $N_{Ht}^N(z)$  is the amount of labor demanded by firm  $z$  and  $Z_{Ht}^N$  is a productivity shock.

Firm  $z$  in the nontradeable sector maximizes its value:

$$E_t \sum_{j=0}^{\infty} M_{Ht,t+j} \left[ P_{H,t+j}^N(z) Y_{H,t+j}^N(z) - W_{H,t+j} N_{H,t+j}^N(z) \right]$$

given demand for its good, which is

$$Y_{Ht}^N(z) = \zeta C_{Ht}^N \left( \frac{P_{Ht}^N(z)}{P_{Ht}^N} \right)^{-\theta}.$$

Firm  $z$  can set its price freely with probability  $1 - \alpha$  as in Calvo (1983). With probability  $\alpha$  the firm must keep its price unchanged.

Analogously to the nontradeable sector, the output of firm  $z$  in the tradeable sector is denoted  $Y_{Ht}^T(z)$ . Its production function is

$$Y_{Ht}^T(z) = Z_{Ht}^T N_{Ht}^T(z),$$

where  $N_{Ht}^T(z)$  is the amount of labor demanded by the firm producing good  $z$  and  $Z_{Ht}^T$  is a productivity shock.

Firm  $z$  in the tradeable sector maximizes its value:

$$E_t \sum_{j=0}^{\infty} M_{Ht,t+j} \left[ P_{H,t+j}^T(z) Y_{H,t+j}^T(z) - W_{H,t+j} N_{H,t+j}^T(z) \right]$$

given demand for its good. Demand in the tradeable sector comes from both the home and foreign regions. Firm  $z$ 's demand is thus given by

$$Y_{Ht}^T(z) = \left( \zeta C_{Ht}^{TH} + (1 - \zeta) C_{Ft}^{TH} \right) \left( \frac{P_{Ht}^T(z)}{P_{Ht}^T} \right)^{-\theta}.$$

The tradeable goods firms also have an opportunity to change their price with probability  $1 - \alpha$  each period and must otherwise keep their prices fixed.

We present the first-order necessary conditions for firm optimization in Online Appendix A.2. The problems of foreign firms are analogous to those of home firms.

3. *Government Policy and Equilibrium.* The federal government operates a common monetary policy for the two regions. This policy takes the form of the following interest rate rule

$$\hat{r}_t^n = \varphi_\pi (\pi_t - \bar{\pi}_t) - \varphi_u (\hat{u}_t - \bar{u}_t) + \varepsilon_{rt},$$

where, as elsewhere in the article, hatted variables denote deviations from a zero-inflation steady state and lowercase

variables are the logs of uppercase variables. Economy-wide inflation  $\pi_t$  is a population-weighted average of inflation in the two regions:  $\pi_t \equiv \zeta \pi_{Ht} + (1 - \zeta) \pi_{Ft}$ , where  $\pi_{Ht} = p_{Ht} - p_{H,t-1}$  is consumer price inflation in the home region and  $\pi_{Ft}$  is defined analogously for the foreign region. In our model, we define unemployment in the home region simply as  $u_{Ht} = 1 - N_{Ht}$ . We define foreign unemployment analogously. This implies that to a first order  $\hat{u}_{Ht} = -\hat{n}_{Ht}$  and  $\hat{u}_{Ft} = -\hat{n}_{Ft}$ . Economy-wide unemployment is a population-weighted average of unemployment in the two regions, so  $\hat{u}_t = \zeta \hat{u}_{Ht} + (1 - \zeta) \hat{u}_{Ft}$ .

Importantly, we allow the monetary authority to have a time-varying inflation target  $\bar{\pi}_t$ .<sup>10</sup> Since the long-run Phillips curve in our model is not vertical, variation in long-run inflation yields variation in long-run unemployment. We assume that the monetary authority targets an unemployment rate that is consistent with its long-run inflation target, that is,  $\bar{u}_t = \frac{(1-\beta)\bar{\pi}_t}{\kappa}$ . We assume that  $\varphi_\pi$  and  $\varphi_u$  obey the Taylor principle, ensuring that the economy has a unique locally bounded equilibrium.  $\varepsilon_{rt}$  is a transitory monetary shock, which we assume follows an exogenous AR(1) process.

For simplicity, the government levies no taxes, engages in no spending, and issues no debt. In other words, there is no fiscal policy. The digital currency issued by the government is in zero net supply. The government's monetary policy has no fiscal implications. An equilibrium in this economy is an allocation that satisfies household optimization, firm optimization, the government's interest rate rule, and market clearing. We focus on the unique locally bounded equilibrium of the model. Implicitly we rule out equilibria in which the inflation rate rises without bound using the trigger strategy argument presented in [Obstfeld and Rogoff \(1983\)](#).

### III.B. Regional and Aggregate Phillips Curves

Taking a log-linear approximation of the model presented in [Section III.A](#) around a zero-inflation steady state with balanced trade yields the following regional Phillips curve for the inflation of nontradeable goods:

$$(11) \quad \pi_{Ht}^N = \beta \mathbf{E}_t \pi_{H,t+1}^N - \kappa \hat{u}_{Ht} - \lambda \hat{p}_{Ht}^N + v_{Ht}^N,$$

10. Prior work that allows for a time-varying inflation target includes [Stock and Watson \(2007\)](#), [Ireland \(2007\)](#), and [Cogley and Sbordone \(2008\)](#).

and aggregate Phillips curve for overall inflation:

$$(12) \quad \pi_t = \beta E_t \pi_{t+1} - \kappa \hat{u}_t + v_t,$$

where  $\pi_{Ht}^N = p_{Ht}^N - p_{H,t-1}^N$  is home nontradeable inflation,  $\hat{p}_{Ht}^N = \frac{p_{Ht}^N}{\bar{p}_{Ht}} - 1$  is the percentage deviation of the home relative price of nontradeables from its steady-state value of one,  $v_{Ht}^N$  is a nontradeable home supply shock,  $v_t$  is a corresponding aggregate supply shock, and the parameter  $\kappa = \lambda \varphi^{-1}$ , where  $\lambda = \frac{(1-\alpha)(1-\alpha\beta)}{\alpha}$ . We provide a detailed derivation of these equations in [Online Appendix A](#).

[Equations \(11\) and \(12\)](#) yield an important result: the slopes of the regional Phillips curve for nontradeables and the aggregate Phillips curve are the same in our model. These slopes are both equal to  $\kappa$ . This result holds for the nontradeable regional Phillips curve but does not carry over to the regional Phillips curve for overall consumer price inflation—which includes both tradeable and nontradeable inflation in the region. As we show in [Online Appendix A.8](#), the slope of the regional Phillips curve for overall consumer price inflation is smaller by a factor equal to the expenditure share on nontradeable goods.

Intuitively, the difference in the slope between the nontradeable and overall regional Phillips curves arises because all regions share the tradeable goods and these goods are priced nationally. The tradeable goods therefore do not contribute to differences in inflation across regions, which means that the regional CPI is made up partly of goods whose regional prices are insensitive to regional variation in unemployment. This makes the regional CPI less sensitive to regional unemployment than the aggregate CPI is to aggregate unemployment.

Our result that the slope of the nontradeable regional Phillips curve is equal to the slope of the aggregate Phillips curve leads us to focus our cross-sectional empirical work on inflation for nontradeable goods. Earlier research that has estimated regional Phillips curves has done so for overall consumer price inflation at the regional level (e.g., [Fitzgerald and Nicolini 2014](#); [McLeay and Tenreyro 2019](#)). Our model suggests that results from such analysis are less directly informative about the slope of the aggregate Phillips curve.

Our assumption that households have GHH preferences helps simplify the derivation of the regional and aggregate Phillips curves in our model—[equations \(11\) and \(12\)](#). GHH preferences

imply that wealth effects on labor supply are zero, which eliminates the dependence of marginal costs on consumption. The absence of a consumption term in the Phillips curve plays a role in the derivation of our result that the nontradeable regional Phillips curve and the aggregate Phillips curve have the same slope. We discuss this point at greater length in [Online Appendix A.9](#). The form of the Phillips curve in our model does not depend on the structure of financial markets. We have assumed complete financial markets across regions, but the Phillips curve is the same in a model with incomplete markets across regions.

An important difference between [equations \(11\)](#) and [\(12\)](#) is the presence of the relative price of nontradeables term  $\lambda \hat{p}_{Ht}^N$  in [equation \(11\)](#). This term implies that inflation in the nontradeables sector will be lower the higher the relative price of nontradeables. Conceptually, this term is very important. It pushes relative prices toward parity in the long run. Also, it implies that even if prices in the economy are very flexible— $\kappa$  is very large—a local boom will not result in unbounded inflation of home nontradeable prices because demand for these goods is affected by their prices relative to other prices in the economy. The mechanical reason this term appears is that the inflation rate for nontradeable goods is driven by variation in the real wage deflated by nontradeable prices. Labor supply in the home region, however, is a function of the real wage deflated by the home CPI. The real marginal-cost variable in the home nontradeable Phillips curve therefore gives rise to an unemployment term and a relative price of nontradeables term.

### III.C. Estimating the Slope of the Phillips Curve with Regional Data

Next we solve the regional Phillips curve—[equation \(11\)](#)—forward to obtain

$$(13) \quad \pi_{Ht}^N = -E_t \sum_{j=0}^{\infty} \beta^j \left( \kappa \tilde{u}_{H,t+j} + \lambda \hat{p}_{H,t+j}^N \right) + E_t \pi_{t+\infty}^N + \omega_{Ht}^N,$$

where  $\tilde{u}_{Ht} = u_{Ht} - E_t u_{H,t+\infty}$  and  $\omega_{Ht}^N = E_t \sum_{j=0}^{\infty} \beta^j v_{H,t+j}^N$ .

A major benefit of estimating the slope of the Phillips curve using regional data from a monetary union is that variation in long-run inflation expectations—the  $E_t \pi_{t+\infty}^N$  term in [equation \(13\)](#)—is constant across regions. This implies that variation

in long-run inflation expectations will be absorbed by time fixed effects in a panel specification. Intuitively, while short-run inflation expectations ( $E_t \pi_{t+1}^N$ ) will differ across regions due to differences in their economic circumstances, long-run inflation expectations ( $E_t \pi_{t+\infty}^N$ ) are independent of the current business cycle. They are determined by beliefs about the long-run monetary regime. In a monetary union like the United States, these beliefs will vary uniformly across regions. This means that these expectations are differenced out in a panel regression with time fixed effects.

The result that long-run inflation expectations are constant across regions (and sectors) in our model relies on productivity and other drivers of real costs having a common trend in the long run. If productivity growth (say) differs across regions even in the long run, this will lead to persistent differences in nontradeable inflation (a Balassa-Samuelson effect). However, if this difference is constant over time, it will be absorbed by region fixed effects in a panel specification.

These observations imply that that we can adopt an empirical specification that replaces the  $E_t \pi_{t+\infty}^N$  term in [equation \(13\)](#) with time and region fixed effects:

$$(14) \quad \pi_{it}^N = -E_t \sum_{j=0}^{\infty} \beta^j \left( \kappa u_{i,t+j} + \lambda \hat{p}_{i,t+j}^N \right) + \alpha_i + \gamma_t + \tilde{\omega}_{it}^N,$$

where  $i$  denotes region,  $\alpha_i$  denotes a set of region fixed effects, and  $\gamma_t$  denotes a set of time fixed effects. Variation in  $E_t \pi_{t+\infty}^N$  in [equation \(13\)](#) that is common across regions will be absorbed by the time fixed effect.<sup>11</sup> Constant differences across regions in  $E_t \pi_{t+\infty}^N$  will be absorbed by the state fixed effects. To the extent that there is remaining variation in  $E_t \pi_{t+\infty}^N$  across regions (e.g., due to changing trends), it will be a part of the error term  $\tilde{\omega}_{it}^N$ .

It is useful to relate [equation \(14\)](#) to the empirical specifications used in the recent regional Phillips curve literature. If we assume that both  $u_{Ht}$  and  $\hat{p}_{Ht}^N$  follow AR(1) processes with autocorrelation coefficients equal to  $\rho_u$  and  $\rho_{pN}$ , respectively, [equation \(14\)](#)

11. The time fixed effects also absorb time variation in the long-run expected unemployment  $E_t u_{t+\infty}$ . We have therefore replaced  $\tilde{u}_{i,t+j}$  in [equation \(13\)](#) with  $u_{i,t+j}$  in [equation \(14\)](#). This equation remains valid if  $\beta = 1$ . The forward sum in the equation is still bounded, because  $u_{it}$  has zero mean conditional on time fixed effects.

simplifies to

$$(15) \quad \pi_{it}^N = -\psi u_{it} - \delta \hat{p}_{it}^N + \alpha_i + \gamma_t + \tilde{\omega}_{it}^N,$$

where  $\psi = \frac{\kappa}{1-\beta\rho_u}$  and  $\delta = \frac{\lambda}{1-\beta\rho_{pN}}$ . This equation is similar to the empirical specification used by much of the recent regional Phillips curve literature. Comparing equations (14) and (15), we see that an important difference between them is that the slope coefficient is not the same. The slope coefficient in equation (14) is  $\kappa$  (which is the same as the slope coefficient in equations (11) and (12)), whereas the slope coefficient in equation (15) is  $\psi = \frac{\kappa}{1-\beta\rho_u}$ . Since unemployment is quite persistent,  $\psi$  is likely to be substantially larger than  $\kappa$ . Note that the AR(1) assumption we use to derive equation (15) is not used in our estimation of  $\kappa$ .

A curious feature of the recent regional Phillips curve literature is that it has tended to yield larger estimates of the slope of the Phillips curve than more traditional estimation strategies based on aggregate data (Fitzgerald and Nicolini 2014; Babb and Detmeister 2017; Hooper, Mishkin, and Sufi 2019; McLeay and Tenreyro 2019). Comparing equations (14) and (15) provides a simple explanation for this discrepancy. The regional Phillips curve literature has been estimating  $\psi$  in equation (15), while the more traditional literature using aggregate variation has typically been estimating  $\kappa$ . Since  $\psi \gg \kappa$ , it is not surprising that the slope of the Phillips curve estimated in the regional literature has seemed large relative to traditional estimates.<sup>12</sup>

The difference between  $\kappa$  and  $\psi$  arises due to the different ways equations (11) and (15) capture the effects of expected future unemployment on current inflation. In equation (11), the effects of expected future unemployment on current inflation are captured by the inflation expectations term  $E_t \pi_{t+1}$  and the coefficient on current unemployment  $\kappa$  only reflects the effect of current

12. This same type of lack of comparability arises in some cases for different estimates based on aggregate data. Some researchers use longer-term inflation expectations, rather than one-period-ahead inflation expectations, to proxy for  $E_t \pi_{t+1}^N$  when estimating the Phillips curve using aggregate data. Our analysis shows, however, that when researchers use data on long-term inflation expectations, they (perhaps inadvertently) end up estimating  $\psi$ , not  $\kappa$ . To compare such estimates with those based on a specification that controls for one-period-ahead expectations, one must translate between the two, for example, by using the formula  $\psi = \frac{\kappa}{1-\beta\rho_u}$  or a version of this formula appropriate for (say) 10-years-ahead inflation expectations.

unemployment on current inflation. In contrast, the slope coefficient in [equation \(15\)](#) captures both the effect of current unemployment and the effect of expected future unemployment into the indefinite future on current inflation—that is, the fact that high unemployment today forecasts high unemployment in future periods.<sup>13</sup>

An advantage of estimating specifications such as [equations \(14\) and \(15\)](#) rather than [equation \(11\)](#) is that identifying the slope coefficient is less sensitive to the exact timing of changes in inflation relative to inflation expectations. In [Figure III](#), we show that the difference between inflation and inflation expectations is quite sensitive to the exact measure of inflation.

We have so far manipulated the Phillips curve under the standard assumption of full-information rational expectations. However, the arguments we make above—solving the Phillips curve forward—rely only on the weaker assumption that the law of iterated expectations holds. We elaborate on this point in [Online Appendix A.10](#), drawing on results from [Adam and Padula \(2011\)](#) and [Coibion, Gorodnichenko, and Kamdar \(2018\)](#).

To derive a tractable empirical specification for the regional Phillips curve in which the coefficient on unemployment is the same as in the aggregate Phillips curve, we have made a number of strong assumptions (perfect labor mobility within a region, no labor mobility across regions, GHH preferences, production linear in labor, etc.). In the world, these assumptions are unlikely to hold exactly, and the empirical specification we estimate is thus unlikely to yield exactly the slope of the aggregate Phillips curve. Deriving an exact analytical mapping is nonetheless useful because it highlights in a transparent way the importance of certain forces (e.g., inflation expectations). In more general models where no exact analytical mapping between the slope of the regional and aggregate Phillips curves exists, our regional slope estimates can be used as empirical targets in a moment-matching exercise. Even aside from the simple case we analyze where the aggregate and regional slopes are equal, these moments are likely to provide

13. [McLeay and Tenreyro \(2019\)](#) control for inflation expectations at the census region level when they estimate the regional Phillips curve. The variation across regions in these inflation expectations data is quite minimal. It may be the case that the variation in this variable is quite attenuated relative to actual variation in inflation expectations across the metropolitan statistical areas that form the regional units in their analysis.



valuable information about the slope of the aggregate Phillips curve (Nakamura and Steinsson 2018; Andrews, Gentzkow, and Shapiro 2020).

#### IV. DATA AND CONSTRUCTION OF STATE-LEVEL PRICE INDICES

The BLS does not publish state-level price indices. Prior work has used metropolitan-level BLS price indices and cost of living estimates from the American Chamber of Commerce Realtors Association (ACCRA) to construct state-level price indices (see Del Negro 1998; Nakamura and Steinsson 2014). An important drawback of this approach is that the BLS imputes missing data using data from other regions. Recent work has used scanner price data to construct state-level price indices (Beraja, Hurst, and Ospina 2019). An important drawback of scanner data is the short sample period available.

We construct new state-level price indices for the United States based on the micro-price data the BLS collects for the purpose of constructing the CPI. Our sample period is 1978 to 2018 (with a 26-month gap in 1986–1988 due to missing micro-data). The micro-data that we base our price indices on are available in the CPI Research Database at the BLS. The data for 1978–1987 were constructed by Nakamura et al. (2018). The micro-price data in the CPI Research Database cover thousands of individual goods and services, constituting about 70% of consumer expenditures. They are collected by BLS employees who visit outlets to record prices. The database does not include the rent prices used to construct the shelter component of the CPI. For this reason, we analyze the behavior of rents separately. Prices are sampled in 87 geographical areas across the United States. In New York, Chicago, and Los Angeles, all prices are collected at a monthly frequency. In other locations, food and energy prices are collected monthly and the prices of other items are collected bimonthly. The CPI Research Database is described in more detail in Nakamura and Steinsson (2008) and Klenow and Kryvtsov (2008).

##### IV.A. *State-Level Price Index Construction*

Our methodology for constructing price indices is a simplified version of the procedure used by the BLS to construct the CPI. One key difference versus the BLS procedure, and a key reason we do not simply use the BLS's own price index software, is that

we do not impute missing price observations using inflation rates calculated for other sectors or regions. We describe our procedure below.

We start by calculating price relatives for individual products. These are the fundamental building blocks of a matched-model price index. For product  $i$  at time  $t$ , the formula we use to calculate the price relative is

$$(16) \quad r_{i,t} = \left( \frac{P_{i,t}}{P_{i,t-\tau}} \right)^{\frac{1}{\tau}},$$

where  $r_{i,t}$  denotes the price relative,  $P_{i,t}$  denote the effective price, and  $\tau$  denotes the number of months since the last time a price was collected for this product. Several details are important. First, it is important to use the effective price rather than the raw collected price. The difference between the collected and effective prices is that the latter adjusts for changes in the number and size of the items being priced (e.g., a 2L bottle of Diet Coke versus a two-pack of 2L bottles of Diet Coke).

Second, we define a product not only by its characteristics (e.g., 2L bottle of Diet Coke), but also by the location in which it is sold. To be precise, in the CPI Research Database, each product is indexed by outlet, quote, and version. The quote is a very narrowly described product, and the version is the exact specification of the item that the price collector identifies in the store. We hold all three of these parameters—outlet, quote, and version—fixed in constructing a product's price relative.

Third, we must decide what to do when prices are missing. Missing prices occur when the product is unavailable because of a temporary stockout, or as a consequence of the bimonthly pricing schedule used by the BLS for most products in most cities. Our procedure is to divide the price change evenly among the periods between successive price observations by taking the  $\tau$ th root of the price change and applying this price relative to all  $\tau$  periods. This implies that  $r_{i,t} = \dots = r_{i,t-\tau+1}$ , where again  $\tau$  is the number of periods between successive price changes. There are several other important details of our index construction procedure that we describe in [Online Appendix B.3](#).

We aggregate the price relatives in several steps. First, we compute an unweighted geometric average of the price relatives within each entry level item (ELI) product category and state. ELIs are relatively narrow product categories such as

“full service meals and snacks” (restaurants) and “motorcycles” defined by the BLS for the purpose of calculating the CPI.<sup>14</sup> We then calculate sectoral state-level price indices by computing a weighted geometric average of the ELI-state indices across the ELIs in that state and sector. We use national weights from the Consumer Expenditure Survey (CEX) for 1998 to perform this aggregation.<sup>15</sup>

Our empirical analysis focuses on nontradeables, but we also construct state-level price indices for tradeables—which we simply define as the complement of nontradeables—and overall state-level price indices. We construct a price index for nontradeables based on our own categorization of BLS’s ELI product categories. In doing this, we try to be conservative in our definition of what constitutes a nontradeable good, since including tradeable goods could lead to attenuation of the slope of the Phillips curve if tradeable goods are priced nationally. In contrast, the main downside of excluding some nontradeable goods is less precise estimates. The goods we classify as nontradeables account for roughly 44% of nonhousing consumer expenditures. Importantly, our index of nontradeables does not include housing services or transportation goods (mainly airline tickets).<sup>16</sup> We estimate regional Phillips curves for housing services separately in [Section V](#) using different data. [Online Appendix B.4](#) provides a detailed list of which ELI categories we classify as nontradeable.

Our method for calculating state-level price indices aims to approximate the nonshelter price index published by the BLS. [Online Appendix Figure C.1](#) illustrates our ability to match the official BLS data by comparing the evolution of 12-month inflation at the aggregate level using our methodology with official CPI inflation excluding housing. The figure shows that we are able to

14. See the appendix to [Nakamura and Steinsson \(2008\)](#) for a list of the ELIs used in the construction of the CPI.

15. Here we follow the BLS in using consumption weights. [Rubbo \(2020\)](#) argues that production networks imply that product-level inflation should be weighted by sales shares.

16. We find that there is much more variability across states in nontradeable inflation than tradeable inflation. For nontradeables, the first principal component of state-level inflation captures only about 37% of the variance in the underlying state-level series. In contrast, for tradeables, the first principal component captures about 71% of the variance in the underlying state-level series. This pattern is consistent with our argument in [Section III.B](#) that many tradeable goods are priced nationally and do not respond to regional marginal costs.

approximate the official BLS data very closely. This is true even for the pre-1988 period when we rely on the micro-data recovered by Nakamura et al. (2018), which likely have greater measurement error.<sup>17</sup>

#### IV.B. Employment Data

The measure of unemployment that we use as our measure of labor market slack in the Phillips curve is the quarterly, seasonally adjusted, state unemployment rate from the Local Area Unemployment Statistics (LAUS) published by the BLS. We also make use of employment data in constructing our tradeable-demand spillovers instrument discussed in Section V. This instrument is a shift-share instrument, similar to the one used in Bartik (1991). It is constructed using employment shares of individual industries at the state level.<sup>18</sup> We seasonally adjust the resulting series by regressing it on an exponentially weighted moving average of its lags and state by quarter-of-year fixed effects. We use the variation not explained by the quarter-of-year dummies as our instrument.<sup>19</sup> We define the tradeable employment share in the same way as Mian and Sufi (2014). Online Appendix B.5 discusses this in more detail.

### V. EMPIRICAL RESULTS

We now turn to our empirical results. We present estimates of the structural parameter  $\kappa$  from equation (14) and  $\psi$  from equation (15). Recall that  $\kappa$  is the structural slope coefficient in the regional Phillips curve for nontradeables from our model, and  $\psi$  is the reduced-form slope coefficient in the type of regional empirical specification often run in prior work. To estimate equation (14), we replace expected future unemployment and relative prices with

17. We drop Arizona because of anomalous trends that we have not been able to investigate due to COVID-19 related access restrictions at the BLS.

18. Industry-state employment data are available from the QCEW, at quarterly frequency for two-digit SIC codes (1975–2000) and three-digit NAICS codes (1990–2017). Before 1990 we use two-digit SIC codes to define industry, whereas after 2000 we use three-digit NAICS codes. For the period 1990–2000, when both the NAICS and SIC code classifications are available, we construct both versions of the instrument and use a simple average of the two.

19. Using the X-11 algorithm for seasonal adjustment yields virtually identical results.

their realized values and an expectation error. We also truncate the infinite sum in [equation \(14\)](#) at  $j = T$ . Doing this yields

$$(17) \quad \pi_{it}^N = \alpha_i + \gamma_t - \kappa \sum_{j=0}^T \beta^j u_{i,t+j} - \lambda \sum_{j=0}^T \beta^j \hat{p}_{i,t+j}^N + \tilde{\omega}_{it}^N + \eta_{it}^N,$$

where  $\eta_{it}^N$  denotes an expectations error that is orthogonal to information known at time  $t$  (and a truncation error). [Equation \(17\)](#) can now be estimated with standard GMM methods, that is, by instrumenting for the two forward sums. We do not attempt to estimate  $\beta$ . Rather, we set it to a standard quarterly value of  $\beta = 0.99$ .

We present results for two approaches to identifying the coefficients  $\kappa$  and  $\lambda$  in [equation \(17\)](#). Our first approach is to instrument for the two forward sums with four-quarter lagged unemployment  $u_{i,t-4}$  and the four-quarter lagged relative price of nontradeables  $\hat{p}_{i,t-4}^N$ . Assuming rational expectations, these lagged variables will be uncorrelated with the expectations error  $\eta_{it}^N$ . The identifying assumption regarding supply shocks is that when one state experiences a boom or bust relative to another state, it does not systematically experience nontradeable supply shocks relative to this other state. For example, when Texas experiences a recession relative to Illinois, this is not systematically correlated with changes in restaurant technology in Texas relative to Illinois. Notice that national supply shocks are absorbed by the time fixed effects, so only regional nontradeable supply shocks are potential confounders.

Our second approach to identification is to construct an instrumental variable that captures variation in demand. The idea behind our instrumental variable is the notion that national variation in demand for specific tradeable goods will differentially affect labor demand for nontradeable goods in states that produce those tradeable goods. For example, an increase in oil prices will differentially affect labor demand in Texas (and other oil-producing states). As a result, wages in Texas will rise, differentially affecting costs of nontradeables in Texas. Building on this idea, we construct a tradeable-demand spillovers instrument as

$$(18) \quad \text{Tradable Demand}_{i,t} = \sum_x \tilde{S}_{x,i} \times \Delta_{3Y} \log S_{-i,x,t},$$

where  $\bar{S}_{x,i}$  is the average employment share of industry  $x$  in state  $i$  over time, and  $\Delta_{3Y} \log S_{-i,x,t}$  is the three-year growth in national employment of industry  $x$  at time  $t$  excluding state  $i$ . This shift-share instrument builds on [Bartik \(1991\)](#) and more closely on [Nguyen \(2014\)](#). The identifying assumption in this case is that there are no supply factors that are correlated with the shifts  $\Delta_{3Y} \log S_{-i,x,t}$  in the time series and correlated with the shares  $\bar{S}_{x,i}$  in the cross section. For example, costs will increase as a result of an increase in oil prices. But if such cost increases are no larger on average for restaurants in Texas than Illinois they will be uncorrelated with our instrument.<sup>20</sup>

Our panel data approach implies that we are relying on cross-state variation in unemployment to identify the slope of the Phillips curve. [Figure IV](#) shows the evolution of the unemployment rate for three states, California, Texas, and Pennsylvania, over our sample period. While there is certainly a great deal of comovement, this figure illustrates well that there is also substantial cross-state variation. One example is that the 1991 and 2007–2009 recessions affected California much more than Texas and Pennsylvania. Another is that Texas experienced a recession in the mid-1980s (partly due to the savings and loan crisis and partly to a large fall in oil prices) while most other states experienced a continued fall in unemployment. Estimates of unemployment at the state level may be plagued by measurement error. Our IV estimation will address this insofar as the measurement error is classical.

The dependent variable in our regressions is  $\pi_{it}^N = p_{it}^N - p_{i,t-4}^N$ , that is, state-level nontradeable inflation over the previous 12 months. Studying inflation over four quarters allows us to reduce measurement error and eliminate seasonality. In [Online Appendix A.11](#), we show that using 12-month inflation as our dependent variable implies that we need to divide our estimates of  $\kappa$  and  $\lambda$  from [equation \(17\)](#) by four to account for the time aggregation. Recall that the inflation rate in our model in [Section III](#) is a quarterly inflation rate.

We truncate the discounted sums on the right-hand side of [equation \(17\)](#) at  $T = 20$  quarters. [Online Appendix Table C.4](#) presents robustness regarding this choice for our main

20. In a related approach, [McLeay and Tenreyro \(2019\)](#) use identified demand shocks from government spending to estimate the slope of the regional Phillips curve.

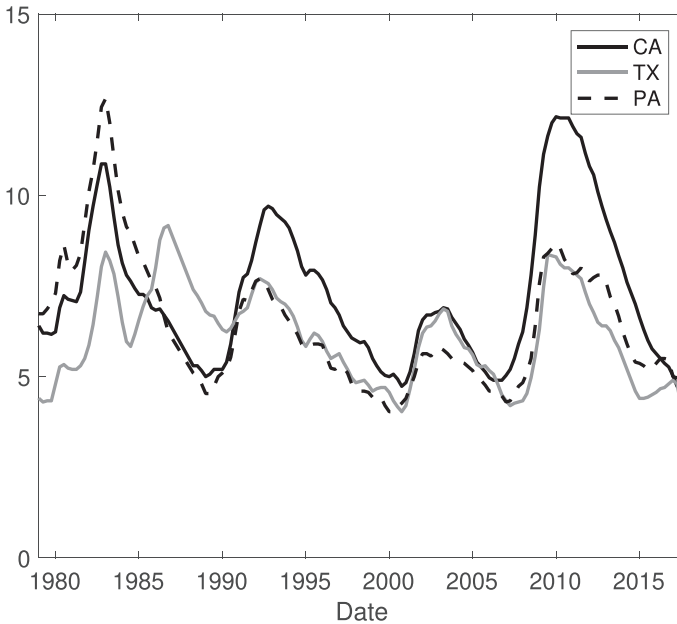


FIGURE IV

## Regional Business Cycles

This figure plots the unemployment rate for California, Pennsylvania, and Texas.

specification. Our results are similar for values between  $T = 20$  and  $T = 40$ . In [Online Appendix A.12](#), we estimate  $\kappa$  using [equation \(17\)](#) with  $T = 20$  on data simulated from our model from [Section III](#). We find that our empirical procedure is able to accurately estimate the true value of  $\kappa$  in this setting for a very wide range of true values of  $\kappa$ .

The forward sums in [equation \(17\)](#) imply that we lose five years of observations at the end of our sample when we set  $T = 20$ . To minimize the effect of this, we use a two-sample two-stage least squares (2SLS) regression. We estimate the first stage on a reduced sample without the last five years and the second stage on the full sample. We cluster standard errors at the state level and apply a correction to our standard errors appropriate for two-sample 2SLS developed by [Chodorow-Reich and Wieland \(2020\)](#).<sup>21</sup>

21. Our tradeable-demand instrument uses all the information in national industry employment growth rates. So our standard errors are not subject to the concerns about inference with shift-share instruments raised by [Adão, Kolesár, and Morales \(2019\)](#).

TABLE I  
SLOPE OF THE REGIONAL PHILLIPS CURVE

	No fixed effects (1)	No time effects (2)	Lagged unempl. (3)	Tradeable demand IV (4)
Panel A: Estimates of $\kappa$ from equation (17)				
$\kappa$	-0.0037 (0.0013)	0.0003 (0.0019)	0.0062 (0.0028)	0.0062 (0.0025)
Panel B: Estimates of $\psi$ from equation (19)				
$\psi$	-0.103 (0.036)	0.017 (0.027)	0.112 (0.057)	0.339 (0.126)
State effects		✓	✓	✓
Time effects			✓	✓

*Notes.* This table presents estimates of  $\kappa$  and  $\psi$  from regression specifications (17) and (19), respectively. The outcome variable is cumulative nontradeable inflation over four quarters, measured in percentage points. We include time and state fixed effects as noted at the bottom of each column. In Panel A, the regressors are the discounted future sum of quarterly state unemployment, in percentage points, and the discounted future sum of the relative price of nontradeables, in  $100 \times \log$  points. For both variables, we truncate the discounted future sum at 20 quarters. In Panel B, the regressors are the fourth lags of quarterly state unemployment, measured in percentage points, and the relative price of nontradeables. In the first three columns we instrument using the fourth lags of quarterly state unemployment and the relative price of nontradeables (this is OLS for  $\psi$ ). In the fourth column, we replace lagged unemployment with our tradeable-demand instrument among the instruments. In all columns, we estimate  $\kappa$  by two-sample 2SLS, and apply the correction to our standard errors from Chodorow-Reich and Wieland (2020). The sample period is 1978–2018. Standard errors are reported in parentheses, clustered by state. All regressions are unweighted. The number of observations is 3,323 in the first three columns of Panel A, with slightly fewer in the last column due to differencing. Likewise, the number of observations is 4,490 in the first three columns of Panel B.

Our empirical specification for estimating  $\psi$  is

$$(19) \quad \pi_{it}^N = \alpha_i + \gamma_t - \psi u_{i,t-4} - \delta p_{i,t-4}^N + \varepsilon_{it}.$$

We use beginning-of-period unemployment and relative price of nontradeables as regressors for consistency with previous studies such as Ball and Mazumder (2019). We present results for two identification approaches analogous to those we use for  $\kappa$ . The first approach is to estimate equation (19) by OLS (i.e., instrumenting for lagged unemployment and the relative price of nontradeables with themselves). The second approach replaces lagged unemployment among the instruments with our tradeable-demand instrument.

#### V.A. Full-Sample Results

Table I presents estimates of  $\kappa$  and  $\psi$  for our full sample period of 1978–2018. Let's start by considering the estimates of  $\kappa$  in Panel A. When we estimate equation (17) without fixed effects, our estimate has the “wrong” sign, that is, higher unemployment is



associated with higher rather than lower inflation ( $\kappa = -0.0037$ ). Adding state fixed effects raises the estimate of  $\kappa$  to 0.0003. Adding time fixed effects further raises the estimate of  $\kappa$  to 0.0062. As we stress throughout the article, time fixed effects eliminate changes in long-run inflation expectations. Finally, using our tradeable-demand instrument as opposed to instrumenting with lagged unemployment yields virtually the same estimate for  $\kappa$  of 0.0062. The fact that our estimate of  $\kappa$  does not change between columns (3) and (4) suggests that the fixed effects we include are sufficient to absorb supply shocks.

Our estimated slope of the Phillips curve is statistically significantly different from zero. In absolute size, however, the slope is small in the sense that it is consistent with the modest response of inflation to changes in unemployment seen in the aggregate time series since 1990. We develop this implication in [Section V.D. Online Appendix](#) Table C.1 presents estimates of the “first-stage” regressions for our IV estimates of [equation \(17\)](#). These first-stage regressions show that our instruments are strong instruments. We separately regress the present value of unemployment and the present value of relative prices on the reduced-form regressors. Lagged unemployment and tradeable demand both strongly predict the present value of unemployment and weakly predict the present value of relative prices. Lagged relative prices strongly predict the present value of relative prices and weakly predict the present value of unemployment.

[Online Appendix](#) Table C.2 reports our estimates of  $\lambda$  for regression specification [\(17\)](#)—the coefficient on the relative price of nontradeables. In our preferred specifications with time and state fixed effects, we estimate values of  $\lambda$  between 0.002 and 0.003. In the model we present in [Section III](#),  $\lambda$  provides an estimate of the degree of nominal rigidities. In a world with flexible prices, our estimate of  $\lambda$  would be large. The fact that our estimate of  $\lambda$  is very small provides further support—over and above our estimate of  $\kappa$ —for the notion that prices are quite rigid in the U.S. economy.

In our baseline results, we calibrate  $\beta = 0.99$ . It may be that firms are considerably less forward looking when they set prices than this calibration implies. Recent work has shown that plausible deviations from full rationality or common knowledge yield a Phillips curve that is less forward looking ([Angeletos and Lian 2018](#); [Gabaix 2020](#)). Also, a model with a combination of sticky information and sticky prices yields a Phillips curve that is less forward looking. [Online Appendix](#) Table C.3 presents estimates of

$\kappa$  where we calibrate  $\beta$  to lower values. As we vary our quarterly calibration of  $\beta$  from 0.99 to 0.9,  $\kappa$  doubles in size. The absolute size of the increase is small because our initial estimate of  $\kappa$  is small.

Our estimates of  $\psi$  in Table I, Panel B have a similar pattern to our estimates of  $\kappa$  discussed above. The estimate without time or state fixed effects is negative, and the estimate increases as we include state and then time fixed effects. An important difference is that the absolute size of our estimates of  $\psi$  are much larger than our estimates of  $\kappa$ . This reflects the fact that in equation (19) the lagged unemployment rate is standing in for the entire future sum in equation (17). Since unemployment is quite persistent, time variation in the future sum is much larger than time variation in the unemployment rate, which results in a much larger coefficient in equation (19) than in equation (17).

Another difference is that  $\psi$  is much larger in column (4) than in column (3), whereas  $\kappa$  is virtually identical. This reflects the fact that the tradeable-demand instrument we use in column (4) is more persistent than the unemployment rate itself. The coefficients in column (4) are therefore identified using more persistent variation, which results in a larger value of  $\psi$  but not a larger value of  $\kappa$ . This highlights an important advantage of estimating  $\kappa$  as opposed to  $\psi$ : estimates of  $\psi$  are hard to interpret because they are sensitive to the persistence of the variation used to identify them. More generally,  $\kappa$  is a structural parameter, and  $\psi$  is not. This implies that  $\psi$  may differ depending on the setting being considered (e.g., may be low in response to a policy change that may be reversed due to a future change in government), while  $\kappa$  is policy invariant.

### V.B. *Subsample Results*

We analyze to what extent the Phillips curve was steeper during the period of the Volcker disinflation than in subsequent years. Table II presents estimates of  $\kappa$  and  $\psi$  for the periods 1978–1990 and 1991–2018. We present these estimates for specifications with and without time fixed effects. All specifications include state fixed effects and control for the relative price of nontradeables.

Consider first the specification without time fixed effects reported in columns (1) and (2). For the pre-1990 sample,  $\kappa$  is estimated to be 0.0278, while  $\psi$  is estimated to be 0.449. In sharp contrast, for the post-1990 sample,  $\kappa$  is estimated to be 0.0002

TABLE II  
HAS THE PHILLIPS CURVE FLATTENED?

	Lagged unempl. IV without time fixed effect		Lagged unempl. IV with time fixed effect		Tradeable-demand IV with time fixed effect	
	Pre-1990 (1)	Post-1990 (2)	Pre-1990 (3)	Post-1990 (4)	Pre-1990 (5)	Post-1990 (6)
Panel A: Estimates of $\kappa$ from <a href="#">equation (17)</a>						
$\kappa$	0.0278 (0.0025)	0.0002 (0.0017)	0.0107 (0.0080)	0.0050 (0.0040)	0.0109 (0.0062)	0.0055 (0.0028)
Panel B: Estimates of $\psi$ from <a href="#">equation (19)</a>						
$\psi$	0.449 (0.063)	0.009 (0.025)	0.198 (0.113)	0.090 (0.057)	0.422 (0.232)	0.332 (0.157)

*Notes.* The table presents estimates of  $\kappa$  and  $\psi$ , before and after 1990. Columns (1), (3), and (5) present results for the sample period 1978–1990; and columns (2), (4), and (6) for the sample period 1991–2018. All specifications include state fixed effects. Specifications in columns (3)–(6) include time fixed effects. The instruments in columns (1)–(4) are the fourth lag of the tradeable-demand instrument and the relative price of nontradeables (i.e., OLS in Panel B). In columns (5) and (6), the instruments are the fourth lag of the tradeable-demand instrument and the relative price of nontradeables. In all columns, we estimate  $\kappa$  by two-sample 2SLS and apply the correction to our standard errors from [Chodorow-Reich and Wieland \(2019\)](#). Standard errors are reported in parentheses, clustered by state. All regressions are unweighted.

and  $\phi$  is estimated to be 0.009. The difference across samples is roughly a factor of 100 for  $\kappa$  and 50 for  $\psi$ . In other words, aggregate inflation became much less sensitive to unemployment after 1990 than it was during the Volcker disinflation.

Contrast this with the results in columns (3)–(4) where time fixed effects are included in the regressions. In this case, the estimated values of  $\kappa$  and  $\psi$  fall only modestly between the early part of the sample and the later part of the sample. For the pre-1990 sample,  $\kappa$  is estimated to be 0.0107 and  $\psi$  is estimated to be 0.198. For the post-1990 sample,  $\kappa$  is estimated to be 0.0050 and  $\psi$  is estimated to be 0.090. The difference across samples is roughly a factor of two and is not statistically significant. The estimates for  $\kappa$  in columns (5) and (6) are very similar to the estimates in columns (3) and (4), and the estimates of  $\psi$  in columns (5) and (6) show an even smaller difference across sample periods.

As we emphasize in [Section II](#), estimates of the Phillips curve based on time-series variation—such as the estimates without time fixed effects in [Table II](#)—are likely to be heavily influenced by time-series variation in long-run inflation expectations  $E_t\pi_{t+\infty}$ . In contrast, the specifications in [Table II](#) that include time fixed effects difference out the influence of long-run inflation expectations. The results in [Table II](#) therefore suggest that the apparent flattening of the Phillips curve in the time series is largely due to

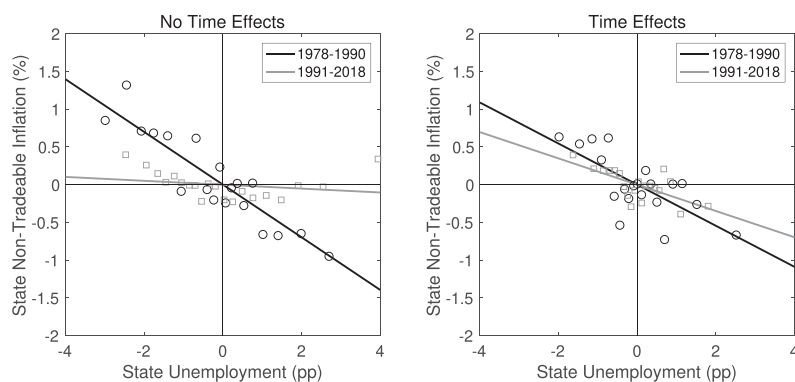


FIGURE V

Scatterplots: Nontradeable Inflation and Unemployment

In the left panel, we residualize state nontradeable inflation and unemployment against state fixed effects and the relative price of nontradeables, before and after 1990. We then plot residualized inflation and unemployment, before and after 1990, grouped by 20 bins of state unemployment. The right panel carries out the same exercise after further residualizing against time fixed effects. In both panels, inflation is cumulated over the previous four quarters and unemployment is lagged by four quarters.

inflation expectations becoming more firmly anchored over time. In the early part of the sample, inflation expectations shifted a great deal and these shifts were negatively correlated with the unemployment rate, which meant that shifts in inflation expectations masqueraded as a steep Phillips curve. The cross-sectional results in [Table II](#), columns (3)–(6) reveal that in fact the Phillips curve has always been quite flat (at least since 1978).

[Figure V](#) provides a visual representation of the results in [Table II](#). In the left panel, we plot a binned scatterplot of state-level nontradeable inflation against state-level unemployment after removing state fixed effects and the effects of the relative price of nontradeables. We plot the data separately for 1978–1990 and 1991–2018. The plot also includes regression lines for each subsample. The data in this panel do not account for time fixed effects and therefore include aggregate time series variation. As a consequence, we see a huge flattening of the Phillips curve in this case.

Contrast this with the right panel in [Figure V](#). This is an analogous figure to the left panel except that we also demean by time fixed effects. These data therefore only reflect regional variation in inflation. In this case, the difference in the slope of

TABLE III  
OUR ESTIMATES COMPARED TO PRIOR WORK

	$\kappa$
Rotemberg and Woodford (1997)	0.019
Galí (2008)	0.085
Nakamura and Steinsson (2014)	0.0077
Our full sample IV estimate	0.0062

*Notes.* We adjust the estimates from Rotemberg and Woodford (1997), Galí (2008), and Nakamura and Steinsson (2014) by the elasticity of output with respect to employment in the model in these papers. For Nakamura and Steinsson (2014), we use the calibration with GHH preferences.

the Phillips curve between the early sample and the late sample is modest. The modest flattening of the Phillips curve that we find over our sample (once we account for time fixed effects) seems consistent with the fact that the frequency of price change in the United States has declined by about 40% as inflation has fallen since the early 1980s (Nakamura et al. 2018).

#### V.C. How Do Our Estimates Compare to Prior Work?

It is instructive to compare our estimate of  $\kappa$  to values of  $\kappa$  arrived at by means of structural estimation or calibration of New Keynesian models. Table III reports three such estimates from Rotemberg and Woodford (1997), Galí (2008), and Nakamura and Steinsson (2014). In all cases, we have adjusted the reported value of  $\kappa$  from these papers by the elasticity of output with respect to employment. As is well known, the value of  $\kappa$  in a New Keynesian model is highly dependent on the degree of nominal and real rigidities assumed. The values for  $\kappa$  used in these papers range from about an order of magnitude larger than our estimated value to a value roughly equal to our estimated value. The main difference between Galí's relatively high value and the much lower values in Rotemberg and Woodford (1997) and Nakamura and Steinsson (2014) lies in the degree of real rigidity that the models used in these papers imply. Galí's model is a relatively simple (textbook) version of the New Keynesian model, which does not incorporate strong sources of real rigidity. Rotemberg and Woodford (1997) and Nakamura and Steinsson (2014) use models with heterogeneous labor markets, which yields a much larger amount of real rigidity. In both cases, the large amount of real rigidity helps these authors match moments that they target in their analysis. Similarly, our estimates imply that the data we have analyzed are also

more consistent with New Keynesian models that incorporate a large amount of real rigidity.

#### *V.D. Aggregate Implications*

A question that naturally arises regarding our cross-sectional estimates of  $\kappa$  is whether they can explain the aggregate time series variation in inflation over our sample. A number of researchers and commentators have suggested that the stability of inflation at the aggregate level in the United States has been surprising over the past 25 years (missing disinflation during the Great Recession and missing reflation during the late 1990s and late 2010s). Some researchers have recently argued that cross-sectional variation suggests a steeper Phillips curve than time series variation for the past few decades. Here, we assess whether this is the case for our estimates.

We start with the solved-forward aggregate Phillips curve—[equation \(6\)](#). In [Section II](#), we made the simplifying assumption that the unemployment rate follows an AR(1). This assumption allowed us to derive a simple aggregate relationship between the discounted future sum of unemployment rates in [equation \(6\)](#) and the current unemployment rate—see [equation \(7\)](#). In reality, the dynamics of the U.S. unemployment rate differ substantially from an AR(1) (see [Neftçi 1984](#); [Sichel 1993](#); [Dupraz, Nakamura, and Steinsson 2020](#)). For this reason, we adopt an approach of estimating a scaling factor  $\zeta$  that relates the current unemployment rate to the discounted future sum in [equation \(6\)](#) using the following regression

$$(20) \quad \sum_{j=0}^T \beta^j \tilde{u}_{t+j} = \zeta \tilde{u}_t + \alpha + \epsilon_t.$$

The series we use for  $\tilde{u}_t$  in this regression is the difference between the aggregate unemployment rate in the United States and the CBO's estimate of the natural rate of unemployment at each point in time.<sup>22</sup> We run this regression for the sample period 1979Q4–2017Q4. This yields an estimate of  $\zeta$  for aggregate variation in the unemployment rate of 6.16 with a Newey-West standard error of

22. We are thus treating the CBO's estimate of the natural rate of unemployment as a forecast of long-run unemployment  $E_t u_{t+\infty}$ .

1.80. Using [equation \(20\)](#), we can rewrite [equation \(6\)](#) as

$$(21) \quad \pi_t - E_t \pi_{t+\infty} = -\kappa \zeta \tilde{u}_t + \omega_t.$$

Our cross-sectional estimates of  $\kappa$  are for nontradeables excluding housing services. As we emphasize in [Section II](#), the treatment of housing services has important implications for the behavior of inflation. [Online Appendix Table C.5](#) presents estimates of  $\kappa$  and  $\psi$  using state-level annual rent inflation data from the American Community Survey for 2001 to 2017. For our baseline specification with state and time fixed effects, we estimate  $\kappa$  to be 0.0243. This estimate of  $\kappa$  is roughly four times larger than our estimate of  $\kappa$  for nonhousing nontradeable goods reported in [Table I](#). We account for this difference below by taking a weighted average of our full-sample  $\kappa$  estimate for nontradeables and this  $\kappa$  estimate for housing services.<sup>23</sup>

[Figure VI](#) plots the left-hand side of [equation \(21\)](#) (black line) against the first term on the right-hand side of [equation \(21\)](#) (gray line) using our estimates of  $\kappa$  and  $\zeta$  from above, which yields  $\kappa \zeta = 0.34$ .<sup>24</sup> We use the 10-year ahead SPF inflation expectations for the CPI as our measure of long-term inflation expectations. The gray line is the demand-induced variation in inflation predicted by our estimates. The figure indicates that the amplitude of inflation fluctuations over the last few business cycles has been roughly in line with what our cross-sectional estimates of  $\kappa$  suggest. In particular, the disinflation during the Great Recession and re-inflation during the 2010s lines up well with what our estimate of  $\kappa$  implies. If the gray line had a larger amplitude than the black line over the business cycle, this would indicate missing disinflation and missing re-inflation. In fact, the amplitude of the gray line is very similar to that of the black line for the Great Recession, the post-Great Recession recovery, and the long 1990s expansion. By this metric, there is thus no missing disinflation or missing re-inflation over this period. These findings echo the results of [Ball and Mazumder \(2019\)](#).<sup>25</sup>

23. We use the shelter and nonshelter expenditure weights in the core CPI. These are 0.42 and 0.58, respectively.

24. The coefficient  $\kappa \zeta$  is calculated as  $4 \times (0.58 \times 0.0062 + 0.42 \times 0.0243) \times 6.16$ , where the factor of 4 accounts for time aggregation to annual inflation.

25. [Online Appendix Figure C.2](#) shows that the modest flattening of the Phillips curve we estimate over our sample period has a minimal effect on the fluctuations in the gray line in [Figure VI](#). [Online Appendix Figure C.3](#) shows that a disproportionate share of the systematic variation in inflation and the fitted

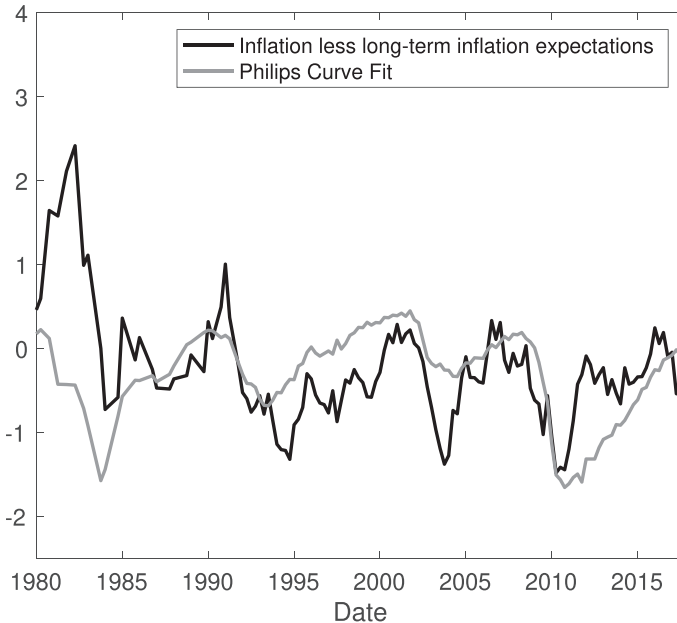


FIGURE VI

## Aggregate Phillips Curve

The figure shows the fit of the aggregate Phillips curve for core inflation. The black line is the difference between core inflation and the 10-years-ahead SPF inflation expectation for the CPI. The gray line plots the first term on the right-hand side of [equation \(21\)](#), which is the demand-induced variation in inflation predicted by our estimates.

The most substantial deviation between the actual and fitted values arises during the Volcker period when actual inflation relative to long-run expectations lies far above the fitted value. While the conventional view is that the Phillips curve has broken down after 1990, we are finding the opposite: a poor fit of our cross-sectional estimate of the Phillips curve when applied to aggregate inflation dynamics over the Volcker period. A natural interpretation of this discrepancy is the presence of adverse

---

value predicted by our model comes from the housing services (rent) component of the CPI. The figure is analogous to [Figure VI](#) except that the black line excludes housing and the fitted value uses only the  $\kappa$  estimate for the nonshelter component of inflation. We see that core inflation excluding housing services varies much less systematically than core inflation including housing services.



supply shocks in the early 1980s, for example, associated with the oil price shocks.

How much of the fall in inflation during the Volcker disinflation can be attributed to the causal effect of higher unemployment working through the slope of the Phillips curve according to our estimates? Unemployment rose by about 5 percentage points between 1979 and 1982. Using a weighted average of our slightly higher pre-1990 nonshelter estimate for  $\kappa$  and our estimate of  $\kappa$  for shelter, we find that this increase in unemployment caused inflation to fall by only about 2 percentage points (see gray line in [Online Appendix Figure C.4](#)). Core CPI inflation first rose from 7% to 10% from 1979 to 1981 and then fell to 4% by 1986. Clearly, the direct causal effect of unemployment working through the slope of the Phillips curve explains only a modest amount of this variation in inflation. Over this same period, long-run inflation expectations first rose from 7% to 8% and then fell to 4%. Our estimates, therefore, suggest that the bulk of the variation in inflation over the early 1980s is due to changes in long-run inflation expectations, with supply shocks also playing an important role.

## VI. CONCLUSION

This article provides new estimates of the slope of the Phillips curve. We estimate that the slope of the Phillips curve is small, and was small even during the Volcker disinflation of the early 1980s. Our results indicate that shifts in expectations about the conduct of monetary policy explain much of the drop of inflation in the early 1980s and more firmly anchored inflation expectations explain the stability of inflation since the mid-1990s. Our estimates are consistent with the insensitivity of inflation to unemployment during both the Great Recession and during the low unemployment periods of the late 1990s and late 2010s.

To reach these conclusions, we estimate the Phillips curve in the cross-section of U.S. states. We use newly constructed state-level price indices for nontradeable goods starting in 1978. We map from our regional estimates to the slope of the aggregate Phillips curve using a multiregion New Keynesian model. The model clarifies that the slope of the aggregate Phillips curve is equal to the slope of the regional Phillips curve for nontradeable goods. We also use the model to show that regional data difference

out the effects of the long-run monetary regime, which otherwise confound estimates of the slope of the Phillips curve. Guided by the model, we show that the conventional empirical specification used to estimate regional Phillips curves must be scaled by a factor relating to the persistence of unemployment fluctuations to yield an estimate of the slope of the Phillips curve. Finally, we develop a new tradeable-demand spillover instrument that allows for flexible patterns of supply shocks at the local level.

An important lesson from our analysis is that when it comes to managing inflation, the elephant in the room is long-run inflation expectations. This view contrasts sharply with the conventional view that managing inflation is about moving up and down a steep Phillips curve. A crucial question for inflation dynamics is why long-run inflation expectations are sometimes so firmly anchored but at other times move sharply? Beliefs about inflation in the long run are governed by beliefs about the long-run behavior of the monetary authority and ultimately the political process that shapes the long-run behavior of the monetary authority. Since this is fundamentally a very low-frequency phenomenon, it is not easily pinned down by half a century or so of data from a single country. While much interesting research has sought to understand the behavior of long-run inflation expectations, we believe it is still not sufficiently well understood and its crucial importance for the conduct of monetary policy implies that even more research should focus on this question.

LONDON SCHOOL OF ECONOMICS, UNITED KINGDOM  
UNIVERSITY OF CALIFORNIA, SAN DIEGO, UNITED STATES  
UNIVERSITY OF CALIFORNIA, BERKELEY, UNITED STATES  
UNIVERSITY OF CALIFORNIA, BERKELEY, UNITED STATES

#### SUPPLEMENTARY MATERIAL

An Online Appendix for this article can be found at *The Quarterly Journal of Economics* online.

#### DATA AVAILABILITY

Code replicating the tables and figures in this article can be found in Hazell et al. (2022) in the Harvard Dataverse, <https://doi.org/10.7910/DVN/OQNZYE>.

## REFERENCES

- Adam, Klaus, and Mario Padula, "Inflation Dynamics and Subjective Expectations in the United States," *Economic Inquiry*, 49 (2011), 13–25.
- Adão, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-Share Designs: Theory and Inference," *Quarterly Journal of Economics*, 134 (2019), 1949–2010.
- Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro, "Transparency in Structural Research," *Journal of Business and Economic Statistics*, 38 (2020), 711–722.
- Angeletos, George-Marios, and Chen Lian, "Forward Guidance without Common Knowledge," *American Economic Review*, 108 (2018), 2477–2512.
- Babb, Nathan, and Alan K. Detmeister, "Nonlinearities in the Phillips Curve for the United States: Evidence Using Metropolitan Data," Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series 2017-070, 2017.
- Ball, Laurence, and Sandeep Mazumder, "Inflation Dynamics and the Great Recession," *Brookings Papers on Economic Activity* (2011), 337–381.
- , "A Phillips Curve with Anchored Expectations and Short-Term Unemployment," *Journal of Money, Credit and Banking*, 51 (2019), 111–137.
- Barnichon, Régis, and Geert Mesters, "Identifying Modern Macro Equations with Old Shocks," *Quarterly Journal of Economics*, 135 (2020), 2255–2298.
- Bartik, Timothy J., "Who Benefits from State and Local Economic Development Policies?," WE Upjohn Institute for Employment Research, 1991.
- Beraja, Martin, Erik Hurst, and Juan Ospina, "The Aggregate Implications of Regional Business Cycles," *Econometrica*, 87 (2019), 1789–1833.
- Bernanke, Ben S., "Inflation Expectations and Inflation Forecasting," Speech given at the NBER Summer Institute, July 2007.
- Bianchi, Francesco, and Cosmin Ilut, "Monetary/Fiscal Policy Mix and Agents' Beliefs," *Review of Economic Dynamics*, 26 (2017), 113–139.
- Blanchard, Olivier, "The Phillips Curve: Back to the '60s?," *American Economic Review: Papers and Proceedings*, 106 (2016), 31–34.
- Calvo, Guillermo A., "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12 (1983), 383–398.
- Carvalho, Carlos, Stefano Eusepi, Emanuel Moench, and Bruce Preston, "Anchored Inflation Expectations," PUC-Rio Working Paper, 2021.
- Chodorow-Reich, Gabriel, and Johannes Wieland, "Secular Labor Reallocation and Business Cycles," *Journal of Political Economy*, 128 (2020), 2245–2287.
- Christiano, Lawrence, Martin Eichenbaum, and Mathias Trabandt, "Understanding the Great Recession," *American Economic Journal: Macroeconomics*, 7 (2015), 110–167.
- Cogley, Timothy, and Argia M. Sbordone, "Trend inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve," *American Economic Review*, 98 (2008), 2101–2126.
- Coibion, Olivier, and Yuriy Gorodnichenko, "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation," *American Economic Journal: Macroeconomics*, 7 (2015), 197–232.
- Coibion, Olivier, Yuriy Gorodnichenko, and Rupal Kamdar, "The Formation of Expectations, Inflation, and the Phillips Curve," *Journal of Economic Literature*, 56 (2018), 1447–1491.
- Crump, Richard K., Stefano Eusepi, Marc Giannoni, and Aysegül Sahin, "A Unified Approach to Measuring  $u^*$ ," *Brookings Papers on Economic Activity* (2019), 143–214.
- Del Negro, Marco, "Aggregate Risk Sharing across US States and across European Countries," Yale University Working Paper, 1998.
- Del Negro, Marco, Marc P. Giannoni, and Frank Schorfheide, "Inflation in the Great Recession and New Keynesian Models," *American Economic Journal: Macroeconomics*, 7 (2015), 168–196.

- Del Negro, Marco, Michele Lenza, Giorgio E. Primiceri, and Andrea Tambalotti, "What's Up with the Phillips Curve," *Brookings Papers on Economic Activity* (2020), 301–373.
- Dupraz, Stephane, Emi Nakamura, and Jón Steinsson, "A Plucking Model of the Business Cycle," University of California, Berkeley, Working Paper, 2020.
- Erceg, Christopher J., and Andrew T. Levin, "Imperfect Credibility and Inflation Persistence," *Journal of Monetary Economics*, 50 (2003), 915–944.
- Fitzgerald, Terry J., Callum Jones, Mariano Kulish, and Juan Pablo Nicolini, "Is There a Stable Relationship between Unemployment and Future Inflation?," Federal Reserve Bank of Minneapolis Working Paper, 2020.
- Fitzgerald, Terry J., and Juan Pablo Nicolini, "Is There a Stable Relationship between Unemployment and Future Inflation? Evidence from US Cities," Federal Reserve Bank of Minneapolis Working Paper, 2014.
- Friedman, Milton, "The Role of Monetary Policy," *American Economic Review*, 58 (1968), 1–17.
- Fuhrer, Jeffrey C., and George Moore, "Inflation Persistence," *Quarterly Journal of Economics*, 110 (1995), 127–159.
- Gabaix, Xavier, "A Behavioral New Keynesian Model," *American Economic Review*, 110 (2020), 2271–2327.
- Galí, Jordi, *Monetary Policy, Inflation, and the Business Cycle* (Princeton, NJ: Princeton University Press, 2008).
- Galí, Jordi, and Mark Gertler, "Inflation Dynamics: A Structural Econometrics Analysis," *Journal of Monetary Economics*, 44 (1999), 195–222.
- Galí, Jordi, Mark Gertler, and J. David Lopez-Salido, "Robustness of the Estimates of the Hybrid New Keynesian Phillips Curve," *Journal of Monetary Economics*, 52 (2005), 1107–1118.
- Gertler, Mark, and John Leahy, "A Phillips Curve with an Ss Foundation," *Journal of Political Economy*, 116 (2008), 533–572.
- Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajšek, "Inflation Dynamics During the Financial Crisis," *American Economic Review*, 107 (2017), 785–823.
- Goodfriend, Marvin, and Robert G. King, "The Incredible Volcker Disinflation," *Journal of Monetary Economics*, 52 (2005), 981–1015.
- Gordon, Robert, "Inflation, Flexible Exchange Rates, and the Natural Rate of Unemployment," In *Workers, Jobs, and Inflation*, Martin N. Baily, ed. (Washington, DC: Brookings Institute, 1982), 89–158.
- Greenwood, Jeremy, Zvi Hercowitz, and Gregory W. Huffman, "Investment, Capacity Utilization, and the Real Business Cycle," *American Economic Review*, 78 (1988), 402–417.
- Hazell, Jonathon, Juan Herreño, Emi Nakamura, and Jón Steinsson, "Replication Data for: 'The Slope of the Phillips Curve: Evidence from U.S. States'," (2022), Harvard Dataverse, <https://doi.org/10.7910/DVN/OQNZYE>.
- Hooper, Peter, Frederic S. Mishkin, and Amir Sufi, "Prospects for Inflation in a High Pressure Economy: Is the Phillips Curve Dead or Is It Just Hibernating?," NBER Working Paper no. 25792, 2019.
- Ireland, Peter N., "Changes in the Federal Reserve's Inflation Target: Causes and Consequences," *Journal of Money, Credit and Banking*, 39 (2007), 1851–1882.
- Jorgensen, Peter Lihn, and Kevin J. Lansing, "Anchored Inflation Expectations and the Flatter Phillips Curve," Federal Reserve Bank of San Francisco Working Paper, 2019.
- Kiley, Michael T., "An Evaluation of the Inflationary Pressure Associated with Short- and Long-Term Unemployment," *Economics Letters*, 137 (2015a), 5–9.
- , "Low Inflation in the United States: A Summary of Recent Research," FEDS Notes, Board of Governors of the Federal Reserve System, 2015b.
- Klenow, Peter J., and Oleksiy Kryvtsov, "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?," *Quarterly Journal of Economics*, 123 (2008), 863–904.

- Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H. Stock, "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve," *Journal of Economic Literature*, 52 (2014), 124–188.
- McLeay, Michael, and Silvana Tenreyro, "Optimal Inflation and the Identification of the Phillips Curve," *NBER Macroeconomics Annual*, 34 (2019), 199–255.
- Mian, Atif, and Amir Sufi, "What Explains the 2007–2009 Drop in Employment?," *Econometrica*, 82 (2014), 2197–2223.
- Mishkin, Frederick S., "Inflation Dynamics," *International Finance*, 10 (2007), 317–334.
- Nakamura, Emi, and Jón Steinsson, "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 123 (2008), 1415–1464.
- , "Fiscal Stimulus in a Monetary Union: Evidence from US Regions," *American Economic Review*, 104 (2014), 753–792.
- , "Identification in Macroeconomics," *Journal of Economic Perspectives*, 32 (2018), 59–86.
- Nakamura, Emi, Jón Steinsson, Patrick Sun, and Daniel Villar, "The Elusive Costs of Inflation: Price Dispersion during the US Great Inflation," *Quarterly Journal of Economics*, 133 (2018), 1933–1980.
- Neftçi, Salih, "Are Economic Time Series Asymmetric over the Business Cycle?," *Journal of Political Economy*, 92 (1984), 307–328.
- Nguyen, Ha, "Demand-Driven Propagation: Evidence from the Great Recession," World Bank Working Paper, 2014.
- Nishizaki, Kenji, and Tsutomu Watanabe, "Output-Inflation Trade-Off at Near-Zero Inflation Rates," *Journal of the Japanese and International Economies*, 14 (2000), 304–326.
- Obstfeld, Maurice, and Kenneth Rogoff, "Speculative Hyperinflations in Maximizing Models: Can We Rule Them Out?," *Journal of Political Economy*, 91 (1983), 675–687.
- Phelps, Edmund S., "Phillips Curves, Expectations of Inflation and Optimal Unemployment over Time," *Economica*, 34 (1967), 254–281.
- Phillips, Alban W., "The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957," *Economica*, 25 (1958), 283–299.
- Roberts, John M., "New Keynesian Economics and the Phillips Curve," *Journal of Money, Credit and Banking*, 27 (1995), 975–984.
- Rotemberg, Julio J., "Sticky Prices in the United States," *Journal of Political Economy*, 60 (1982), 1187–1211.
- Rotemberg, Julio J., and Michael Woodford, "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual*, 12 (1997), 297–346.
- Rubbo, Elisa, "Networks, Phillips Curves, and Monetary Policy," Harvard University Working Paper, 2020.
- Rudd, Jeremy, and Karl Whelan, "New Tests of the New-Keynesian Phillips Curve," *Journal of Monetary Economics*, 52 (2005), 1167–1181.
- Samuelson, Paul, and Robert Solow, "Analytical Aspects of Anti-Inflation Policy," *American Economic Review*, 50 (1960), 177–194.
- Sargent, Thomas J., "The Ends of Four Big Inflation," in *Inflation: Causes and effects*, Robert E. Hall, ed. (Chicago: University of Chicago Press, 1982), 41–98.
- Sbordone, Argia M., "Prices and Unit Labor Costs: A New Test of Price Stickiness," *Journal of Monetary Economics*, 49 (2002), 265–292.
- Sichel, Daniel, "Business Cycle Asymmetry: A Deeper Look," *Economic Inquiry*, 31 (1993), 224–236.
- Stock, James H., and Mark W. Watson, "Why Has US Inflation Become Harder to Forecast?," *Journal of Money, Credit and Banking*, 39 (2007), 3–33.
- , "Slack and Cyclically Sensitive Inflation," NBER Working Paper no. 25987, 2019.
- Taylor, John B., "Staggered Contracts in a Macro Model," *American Economic Review*, 69 (1979), 108–113.

- , “Aggregate Dynamics and Staggered Contracts,” *Journal of Political Economy*, 88 (1980), 1–24.
- Woodford, Michael, “Doing without Money: Controlling Inflation in a Post-Monetary World,” *Review of Economic Dynamics*, 1 (1998), 173–219.
- , *Interest and Prices* (Princeton, NJ: Princeton University Press, 2003).