

NBER WORKING PAPER SERIES

THE ACTIVE ROLE OF THE NATURAL RATE OF UNEMPLOYMENT

Robert E. Hall
Marianna Kudlyak

Working Paper 31848
<http://www.nber.org/papers/w31848>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2023, revised April 2024

Kudlyak presented some material from this paper as a part of her Alejandro Justiniano memorial keynote address at the annual Federal Reserve System Macroeconomics meeting, November 9 and 10, 2021. We thank Jordi Galí for providing an extended version of his estimates of the natural rate of unemployment. We thank Yuriy Gorodnichenko, Robert Hetzel and Jeff Lacker for insightful discussions. Hall's research was supported financially by the Hoover Institution. The opinions expressed are those of the authors and do not reflect those of the National Bureau of Economic Research, the Federal Reserve Bank of San Francisco, the Federal Reserve System, or any other institution with which authors are affiliated.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Robert E. Hall and Marianna Kudlyak. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Active Role of the Natural Rate of Unemployment
Robert E. Hall and Marianna Kudlyak
NBER Working Paper No. 31848
November 2023, revised April 2024
JEL No. E32,J63,J64

ABSTRACT

We propose that the natural rate of unemployment may have an active role in the business cycle, in contrast to a widespread view that the rate is fairly smooth and at most only weakly cyclical. We demonstrate that the tendency to treat the natural rate as near-constant would explain the surprisingly low slope of the Phillips curve. We observe that evidence is weak about this basic point—the evidence neither comes close to rejecting the conventional view nor does it reject a very different view in which fluctuations in the natural rate are associated with a substantial fraction of cyclical volatility. We show that the natural rate may have closely tracked the actual rate during the long recovery that began in 2009 and ended in 2019. We explain how the common finding of research in the Phillips-curve framework of low—often extremely low—response of inflation to unemployment could be the result of fairly close tracking of the natural rate and the actual rate in recoveries. Our interpretation of the data contrasts to that of many Phillips-curve studies, that conclude that inflation has little relation to unemployment. But we also call attention to a recently growing body of research along New Keynesian lines that supports the credibility of the hypothesis of a strong relation between movements of the natural rate and actual rate of unemployment.

Robert E. Hall
Hoover Institution
Stanford University
Stanford, CA 94305-6010
and NBER
rehall@gmail.com

Marianna Kudlyak
101 Market St
Federal Reserve Bank of San Francisco
San Francisco, CA 94105
mkudlyak@gmail.com

1 Introduction

Following its introduction in Friedman (1968), the natural rate of unemployment became an essential element of thinking about the business cycle, about inflation, and about monetary policy. In the New Keynesian analysis that emerged from Friedman’s insights, the natural rate of unemployment is the unemployment rate that would hold if the economy were in equilibrium, free of transitory forces that raise inflation above its normal level or depress inflation below that level.

We investigate the relation between the actual rate of unemployment and the natural rate. We start by describing the relation between inflation and its two determinants in the Phillips curve. One is the inflation anchor, the equilibrium value described above. The other is the gap between the actual unemployment rate and the natural unemployment rate.

We show that the basic Phillips curve model is not identified—observations of the inflation gap and unemployment do not pin down the corresponding values of the underlying parameters, the natural unemployment rate and the slope of the Phillips curve. An inverse relation between the slope and the unemployment gap implies that if the natural rate is close to the actual rate, the slope is large and the Phillips curve is steep, and vice versa. Faced with this problem, Phillips-curve researchers have come up with potential solutions. These methods deliver vastly different time series of the natural rate in terms of its correlation with the actual unemployment rate and other aspects.

We add to the stock of proposed solutions with what we call the anchored-inflation method, in which the natural rate is identified from the actual unemployment rate for months when inflation is equal to its anchored value. We study the cyclical recovery beginning in 2009 and ending in early 2020. During that recovery, inflation stayed close to the Fed’s target of two percent per year. We posit that the inflation anchor remained close to constant at two percent. Accordingly, the unemployment rate was close to the natural rate. Observed unemployment glided down from 10.0 percent at its maximum to 3.5 percent at its minimum. Our conclusion is that the natural rate declined along fairly close to the same path. That is, the natural unemployment rate stayed close to the actual rate during the recovery.

We describe three categories of existing methods for constructing estimates of the natural rate of unemployment. First, the widely-used unemployment-trend method is based on the assumption that the natural unemployment rate is constant or a slow-moving index of demographic change. Second, the statistical-submodel method harnesses multivariate statistical methods to form an index of the natural unemployment rate. With this method, inflation pressure typically follows an AR-process and the natural rate follows another random walk. This approach yields a path that is more volatile and correlated with the actual

unemployment rate. Third, the macro-model approach uses a general-equilibrium model to calculate the natural rate as the unemployment rate in a counterfactual equilibrium of the model without sticky prices and wages. This approach also finds that the natural rate tracks the actual rate fairly closely.

We demonstrate that if the true natural rate is substantially positively correlated with the actual rate, a Phillips curves estimated with a constant in place of the true natural rate will likely appear to be close to flat. The same conclusion follows if the Phillips-curve estimation is based on a time series for the natural rate that is uncorrelated or only weakly correlated with the actual rate.

Historically, most investigators identified the slope of the Phillips curve by assuming that the natural rate of unemployment is constant over time or that it moves slowly on a path determined by changes in demography. Under that type of assumption, the Phillips curve is found to be quite flat, especially after 2000.

The resolution of the Phillips curve identification problem has profound implications for a key question in modern macroeconomics—is inflation sticky or flexible? As we noted above, a large branch of the Phillips-curve literature posits that the natural rate of unemployment is constant or weakly correlated with actual unemployment. Under that view, in recoveries with anchored inflation, unemployment must be highly variable relative to the natural rate, because unemployment falls so far during recoveries. In those times, the slope of the Phillips curve must be close to zero. This is a sticky view of inflation.

The differentiating factor between the two views is flexibility of prices. According to the flexible view of inflation, the negative slope of the Phillips curve is substantial, while the gap between the actual and natural unemployment rate is small and transitory. A frequently used name for this literature is the “real business cycle model”, a name that calls attention to the important role of flexible inflation in limiting the effects of monetary policy on output and employment.

Under the sticky view, widely present in the literature, the Phillips-curve slope is small. An implication of the sticky view is that the unemployment gap can be large and persistent.

With respect to the measurement of the driving force of inflation, we find that, during recoveries, unemployment by itself is not the proper measure. Rather, that measure is the gap between the actual unemployment rate and the natural rate. This finding explains why recoveries can continue over extended periods, as long as a decade, when unemployment is gliding downward by many percentage points, while inflation is close to constant. Another way to express the same finding is that the labor market can gradually tighten in the sense of the Diamond-Mortensen-Pissarides model’s measure of labor-market tightness, the vacancy/unemployment ratio, while inflation remains at a constant low level.

This paper is mainly about the labor market during cyclical recoveries. Most of the time, the labor market evolves smoothly during recoveries, which last a few years in some cases and more than a decade at the longest. Each recovery ends with an important adverse aggregate shock. During the ensuing recessions, unemployment rises rapidly. The paper has relatively little to say about recessions.

In a related earlier paper, Hall and Kudlyak (2022a), we find that actual unemployment behaves much the same way across recoveries. Rather than vibrating around a fixed natural rate, the observed behavior of unemployment comprises (1) occasional sharp upward movements in times of economic crisis (recessions), and (2) patterns of inexorable downward glide at a low but reliable proportional rate. The glide continues until unemployment reaches a low barrier of approximately 3.5 percent or until another economic crisis interrupts the glide. The fact that actual unemployment behaves in much the same way in all recoveries stimulated our interest in the hypothesis that the natural rate behaves the same way. In another related paper, Hall and Kudlyak (2022b), we investigate a variety of mechanisms that account for this finding.

This paper is organized as follows. Section 2 describes inflation's relation to the inflation anchor and the unemployment gap through the Phillips curve and sets forth the definition of the natural rate of unemployment used in this paper, derived from Friedman's pioneering analysis. Section 3 discusses the failure of identification of the slope parameter of the basic Phillips curve. Section 4 describes a method of inferring the natural rate of unemployment during periods of anchored inflation. Section 5 describes other existing methods of modelling and calculating the natural rate. Section 6 deals with the bias toward understatement of the slope of the Phillips curve that occurs when the slope is estimated without consideration of the correlation of the actual and natural unemployment rates. This bias may help understanding the tendency for Phillips curves estimated during the long recovery from 2009 to 2019 to be unexpectedly flat. Section 7 reviews the evidence in favor the hypothesis that inflation is sticky or the hypothesis that it is flexible. Section 8 discusses the effect of the pandemic on the Phillips curve.

2 The Inflation Process and the Phillips Curve

2.1 The sticky-price model of the inflation process

Modern macroeconomics combines a longer-term flexible-price model of employment, output, and the price level with a short-term model of deviations from that model. In the longer run, the price level adjusts so supply equals demand in the output and labor markets. market and similarly for the wage in the labor market. The government operates a central bank that

issues interest-bearing reserves. If the bank cuts the rate paid on reserves, or, equivalently, increases the volume of reserves, the price and wage levels increase, so one dollar buys a smaller amount of output. The price level adjusts as needed to maintain equilibrium. Inflation is the rate of change in the price level. Under the widely accepted New Keynesian paradigm, for many goods and services, prices do not clear markets quickly. Wages may be even stickier. Accordingly, the model adds price-setting inertia—sellers face lags before they can change their prices.

The New Keynesian inflation model views sellers as considering two forces in deciding how to set a price today that remains in force some time into the future. One is that inflation has a component that reflects the success of monetary policy in stabilizing inflation at a low and reasonably constant level. We call this component the inflation anchor. It is the result of an inference that a price-setter makes about how to set a price that will remain in effect for some time into the future. The basis for the inference includes information about the likely success of the central bank in stabilizing inflation in the near future, recent actual inflation, and the sources of that inflation. If monetary policy loses its grip and high inflation sets in, the anchor rises to reflect that development.

The second force operates at business-cycle frequencies. It captures the negative relation between inflation and unemployment. In the New Keynesian model, this force is measured by the gap between the actual unemployment rate and the natural unemployment rate. The level of downward pressure from unemployment is not observed directly because it depends on the unemployment rate minus the natural rate, and the natural rate is not observed.

The source of the negative relation between the unemployment gap and the rate of inflation in the basic New Keynesian model is the following: If unemployment is below the natural rate, it means that the price level is below its equilibrium value. The reason is that the sticky price level is below equilibrium and the consumers and firms, therefore, demand more than the equilibrium levels of output. Their demands are automatically satisfied by producers because, according to a basic Keynesian principle, demand determines output. Higher output implies lower unemployment. As time passes, the previously stuck prices become unstuck, prices free up, the price level rises from their low levels back toward equilibrium. Inflation continues higher than the anchored rate during this process.

Sellers choose the amount of inertia in pricing by varying the time between repricing events: more frequent repricing reflects greater response to the pressure. Sellers can also vary the size of the repricing change.




Figure 1: The New Keynesian Model in a Phase Diagram

2.2 Inflation and unemployment in the basic model

Figure 1 displays the basic model in a phase diagram. The unemployment rate u is on the horizontal axis and the inflation rate π is on the vertical axis. The natural unemployment rate, u^* , is marked on the u axis and the anchored inflation rate, π^* , is marked on the π axis. The natural rate and the anchored rate are two key parameters of the model.

A conspicuous dot with coordinates $[u^*, \pi^*]$ marks the resting point of the economy, with unchanging unemployment and inflation. A line rising from the high-unemployment, low-inflation region in the lower right describes the upward convergence toward the point of rest. A line declining from the low-unemployment, high inflation region describes convergence from that region. Both slopes are negative. Jointly, they trace out the Phillips curve of the model. Its slope is the inflation change divided by the unemployment gap change. That ratio, designated ϕ , is the slope of the Phillips curve, the third key parameter of the model.

2.3 The Phillips curve

The model laid out above defines the Phillips curve, an equation relating inflation to an inertial term and to a term involving economic activity. The Phillips curve is a key component of the New Keynesian class of macroeconomic models—see Woodford (2003) for a detailed exposition and Chapters 6 and 7 in Romer (2019) for a recent advanced textbook treatment of New Keynesian macroeconomics.

We consider a setup that embodies the properties described above:

$$\pi_t - \pi_t^* = -\phi \cdot (u_t - u_t^*). \quad (1)$$

Variables subscripted with t are time series. Variables without subscripts are parameters.

This equation is the Phillips curve. π_t is the actual rate of inflation; π_t^* is the inflation anchor; and ϕ is a non-negative coefficient governing the strength of the response of inflation to the gap between unemployment and the natural rate, $u_t - u_t^*$. If ϕ is small, prices are quite sticky and movements of $u_t - u_t^*$ are large and persistent. If ϕ is large, prices are flexible and $u_t - u_t^*$ returns quickly to its normal value of zero.

The Phillips curve has the property that $u_t = u_t^*$ when $\pi_t = \pi_t^*$, that is, when inflation is at its anchored level, unemployment is at its natural rate.

Inflation also fluctuates for reasons apart from unemployment, notably from fluctuations in the supply of energy and agricultural products, and, more recently, in the pandemic, products with supplies cut back by the shutdowns. These fluctuations could be included in an extended model. In most empirical Phillips curves, they enter as separate additive terms.

Although it is conventional to display unemployment on the horizontal axis of a Phillips curve and to treat the unemployment gap informally as an exogenous determinant of inflation, the variables under discussion here are obviously jointly determined.

If ϕ is large, the Phillips curve is nearly vertical; even small values of the unemployment gap go with large effects on inflation. If ϕ is small, the Phillips curve is nearly flat.

2.4 Sticky prices and sticky inflation

We use the terms *sticky prices* and *sticky inflation* as equivalent. In discrete time, the change in the price level in a given month incorporates the level of inflation. We always use π to denote the level of inflation. We could also write $\frac{p_t - p_{t-1}}{p_{t-1}}$ but it would not add anything. In continuous time, on the other hand, there can be a difference—the price level can jump. We avoid the issue by using discrete time.

2.5 Our definition of the natural rate

The natural rate of unemployment, in our definition, is the unemployment rate at the point of rest in Figure 1. At this point, inflation equals the inflation anchor. The natural rate of unemployment also goes by the name *non-accelerating-inflation rate of unemployment* or NAIRU, which we take to be a synonym for the natural rate. The name captures a key property of that rate: Periods of stable inflation are times when the unemployment rate is at its natural level.

Our definition of the natural rate may differ from others based on a hypothesized absence of frictions in price and wage formation. Our definition recognizes that positive unemployment prevails when the economy is at rest, owing to normal turnover in the labor market. Defining the natural rate of unemployment is essential, given the myriad of definitions in the literature as summarized by Rogerson (1997).

2.6 Price flexibility

We can rewrite the Phillips curve in aggregate-supply form as

$$u_t = u_t^* - \frac{\pi_t - \pi_t^*}{\phi}. \quad (2)$$

The coefficient ϕ , the downward slope of the Phillips curve, controls the influence of inflation on real activity, as measured by unemployment. Higher values of ϕ make the model more like the real-business-cycle model, where real activity is not influenced by monetary factors such as inflation. In the polar case of full monetary neutrality, ϕ is infinite and unemployment tracks the natural rate all the time. At the other end, ϕ is zero, and the Phillips curve is flat—large movements of unemployment are paired with a fixed value of inflation.

The parameter ϕ captures the price-flexibility of the economy. If high, the economy has flexible prices and fits the real business cycle paradigm. If close to zero, prices are somewhat or fully sticky, and monetary factors have important involvement in the determination of unemployment and other real variables.

3 Estimation of the Slope of the Phillips Curve by Regression

A natural way to begin econometric analysis of the Phillips curve is to study the regression of $\pi_t - \pi_t^*$ on unemployment u_t . We assume the availability of accurate data for π_t , π_t^* , and u_t . See, for example, Jorgensen and Lansing (2019) for a discussion of measurement of π_t^* .

3.1 The regression of the inflation gap on the unemployment rate

We treat the natural rate u_t^* as unobservable. We denote the observed coefficient of the regression of the negative of the inflation gap on the observed unemployment rate as $\hat{\phi}$. The coefficient is

$$\hat{\phi} = \frac{\text{Cov}(-(\pi_t - \pi_t^*), u_t)}{V(u_t)}. \quad (3)$$

We substitute the model's value for $\pi_t - \pi_t^*$:

$$\hat{\phi} = \frac{\text{Cov}(\phi(u_t - u_t^*), u_t)}{V(u_t)}. \quad (4)$$

Now let C be the unobservable regression coefficient of u_t^* on u_t . The relation between $\hat{\phi}$ and ϕ simplifies to

$$\hat{\phi} = (1 - C) \phi. \quad (5)$$

If C is zero, the regression gives the true slope of the Phillip's curve, ϕ . If the natural rate is a positive component of total unemployment, C is positive and the regression coefficient $\hat{\phi}$ understates the true relation between the inflation gap and the unemployment gap.

This property is important because many estimates of the slope of the Phillips curve, ϕ , turned surprisingly low in the recovery from 2009 to 2019. Rising values of C may be a cause—the decline in estimated ϕ may be an illusion.

3.2 Failure of identification

Advance knowledge of the parameter C is useful for any value of C , not just zero. Regress to get $\hat{\phi}$, then calculate $\phi = \hat{\phi}/(1 - C)$.

The explicit or implicit reliance on a belief that $C = 0$ is close to universal in research based on a regression of the inflation gap on unemployment.

The Phillips curve regression reviewed in this section rests on the identifying assumption that C has a known value. Absent a persuasive reason to believe this assumption, regression yields no usable information about the slope coefficient, ϕ , or the natural rate, u^* .

We can state this conclusion in terms of the notion of *identification*: A model is identified if there is a unique mapping from its observables to its parameters. Frequently, a model achieves identification through the addition of assumptions. These *identifying assumptions* capture plausible properties of a model that are not tested.

The basic regression model here is not identified—observations of the inflation gap and unemployment do not pin down the corresponding values of the underlying parameters: the slope of the Phillips curve, ϕ , and the coefficient, C , relating the natural rate to the observed unemployment rate.

3.3 Nonlinear Phillips curves

There is growing interest in exploring curvature in the Phillips curve, so the Phillips curve takes the form,

$$\pi_t - \pi_t^* = -\phi(u_t - u_t^*). \quad (6)$$

Here $\phi()$ is an increasing smooth function of the unemployment gap, $u - u^*$. It can be approximated as a spline function—see Leduc, Marti and Wilson (2019).

Nonlinearity by itself does not lift the identification barrier. On the other hand, although the spline may add quite a few more parameters, there is no added burden of identification. We discuss methods that have emerged that may achieve identification, explicitly or implicitly. In the section below, we discuss our anchored-inflation method, which offers partial identification. In the following section, we discuss three existing methods. The methods deliver vastly different time series for the natural rate.

4 Inferring the Natural Rate of Unemployment during Periods of Anchored Inflation

This section pursues the idea that there is information bearing on the Phillips curve that is not captured in the regression approach. According to the ideas launched by Milton Friedman and rapidly adopted in Phillips-curve discourse, a time when unemployment is at its natural level is also a time when inflation is at its anchored level.

To implement this idea, we solve the Phillips curve for the natural rate of unemployment:

$$u_t^* = u_t + \frac{1}{\phi}(\pi_t - \pi_t^*). \quad (7)$$

Thus the natural rate is the actual unemployment rate with an adjustment for the effect on unemployment inferred from the inflation pressure, $\pi_t - \pi_t^*$.

There is a special case where outside information is arguably available that permits calculation of u_t^* from equation (7). We propose a simple anchored-inflation method, which exploits the principle that the natural rate of unemployment equals the observed rate when inflation equals its anchored value. Only observations with inflation at the anchored level are considered: $\pi_t = \pi_t^*$. Identification is partial. We assume that the slope, ϕ , is strictly positive. We also assume that the anchored rate is known.

Inserting this restriction and these assumptions into equation (1) yields

$$u_t^* = u_t. \quad (8)$$

Thus, for a subset of the observations, we can establish that unemployment u_t reveals the value of the natural rate, u_t^* . The current rate of unemployment reveals the natural rate. We have partial identification of dates when $\pi_t = \pi_t^*$ and $\phi > 0$.

The issue becomes, what configuration of theory and data would make a finding plausible that actual inflation was equal to anchored inflation at a particular time? Friedman himself associated the source of inertia in inflation with expectations of future inflation, and that

theme has resonated in most discussions of inflation dynamics ever since. One determinant of expected inflation is the central bank’s target rate of inflation. Another is the success of the bank in achieving its target. We look for a moderately lengthy historical episode where inflation was close to constant at a low level consistent with the central bank’s stated objective. During such an episode, the observed unemployment rate is close to the natural rate.

The recovery in the US economy starting in November 2009 and lasting until the pandemic terminated the recovery partway through March 2020, deserves consideration as an episode when inflation was close to the target rate of 2 percent, adopted formally by the Fed in 2012. The anchored rate coincided fairly closely with the target rate.

Under these two assumptions about the US labor market in 2009 to 2019, we construct Figure 2 to demonstrate the resulting inferences about the natural rate during that period. The figure plots the actual unemployment rate and indicates with red dots the months when inflation was close to the anchored rate of two percent. Specifically, it indicates the months when the year-over-year percentage change of the price of personal consumption expenditures was between 1.5 and 2.5 percent. The figure shows that during the 2009-2019 period, there were numerous months when the natural rate closely tracked actual unemployment, according to the logic developed in this section.

In practice, we designate an interval of values such as $[\pi_t - \frac{\epsilon}{2}, \pi_t + \frac{\epsilon}{2}]$ for these calculations, where ϵ is a fairly small number such as 0.01, its value in the figure. Equation (7) shows that the corresponding interval of values of the natural rate of unemployment is

$$\frac{1}{\phi} \left[\pi_t - \frac{\epsilon}{2}, \pi_t + \frac{\epsilon}{2} \right] \quad (9)$$

The division by ϕ is crucial. If ϕ is fairly small, as many investigators believe, the width of the interval could be substantial. The case that actual unemployment reveals natural unemployment is not overwhelming, given the lack of consensus on the slope of the Phillips curve.

Note that this method applies only to the one recovery with exceptionally stable inflation. In the other recoveries, actual inflation was sufficiently variable that we cannot make a similar inference. And there is no case of a recession with constant inflation, so there is no direct information about the natural rate during recessions, using this approach.

In general, equation (7) implies that systematic movements of actual unemployment resemble the corresponding movements of the natural rate unless the $\frac{1}{\phi} \cdot (\pi_t - \pi_t^*)$ component is large, which happens when inflation exhibits substantial deviations from the anchor, or when the coefficient ϕ is small.




Figure 2: The Actual Unemployment Rate and the Estimated Natural Rate during the 2009-2019 Recovery

Note: Authors' calculations using monthly data from the CPS and the National Income and Product Accounts. October 2010—January 2020.

5 Existing Methods for Inferring the Natural Rate of Unemployment

In this section, we consider three existing methods for extracting values of ϕ and u_t^* from available data, on the assumption that π_t , π_t^* and u_t are observed with reasonable accuracy. These are:

- The *unemployment-trend* method defines the natural rate of unemployment as a constant or a slow-moving index of demographic change derived from the actual lower-frequency movements of unemployment and its demographic components
- The *statistical-submodel* method, forms a natural-rate index under identifying assumptions about its response to other variables and the role of unobserved shocks
- The *macro-model* method, uses a general-equilibrium model to form a counterfactual stickiness-free equilibrium whose equilibrium values of unemployment play the role of the natural rate.

5.1 The unemployment-trend method—the natural rate as the low-frequency trend in actual unemployment

The most common identifying method uses values of time series from the unemployment survey. In many cases, the values trace out a slow-moving trend related to the composition of the labor force. The Congressional Budget Office publishes a frequently updated estimate of the time path of the natural unemployment rate with demographic adjustment, and many Phillips-curve studies have adopted the CBO’s path. That path rose gradually to a maximum in the 1970s and declined thereafter—see Figure 3.

In 2021, the CBO changed the name of the path, to the *noncyclical rate of unemployment*. The new name clarifies that the counterfactual underlying the calculation is the absence of cyclical movements, not the absence of all high and medium-frequency fluctuations. The clarification differentiates the CBO’s estimate from those calculating the natural rate as the result of passing the actual rate through a band-pass filter, such as Hodrick-Prescott, and retaining only the low-frequency component. Other recent estimates of the long-run trend in unemployment are Barnichon and Matthes (2017), Tasci (2018), Barnichon and Mesters (2018), and Hornstein and Kudlyak (2019).

A related method uses a search and matching model to define the trend. See Daly, Hobijn, Sahin and Valletta (2012), who estimate that during the 2007-09 recession the natural rate of unemployment rose to a value in the range of 5.5 to 6.6 percent.




Figure 3: The Actual Unemployment Rate and the Noncyclical Rate of Unemployment from the Congressional Budget Office

Figure 3 shows that the unemployment-trend method classifies almost all of the movements of the unemployment rate as cyclical, leaving the natural rate almost no role in the evolution of unemployment. Compare the path traced out by the CBO’s reckoning of the natural rate during the long recovery from 2009 to 2019, shown in Figure 3, to the natural rate for the same period based on our definition of the natural rate, shown in Figure 2. Figure 3 shows a long, smooth decline in unemployment, almost all of which is characterized as a decline in inflationary pressure. Figure 2, for the same period, shows that close to all of the decline is characterized as a decline in the natural rate. During the period, inflation was stable. Figure 3 says that the Phillips curve was flat during the period—inflation pressure grew substantially, but inflation remained constant.

Unemployment-trend methods will be appropriate if the volatility of the true natural rate is low, so it matches the volatility of the smooth time series taken as the natural rate. But if the true value of the natural rate has volatility closer to actual unemployment, the unemployment-trend method will understate the true downward slope of the Phillips curve.

5.2 The statistical sub-model method—the natural rate as a latent variable that follows a specified stochastic process

The statistical sub-model method carries out joint statistical estimation of the Phillips curve and the natural rate. The rate is a latent variable in the joint system and the slope of the Phillips curve is a latent parameter of the system.

Gordon (1997) estimated a time-varying natural rate from a statistical model comprising an inflation equation with the inflation pressure and an equation for the natural rate, which follows a random walk. From 1955 through 1995, the estimated natural rate varied between 5.4 and 6.5 percent—see the series under the author’s preferred smoothness parameter in his Figure 2. It declined by a percentage point between the mid-1980s and mid-1990s. See also Gordon (1998). Gordon’s statistical sub-model method yields an estimated path of the natural rate that resembles the path from the unemployment-trend method. It shares the potential of that method to understate the slope of the true Phillips curve if the true natural rate accounts for a substantial fraction of the movements of unemployment.

King and Morley (2007) define the natural rate of unemployment as the long-run equilibrium unemployment rate. They develop a vector-autoregression model of real GDP, inflation, and the unemployment rate. They impose long-run identifying restrictions on the relationship between the observed series and the unobserved structural shocks. Specifically, they require that the first structural shock has no long-run effect on the unemployment rate and that the second structural shock has no long-run effect on output and the unemployment rate. They suggest that their third shock, which is unrestricted, may be thought of as the natural rate shock, as it is the only one that is allowed to affect the unemployment rate in the long run. The results in their Figure 2 show that the estimates of the natural rate of unemployment closely follow the actual unemployment rate. Consequently, their version of the natural rate implies that movements along the slope, ϕ , play a large role in the volatility of the natural rate.

Using a state-space model, Crump, Eusepi, Giannoni and Sahin (2019) take the natural rate of unemployment as the sum of a secular trend component and a cyclical component. Having estimated the trend, they estimate the cyclical component from a forward-looking Phillips curve model under the assumption of an AR(2) process for the unemployment gap. Bok, Crump, Nekarda and Petrosky-Nadeau (2023) (BCNP) estimate a model that combines a Phillips curve and the extraction of the natural rate of unemployment in a three-equation state-space model. The model draws upon earlier work by Laubach (2001); Crump et al. (2019); and Crump, Eusepi, Giannoni and Sahin (2022). In BCNP, inflation pressure follows an AR(1) process and the natural rate follows another random walk. Figure 2 of BCNP shows the unemployment rate and the estimated natural rate, labeled as the “preferred

stable-price unemployment rate”. Their estimated natural rate is substantially positively correlated with the actual rate as compared to the CBO’s natural rate.

The statistical sub-model methods tend to find moderate to large roles for the natural-rate component of the unemployment rate.

5.3 The macro-model method—the natural rate defined by a counterfactual solution of a general-equilibrium macro model

General-equilibrium models provide a way to determine the path of the natural rate of unemployment. The idea is to construct a counterfactual solution to a version of the model that describes an economy satisfying a definition of the natural rate of unemployment. For example, the version could impose the condition that the actual rate of inflation equals the anchored rate and both equal the central bank’s inflation target.

Galí, Smets and Wouters (2011) (GSW) is a leading example of this approach to extracting the natural rate of unemployment from a general-equilibrium model. In GSW’s case, the model is a widely used New Keynesian model developed by Smets and Wouters. The model includes a sub-model that deals explicitly with unemployment along the lines of the DMP framework. The calculation runs as follows: “We [GSW] ...assess the role of wage rigidities as a factor underlying observed unemployment fluctuations by comparing the observed unemployment rate to its estimated natural counterpart, where the latter is defined as the unemployment rate that would be observed in the absence of nominal wage rigidities...” (pages 348-349). Their counterfactual is somewhat different from ours, but its results are indicative of one based on ours or other definitions of the natural rate.

Figure 4 plots the resulting calculated natural rate from 1966, first quarter, through 2015, third quarter.¹ The path of the natural rate captures the bulge of unemployment in the 1970s, and it also moves substantially in harmony with actual unemployment at other times, especially after the serious recessions starting in 1981 and 2007.

Furlanetto and Groshenny (2016), using an estimated medium-scale dynamic stochastic general equilibrium model with sticky prices and equilibrium search unemployment that features a shock to matching efficiency, find that this shock contributed to raising the actual unemployment rate by around 1.3 percentage points and the natural rate by around 2 percentage points during the 2007-09 recession. Although this effect is considerably smaller than the effect found by GSW, it refers to the effect of the matching-efficiency shock rather than the shock-regime measured by GSW.

¹This figure displays data supplied to us by Galí covering three more years than the published Figure 10.




Figure 4: The Actual Unemployment Rate and the Natural Rate based on Galí et al. (2011)

Note: Data provided to the authors by Galí; quarterly, Q1 1966—Q3 2015.

6 The Natural Rate and the Flat Phillips Curve

Our discussion above makes it clear that determining the path of the natural unemployment is a challenge. That path is an intrinsic element of the Phillips curve—every study of the Phillips curve rests on a specification of the natural rate.

As we noted earlier, a simple specification, adopted by many authors, is to take the natural rate to be a constant, or, equivalently, to omit the natural rate altogether. Another specification takes a long-run trend in the actual unemployment rate as the natural rate of unemployment. Such measures, even if time-varying, do not capture the cyclical variation in the natural rate described by the anchored-inflation method for the recovery from the 2007-2009 recession.

In this section, we discuss the potential bias toward understatement of the slope of the Phillips curve resulting from the absence of cyclical variation over time in the natural rate of unemployment. We show that if the true natural rate of unemployment is highly correlated with the actual rate, Phillips curves estimated with constant natural rates or natural rates uncorrelated with the actual unemployment rate will inevitably be close to flat.

6.1 Implications of mis-specifying the natural rate as uncorrelated with actual unemployment

We are studying the Phillips curve from equation (1), replicated here

$$\pi_t - \pi_t^* = -\phi \cdot (u_t - u_t^*).$$

As before, we presume that we have solved the problem of measuring π_t^* , and focus on the unobserved u_t^* .

We now demonstrate that taking the natural rate of unemployment as constant or near-constant has profound implications for estimation of the Phillips curve if the true covariance between the natural and actual rates is materially positive.

Consider a Phillips-curve regression that does not include the natural unemployment rate:

$$\pi_t - \pi_t^* = -\hat{\phi} u_t. \tag{10}$$

In Section 3, we demonstrated that the regression coefficient is

$$\hat{\phi} = (1 - C) \phi.$$

C is the regression coefficient of u_t^* on u_t . It is an index of the relevance of the natural rate. If $C = 0$, the natural rate is irrelevant and the regression coefficient $\hat{\phi}$ will be an unbiased estimate of the Phillips-curve slope, ϕ . If $C = 1$, then $\hat{\phi} = 0$ —the Phillips curve appears to be totally flat, even if the true slope is robustly positive. In that case, the natural rate is highly relevant.

Note that C is not sensitive to the overall level of the natural rate, because the constant part of C is absorbed by the constant that would normally be part of the functional form of the Phillips curve. C is sensitive to the co-variation of the natural rate and the actual unemployment rate.

Our discussion of the omission of a time-varying natural rate from a regression for the slope of the Phillips curve, is an application of the standard analysis of the bias from an omitted right-hand variable.

The denominator in C , $V(u_t)$, is observed directly and is robustly positive because unemployment rises briskly in recessions and falls reliably in recoveries. The big question is the covariance of the natural rate u_t^* with actual unemployment. If the covariance is zero—possibly because u_t^* is constant over time—there is no bias. If u_t tracks the natural rate u_t^* almost perfectly, C will be almost 1, and the estimate of $\hat{\phi}$ will be essentially zero, even if the true value of ϕ is quite positive.

One particularly salient conclusion from this analysis is the following: If the true natural rate is highly correlated with the actual rate, Phillips curves estimated with constant

or nearly constant natural rates of unemployment uncorrelated with the actual rate will inevitably be close to flat.

6.2 The relevance of the natural rate of unemployment for actual unemployment under various approaches to identifying the natural rate

Our finding that in the long expansion of 2009 through 2019, actual unemployment tracked the natural rate closely, shows that period C for that period is very close to one.

We can use the results from Galí et al. (2011) to illustrate our analysis of the bias from failing to consider the movements of the natural rate. Suppose that we studied the Phillips curve in the GSW model by regressing an appropriate version of $\pi_t - \pi_t^*$ on u_t . The regression coefficient for u_t^* on u_t is $C = 0.60$ in the 1966-2015 sample, so the estimated slope of the Phillips curve is depressed to $1 - C = 0.40$ times its true value during that period.

Using BCNP's estimate of the natural rate over 1985 through 2019, we find the regression coefficient for u_t^* on u_t to be $C = 0.47$ over that period. The existing state-space statistical models of the inflation pressure or the natural rate may overstate the pressure following recessions. As we suggest in Section 4, the natural rate may jump upward to come close to matching the high level of unemployment coming out of a crisis. Then, for a protracted period, the actual unemployment rate and the natural rate glide down together, implying that the gap is small throughout the recovery. Accommodating such possibilities for the natural rate requires relaxing the assumption of an AR process with a constant-variance error for the gap. We note that during the period for which we have both GSW's and BCNP's estimates of the natural rate, from Q1 1985 to Q3 2015, we find C to be 0.46 and 0.45, respectively.

The natural rate does not account for all of the movement of actual unemployment—inflation pressure accounts for some of the cyclical movements. The key conclusion is that the movements of the natural rate constitute a central factor in the economics of the Phillips curve and employment volatility. The natural rate and the actual rate of unemployment move together—in some episodes, notably in the recovery from the financial crisis, the regression coefficient is close to 1.

6.3 Changes over time in the relevance of the natural rate of unemployment

Increases over time in the relevance of the natural rate might account for the apparent flattening of the Phillips curve. To check this possibility in the Galí et al. (2011) results,




Figure 5: Estimated Values of C , the Relevance of the Natural Rate of Unemployment for Actual Unemployment, 20 Quarters Forward

Note: The figure shows the coefficients from rolling regressions covering 20 forward quarters from the date shown on the horizontal axis, Q1 1966—Q4 2010.

we replicated the estimates in Figure 4 in the form of coefficients from rolling regressions covering 20 forward quarters from the date shown on the horizontal axis. Figure 5 shows the movements of C calculated in this way. In general, C , is cyclical but untrended over the long run. There is a sign of an increase in C that would explain the flattening of the Phillips curve over the period from 1994 through 2010.

Another approach to the issue is based on del Negro, Lenza, Primiceri and Tambalotti (2020) (DLPT), a recent careful empirical study of the response of inflation and the unemployment rate to what they call an unemployment shock, in an 8-variable vector autoregression. The authors argue that the Phillips curve has flattened since 1990 and seek to explain the flattening. Lines 1, 2, and 3 of Table 1 display their results. Line 1 shows the response of unemployment to the shock for the first half of the sample through 1989, and for

<i>Line</i>	<i>Variable</i>	<i>Before 1990</i>	<i>After 1990</i>	<i>Source</i>
1	Unemployment (positive)	0.3	0.3	DLPT
2	Inflation (negative)	0.3	0.1	DLPT
3	Slope of Phillips curve regression	1.0	0.3	row 2/row 1
4	Relevance of natural rate	0.6	0.9	GSW (see text)
5	True slope of Phillips curve	2.5	2.5	row 3/(1-row 4)

Table 1: Calculations Relating to the Downward Bias in the Estimated Slope of the Phillips Curve, based on the estimates from del Negro et al. (2020) (DLPT) and Galí et al. (2011) (GSW)

the second half, starting in 1990. The response is measured four quarters after the shock. Line 2 shows the responses of inflation, and line 3 shows ratio of the price response to the unemployment response. We interpret the ratio as the slope of the Phillips curve, in units of the percentage-point decline in inflation per percentage point increase in unemployment. According to these results, the estimated slope of the Phillips curve was 1.0 in the earlier period, and declined to 0.3 in the later period. These findings are in line with the substantial literature, which is surveyed extensively in the DLPT paper.

DLPT include CBO’s time series for the natural rate of unemployment as one of its 8 basic indicators and as their measure of natural unemployment, but does not report the functions for the responses to the natural rate impulses. It does report the impulse response function from unemployment to the CBO series for the natural rate, which suggests almost no volatility.

We use our earlier derivations of the relevance of the natural rate based on leaving out the natural rate and the estimates of the natural rate series from GSW. Lines 4 and 5 of Table 1 provide information about the possibility that the econometric downward bias in the slope of the Phillips curve, caused by omitting data on the natural rate, accounts for the large decline in the slope of the estimated Phillips curve from the pre-1990 to post-1990 period. Line 4 reports the relevance statistic C , which we derived from GSW, as potentially applicable to the earlier period in DLPT’s results. Line 5 shows the implied value of the slope of the Phillips curve adjusted for the given C . For the earlier period, the adjusted estimate, ϕ , is $2.5 = 1.0/(1 - 0.6)$.

We then calculate what the value of C should be for the later period so that the adjusted Phillips curve slopes are the same in both periods. With $C = 0.9$, the slope in the later period $\phi = 2.5 = 0.3/(1 - 0.9)$ equals the slope in the earlier period. While we do not know

the increase in the value of C , the hypothesis that the relevance of the natural rate of unemployment for actual unemployment has increased in the later period has some support in Figure 5. It shows that during the long recovery from the 2007-09 recession, C , calculated from GSW data, reached its all-time maximum.

We propose that the US labor market experienced an increase in the relevance of the natural rate as measured by C , the regression coefficient for the natural rate on the actual rate of unemployment. This supports our hypothesis that the scenario in Table 1 accounts for at least some of the apparent decline in the slope of the Phillips curve after 1990.

6.4 Other related literature

6.4.1 Estimates of the Phillips curve

As we noted earlier, a substantial literature starts from the assumption that the natural rate of unemployment is constant or only slow-moving, reflecting movements in a long-run trend in actual unemployment, and, therefore, uncorrelated or only weakly correlated with actual unemployment. Another branch of the literature notes that there have been several instances when large movements in the unemployment rate have coincided with small changes in the inflation rate. Using a smooth trend as the measure of the natural rate, studies typically find a flat Phillips curve, especially in recent decades. Some studies have invoked nonlinearities, whereby the slope differs when unemployment is high versus low, or time-variation of the Phillips curve across time-periods.

Hooper, Mishkin and Sufi (2020), review earlier research and contribute their own national and cross-state evidence on the flattening of the Phillips curve. They conclude that “national data going back to the 1950s and 60s yield strong evidence of negative slopes and significant nonlinearity in those slopes, with slopes much steeper in tight labor markets than in easy labor markets. The evidence of both slope and nonlinearity weakens dramatically based on macro data since the 1980s for the price Phillips curve.” For the aggregate analysis, they use the CBO’s measure of the national natural unemployment rate. For the state-level results, they use bins of actual state data because no estimates of the state-level natural rate were available at the time of their work.

Hazell, Herreno, Nakamura and Steinsson (2022) is a recent contribution that finds very flat state-level Phillips curves for non-tradeable goods and services at the state level, using state-level unemployment as the measure of the state-level gap and the CBO’s natural-rate measure for the aggregate analysis.

Most recently, using the CBO’s measure in the aggregate analysis and the unemployment rate rather than the unemployment gap in the MSA-level analysis (because “there are no

estimates of the natural rate of unemployment...for the city-level unemployment data”), Smith, Timmermann and Wright (2023) study time-variation in the slope of the Phillips curve and find two regime changes: prior to 1972, the estimated slope is 0.51 units of price decline per unit of unemployment increase; this slope steepens 1972–2001 to 0.87; after, 2001 break, the slope of the Phillips curve becomes essentially zero.

Leduc and Wilson (2017) measure the city-level natural rate of unemployment as a 10-year trailing average of the city’s actual unemployment. They estimate the slope of the cross-city wage Phillips curve using seven-year rolling regressions and find a steady flattening of the Phillips curve slope starting with the 2001–2007 sample.

Dotsey, Fujita and Stark (2018), measuring the inflation pressure by the deviations of actual unemployment from its Hodrick–Prescott trend, conclude that “using the Phillips curve may add value to the monetary policy process during downturns...We find no evidence for relying on the Phillips curve during normal times, such as those currently facing the U.S. economy.” (p.90) Ashley and Verbrugge (2023), using the long-run trend of unemployment from Tasci (2018) and from CBO as the measures of the natural rate, estimate what they call a “persistence-dependent” version of the Phillips curve that varies across three phases of the business cycle and find that in the recovery phase, inflation is unrelated to the unemployment gap. Using the CBO’s measure, Doser, Nunes, Rao and Sheremirov (2023) estimate a piecewise-linear specification and document that the data favor a model with two regions, with the response of inflation to an increase in unemployment slower in the region where unemployment is already high. Barnes and Olivei (2003) estimate a piecewise linear specification of the Phillips curve. Using state fixed effects, Leduc et al. (2019) estimate nonlinear wage Phillips curve in the state-level unemployment data and find some evidence of steeper Phillips curve in hot labor markets. Other recent studies of the Phillips curve include Ball and Mazumder (2011), Stella and Stock (2013), Coibion and Gorodnichenko (2015), Cecchetti, Feroli, Hooper, Kashyap and Schoenholtz (2017).

Stock and Watson (2010) present evidence consistent with our conclusion that, in recoveries, the natural rate follows the smoothly declining path of the actual rate. They show that inflation takes a step downward early in a recession, but then remains unrelated to unemployment changes as the business cycle progresses through recovery—see their Figure 2. Constant inflation with declining unemployment suggests that the natural rate of unemployment is declining in parallel with actual unemployment, according to equation (7).

6.4.2 Policymakers’ views

Policymakers have faced a great deal of uncertainty about the natural rate of unemployment. Orphanides (2002) and Orphanides and Williams (2013), discuss the magnitude of

informational problems and disagreement over the natural rate of unemployment among policymakers in real time. Staiger, Stock and Watson (1997) find that it is measured quite imprecisely. A similar conclusion is reached in a comment by Davis (2019). Crump, Nekarda and Petrosky-Nadeau (2020) discuss the range of the natural unemployment rate benchmarks used by policymakers. Hetzel (2022), in a book on the history of the Federal Reserve system, expresses skepticism about the role of the Phillips curve as a structural model of the economy to guide the choice of monetary policy, evoking the Lucas critique (Lucas (1976)) and the lack of certainty about the natural rate of unemployment. Hetzel (2024) discusses articulation of the monetary policy regime along the lines of what the FOMC controls and how it exercises that control.

In a departure from the prevailing view among policymakers and other observers, that the natural rate moves slowly and does not track the cycle, Lacker (2012) took the opposite view: “There is a reference unemployment rate to which it’s most appropriate to compare the current unemployment rate for the purposes of assessing current policy...The most common term for this reference rate is “the natural rate” of unemployment...There is a clear intuition for having the unemployment yardstick for monetary policy vary with economic conditions...Estimates of [the natural rate] invariably impose the assumption that it varies only slowly and does not respond to many transitory shocks.

7 Is Inflation Sticky or Flexible?

Whether inflation is sticky or flexible is a key question of macroeconomics. There is a broad range of opinions in the macro profession. One view, widely present in the literature, is that the Phillips-curve slope, ϕ , is small. The profession call this the sticky view of the slope of the Philips curve. An implication of the sticky view is that the unemployment gap, $u_t - u_t^*$, can be large and persistent.

According to the flexible view of inflation, the slope of the Phillips curve, ϕ , is substantial, while the unemployment gap, $u_t - u_t^*$, is small and transitory. A frequently used name in the literature is the “real business cycle model”, a name that calls attention to the important role of flexible inflation in limiting the effects of monetary policy on output and employment.

In the flexible view of the Phillips curve, low unemployment does not necessarily signal a high unemployment gap. During recoveries, the pressure, $u_t - u_t^*$, is close to zero, based on the evidence that during recoveries natural rate of unemployment glides down together with actual unemployment and they are likely close to each other. Our finding means that the unemployment gap is close to zero and, therefore, inflation pressure is weak. We do not take a position on the gap during contractions when unemployment is rising rapidly.

Under the flexible view, during recoveries, the pressure is zero. That is, during recoveries, an economy resembles a real- business-cycles economy, with $1/\phi$ being relatively low and the Phillips curve being steep. The flexibility of prices is the key differentiating factor. Our view requires that prices are somewhat flexible, so the Phillips curve is reasonably steep, whereas the opposing low and sticky view posits stickier prices and a flatter Phillips curve.

By far the most important difference between New Keynesian macro models and RBC macro models is the Phillips curve. Changing the single parameter we call ϕ in an NK model to a large value, converts the NK model to a RBC model.

The Phillips curve sets the excess of the rate of inflation over its anchor to the product of the slope parameter and the inflation pressure. Absent a convincing identifying condition, the available information identifies a range of paths of ϕ paired with the corresponding paths of u_t^* , that satisfy the Phillips curve, given the observed $\pi_t - \pi_t^*$ and u_t . Near one end, ϕ is close to zero and $\pi_t - \pi_t^*$ is correspondingly large. Near the other end, ϕ is large and the unemployment gap $u_t - u_t^*$ is small. Both views fit the specified data. Additional information helping to reveal u_t^* or ϕ would be needed to determine which view is correct.

We can compare the two views during a period in which the difference is particularly clear, namely the recovery from 2009 to 2019, previously discussed in Section 4. During that recovery, unemployment declined along a smooth path while inflation was level at close to the two-percent Fed target. Believers in sticky inflation attribute the near-constant inflation to a flat Phillips curve. The decline in unemployment was the result of gradually rising demand that moved the economy to the left along the flat Phillips curve. Believers in flexible inflation hold that the Fed adjusted monetary policy to peg inflation at two percent, with no effects on unemployment or other real variables.

The flexible-believer cites the fact that quite a few months during the long recovery had inflation close to the Fed’s target, which implies that unemployment in those months is close to the natural rate. But the sticky-believer counters with the point that this finding rests on the assumption that those months had Phillips-curve slopes that were definitely positive, and the sticky-believer questions that assumption.

The sticky-believer cites regression evidence that the CBO’s time series has a statistically insignificant coefficient in a Phillips-curve estimated with data from 2009 to the present. The flexible-believer dismisses that evidence on the grounds that it is contaminated by likely understatement of the covariance of the CBO’s measure with the true natural unemployment rate.

The most radical potential conclusion about the relation between the observed rate of unemployment and the natural rate is that there is no difference—observed unemployment is at its natural level all the time. This conclusion would cut the heart out of the Phillips curve

and the distinctive features of the New Keynesian model. It would deny unemployment any role as a measure of inflation pressure. That conclusion goes beyond the evidence, however. We have relatively few observations of stable inflation in times of rising or really high unemployment. We believe that a reasonable interpretation of the evidence is that, during long, slow, reliable recoveries with gradually declining unemployment, unemployment is close to its natural rate and is not a measure of inflation pressure. Under those conditions, there is no meaningful unemployment gap.

8 The Effect of the Pandemic on the Phillips Curve

During the long recovery from 2009 to 2019, inflation became powerfully anchored at just below 2 percent per year. In this stable environment, sellers adapted their price-setting procedures to stability. The Phillips curve became relatively flat because sellers tended to leave prices unchanged for extended periods—relatively few sellers responded to change each month.

The pandemic created a completely different environment for pricing decisions, with rapid new developments. A quick response to each new development was required. Policy responses to the pandemic included expansionary monetary and fiscal policy. Another important development was a reduction in output supply, due to idling of important sectors of the economy, notably hospitality. The turbulence that the pandemic brought to seller's economic situations made it necessary to make more frequent prices changes than in the tranquil pre-pandemic times.

In the Phillips curve framework, an increase in turbulence represents a regime change—the Phillips curve becomes more sensitive to changes in unemployment. That is, in tranquil times the Phillips curve is relatively flat—any shifts in demand show up mostly as quantity changes, while in turbulent times, the shifts in demand have large effects on inflation. One possible explanation behind the post-pandemic inflation is an increased volatility of the inflation anchor— π_t^* in the Phillips curve. As discussed above, the inflation anchor also depends on the current speed of adjustment of prices—if information becomes more volatile, sellers will choose to change prices more frequently.

Hall (2023) suggests that the New Keynesian Phillips curve became steeper in the volatile conditions of the pandemic. He establishes that sectors with price stickiness are prone to rapid transitions from stickiness to flexibility, as sellers elect to reset their prices and abandon anchoring. He argues that the logic of the New Keynesian model of the Phillips curve links inflation to volatility, because a larger fraction of sellers are pushed out of their regions of

inaction when volatility is elevated. See also Blanco, Boar, Jones and Midrigan (2023) and Cavallo, Lippi and Miyahara (2023), and work cited there.

With respect to the pandemic, this line of thought implies that the prices that we thought were sticky turned much more flexible. The pandemic economy is closer to a flexible-price economy, an economy with significantly less price inertia, compared to the sticky-price pre-pandemic economy.

As the environment changed toward lower inflation, because the constraints from the pandemic relaxed and the Fed adopted a contractionary policy to combat inflation, the steep slope of the Phillips curve became an advantage. Disinflation is occurring with a materially smaller bulge of unemployment than would occur under the pre-pandemic, more stickier-price economy.

Sargent (1982) studies the disinflations in four economies as they overcame high inflation rates and achieved price stability with little dislocation of economic activity. Their Phillips curves became vertical. The steeper is the Phillips curve, the less the cost of disinflation. In times of rapid change, especially those involving fiscal or monetary reforms, the real cost of disinflation can be smaller than in more tranquil times.

9 Concluding Remarks

For decades, Friedman's invention, the natural rate of unemployment, languished as an unimportant constant in the Phillips curve, or as a slow-moving trend. But puzzling evidence accumulated about the role of the resulting Phillips curve in the behavior of inflation and unemployment. In particular, the decade-long expansion starting in 2009 combined near-constant inflation with continuing declines in unemployment from 10 percent to 3.5 percent.

Phillips curves constructed with constant natural rates and constant slopes became untenable as this process unfolded. A few investigators reconsidered constancy of the natural rate in favor of a decline, but many concluded that it was the slope of the Phillips curve that had declined. Some estimates of the slope by the end of the recovery were close to zero.

Our investigation of the recovery starting in 2009 concludes in favor of a declining natural rate. The logical basis for this conclusion is that the anchored inflation rate converged to the Fed's target rate of two percent over such a long period of stable inflation so close to that target. A bedrock principle of the New Keynesian model is that in an economy with actual inflation equal to its anchor (sometimes called the expected rate), the observed unemployment rate equals the natural rate.

Based on our earlier work on the behavior of unemployment in cyclical recoveries in the 10 recessions since 1948, we extend our conclusion, but with more wiggle room, to all of those

recessions, because all of them share the pattern of the 2009-19 recovery. Unemployment starts at an elevated level and glides downward until interrupted by the next recession or unemployment hits bottom at around 3.5 percent.

We provide some surrounding analysis and evidence for our conclusion. The finding of declining estimated values of the slope of the Phillips curve could be an artifact of the omission of the natural rate from the regression or the use of a proxy for the natural rate that lacks the true correlation with the unemployment rate. Adjusting for this bias could disclose that the true slope of the Phillips curve remained the same or even increased as unemployment declined.

Another important conclusion is that the labor market may change dramatically upon changes in the macroeconomic environment. There may have been a substantial loosening so that prices rose instead of unemployment falling in 2021 and early 2022 followed by disinflation currently instead of rising unemployment. Prices may have had a bigger role in stabilizing output than is implied by sticky-price models.

References

- Ashley, Richard and Randal J. Verbrugge, “The Intermittent Phillips Curve: Finding a Stable (But Persistence-Dependent) Phillips Curve Model Specification,” 2023. Federal Reserve Bank of Cleveland Working Paper 19-09R2.
- Ball, Laurence and Sandeep Mazumder, “Inflation Dynamics and the Great Recession,” *Brookings Papers on Economic Activity*, 2011, 42 (1), 337–405.
- Barnes, Michelle L. and Giovanni P. Olivei, “Inside and Outside the Bounds: Threshold Estimates of the Phillips Curve,” *Federal Reserve Bank of Boston New England Economic Review*, 2003, pp. 3–18.
- Barnichon, Regis and Christian Matthes, “The Natural Rate of Unemployment over the Past 100 Years,” *FRBSF Economic Letter*, 2017, 23.
- and Geert Mesters, “On the Demographic Adjustment of Unemployment,” *Review of Economics and Statistics*, 2018, 100, 219–231.
- Blanco, Andres, Corina Boar, Callum Jones, and Virgiliu Midrigan, “Nonlinear Inflation Dynamics in Menu Cost Economies,” Working Paper, NYU 2023.

- Bok, Brandyn, Richard K. Crump, Christopher J. Nekarda, and Nicolas Petrosky-Nadeau, “Estimating Natural Rates of Unemployment: A Primer,” Working Paper 2023-25, Federal Reserve Bank of San Francisco 2023.
- Cavallo, Alberto, Francesco Lippi, and Ken Miyahara, “Large Shocks Travel Fast,” Working Paper 31659, National Bureau of Economic Research 2023.
- Cecchetti, Stephen, Michael Feroli, Peter Hooper, Anil Kashyap, and Kermit L. Schoenholtz, “Deflating Inflation Expectations: The Implications of Inflation’s Simple Dynamics,” Discussion Paper 11925, CEPR 2017.
- Coibion, Olivier and Yuriy Gorodnichenko, “Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation,” *American Economic Journal: Macroeconomics*, 2015, 7 (1), 197–232.
- Crump, Richard K., Christopher J. Nekarda, and Nicolas Petrosky-Nadeau, “Unemployment Rate Benchmarks,” Working Paper 2020-072, Finance and Economics Discussion Series, Federal Reserve Board 2020.
- , Stefano Eusepi, Marc Giannoni, and Aysegul Sahin, “A Unified Approach to Measuring u^* ,” *Brookings Papers on Economic Activity*, 2019, pp. 143–214.
- , —, —, and —, “The Unemployment-Inflation Trade-Off Revisited: The Phillips Curve in COVID Times,” Working Paper 29785, National Bureau of Economic Research 2022.
- Daly, Mary C., Bart Hobijn, Aysegul Sahin, and Robert G. Valletta, “A Search and Matching Approach to Labor Markets: Did the Natural Rate of Unemployment Rise?,” *Journal of Economic Perspectives*, 2012, 26, 3–26.
- Davis, Steven J., “Comment on “A Unified Approach to Measuring u^* ” by Richard K. Crump, Stefana Eusepi, Marc Giannoni and Aysegul Sahin,” *Brookings Papers on Economic Activity*, 2019, *Spring*, 215–223.
- del Negro, Marco, Michele Lenza, Giorgio Primiceri, and Andrea Tambalotti, “What’s Up with the Phillips Curve?,” *Brookings Papers on Economic Activity*, 2020, pp. 301–357.
- Doser, Alexander, Ricardo Nunes, Nikhil Rao, and Viacheslav Sheremirov, “Inflation Expectations and Nonlinearities in the Phillips Curve,” *Journal of Applied Econometrics*, 2023, 38, 453–471.

- Dotsey, Michael, Shigeru Fujita, and Tom Stark, “Do Phillips Curves Conditionally Help to Forecast Inflation?,” *International Journal of Central Banking*, 2018, 14, 43–92.
- Friedman, Milton, “The Role of Monetary Policy. Presidential address delivered at the 80th Annual Meeting of the American Economic Association,” *American Economic Review*, 1968, 58 (1), 1 – 17.
- Furlanetto, Francesco and Nicolas Groshenny, “Mismatch Shocks and Unemployment during the Great Recession,” *Journal of Applied Econometrics*, 2016, 31, 1197–1214.
- Galí, Jordi, Frank Smets, and Rafael Wouters, “Unemployment in an Estimated New Keynesian Model,” *NBER Macroeconomics Annual*, 2011, 26, 329–360.
- Gordon, Robert J., “The Time-Varying NAIRU and Its Implications for Economic Policy,” *Journal of Economic Perspectives*, 1997, 11 (1), 11–32.
- , “Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU,” *Brookings Papers on Economic Activity*, 1998, (2), 297–346.
- Hall, Robert E., “A Major Shock Makes Prices More Flexible and May Result in a Burst of Inflation or Deflation,” Working Paper 31025, National Bureau of Economic Research 2023.
- and Marianna Kudlyak, “The Inexorable Recoveries of Unemployment,” *Journal of Monetary Economics*, 2022, 131, 15–25.
- and —, “Why Has the US Economy Recovered So Consistently from Every Recession in the Past 70 Years?,” *NBER Macro Annual*, 2022, 36.
- Hazell, Jonathon, Juan Herreno, Emi Nakamura, and Jon Steinsson, “The Slope of the Phillips Curve: Evidence from U.S. States,” *Quarterly Journal of Economics*, 2022, 137 (3), 1299–1344.
- Hetzel, Robert L., *The Federal Reserve System. A New History*, University of Chicago Press, 2022.
- , “Learning What Is the Optimal Monetary Policy Regime,” Mercatus Research Working Paper, Mercatus Center at George Mason University 2024.
- 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?,” *Research in Economics*, 2020, 74, 26–62.

- Hornstein, Andreas and Marianna Kudlyak, “Aggregate Labor Force Participation and Unemployment and Demographic Trends,” Working Paper 2019-07, Federal Reserve Bank of San Francisco 2019.
- Jorgensen, Peter Lihn and Kevin J. Lansing, “Anchored Inflation Expectations and the Slope of the Phillips Curve,” 2019. Federal Reserve Bank of San Francisco, Working Paper No. 2019-27.
- King, Thomas B. and James Morley, “In Search of the Natural Rate of Unemployment,” *Journal of Monetary Economics*, 2007, 54, 550–564.
- Lacker, Jeffrey M., “Maximum Employment and Monetary Policy,” Speech, The Money Marketeers of New York University, New York 2012.
- Laubach, Thomas, “Measuring the NAIRU: Evidence from Seven Economies,” *Review of Economics and Statistics*, 2001, 83 (2), 218–231.
- Leduc, Sylvain and Daniel J. Wilson, “Has the Wage Phillips Curve Gone Dormant?,” *Federal Reserve Bank of San Francisco Economic Letter*, 2017, (2017-30).
- , Chitra Marti, and Daniel J. Wilson, “Does Ultra-Low Unemployment Spur Rapid Wage Growth?,” *Federal Reserve Bank of San Francisco Economic Letter*, 2019, (2019-02).
- Lucas, Robert E., “Econometric Policy Evaluation: A Critique,” *Carnegie-Rochester Conference Series on Public Policy*, 1976, 1, 19–46.
- Orphanides, Athanasios, “Monetary Policy Rules Based on Real-Time Data,” *American Economic Review*, 2002, 91 (4), 964–985.
- and John C. Williams, “Monetary Policy Mistakes and the Evolution of Inflation Expectations,” in *The Great Inflation: The Rebirth of Modern Central Banking*, Bordo and Orphanides (Eds.), 2013.
- Rogerson, Richard, “Theory Ahead of Language in the Economics of Unemployment,” *Journal of Economic Perspectives*, 1997, 11 (1), 73–92.
- Romer, David, *Advanced Macroeconomics*, McGraw-Hill Education; 5th edition, 2019.
- Sargent, Thomas J., “The Ends of Four Big Inflations,” *NBER Chapters*, in: *Inflation: Causes and Effects*, National Bureau of Economic Research, 1982, pp. 41–98.

- Smith, Simon C., Allan Timmermann, and Jonathan H. Wright, “Breaks in the Phillips Curve: Evidence from Panel Data,” Working Paper 2023-015, Federal Reserve Board, Washington, D.C., Finance and Economics Discussion Series 2023.
- Staiger, Douglas O., James H. Stock, and Mark W. Watson, “How Precise Are Estimates of the Natural Rate of Unemployment?,” *In Reducing Inflation: Motivation and Strategy*, Christina D. Romer and David H. Romer, Editors, 1997, pp. 195 – 246.
- Stella, Andrea and James H. Stock, “A State-Dependent Model for Inflation Forecasting,” 2013. Board of Governors of the Federal Reserve System, International Finance Discussion Papers No. 1062r.
- Stock, James H. and Mark W. Watson, “Modeling Inflation after the Crisis,” 2010. Prepared for the Federal Reserve Bank of Kansas City Symposium, “Macroeconomic Policy: Post-Crisis and Risks Ahead,” Jackson Hole, Wyoming, August 26-28 2010.
- Tasci, Murat, “The Ins and Outs of Unemployment in the Long Run: Unemployment Flows and the Natural Rate,” Working Paper 12-24, Federal Reserve Bank of Cleveland 2018.
- Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton, NJ: Princeton Univ. Press, 2003.