

NBER WORKING PAPER SERIES

THE INEXORABLE RECOVERIES OF US UNEMPLOYMENT

Robert E. Hall
Marianna Kudlyak

Working Paper 28111
<http://www.nber.org/papers/w28111>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2020

The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research, the Federal Reserve Bank of San Francisco or the Federal Reserve System.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 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 Inexorable Recoveries of US Unemployment
Robert E. Hall and Marianna Kudlyak
NBER Working Paper No. 28111
November 2020
JEL No. E32,J64

ABSTRACT

Unemployment recoveries in the US have been inexorable. It is a remarkable fact that, prior to 2020, after unemployment reached its peak in a recession, and a recovery began, the annual reduction in the unemployment rate was stable at around 0.1 log points per year. The economy seems to have an irresistible force toward restoring full employment. Unless another crisis intervenes, unemployment continues to glide down to its minimum level of approximately 3.5 percentage points. The observed behavior of unemployment casts doubt on the common assumption that there is a constant natural rate of unemployment around which unemployment oscillates. Instead, the natural rate of unemployment during recoveries tracks actual unemployment on its downward path.

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
marianna.kudlyak@sfrb.org

Contents

1	Business Cycle Measurement	5
1.1	Our measure of the business cycle	5
1.2	Econometrics of business-cycle measurement	6
2	Estimation Methods	7
2.1	Estimation based on chronologies	7
2.2	Estimation using the Hidden Markov Approach	9
2.3	Sampling distributions of the estimators of the recovery rate	10
3	Estimates of the Unemployment Recovery Rate	10
3.1	Unemployment recovery rate	10
3.2	Estimates by recovery	12
3.3	Recovery rates by demographic group	13
3.4	Results for the hidden Markov approach	13
3.5	Comparison of the two approaches	14
4	The Natural Rate of Unemployment	14
4.1	Inferring the natural rate of unemployment from periods of stable inflation .	14
4.2	Existing evidence that the natural rate of unemployment closely tracks the actual rate during recoveries	15
4.3	Takeaway about the natural behavior of unemployment	17
5	Concluding Remarks	18

We undertake a close examination of the behavior of unemployment over the period from 1948 to 2019 using data from the Current Population Survey (CPS). We find that during this period unemployment has shot upward 10 times as the economy has experienced economic crises. Following a crisis, the unemployment rate glides downward on a predictable but slow recovery path. In some cases, the path ends with unemployment still above its minimal level, and in others, such as the longest recovery, from 2009 to February 2020, unemployment reached 3.5 percent, which may be the current minimal level. Unemployment reached its historical minimal level over the entire period in the early 1950s, at 2.5 percent.

This paper is empirical, and we do not venture outside our sample data. We are not informed about recoveries lasting longer than the one that ended in February 2020. We are only partly informed about the current recovery following the pandemic crisis of the spring of 2020. Further, we do not enter the thicket of general equilibrium models or Phillips curves. Rather, we study the behavior of unemployment in completed recoveries recorded in the CPS. After documenting a regularity of the process, we discuss implications for inferring the path of the natural rate from the observed behavior of unemployment.

We find that rather than vibrating around a fixed natural rate, the observed behavior of unemployment comprises (1) occasional sharp upward movements in times of economic crisis, and (2) an inexorable downward glide at a low but reliable proportional rate at all other times. The glide continues until unemployment reaches a low barrier of approximately 3.5 percent or until another economic crisis interrupts the glide.

We focus on recoveries. Our measurement starts in an economy that has been hit recently by an adverse shock that triggered a recession. These shocks have heterogeneous sources. The major recession that began in 1981 is generally viewed as the result of a sharp monetary contraction, while the major recession that began at the end of 2007 got much of its strength from the financial crisis of September 2008. This paper recognizes that the shocks that propel unemployment sharply upward are heterogeneous. It is about the homogeneity of historical recoveries.

Figure 1 shows our main evidence. It displays the log of the unemployment rate during the 10 completed recoveries since 1948, with the recession spells of sharply rising unemployment left blank. The key fact about recoveries is apparent in the figure: Unemployment declines smoothly but slowly throughout most recoveries most of the time, at close to the same proportional rate. In the log plot, the recoveries appear as impressively close to straight lines. In terms of levels rather than logs, this behavior implies that unemployment falls in a year by one tenth of its level at the beginning of the year. For example, in a year starting with 7 percent unemployment, the rate falls by 0.7 percentage points during the year. We document this regularity within the two main statistical approaches to business-




Figure 1: The Paths of Log-Unemployment During Recoveries

cycle analysis and measurement: (1) construction of a chronology of turning points, and (2) estimation of a Markov regime-switching model.

Based on our findings, we discuss what can be learned about the natural rate of unemployment from the observed behavior of unemployment during recoveries. We take the *non-accelerating-inflation rate of unemployment* (NAIRU) to be a synonym for the natural rate. It captures a key property of that rate: Times of stable inflation are times when the unemployment rate is at its natural level. Most commentators have taken the natural rate to follow a smooth, low-volatility path determined by demographics. In recoveries, by contrast, we find that the natural rate follows the smoothly *declining* path of the actual rate.

We cite work by Galí, Smets and Wouters (2011) that finds an estimated path for the natural rate with high volatility, tracking the actual unemployment rate, including during recoveries. We also present simple non-parametric evidence of systematic cyclical movements of the natural rate, closely tracking the observed rate during recoveries. We conclude that our findings about the actual rate carry over closely to the natural rate. In other words, recoveries are times when the gap between the actual unemployment rate and the natural rate is small. We do not take a position on the gap during contractions when unemployment is rising rapidly.

As we write, the United States and the world are in the throes of a major pandemic and resulting deep slump. The recovery of the US unemployment rate has been vastly speedier so far than its low historical value, dropping from its maximum of 14.7 percent in April to

6.9 percent in October. In a separate paper (Hall and Kudlyak (2020a)), we discuss how recalls of a completely unprecedented volume of unemployed individuals to their existing jobs accounts for the highly unusual rate of decline of unemployment, and why it is likely that the normal pattern of low but reliable decline of unemployment will resume once those individuals are back at work. Also see Fujita and Moscarini (2017) on recalls in general and Gregory, Menzio and Wiczer (2020) on the role of recalls in the recovery from the pandemic.

We are not the first to study the time-series properties of US unemployment. The basic asymmetry between the sharp rise of unemployment in contractions and the slow pace of expansions is well known, and studied carefully with new results and a thorough discussion of the earlier literature in Dupraz, Nakamura and Steinsson (2019). We note that a well-documented property of the unemployment rate—most recently confirmed by Dupraz and co-authors—is that unemployment rises rapidly in response to a significant aggregate adverse shock and then gradually recovers to a level of 3 to 5 percent of the labor force. Like fuel prices, unemployment rises like a rocket and falls like a feather. Our addition to this literature is our demonstration of the astonishing reliability of the recovery process. We measure the rate of recovery of unemployment from recession-highs and demonstrate how uniform the rate is. Based on our findings, we discuss what can be learned about the natural rate of unemployment.

In a separate paper (Hall and Kudlyak (2020b)), we point out the puzzle of *slow decline* of unemployment in recoveries. Initially pointed out by Cole and Rogerson (1999), the puzzle is that unemployment declines much slower than the measured individual job finding rates would seem to indicate. In that paper, we also discuss models in the Diamond-Mortensen-Pissarides tradition that can account for this observation.

Recently, the Federal Reserve Board announced that, in future expansions, policy would not lean against a glide path that brought unemployment below 4 percent until there were clear signs of rising inflation (Powell (2020)). The Fed’s new policy of not resisting the downward glide in unemployment during periods of calm is consistent with our conclusions.

1 Business Cycle Measurement

1.1 Our measure of the business cycle

To study recoveries, we need a measure of the business cycle. Romer and Romer (2019) discuss cycle measures in detail. They conclude that the preferred defining characteristic of the measure is its ability to capture unused resources. In current business-cycle research, the primary alternative definition is to extract a higher-frequency component from real GDP or other output measure. That component is the higher-frequency series from the Hodrick-

Prescott or other bandpass filter. We agree with the Romers that tying the business cycle to unused resources is conceptually superior to tying it to higher-frequency movements.

Our view further adopts the Romers’ conclusion that the unemployment rate, or a measure derived from the unemployment data from the Current Population Survey, is the best available measure of the cycle. The unemployment rate appears to contain almost no movements associated with productivity or similar forces that would call for filtering out. A modest slow-moving demographic component of the unemployment rate is present—see Hornstein and Kudlyak (2019) and Crump, Eusepi, Giannoni and Sahin (2019).

1.2 Econometrics of business-cycle measurement

We treat log-unemployment as the sum of a latent trend component and a latent stationary component capturing survey sampling errors and other deviations from the trend. In a crisis, the trend is fairly sharply upward. During a recovery, the trend is modestly downward. Our objective is to measure the central tendency and dispersion of the rate of decline of the latent systematic component of the monthly change of log-unemployment rate during recoveries.

We consider the general class of models

$$f(u_t) = x_t + \epsilon_t, \tag{1}$$

where $f(\cdot)$ is a monotonic transformation, x_t is the systematic trend component capturing the business cycle, and ϵ_t is the random unsystematic component, taken to be uncorrelated with x_t . The model for x_t describes the random arrival of peaks and troughs with linear paths connecting these turning points. We use the log transformation because it fits the data quite a bit better than a model where the level of unemployment is the dependent variable. In Figure 1, the implied linearity of the recovery paths of log unemployment is quite clearly confirmed.

In the specification with $\log u_t$ on the left-hand side, the slope with respect to x_t is measured in *log points*, that is, percent changes in unemployment per unit of x . Where possible, we avoid stating the results in the potentially confusing terms of percents of percents, but that is the actual implication of the specification. We use the term log points and state them as decimals. For example, a typical finding is that unemployment declines during a recovery by 0.1 log points per year, which is 0.7 percentage points if the unemployment rate starts at 7 percent of the labor force.

The literature has focused on two general classes of specifications for the systematic component. One is *chronology-based* and proceeds by assigning turning points—dates when recessions end and recoveries begin and dates when recoveries end and recessions begin.

Chronologies are available from published sources, notably the National Bureau of Economic Research, which identifies monthly dates of turning points in a latent measure called economic activity. They can be created for a particular time series, such as the unemployment rate, as an exercise in human pattern recognition. And chronologies can be created by algorithms, such as the one described in Dupraz et al. (2019).

Given a chronology, one can approximate the systematic component x_t by interpolating between the turning points and measuring the noise ϵ_t as the residual between $\log u_t$ and x_t . The systematic trend component x_t is a smooth function of t . We take it to be a straight line between the turning points of the series, so x_t has equal increments over time, between the turning points. Overall, the trend component is a linear spline.

The other class of models focuses on *regime switching*, where the systematic component modeled as a statistical time series that obeys one model in contractions and another in recessions. Hamilton (1989) launched the econometric literature on Markov-switching models in this class.

The key difference between these classes is that turning points are latent unobserved events in regime-switching models. The models yield a probability that a given month is a turning point, not an unambiguous turning-point date.

2 Estimation Methods

2.1 Estimation based on chronologies

We consider three monthly business-cycle chronologies:

1. NBER: The chronology maintained by the National Bureau of Economic Research identifying turning points in economic activity, as described in detail at NBER.org
2. DNS: The chronology produced by the DNS algorithm based on US unemployment from January 1948 through February 2020, with size parameter 1.5
3. HK: The chronology produced by the authors based on observed business cycle peaks and troughs, using the same data as DNS

DNS developed an algorithm that maps a time series into another time series taking on three discrete values: trough, peak, and neither. For unemployment, most months are classified as neither a trough nor a peak, but rather a continuation of a previous trend. The DNS algorithm is based on judgment about how to extract turning points from time-series data, but its application banishes human judgment from the actual determination. The algorithm is a filter that applies prior beliefs embodied in the algorithm to determine turning




Figure 2: Three Chronologies for the US Unemployment Rate

points. Because the algorithm makes it cheap to extract a chronology from hypothetical data, and because it is a function, producing a single chronology from any particular input, it is a suitable basis for experimenting with the use of a chronology in a situation where noise partly obscures an underlying true chronology. We have subjected the algorithm to thousands of experimental paths of unemployment in this endeavor.

Our procedure (HK) delivers turning points similar to the DNS algorithm. However, we pick the latest points for peaks and troughs, consistently with our definition of the recovery. The results based on the HK chronology are quite similar to those based on the DNS chronology.

Figure 2 shows the three chronologies. One disagreement is immediately apparent—the NBER chronology has a recovery beginning in July 1980 and ending 12 months later in July 1981. There is no comparable recovery in the other chronologies. In general, DNS and HK are similar to one another and differ from NBER. The reason is that DNS and HK are chronologies for unemployment alone, while NBER is a chronology for latent economic activity. For the dates in the table starting with January 1980, NBER.org has published explanations of the various indicators that form the basis for the determination of the dates.

We measure the recovery rate as the average annual decline in log unemployment over the entire set of recoveries and over individual recoveries. Although the mean is the natural estimator irrespective of the time-series process of the shocks ϵ , the process does matter for the standard error of the sample mean.

Given a recovery running from an initial high point of unemployment, which we number as zero, to the following low point, which we number as T , our model for a single recovery is

$$12(\log u_T - \log u_0) = -\beta T + \epsilon_T. \quad (2)$$

We include the 12 so that the recovery rate β is in log points per year. We use the estimator

$$\hat{\beta} = -\frac{12(\log u_T - \log u_0)}{T}. \quad (3)$$

We use a modified bootstrap procedure, described later in the paper, to approximate the sampling distribution of $\hat{\beta}$.

2.2 Estimation using the Hidden Markov Approach

Our second approach to modeling business cycles posits the same basic cyclical structure,

$$\log u_t = x_t + \epsilon_t, \quad (4)$$

where x_t is unobserved, but hypothesized to be piecewise linear, with switching between positive and negatively sloped segments at random, according to a Markov process. Under the assumption that the disturbance is a random walk, $\Delta\epsilon_t = \eta_t$, with η_t being white noise, the model becomes

$$\Delta \log u_t = \beta_i + \eta_t. \quad (5)$$

The monthly increment, β_i , $i \in \{1, 2\}$, shifts back and forth between $i = 1$ for recessions and $i = 2$ for recoveries. We focus on β_2 , the log-decline in unemployment during recoveries. James Hamilton pioneered the econometric analysis of this class of models. He derived the likelihood function in a computationally convenient form (Hamilton (1989)).

We note that the assumption that ϵ is a random walk is essential to our application of the hidden Markov model. We need the assumption to justify taking first-differences, which has the effect of isolating β_i on the right-hand side of the equation. This step also puts the iid innovation η_t on the right-hand side, a property that is the starting point for the regime-change class of models. However, we know that ϵ has a high AR(1) parameter ρ but somewhat less than one. Thus first-differencing does not yield the true innovation η , but only something close to it.

Because Hamilton’s approach is an application of maximum likelihood, the information matrix is the basis of an estimator of the covariance matrix of the estimated parameters. But this principle only applies in the absence of specification error. Our application involves the specification error just mentioned. Consequently, we estimate the model as if ϵ were a random walk, and then validate the estimator by bootstrap. Marcelo Perlin provided the Matlab package for estimating hidden Markov models that we used (Perlin (2015)). Our validation efforts supported the accuracy of the results of his package but found that the standard error from the information matrix was considerably smaller than from the bootstrap. We report both.

2.3 Sampling distributions of the estimators of the recovery rate

We use a modified bootstrap procedure to approximate the sampling distribution of the estimates of the recovery rate and its standard deviation across recoveries. We resample the innovations η from the HK chronology, form the AR(1) values of the disturbances using $\rho = 0.92$, and construct simulated chronologies using the transition probabilities from the HK chronology. Next we calculate a time series of corresponding values of the unemployment rate as the sum of the calculated disturbances and the fitted values from the HK chronology. We then pass the simulated unemployment numbers through the DNS chronology software. Using the resulting chronology, we calculate the recovery rates by recovery and the averages for the full sample and the later sample. We also re-estimate the hidden-Markov model. We repeat the process 10,000 times for the chronology-based estimates and 500 times for the hidden Markov calculations, which are much more computationally intensive.

This procedure captures the variation in the estimates arising from the disturbances, from the randomness of the underlying constructed chronology, and from the blurring of the chronology from the added disturbances. It understates the resulting standard errors on account of omitting the variation arising from the estimation of ρ and from estimating the transition probabilities. We believe the understatement is small.

3 Estimates of the Unemployment Recovery Rate

3.1 Unemployment recovery rate

Table 1 shows our statistical results for both approaches. The upper panel displays results for the entire sample, starting in 1949, and the lower panel displays the same statistics for the post-1959 sample. It shows the estimates of the two key statistics in this study: the annual recovery rate in log points, β , and the standard deviation of β_r across the recoveries,

	Chronology			Hidden Markov	Difference
	NBER	Dupraz-Nakamura-Steinsson	Hall-Kudlyak		
Full sample					
Annual recovery rate, log points	0.087	0.132	0.129	0.066	
Bootstrap standard error	(0.016)	(0.017)	(0.017)	(0.042)	
Information matrix standard error				(0.015)	
Standard deviation of recovery rate across recoveries					
Bootstrap standard error	0.076 (0.119)	0.084 (0.117)	0.084 (0.115)		
Difference between HK estimate and hidden-Markov estimate					0.066
Bootstrap standard error					(0.015)
After 1959					
Annual recovery rate, log points	0.067	0.106	0.103	0.070	
Bootstrap standard error	(0.013)	(0.007)	(0.007)	(0.039)	
Information matrix standard error				(0.014)	
Standard deviation of recovery rate across recoveries					
Bootstrap standard error	0.025 (0.038)	0.011 (0.039)	0.016 (0.038)		
Difference between HK estimate and hidden-Markov estimate					0.036
Bootstrap standard error					(0.012)

Table 1: Statistical Results

indexed by r . The upper panel covers the entire sample period, from October 1949 through February 2020. The left three columns show the recovery rate β in log points per year using the chronology approach, along with the standard deviation of the recovery rate, both with bootstrap standard errors.

For the NBER chronology with the full sample, the decline rate pooled across recoveries is estimated as 0.087. Recovery rates for the DNS and HK chronologies are similar to each other and are well above the NBER level, at 0.129 and 0.132. The DNS and HK chronologies, constructed from unemployment alone, are more successful at capturing the movements of unemployment during recoveries because they are better synchronized with the actual movements. Of course, DNS and HK would be correspondingly poorer at tracking economic activity, the concept behind the NBER chronology.

We illustrate the interpretation of the annual decline figures in the table with an example from the recovery rates based on the NBER chronology. Consider the situation just after a

severe recession, with the unemployment rate starting at 10 percent. The expected unemployment rate a year later is $10 \exp(-0.087) = 9.2$ percent. With the higher recovery rate based on DNS, the rate a year later would be 8.8 percent, essentially the same as with the HK chronology. According to the DNS rate, starting from 6 percent, the unemployment rate a year later would be 5.3 percent.

The recovery rates for the sample running from May 1961 through February 2020, shown in the lower panel of Table 1, are somewhat lower, because, as shown in Figure 3, the first three recoveries in the full sample had substantially higher recovery rates than did any of the later recoveries.

3.2 Estimates by recovery

Figure 3 shows the results of estimating each of the 10 recoveries in the HK chronology separately. The separate rates from 1961 onward cluster close to 0.10. Over that 60-year period, with 7 recessions and recoveries, some mild and two quite severe, the recovery rates are remarkably similar. Table 1 summarizes this finding in terms of the standard deviations across either the 10 recoveries of the full sample or the 7 recoveries of the later sample.

The standard deviations across the individual recoveries over the full sample, shown in the upper panel, are all around 0.08 with standard errors around 0.12. As the figure shows, the first three recoveries have rather higher recovery rates than the later seven. The hypothesis of equality, that is, a standard deviation of zero, is easily not rejected, but the confidence interval is fairly wide, for all three chronologies. On the other hand, the standard deviations of the seven later recoveries have point estimates of 0.025 or under confirming the visual evidence of the figure. And the confidence intervals, constructed with standard errors of 0.39 or 0.38, are tight. These results nail down the primary thesis of our study—the uniformity of recovery rates over the past 60 years.

Why is there a widespread impression that the recovery from the 2007 recession and financial crisis of 2008 was slower than previous recoveries? The answer is that recoveries tend to be judged in terms of output. Both actual growth of real GDP and growth of potential GDP were lower for a number of reasons, including especially the decline in the rate of productivity growth—see Fernald, Hall, Stock and Watson (2017). The facts are that output growth was substandard during the recovery but the decline in unemployment was at the normal rate for recoveries after 1960.




Figure 3: Estimated Recovery Rates by Recovery

3.3 Recovery rates by demographic group

We repeated the chronology-based estimation for the unemployment rates of four different demographic groups, using the DNS chronology function and the post-1959 sample. Recovery rates are, for young people (age 16 to 25), 0.091; women (age 16+), 0.095; non-white (age 16+), 0.110; prime age (25 through 64), 0.125; total (age 16+), 0.106, as in the lower panel and DNS column of Table 1. The uniformity across demographic groups confirms the robustness of our findings and conclusions.

3.4 Results for the hidden Markov approach

The column of Table 1 labeled *Hidden Markov* shows results from estimating the hidden Markov model using the full sample and the sample starting in 1961. For the full sample, the estimated recovery rate is 0.066, well below the earlier results for the chronology-based estimates, especially in the case of DNS and HK. Estimates for the post-1960 sample are closer, at 0.070 for hidden-Markov and 0.103 for DNS. The NBER estimate concurs with the hidden-Markov one, but we have good reason to believe that it is downward biased relative to estimates that are based only on the unemployment rate, which include hidden-Markov.

3.5 Comparison of the two approaches

In the lower right-hand corner of both panels of Table 1, we report the difference between the HK variant of the chronology-based estimate of the pooled recovery rate and the hidden-Markov-based estimate. Below each estimate, we give the bootstrap standard error of the difference, based on the difference in 500 replications of the bootstrap simulations. In both cases, we reject the hypothesis that the difference arises from sampling variation alone. The rejection is stronger for the full sample than for the later sample.

One reason for the disagreement between the two estimators is that the theory of the application of the hidden-Markov setup to our problem requires the assumption that the disturbance is a random walk, whereas it is actually an AR(1) process with coefficient somewhat less than one. The disagreement between the information-matrix estimate of the hidden-Markov standard error and the bootstrap may also result from this conflict.

The chronology-based estimator can be considered an application of Bayesian thinking, in that it imposes prior beliefs about the process. The posterior, so to speak, may involve a higher implied value of the recovery rate because the prior belief pushes the likelihood in that direction.

In any case, the two statistical approaches agree that the basic structure that they share is supported in the data. It really is true and not the result of an optical illusion that unemployment rises rapidly and then declines slowly and uniformly over the business cycle. Figure 4 shows how well the fitted value from the HK estimation with different recovery rates for each recovery (those shown in Figure 3) fits the data. The R^2 of the fit is 0.91.

4 The Natural Rate of Unemployment

Milton Friedman introduced the natural rate of unemployment (Friedman (1968)), a concept that came to dominate macroeconomics under that name or its synonym, the non-accelerating inflation rate of unemployment rate or NAIRU. Friedman's analysis implies that years of stable inflation are years when the unemployment rate is at its natural level.

4.1 Inferring the natural rate of unemployment from periods of stable inflation

Figure 5 takes a non-parametric, minimal-assumption approach to inferring the natural rate of unemployment. The dots in the figure identify calendar quarters when the acceleration of the GDP price deflator is within a narrow band around zero. To identify those quarters, we first form the inflation rate in the normal way as the percentage increase of the index in one




Figure 4: Fit of the Chronology Model

quarter over its value 4 quarters earlier. Using the 4-quarter span dramatically reduces the role of transitory influences. Second, we calculate acceleration as the 4-quarter change in the inflation rate. Finally, a quarter has price-level acceleration close to zero, and appears as a dot in the figure, if the absolute value of its acceleration is less than 0.20 percentage points.

In the most recent recovery, there were 7 quarters of stable inflation by our measure. In the first four, unemployment was high—well above the five-percent level of most estimates of the natural rate based on demographics, which we discuss below. The figure shows that quite a wide range of unemployment levels are identified as time-specific values of the natural rate of employment, but no value above 8 percent received that designation.

4.2 Existing evidence that the natural rate of unemployment closely tracks the actual rate during recoveries

A large literature has developed that tracks the evolution of the natural rate arising from changes in the demographic composition of the unemployed. Estimates of the long-run demographic trends in unemployment are often taken as the estimate of the natural rate. For example, the Congressional Budget Office maintains estimates of the path of the natural rate which show it rising gradually to a maximum in the 1970s, on account of a bulge of young workers more prone to unemployment, then declining as that cohort aged and enjoyed lower unemployment. Recent estimates of the long-run demographic trend in unemploy-




Figure 5: Quarterly Unemployment with Quarters Marked when Price-Level Acceleration Was Close to Zero So Unemployment Was Presumptively at Its Natural Level

ment are available from Crump et al. (2019), Hornstein and Kudlyak (2019), Barnichon and Mesters (2018) and Barnichon and Matthes (2017). These recent studies and their numerous predecessors consider only demographic change.

Capturing the other forces that affect the natural rate of unemployment is a challenge. In principle, if we knew the slope of the Phillips curve, we could back out the implied natural rate from the observed relation between inflation and unemployment. Few investigators have attempted this approach, because they need outside evidence on the natural rate to identify the slope of the Phillips curve. Studies of the Phillips curve often use a smooth, low-volatility path determined by demographics as the natural rate. Recent influential studies of the Phillips curve include Hazell, Herreno, Nakamura and Steinsson (2020), Jorgensen and Lansing (2019), Coibion and Gorodnichenko (2015), Stella and Stock (2013), Meier (2010), Laubach (2001), Staiger, Stock and Watson (1997). Crump, Nekarda and Petrosky-Nadeau (2020) surveys this literature.

Galí et al. (2011) undertake the challenge of backing out the natural rate from an economic model. That paper employs a New Keynesian model with a full sub-model that deals explicitly with unemployment. It solves the model to find the implied value of the natural rate. Figure 10 in that paper plots the calculated natural rate from 1965 through 2011. The path of the natural rate is completely different from the paths calculated with demographic adjustments. It does capture the bulge of unemployment in the 1970s, but it also rises

substantially in harmony with actual unemployment, especially after the serious recessions of 1981 and 2007-2009. The natural rate does not account for all of the cyclical movement of actual unemployment—the gap between actual and natural unemployment accounts for some of the cyclical movements. The authors reach their conclusion in the context of a particular model that embodies many prior beliefs and parametric restrictions. But the key conclusion is similar to the one reached from our non-parametric, minimal-assumption approach in Figure 5.

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.

A number of studies based on the Phillips curve note that there have been several instances when large movements in the unemployment rate have coincided with small changes in the inflation rate. Taking a smooth demographic trend as the measure of the natural rate, the studies have invoked non-linearities in the Phillips curve where the coefficient on the unemployment gap in recoveries is close to zero. These studies include Barnes and Olivei (2003), Dotsey, Fujita and Stark (2018), Ashley and Verbrugge (2020), Doser, Nunes, Rao and Sheremirov (2017), and Hooper, Mishkin and Sufi (2020). We view these studies as performing a variety of parametric extensions, in the Phillips curve framework, of the facts revealed in our Figure 5.

4.3 Takeaway about the natural behavior of unemployment

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 inflationary 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 at the natural rate and is not a measure of inflationary pressure. Under those conditions, there is no meaningful unemployment gap.

5 Concluding Remarks

We have developed a parsimonious statistical model of the behavior of observed unemployment. It describes: (1) occasional sharp upward movements in unemployment in times of economic crisis, and (2) an inexorable downward glide at a low but reliable proportional rate at all other times. The glide continues until unemployment reaches approximately 3.5 percent or until another economic crisis interrupts the glide.

Using the premise that times of stable inflation are times when the unemployment rate is at its natural level, we use our findings on the observed behavior of the unemployment rate to infer the behavior of the natural rate of unemployment. Our evidence from a simple non-parametric exercise shows that instead of oscillating around a constant or slow-moving time-varying trend, the natural rate follows the smoothly *declining* path of the actual rate during recoveries. A similar conclusion follows from the general equilibrium model of Galí et al. (2011).

As for policy, the Fed's new policy of not resisting the downward glide in unemployment during periods of calm is consistent with our conclusions.

References

- Ashley, Richard and Randal J. Verbrugge, “Finding a Stable Phillips Curve Relationship: A Persistence-Dependent Regression Model,” 2020. Federal Reserve Bank of Cleveland Working Paper 19-09R.
- 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.
- 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.
- Cole, Harold L. and Richard Rogerson, “Can the Mortensen-Pissarides Matching Model Match the Business-Cycle Facts?,” *International Economic Review*, 1999, 40 (4), 933–959.
- Crump, Richard K., Christopher J. Nekarda, and Nicolas Petrosky-Nadeau, “Unemployment Rate Benchmarks,” 2020. Finance and Economics Discussion Series, Federal Reserve Board, Working Paper No. 2020-072.
- , Stefano Eusepi, Marc Giannoni, and Aysegul Sahin, “A Unified Approach to Measuring u^* ,” 2019. Paper presented at the Brookings Papers on Economic Activity Conference, held at the Brookings Institution, Washington, March 7–8.
- Doser, Alexander, Ricardo Nunes, Nikhil Rao, and Viacheslav Sheremirov, “Inflation Expectations and Nonlinearities in the Phillips Curve,” 2017. Federal Reserve of Boston Working Paper No. 17-11.
- Dotsey, Michael, Shigeru Fujita, and Tom Stark, “Do Phillips Curves Conditionally Help to Forecast Inflation?,” *International Journal of Central Banking*, 2018, 14, 43–92.
- Dupraz, Stephane, Emi Nakamura, and Jon Steinsson, “A Plucking Model of Business Cycles,” Working Paper No. 26351, National Bureau of Economic Research 2019.

- Fernald, John., Robert E. Hall, James H. Stock, and Mark W. Watson, “The Disappointing Recovery of Output after 2009,” *Brookings Papers on Economic Activity*, 2017, *Spring*.
- 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.
- Fujita, Shigeru and Giuseppe Moscarini, “Recall and Unemployment,” *American Economic Review*, 2017, *107* (12), 3875–3916.
- Galí, Jordi, Frank Smets, and Rafael Wouters, “Unemployment in an Estimated New Keynesian Model,” *NBER Macroeconomics Annual*, 2011, *26*, 329–360.
- Gregory, Victoria, Guido Menzio, and David Wiczer, “Pandemic Recession: L or V-Shaped?,” 2020. New York University, Working Paper.
- Hall, Robert E. and Marianna Kudlyak, “Unemployed With Jobs and Without Jobs,” Working Paper No. 27886, National Bureau of Economic Research 2020.
- and — , “Why Has the US Economy Recovered So Consistently from Every Recession in the Past 70 Years?,” Working Paper No. 27234, National Bureau of Economic Research 2020.
- Hamilton, James D., “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle,” *Econometrica*, 1989, *57*, 357–384.
- Hazell, Jonathon, Juan Herreno, Emi Nakamura, and Jon Steinsson, “The Slope of the Phillips Curve: Evidence from U.S. States,” 2020. University of California, Berkeley, Working Paper.
- 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,” 2019. Federal Reserve Bank of San Francisco, Working Paper No. 2019-07.
- Jorgensen, Peter Lihn and Kevin J. Lansing, “Anchored Inflation Expectations and the Flatter Phillips Curve,” 2019. Federal Reserve Bank of San Francisco, Working Paper No. 2019-27.

- Laubach, Thomas, “Measuring the Nairu: Evidence from Seven Economies,” *Review of Economics and Statistics*, 2001, 83 (2), 218–231.
- Meier, Andre, “Still Minding the Gap—Inflation Dynamics during Episodes of Persistent Large Output Gaps,” 2010. IMF Working Paper No. 10/189.
- Perlin, Marcelo, “MS Regress - The MATLAB Package for Markov Regime Switching Models,” 2015. Available at SSRN.
- Powell, Jerome H., “New Economic Challenges and the Fed’s Monetary Policy Review,” Technical Report, Remarks by Chair Board of Governors of the Federal Reserve System at “Navigating the Decade Ahead: Implications for Monetary Policy,” an economic policy symposium sponsored by the Federal Reserve Bank of Kansas City Jackson Hole, Wyoming August 2020.
- Romer, Christina D. and David H. Romer, “NBER Business Cycle Dating: History and Prospect,” December 2019. University of California, Berkeley.
- 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.