

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

The Inexorable Recoveries of US Unemployment

Robert E. Hall
Hoover Institution and Department of Economics
Stanford University

Marianna Kudlyak
Federal Reserve Bank of San Francisco
Hoover Institution, IZA, CEPR

June 2022

Working Paper 2021-20

<https://www.frbsf.org/economic-research/publications/working-papers/2021/20/>

Suggested citation:

Hall, Robert E., Marianna Kudlyak. 2022 “The Inexorable Recoveries of US Unemployment,” Federal Reserve Bank of San Francisco Working Paper 2021-20.
<https://doi.org/10.24148/wp2021-20>

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Board of Governors of the Federal Reserve System.

The Inexorable Recoveries of Unemployment *

Robert E. Hall

Hoover Institution and Department of Economics

Stanford University

rehall@stanford.edu; stanford.edu/~rehall

Marianna Kudlyak

Federal Reserve Bank of San Francisco

Hoover Institution, IZA, CEPR

marianna.kudlyak@sf.frb.org; www.mariannakudlyak.com

June 20, 2022

Abstract

Unemployment recoveries in the US have been inexorable. In the aftermath of a recession, unless another crisis intervenes, unemployment continues to glide down. Between 1948 and 2019, the annual reduction in the unemployment rate during cyclical recoveries was distributed around 0.1 log points per year. The economy seems to have an irresistible force toward restoring full employment. Occasionally, unemployment rises rapidly during an economic crisis, while most of the time, unemployment declines slowly and smoothly at a near-constant proportional rate. Similar properties hold for other measures of the US unemployment rate and for unemployment in emerging and advanced countries.

JEL: E32, J63, J64.

Keywords: Business cycle, Recovery, Unemployment, Recession.

*Corresponding author: Robert E. Hall, Hoover Institution, Stanford University, Stanford, California 94305 USA. +1 650 723 2215, rehall@stanford.edu. Hall's research was supported by the Hoover Institution. The opinions expressed are those of the authors and do not reflect those of the Federal Reserve Bank of San Francisco, the Federal Reserve System, or the National Bureau of Economic Research. We thank Stephane Dupraz, Emi Nakamura, and Jón Steinsson for providing the code for their business-cycle chronology software. We thank Markus Brunnermeier, Peter Diamond, and the referee and editor for comments. Backup data and code are available at Kudlyak's website.

1 Introduction

A close examination of the behavior of US unemployment during cyclical recoveries, over the period from January 1948 to February 2020, using data from the Current Population Survey (CPS), reveals that unemployment shot upward 10 times as the economy experienced economic crises. Following each crisis, the unemployment rate glided downward on a predictable but slow recovery path. In the longest recovery, from October 2009 to February 2020, unemployment reached 3.5 percent. Unemployment reached its historical minimal level over the entire period in the early 1950s, at 2.5 percent.

This paper is empirical and limited to the period from the beginning of modern unemployment measurement, in January 1948, to the end of the last completed recovery, in February 2020. Further, it does not enter the thicket of general equilibrium models or Phillips curves. Rather, it studies the behavior of unemployment in completed recoveries recorded in the CPS. The behavior 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 other times.

The paper focuses on recoveries. Our measurement starts in an economy that has been hit recently by an adverse shock that triggered a recession. 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. Shocks that propel unemployment sharply upward have heterogeneous sources. The paper is about the homogeneity of historical recoveries.

Figure 1 shows the main evidence. It displays the log of the unemployment rate during the 10 completed recoveries since 1948, with recession spells of sharply rising unemployment left blank. The key facts about recoveries are 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. The paper documents this regularity using a standard statistical approach to business-cycle analysis and measurement—construction of a chronology of turning points. It also shows that measures of US unemployment extended to include discouraged workers and others, not counted in the labor force, display the same consistent pattern as the standard unemployment rate. And it shows that unemployment in other advanced and emerging countries behaves in much the same way as in the US.

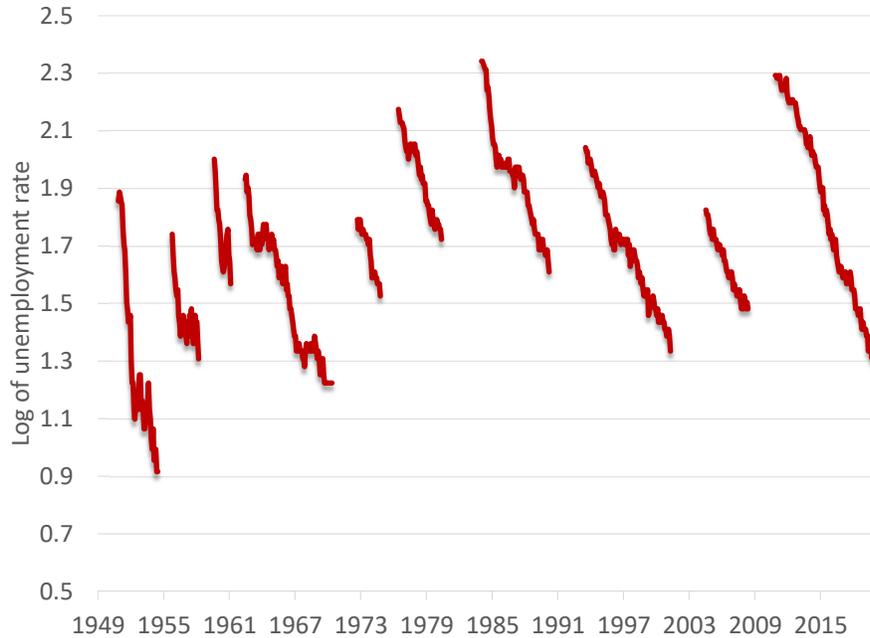


Figure 1: The Paths of Log-Unemployment during Recoveries

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 in US unemployment is well known. It is studied carefully with new results and a thorough discussion of the earlier literature in Dupraz, Nakamura and Steinsson (2022). A well-documented property of the unemployment rate—most recently confirmed by those authors—is that unemployment rises rapidly in response to a significant aggregate adverse shock and then gradually recovers. Like fuel prices, unemployment rises like a rocket and falls like a feather. What is new here is a demonstration of the reliability of the recovery process. This involves documenting the slow rate of recovery of unemployment from recession-highs and showing how uniform the rate is.

A companion paper, Hall and Kudlyak (2022), considers resolutions of the puzzle of slow decline of unemployment in recoveries. Initially pointed out by Cole and Rogerson (1999), the puzzle is that unemployment declines much more slowly than the measured individual job finding rates would seem to indicate. The companion paper discusses models in the Diamond-Mortensen-Pissarides tradition that can account for the puzzle.

2 Business Cycle Measurement

The study of recoveries requires a measure of the business cycle. Romer and Romer (2020) 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 based on extracting 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.

This paper further adopts the Romers’ conclusion that the unemployment rate, or a similar 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).

We model log-unemployment in recoveries as the sum of a latent declining path component and a latent stationary component capturing survey sampling errors and other deviations from the path. The path is modestly downward. The 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.

The formal model used here is

$$\log u_t = \alpha - \beta t + \epsilon_t, \tag{1}$$

$\alpha - \beta t$ is the systematic path component capturing the recovery phase of the business cycle, and ϵ_t is the random unsystematic component, taken to be uncorrelated with t .

In Figure 1, the approximate linearity of the recovery paths of log unemployment is plainly visible. A later section of the paper considers alternative parametrizations of the time path. The conclusion is that the log specification is a reasonable approximation.

With $\log u_t$ on the left-hand side, the downward slope β is measured in log points, that is, percent declines in unemployment per unit of t . Where possible, the description avoids stating the results in the potentially confusing terms of percents of percents, but that is the actual implication of the specification. Our explanation uses 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.

A general class of specifications for the systematic component is based on a chronology 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. Chronologies 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. (2022). Given a chronology, we estimate the systematic component $\alpha - \beta t$ by standard econometric methods.

A second approach to modeling business cycles draws on a class of models launched in Hamilton (1989). Results based on this approach appear in Hall and Kudlyak (2020a), a working paper version of this paper. Marcelo Perlin provided the Matlab package for estimating hidden Markov models that we used (Perlin (2015)).

3 Estimation

3.1 Chronologies of US business-cycle recoveries

We consider three monthly business-cycle chronologies: First is NBER, the chronology maintained by the National Bureau of Economic Research identifying turning points in economic activity, as described in detail at NBER.org. Second is DNS, the chronology described in Dupraz et al. (2022). These authors developed an algorithm to generate a chronology from a time series. The third chronology, HK, is one produced by us, using the same data as DNS.

The NBER employs a committee of business-cycle specialists to construct its chronology. The committee has developed a definition of economic activity, based primarily on real output and real income, together with definitions of contractions and recoveries. The chronology comprises turning points—activity reaches a trough in the transition month from contraction to recovery and it reaches a peak in the transition month from recovery to contraction. The committee operates in delayed real time—announcements of the determination of a turning point occurs around 6 to 18 months after the event.

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 path. The DNS algorithm is based on its creators’ judgments 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 points. The algorithm makes it cheap to extract a chronology, and, because it is a function, producing a unique chronology from any particular input, it is a suitable basis for experimenting with the use of a chronology.

Our procedure (HK) uses informed human judgment to identify turning points based on criteria similar to those of the NBER, but using data from just one cyclical indicator, the unemployment rate compiled by the Bureau of Labor Statistics. Specifically, it is the

standard monthly unemployment rate, called U-3 in the BLS’s nomenclature. The main purpose of the HK chronology is to validate our use of the pure algorithm of DNS as a proxy for human judgment.

The HK chronology has turning points similar to the DNS algorithm—the validation is successful. However, we pick some later dates for peaks and troughs. The statistical results based on the HK chronology are quite similar to those based on the DNS chronology. Accordingly, we are confident that our application of the DNS algorithm to other measures of US unemployment, and to unemployment data from other countries, generates chronologies that resemble reasonably closely those might have been generated by human judgment.

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 two chronologies. This disagreement reflects the small magnitude of the rise and fall of unemployment in the 1980 cycle. 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.

Although the NBER has determined that April 2020 was a turning point from recession to recovery, that recovery is not included here because it is incomplete as of this writing, and because of the explosion of temporary-layoff unemployment, discussed briefly later in this paper.

3.2 Modeling recoveries using a chronology

We develop a separate model with its own parameters for each recovery in a chronology. For a recovery running from an initial high point of unemployment, which we number as $t = 0$, to the following low point, which we number as T , the model for a single recovery is

$$\log u_t = \alpha - \beta t + \epsilon_t. \tag{2}$$

The residual, ϵ_t , follows an AR(1) process,

$$\epsilon_t = \rho\epsilon_{t-1} + \eta_t. \tag{3}$$

The innovation η_t is white noise. An autoregressive transformation yields

$$\log u_t = \rho \log u_{t-1} + (1 - \rho)\alpha - \beta[t - \rho \cdot (t - 1)] + \eta_t. \tag{4}$$

A re-parametrization in terms of $\kappa = (1 - \rho)\alpha - \rho\beta$ and $\gamma = (1 - \rho)\beta$ yields

$$\log u_t = \rho \log u_{t-1} + \kappa - \gamma t + \eta_t. \tag{5}$$

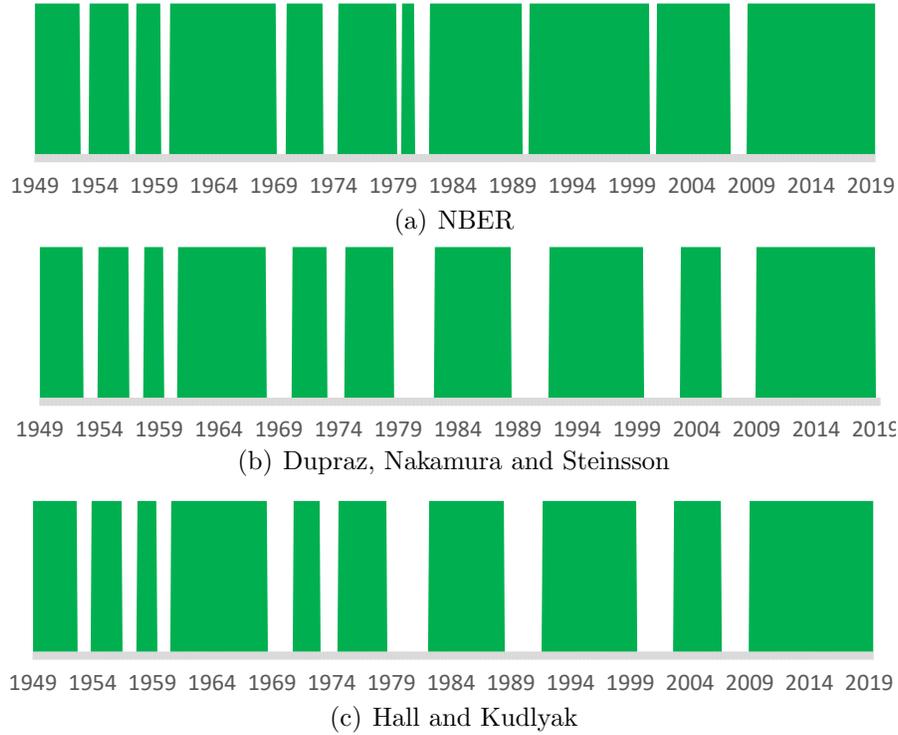


Figure 2: Three Chronologies for the US Unemployment Rate

The implied value of the recovery-rate parameter is

$$\beta = \frac{\gamma}{1 - \rho}. \quad (6)$$

The autoregressive parameter ρ indexes a range of estimators of the key parameter β , the recovery rate. If the random part of the unemployment path is serially uncorrelated, with $\rho = 0$, the equation conditional on ρ is

$$\log u_t = \alpha - \beta t + \eta_t, \quad (7)$$

and the estimated recovery rate is the coefficient of $-t$ in a simple regression.

If the random part of the unemployment path is a random walk, with $\rho = 1$, the conditional equation is

$$\log u_t - \log u_{t-1} = -\beta + \eta_t, \quad (8)$$

and the estimated recovery rate is the mean decline in log unemployment over the recovery, which can also be expressed as estimated $\beta = (\log u_0 - \log u_T)/T$. Thus, the parameter of interest, the recovery rate β , is identified when unemployment is a random walk, but the constant, α , is not.

The estimates of ρ are generally positive and, in some cases, close to 1. Equation (6) shows that the estimator based on $\rho < 1$ blows up if ρ is really close to 1. We proceed by switching to the $\rho = 1$ estimator of β if our estimate of ρ is above a threshold value, which we take to be 0.9.

Although the key parameter in the model is identified even in the random-walk case, there is some interest in the special problems that arise in that case. In particular, as equation (6) shows, the lack of identification of α rules out separation of that parameter from the initial level of unemployment. And testing the null hypothesis $\rho = 1$ uses the standard t statistic, but its distribution is distinctly more dispersed than the t distribution.

Our specification fits the analysis of Dickey and Fuller (1979) as explained in Hamilton (1994), pages 497-502 (case 4), where there is a constant and a trend on the right side of the equation, as in equation (5). In the range of sample sizes in our work, the critical value of the Dickey-Fuller test statistic at the 0.05 level, is 3.7. The statistic is the ratio of the estimated value of $1 - \rho$ from OLS applied to equation (5) divided by the OLS standard error (the "t"-statistic).

To approximate the standard error of the rate of decline of unemployment, β , we carry out a bootstrap-style simulation. For each of the historical recoveries, 100,000 vectors of values of the innovation η_t are generated by re-sampling the innovations from our estimation, with replacement. Then we calculate the implied vectors of $\log u_t$ from equation (3) and equation (2), using the estimated values of the parameters. Finally, we re-estimate β from the bootstrapped data following the same steps as for the earlier estimation from the actual data. The bootstrap standard error is the standard deviation of the re-estimated values of β .

Our procedure includes the switch to the first-difference estimator when the estimate of ρ exceeds 0.9 in the bootstrap. Absent this precaution, the bootstraps sometimes involve a small fraction of values of ρ almost infinitesimally close to 1, which result in substantial exaggeration of the sampling dispersion of the estimate of β .

The dispersion measures are conditional on the chronologies; that is, they presume knowledge of the turning-point dates when in fact the dates are subject to sampling variation. A sufficiently large realization of η_t can bump the assigned turning point date by a month away from its true value. The bootstrap procedure reveals the contribution of the random innovations η_t to sampling errors in β , but not the contribution of the sampling error in the locations of the turning points in the chronology. The finding of essentially identical results from the HK and DNS chronologies suggests that noise in the dates is not an important source of sampling error in the estimates of β . Even the large differences in cycle dates between the NBER chronology on the one hand and the HK and DNS chronologies on the

	<i>Chronology</i>		
	<i>NBER</i>	<i>Dupraz- Nakamura- Steinsson</i>	<i>Hall- Kudlyak</i>
Annual recovery rate, log points	0.089	0.108	0.107
(Average of bootstrap standard errors)	(0.028)	(0.024)	(0.023)
Standard deviation of recovery rate across all recoveries	0.040	0.033	0.034
Standard deviation of recovery rate omitting the first 3 recoveries	0.025	0.015	0.020

Table 1: Summary of the Estimation Results

other has little effect on the sampling error and a moderate downward bias in the recovery rates estimated from the NBER chronology.

True bootstrap estimation of the standard errors of the estimates of β is not available with time series, because the observations are not independent, but re-sampling estimated innovations is a recognized approximation—see Chernick (2008), page 99.

4 Estimates of the Unemployment Recovery Rate

4.1 Main results

Table 1 summarizes the main statistical results—estimates of the recovery rate and its dispersion across the recoveries, for each of the three chronologies. Here and in the rest of the paper, recoveries are reported at annual rates, 12 times the monthly rates from the estimation.

The table shows the estimates of the key results in this study: the annual recovery rate in log points per, β , its statical standard error, and its standard deviation, describing the variation of the rate across recoveries. The sample period is October 1949 through February 2020 (data are available starting in January 1948 but do not contain any recoveries known to be complete).

The top of Table 1 reports the equally weighted average recovery rate across the ten (in the case of DNS or HK chronologies) or eleven (in the case of the NBER chronology) recoveries,

along with the average of their bootstrap standard errors. Note that it is the average of the standard errors, not the standard error of the average. The recovery-specific rates are not draws from a homogeneous population, so the entries are not the result of pooling across the recoveries. Most of the recoveries are long enough and have sufficiently small variances of the innovations that the sampling dispersion for β is quite small for individual recoveries.

The lower part of the table reports two measures of the variability of the estimated recovery rate across recoveries. The first is the standard deviation across all of the 10 or 11 recoveries. The second is the standard deviation omitting the first three recoveries—those occurring before 1960, when unemployment volatility was generally higher.

For the NBER chronology, the estimated average decline rate across recoveries is 0.089 log points per year. Recovery rates for the DNS and HK chronologies are similar to each other and are above the NBER level, at 0.108 and 0.107. 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.

An example from the recovery rates based on the HK chronology illustrates the interpretation of the annual decline figures in the table. Consider the situation just after a severe recession, with the unemployment rate starting at 10 percent. The expected unemployment rate a year later is $10 \exp(-0.107) = 9.0$ percent. With the recovery rate based on DNS, the rate a year later would be essentially the same as with the HK chronology.

The lower part of Table 1 summarizes findings about the similarity of recovery rates across recoveries, as captured by the standard deviation across the estimated rates. The first line shows the standard deviations across all 10 or 11 estimates. These are in the range of 0.03 to 0.04. The second line omits the first three recoveries and demonstrates that unemployment recoveries have been quite a bit more homogeneous in the years since 1960.

4.2 Estimates by recovery

Figure 3 shows the separate results for the ten recoveries in the HK chronology. Estimated rates for the first three recoveries are more variable than the remaining seven recoveries and have much higher statistical standard errors. The estimates from 1961 to 2020 cluster around 0.10 with small standard errors. Over that 60-year period, with seven recoveries, some mild and two quite severe, the recovery rates are similar.

The standard deviation of the ten annual recovery rates over the entire sample period is 0.034 log points, while the standard deviation for the seven recoveries starting in the 1960s is 0.020 log points.

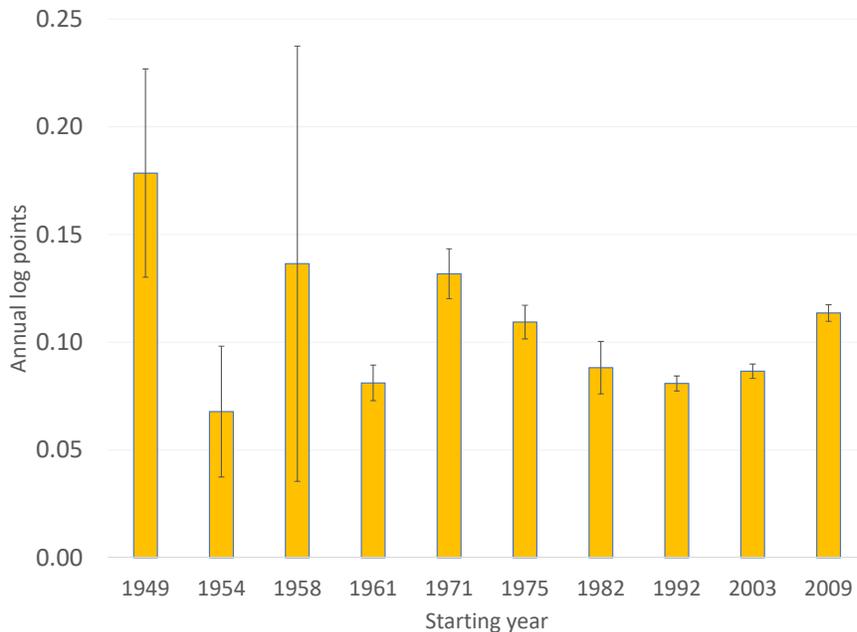


Figure 3: Estimated Recovery Rates by Recovery, HK chronology

These results nail down the primary conclusions of the study—the uniformity of recovery rates over the past 60 years and their remarkably low levels.

The results reveal accurately—as is evident from the small standard errors—that the values of β from the seven later recoveries are similar. Figure 4 and Figure 5 confirm that the two other chronologies generate similar patterns of recovery rates across recoveries.

Note that it is obvious that a standard statistical test would reject the hypothesis that the values of β are uniform across the seven later recoveries. For the HK chronology, the p value for the hypothesis that the 10 values of the recovery rate β are equal is 0.00000095. The hypothesis is rejected with similarly tiny p values for the other two chronologies. The choice not to pool across recoveries is powerfully supported.

Notice that it is simultaneously true that the recovery rates are reasonably similar across recoveries and overwhelmingly true that they are not identical. There is sufficiently little noise in the data to make the test of homogeneity sensitive to small deviations from homogeneity.

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

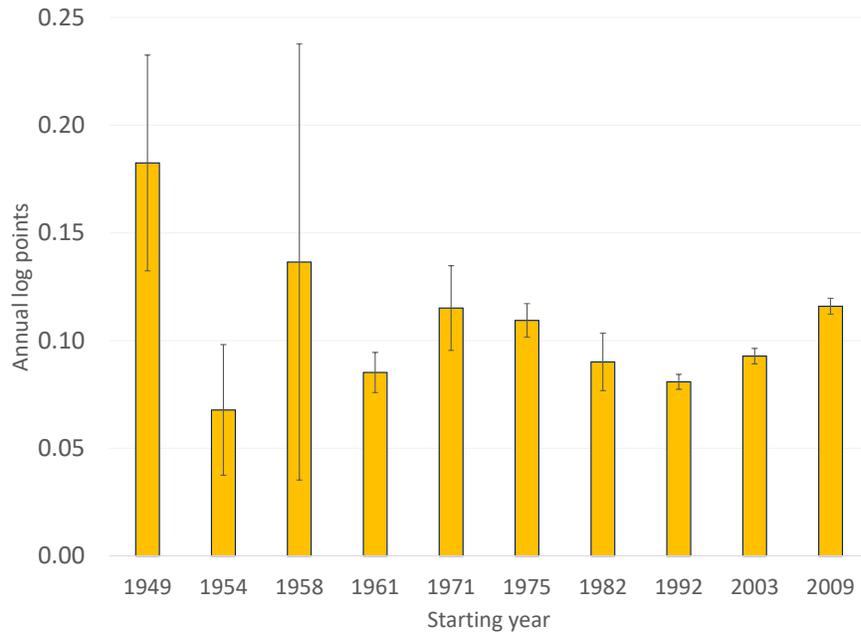


Figure 4: Estimated Recovery Rates by Recovery, DNS chronology

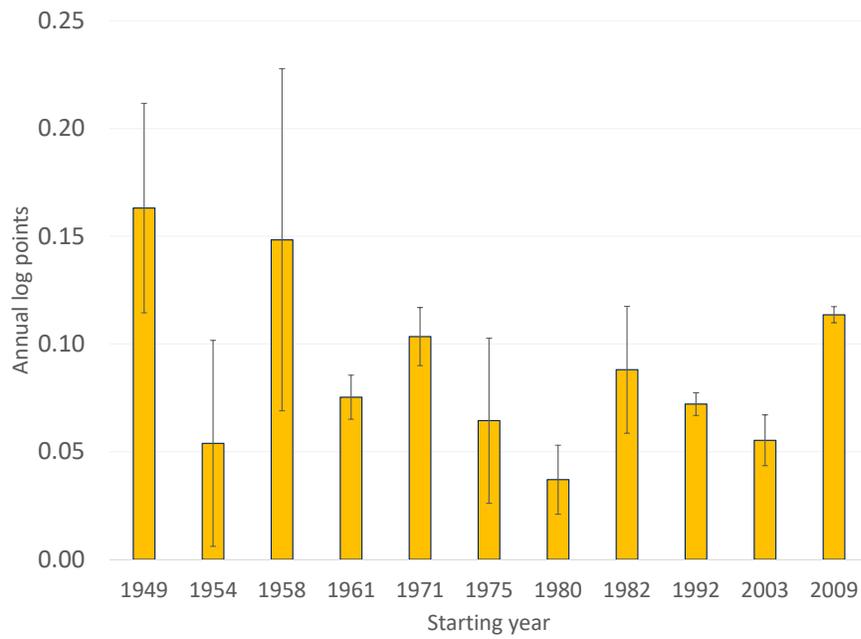


Figure 5: Estimated Recovery Rates by Recovery, NBER chronology

1	Starting month	Oct-49	Sep-54	Jul-58	May-61	Aug-71	May-75	Dec-82	Jun-92	Jun-03	Oct-09
2	Ending month	Jun-53	Mar-57	Feb-60	May-69	Oct-73	May-79	Mar-89	Apr-00	May-07	Feb-20
3	β , annual rate of recovery	0.179	0.068	0.136	0.081	0.132	0.109	0.088	0.081	0.087	0.114
4	Standard error	(0.048)	(0.030)	(0.101)	(0.008)	(0.012)	(0.008)	(0.012)	(0.003)	(0.003)	(0.004)
5	ρ , serial correlation coefficient of random component of recovery path	0.834	0.672	0.813	0.848	0.473	0.664	0.880	0.705	0.369	0.823
6	Standard error	(0.069)	(0.103)	(0.127)	(0.053)	(0.186)	(0.110)	(0.047)	(0.073)	(0.136)	(0.052)
7	Dickey-Fuller statistic	2.41	3.20	1.48	2.85	2.84	3.07	2.53	4.05	4.63	3.43
8	p -value for curvature	0.487	0.890	0.910	0.775	0.261	0.449	0.362	0.414	0.108	0.155
9	Box-Cox λ	-0.349	6.098	2.019	-0.662	1.383	1.145	0.163	0.429	-0.383	0.386
10	p -value for log	0.358	0.000	0.220	0.078	0.443	0.185	0.700	0.415	0.628	0.084

Table 2: Additional Results by Recovery, Based on the HK chronology

output growth was substandard during the recovery but the decline in unemployment was at the normal rate for recoveries.

4.3 Additional results

Table 2 provides additional information about the ten recoveries in the HK chronology. Rows 1 and 2 show the starting and ending months of the recoveries. Rows 3 and 4 give the estimated annual recovery rates and their bootstrap standard errors, as graphed in Figure 3. Rows 5 and 6 show the monthly serial correlation parameters ρ , and their conventional regression standard errors. The parameters are all unambiguously positive and half of them are over 0.8. None of them exceeds the threshold value, 0.9, so the procedure of switching to the random-walk estimator has no bearing on the main estimated β values. It only improves the performance of the bootstrap.

Row 7 in the table deals with testing the hypothesis that log unemployment follows a random walk within a given recovery. For this purpose, we use the Dickey-Fuller statistic, calculated as the usual t -statistic for the hypothesis that a coefficient is one, but with a critical value at the 0.05 level of 3.7 rather than the value for the actual t distribution of about 2.0. The random-walk hypothesis is rejected only for recoveries 8 and 9. However, the test statistic exceeds 3 for three other recoveries, indicating stationarity with considerable but not overwhelming confidence. Recall that stationarity is not an assumption of our approach, so a finding of high serial correlation or a random walk is not evidence against the main conclusions of the study.

Rows 8, 9 and 10 investigate evidence of departures from the log-linear specification. Line 8 reports p -values for the hypothesis that the path is actually log-linear within a setting

where it could be log-quadratic. A small p would signal that the restriction to a straight patch of $\log u$ is unlikely and a curving line is probable. The test adds a curving quadratic term to the model, controlled by a parameter χ , so the model for a recovery becomes

$$\log u_t = \alpha - \beta t - \chi t^2 + \epsilon_t. \quad (9)$$

Most of the p values are moderate and indicate an absence of strong evidence of curvature. None of the ten p -values fall short of 0.1. The evidence gives little support to the idea that there is systematic curvature in the same direction, as four of the estimates of the χ -parameters are positive and the other six are negative.

Rows 9 and 10 of Table 2 use the Box-Cox Box and Cox (1964) approach to evaluate the log specification. In this approach, the model is generalized to encompass linear and other specifications as well as the log. The generalized model is

$$\tilde{u}_t = \alpha - \beta t + \epsilon_t, \quad (10)$$

where \tilde{u} is the Box-Cox transformation of unemployment,

$$\tilde{u} = \frac{u^\lambda - 1}{\lambda}. \quad (11)$$

If $\lambda = 0$, \tilde{u} is $\log u$, while if $\lambda = 1$, \tilde{u} is $u - 1$, linear.

The results make it clear that layering the flexibility of the Box-Cox transformation on top of the flexibility of the baseline specification is asking too much of the data. The excess flexibility results in wild parameter estimates with large standard errors. We believe that the log specification comes out somewhat ahead of the linear one. Only one of the p -values is small enough to meet the conventional standard for an unambiguous value of p in the range of 0.05 or less. We conclude that the evidence confirms the visual impression in Figure 1 that unemployment in recoveries basically follows a log-linear path with small but detectable deviations. That said, none of the substantive conclusions would be affected if the baseline specification were linear rather than log-linear.

5 Alternative Measures of US Unemployment

For the US, the BLS provides compilations of alternative measures of unemployment. This section investigates the behavior of some of these measures in recoveries, using the framework of this paper. We form chronologies using the DNS software and then estimate recovery rates as described earlier.

<i>Unemployment measure</i>	<i>Number of recoveries</i>	<i>Average annual recovery rates with average standard errors</i>	<i>Standard deviation across recovery rates</i>
Standard unemployment for last 3 recoveries	3	0.094 (0.004)	0.017
Standard unemployment plus discouraged workers	3	0.094 (0.003)	0.016
Above plus marginally attached to labor force	3	0.089 (0.003)	0.007
Above plus part time for economic reasons	3	0.084 (0.003)	0.023
Jobless unemployment--standard unemployment minus temporary layoffs	5	0.096 (0.007)	0.016

Table 3: Results for Other Measures of Unemployment

5.1 More inclusive overall measures of unemployment

Some of the alternative measures include more individuals than does the standard unemployment rate. We report results for the three extended unemployment rates, called U-4, U-5, and U-6, over the period of publication, which began in 1994, as part of a comprehensive revision of the CPS. The results are quite similar to those graphed in Figure 3. We believe that this evidence demonstrates that the findings are robust across measures of unemployment and are not artifacts of the specific choices embodied in the standard unemployment rate.

Table 3 shows the results for four measures of unemployment with ascending scope. The table gives the average annual recovery rates in log points and the averages of their standard errors. The first row shows the average results for the standard unemployment rate, for the last three recoveries. They are similar to but more precise than the averages over the full span of 10 recoveries reported earlier in Table 1.

The second row of the table adds people to the standard unemployment count who do not satisfy the requirement of active job search in the four weeks prior to the survey, because they have become discouraged, but who indicate a desire to work and are available to work. The BLS calls this group U-4. The recovery rates and their standard errors remain almost exactly the same.

Adding another group, called marginally attached workers, whose lack of recent job-seeking effort arises from a variety of other causes, results in a slightly smaller recovery rate of the measure called U-5. Another small decrement in the recovery rate applies to group U-6, shown in the fourth line in the table. This group adds those working part-time for

economic reasons. These people are working fewer than their desired hours as a result of adverse economic conditions.

As a general matter, the recovery rates for extended measure of unemployment are quite similar to those classified in the standard way. The results of this paper are robust to inclusion of unemployed individuals with lower job-seeking efforts or other reasons that they are not included in the standard measure.

5.2 Temporary-layoff unemployment

At this writing, the United States is in the last stages of recovering from a major pandemic and resulting sharp recession. The recovery of the US unemployment rate has been speedier than its low historical value. The rate dropped from its maximum of 14.7 percent in April 2020 to 3.6 percent in May 2022. A separate paper (Hall and Kudlyak (2020b)), discusses how a completely unprecedented volume of temporary layoffs accounts for the highly unusual rate of decline of unemployment. A substantial fraction of workers on temporary layoff are recalled to their previous positions—in effect, these individuals are on leave from jobs that they continue to hold. 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.

Starting in 1967, the CPS included questions that identify workers on temporary layoff. In recessions in the 1970s and 1980s, temporary layoffs spiked to just over two percent of the labor force, but subsequently declined to under one percent in recent decades and around 0.5 percentage points during the long recovery after 2009. In related research, we make the case that unemployment analysis should distinguish temporary-layoff unemployment from what we call jobless unemployment—individuals who are searching actively and do not hold existing jobs. Accordingly, we have repeated our measurement of recovery rates using data starting in 1967 that excludes workers classified as temporarily laid off. The results appear at the bottom of Table 3. The recoveries during 1967 through 2019 of the jobless unemployment were at the same rate of about 0.1 log points per year as found here for standard unemployment.

6 Results for Advanced and Emerging Economies

We also perform our basic analysis on the unemployment rates from a considerable number of advanced and emerging countries. Again, we form chronologies using the DNS software and then estimate recovery rates and their standard errors as described earlier.

The Organisation for Economic Co-operation and Development compiles harmonized unemployment data for many countries. Table 4 reports results for these countries, in the

<i>Country</i>	<i>Initial year</i>	<i>Average annual recovery rate</i>	<i>Average of standard errors</i>	<i>Standard deviation of recovery rates</i>	<i>p-value for uniformity</i>
Canada	1955	0.139	(0.015)	0.009	0.000
Chile	1986	0.184	(0.023)	0.016	0.000
Estonia	1983	0.354	(0.066)	0.024	0.000
Belgium	1983	0.125	(0.026)	0.004	0.004
Finland	1991	0.145	(0.014)	0.011	0.000
Netherlands	1990	0.152	(0.012)	0.005	0.000
Austria	1993	0.113	(0.021)	0.001	0.560
Czech Republic	1993	0.165	(0.038)	0.006	0.001
Denmark	1991	0.112	(0.017)	0.003	0.000
France	1987	0.047	(0.006)	0.001	0.007
Greece	2000	0.072	(0.006)	0.002	0.001
Ireland	1991	0.138	(0.009)	0.006	0.000
Italy	1999	0.054	(0.012)	0.001	0.020
Japan	1996	0.074	(0.005)	0.001	0.002
Norway	1999	0.191	(0.033)	0.008	0.000
Poland	1983	0.178	(0.010)	0.003	0.000
Portugal	1999	0.137	(0.020)	0.002	0.071
Slovak Republic	1987	0.109	(0.018)	0.005	0.000
Spain	1989	0.082	(0.010)	0.001	0.248
Sweden	2005	0.116	(0.006)	0.007	0.000
Costa Rica	2011	0.017	(0.011)	0.002	1.000
Germany	1997	0.085	(0.007)	0.001	0.703
Hungary	1988	0.143	(0.015)	0.004	0.423
Korea	1983	0.141	(0.024)	0.010	0.000
Lithuania	2012	0.177	(0.019)	0.005	0.000
Luxemburg	1983	0.101	(0.020)	0.001	0.498
Mexico	1955	0.154	(0.009)	0.011	0.000
Slovenia	1983	0.107	(0.029)	0.005	0.330
Turkey	1997	0.196	(0.014)	0.003	0.081
United Kingdom	1983	0.087	(0.013)	0.002	0.026
Colombia	2007	0.058	(0.003)		
Iceland	1983	0.123	(0.016)	Single recovery	
Israel	1983	0.088	(0.004)		
Latvia	2003	0.129	(0.022)		

Table 4: Results for Other Countries Based on OECD Harmonized Data

framework of this paper. The countries are sorted in the order of the number of recoveries recorded for each. The right-hand column shows the average annual recovery rate and the average annual standard errors of the underlying estimates.

Only Canada has a record almost as long as the US, starting in 1955. Its average recovery rate is somewhat above the finding for the US. It is estimated with good precision. Its recovery rates are substantially more homogeneous than found for the US. The standard deviation of the rates is 0.009 compared to 0.020 and 0.034 for the US.

Note that our general application of the chronology approach, of treating each recovery as a separate object, implies that countries that only recently began data collection can achieve high precision. Among the four countries that have only a single recovery, the distribution of standard errors is similar to the distribution in the countries with multiple recoveries.

The unemployment recovery rates for advanced and emerging economies cluster in the range of the US rates for more recent recoveries, around 0.1 log points per year. Slow but sure is not limited to the US. And the cross-recovery standard deviations share the same low general values as found earlier, confirming the finding of uniformity for these in the US.

We conclude that the findings of this paper are robust to inclusion of a great many other countries.

7 Concluding Remarks

We have developed a parsimonious statistical model of the behavior of unemployment in cyclical recoveries. In economies subject to occasional major negative shocks, it describes an inexorable downward glide at a low but reliable proportional rate during quiescent times. That rate is around 0.1 log points per year. The model describes: first, occasional sharp upward movements in unemployment in times of economic crisis, and second, an inexorable downward glide at a low but reliable proportional rate at all other times. The glide continues until another economic crisis interrupts the glide. The behavior of unemployment has a similar character in alternative measures of US unemployment and in many other advanced and emerging countries.

The companion paper, Hall and Kudlyak (2022), considers explanations of the mechanisms that underlie the movements we document in this paper. We show that the immediate victims of job loss in a crisis tend to have downstream unemployment lasting several years, but not long enough to account for more than a fraction of the persistence documented in this paper. And the evidence shows that the long bulge in unemployment following a crisis involves recruitment of additional victims who did not lose jobs in the crisis itself.

In view of these findings, the companion paper seeks a mechanism that delivers consistent but slow recoveries of unemployment during the last seven decades, in the US and other advanced economies. We argue that such a mechanism generates self-recovery in the labor market. Self-recovery is present in the standard Diamond-Mortensen-Pissarides model of unemployment, but it is faster than in the data. A mechanism whereby a negative feedback from high unemployment to job creation early in the recovery generates reliable but slow recoveries, as in the data. Models of congestion are a leading example of the mechanisms we discuss.

References

- Box, George E.P. and David R. Cox, “An Analysis of Transformations,” *Journal of the Royal Statistical Society: Series B (Methodological)*, 1964, 26 (2), 211–243.
- Chernick, Michael R., *Bootstrap Methods: A Guide for Practitioners and Researchers*, Wiley, 2008.
- 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., Stefano Eusepi, Marc Giannoni, and Aysegul Sahin, “A Unified Approach to Measuring u^* ,” *Brookings Papers on Economic Activity*, 2019, pp. 143–214.
- Dickey, David and Wayne Fuller, “Distribution of the Estimators for Autoregressive Time Series with a Unit Root,” *Journal of the American Statistical Association*, 1979, 74 (366), 427–431.
- Dupraz, Stephane, Emi Nakamura, and Jon Steinsson, “A Plucking Model of Business Cycles,” 2022. University of California, Berkeley.
- 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.
- Fujita, Shigeru and Giuseppe Moscarini, “Recall and Unemployment,” *American Economic Review*, 2017, 107 (12), 3875–3916.
- Gregory, Victoria, Guido Menzio, and David Wiczer, “Pandemic Recession: L or V-Shaped?,” *Quarterly Review*, 2020, (4011). Federal Reserve Bank of Minneapolis, May 27.
- Hall, Robert E. and Marianna Kudlyak, “The Inexorable Recoveries of US Unemployment,” Working Paper 28111, National Bureau of Economic Research 2020.
- ___ and ___, “The 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?,” *NBER Macroeconomics Annual 2021*, 2022, 36, 1–55.

Hamilton, James D., “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle,” *Econometrica*, 1989, 57, 357–384.

—, *Time Series Analysis*, Princeton University Press, 1994.

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.

Perlin, Marcelo, “MS Regress - The MATLAB Package for Markov Regime Switching Models,” 2015. Available at SSRN.

Romer, Christina D. and David H. Romer, “NBER Recession Dates: Strengths, Weaknesses, and a Modern Upgrade,” 2020. University of California, Berkeley.

Figure 1: Log-unemployment declines slowly but surely during recoveries

Figure 2. Three business-cycle chronologies. The solid bars show periods when the economy is recovering and unemployment is on its downward glide path, according to 3 different chronologies. The NBER tracks overall economic activity, while the DNS and HK chronologies consider only unemployment.

Table 1. The 3 chronologies are close to agreeing that the annual rate of decline of unemployment is about about 10 percent per year during recoveries. The standard deviation of the recovery rate across recdoveries is fairly low, meaning that the rates are similar to each other, especially in more recent recoveries.

Figure 3. The height of each bar depicts the measured rate of decline of unemployment within each recovery, using the HK chronology. The error bars show that the measurement procedure has small statistical errors are small starting in the 1960s.

Figure 4. Measured recovery rates in log unemployment are quite similar using the DNS chrnology as for the HK chronology.

Figure 5. The NBER chronology results in somewhat different recovery rates, though the general levels and pattern are similar to those for the other two chronologies.

Table 2. Lines 1 and 2 show the dates of the recoveries. Lines 3 and 4 provided the recovery rates and their statistical measures of uncertainty. Lines 5 and 6 give the serial correlation of the random components of log unemployment. Line 7 shows that the random components could be random walks, though the evidence is far from definitive. Lines 8, 9, and 10 shows that the focus on log unemployment is reasonably consistent with the data, but unemployment itself would work almost as well.

Table 3. Results based on more inclusive measures of unemployment yield quite similar conclusions as those based on the standard measure. Those added are not looking actively for work because they have found it difficult to land jobs, people who have productive alternatives to work in the market, and workers who have been put on part-time because of lack of demand. Removing workers from the data who have jobs but are not working also has little effect on the results.

Table 4. Canada is the only country close to the US in terms of the starting date of unemployment measurement. Canada's unemployment recovery rate is noticeably higher than the US rate and its homogeneity of rates across recoveries is even higher than the US. In general, many of the advanced and emerging countries show the same general behavior, with recovery rates between 0.05 and 0.15 with even more homogeneity of rates across recoveries than in the US.