The Consistent Recovery of US Economy from Every Previous Recession in the Past 70 Years *

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
marianna.kudlyak@sf.frb.org; sites.google.com/site/mariannakudlyak/

August 27, 2020

Abstract

It is a remarkable fact about the historical US business cycle that, after unemployment reached its peak in a recession, and a recovery began, the annual reduction in the unemployment rate was stable at close to one tenth of the current level of unemployment. For example, when the unemployment rate was 7 percent at the beginning of a year, the unemployment rate fell by 0.7 percentage points during the year. The economy seems to have had an irresistible force toward restoring full employment. There was high variation in monetary and fiscal policy, and in productivity and labor-force growth during the recoveries, but little variation in the rate of decline of unemployment. We explore models of the labor market's self-recovery that imply gradual working off of unemployment following a recession shock. These models explain why the recovery of market-wide unemployment is so much slower than the rate at which individual unemployed workers find new jobs. The reasons include the fact that the path that individual job-losers follow back to stable employment often includes several brief interim jobs, sometimes separated by time out of the labor force. We show that the evolution of the labor market involves more than the direct effect of persistent unemployment of job-losers from the recession shock—unemployment during the recovery is elevated for people who did not lose jobs during the recession.

*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 Jon Steinsson for providing the code for their business-cycle chronology function.

We study data from the labor market during recoveries from past recessions, not including the nascent recovery from the pandemic recession of 2020. We find that the recovery phase of the US business cycle has invariably been slow but irresistible. Each year, unemployment falls by one tenth of the level of unemployment—say by 0.7 percentage points in a year with 7 percent unemployment. We document this regularity within the two main statistical approaches to business-cycle analysis and measurement: (1) construction of a chronology of turning points, and (2) estimation of a Markov regime-switching model.

As we write, the United States and the world are in the throes of a major pandemic and resulting deep recession. We believe that an investigation of historical recoveries may have a bearing on understanding the recovery from the current recession. We plan to update our analysis in the coming year to include a comparison of the current recovery to past recoveries.

Following Romer and Romer (2019), we start from the premise that the business cycle involves variations in aggregate slack, and that the best measure of slack is the unemployment rate. We note that a well-documented property of the unemployment rate—most recently confirmed by Dupraz, Nakamura and Steinsson (2019)—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 rock and falls like a feather. In March and April of 2020, unemployment rocketed upward.

We believe our findings are relevant for the recovery that is expected to follow the deep recession triggered by the pandemic, but we do not try to predict how the recovery will evolve. In particular our work should not be interpreted as a forecast that unemployment will decline as slowly in the coming recovery as it has in past recoveries. Our work is complementary to Gregory, Menzio and Wiczer (2020) in this respect.

We focus on recoveries. Our measurement starts in an economy that has just been hit 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. And the current shock is the result of a global pandemic. This paper recognizes that the shocks that propel unemployment sharply upward are heterogeneous. Historical recoveries have been much more homogeneous.

We demonstrate that a steady decline in unemployment following a recession shock occurred reliably in the recoveries in the 70 years we include—years when unemployment has been measured reliably and consistently over time in the Current Population Survey. We point out the puzzle of *slow decline* of unemployment. Cole and Rogerson (1999) called attention to the puzzle—unemployment declines much more slowly than the measured exit rates from unemployment among individuals would seem to indicate.

Our thesis is that the economy has a powerful tendency to recover from serious adverse shocks, but recovery takes time. We show that other potential driving forces do not appear to generate deviations from the downward path of unemployment during recoveries. These forces include monetary and fiscal initiatives and variations in productivity and labor-force growth.

Our results suggest a rather different view of cyclical fluctuations from most current thinking. In our view, a variety of serious shocks can trigger a recession. The differences and similarities between the current pandemic shock and its immediate predecessor, the financial crisis of 2008-9, are illustrative. Millions of workers lose jobs they have held for some time and where they have accumulated job-specific human capital. In the following recovery, rebuilding gradually and inevitably occurs. Until the recovery is complete, employment is depressed because the recovery process involves higher unemployment. This unemployment is effectively a reduction in labor supply. In contrast, the dominant current view emphasizes slow recoveries in labor demand and treats the elevated unemployment as the result of a shortfall of demand relative to supply in the labor market.

We also emphasize a parallel role for the adverse effect of shocks on relations between employers and suppliers of credit. Lending standards spike along with unemployment when a recession strikes and breaks up some previously durable relations. Relations between firms and workers gradually rebuild. There is an interaction between events in the credit market and in the labor market—higher lending standards reduce incentives to hire workers and drive unemployment upward.

Our view of the labor market has points in common with Pries (2004) and can be seen as responding to the challenge of Cole and Rogerson (1999) to explain how aggregate unemployment recovers much more slowly than does an individual spell of unemployment. Like these authors, we build models within the framework of the dominant modern view of unemployment, associated with Diamond, Mortensen, and Pissarides.

We proceed in the following steps:

First, we develop an econometric framework for measuring the speed of recoveries. We demonstrate the a framework that explicitly detects turning points and creates a cyclical chronology yields more reliable estimates of recovery speeds than a framework that treats the chronology as a latent object.

Second, we study the job loss that occurs when a crisis launches a recession. A spike in job loss is visible in a variety of data sources, measuring layoffs, job destruction, displacement, and unemployment insurance claims. But the spike in job loss in these measures is far smaller than the total increase in the number of unemployed workers.

Third, we ask whether the volume of job losers and their likely speed of finding long-term replacement jobs is enough to explain the long-lasting bulge of total unemployment that is only gradually worked off during even a long recovery like the one that ended in early 2020. We conclude, from data on displaced workers collected every two years in the Current Population Survey, that the number of workers displaced even in the severe recession starting in 2007 was not enough to explain the volume of excess unemployment present in the US economy during the period from 2009 through 2014. Something happened in the labor market during that period that caused elevated unemployment among workers who were not displaced around 2009.

Fourth, we examine the puzzle of low recovery speed in the framework of the DMP model. We calculate the effective exit rate from unemployment, which is lower than the exit rate for individuals from one month to the next. Those individual exits are frequently temporary departures from the labor force or short-term jobs, and are then followed by additional spells of unemployment, as described in Hall and Kudlyak (2019). In the DMP equilibrium, the unemployment rate falls more gradually, so it accounts for some of the puzzle of low recovery speed.

We build another model that explains slow but sure recoveries with a congestion externality. When unemployment spikes, employers' costs of recruiting rise. According to standard DMP principles, higher costs of filling vacancies discourage job creation and raises equilibrium unemployment. This model generates a generally slow decline of unemployment during a recovery, but appears to fall short of offering a full explanation of the slow but reliable process of recovery in the US labor market.

We also develop a model of the disruption of credit markets that occurs in financial crises. The operation of markets for publicly traded securities is impeded during and after a crisis, but these markets tend to return to normal faster than does the labor market. We develop evidence that traditional bank lending based on continuing bilateral relations suffers disruptions that heal at the same low but steady speed as do similar relations in the labor market. The disruptions put a wedge between the cost of capital in open markets and the effective cost to credit intermediated through banks. The wedge is visible in tougher lending standards for an extended period following a crisis. The wedge erodes the incentive to create jobs and raises the unemployment rate through the principles of the DMP model. This model comes closer to matching the observed path of unemployment.

Fifth and finally, we show that recoveries have had a wide variety of accompanying movements of policy instruments and arguably exogenous aggregate driving forces. Unemployment declines smoothly despite irregular movements of possible driving forces, including

fiscal and monetary policy, productivity growth, labor-force growth, and variations in stock-market discounts.

We conclude that the economy includes a strong internal force toward recovery that operates apart from policy instruments and apart from productivity growth and financial developments revealed in the stock market. We make the case that the internal force originates in the labor market. After a negative shock, employers gradually find it profitable to hire more aggressively. Unemployment falls as the unemployed are put back to work. Rather than a pull from expansionary policy, the growth in employment arises from a push toward lower unemployment.

1 Recovery Following a Recession

1.1 Measuring recoveries

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 slack or 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 filter or from a bandpass filter. We agree with the Romers that tying the business cycle to slack is conceptually superior to tying it to higher-frequency movements. In any case, this paper concentrates on the behavior of unemployment in recoveries, and should be of interest to investigators who believe that unemployment is only one of several important business-cycle indicators.

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. It appears to contain almost no movements associated with productivity or similar forces that would call for filtering out. Fortunately, the movements of the CBO's measure and the unemployment measure are quite similar. Figure 1 compares the CBO's real GDP ratio to a measure based on the labor market, constructed as 100 minus the unemployment rate.

The magnitudes of cyclical fluctuations in the GDP ratio is greater than the magnitude for the unemployment-based measure, especially in the years before 1990. The chronology of peaks and troughs is similar for the two measures. We will use unemployment in what follows, but the use of the CBO's GDP ratio measure would give much the same results.

Movements of some cyclical variables, notably real GDP, contain components such as productivity growth that need to be filtered out. The unemployment series from the Current

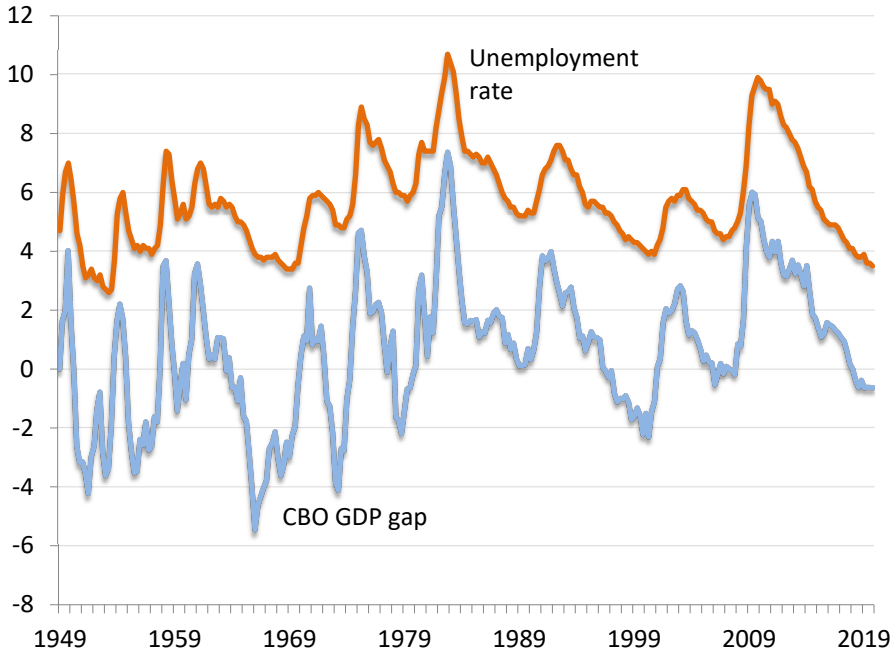


Figure 1: Comparison of Cycle Measure Based on the GDP Ratio to a Measure Based on the Unemployment Rate

Population Survey appears to contain almost no movements associated with productivity or similar forces that would call for filtering out. On the other hand, real GDP has non-cyclical movements associated with the movements of productivity, labor-force size, and other forces not usually associated with the business cycle. The concept of potential real GDP is aimed at accomplishing that filtering. The ratio of actual to potential real GDP filters out all influences captured by potential GDP. The Congressional Budget Office produces a widely used and respected measure of potential real GDP.

2 Econometrics of Business-Cycle Measurement

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 the unemployment rate during recoveries. We consider the general class of models

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

where $f(\cdot)$ is a monotonic transformation, x_t is the systematic 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 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 0.10 log points per year, which is 0.7 percentage points if the unemployment rate is 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 by Dupraz et al. (2019).

The other class of models focuses on *regime switching*, where the systematic component is modeled as a statistical time series that obeys one model in contractions and another in recessions. A number of authors, notably James Hamilton, have created Markov-switching models in this class.

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 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. Chronologies appear to overstate the precision of the information actually available. We investigate this issue in the context of the algorithm of Dupraz et al. (2019) (DNS).

2.1 Dupraz, Nakamura, and Steinsson's chronology algorithm

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. 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 hundreds of thousands of experimental paths of unemployment in this endeavor.

2.2 Measuring the effect of random noise in obscuring an underlying chronology

Here we consider a basic question about extracting the underlying chronology from noisy data. In the model $\log u_t = x_t + \epsilon_t$, we suppose that the DNS algorithm would extract the chronology perfectly from x_t, absent the noise ϵ_t. We ask how much distortion of the chronology for x_t results from applying the algorithm to the noisy data. Some blurring of the chronology is bound to occur from the noise. For example, if there is actually a trough in December 2007, but a substantial amount of noise happens to occur in that month, the algorithm may find a false trough in November 2007 or January 2008.

For the purpose of investigating this question, we assume that x_t is a linear spline, composed of linear segments connecting the turning points of the unemployment series.

We create a reference chronology over the 866 months from January 1948 to February 2020. We intend that it capture the likely properties of unemployment chronologies, so that the results of our experiments are relevant to the actual history of unemployment. We create values $\hat{x}_t$ as the straight line connecting the observed values of $\log u_t$ at the turning points. Then we construct residuals $\hat{\epsilon}_t = \log u_t - \hat{x}_t$ to serve as the reference values of the noise ϵ_t. We fit an $AR(1)$ process to $\hat{\epsilon}_t$ to provide a reference value of the serial correlation ρ of the noise and a reference value of the standard deviation σ of the innovation in $\hat{\epsilon}_t$.

We are concerned primarily with estimating the average recovery rate of unemployment and the standard deviation of the average recovery rate across recoveries. We define the recovery rate of unemployment, $\mathbb{R}_r$, as the negative of the slope of the straight line segment connecting a peak in unemployment to the next trough; r is the recovery number in the chronology. We let $\mathbb{R}$ be the average of $\mathbb{R}_r$ across all recoveries, weighted by the number of months. As with the other statistics, the absence of a hat refers to the reference chronology and the presence of a hat refers to the chronology inferred in the presence of noise. Similarly, we define $\mathbb{S}$ be the weighted average of $\mathbb{S}_r$, the standard deviation of $\mathbb{R}_r$ across all recoveries.

The plan is to calculate the values of these statistics under accurate information about the systematic component x, which is a non-stochastic function of the chronology. The result is the non-stochastic pair of values $\mathbb{R}$ and $\mathbb{S}$, based on the reference chronology. We

	Reference	*With noise*
Recovery rate, log points per year	0.123	0.140
		(0.0001)
Standard deviation of recovery rate	0.071	0.078
		(0.0002)

Table 1: How Noise in Unemployment Biases Recovery Rates

then construct 10,000 time-series paths for unemployment as the sum of the systematic component and a randomly drawn normally distributed AR(1) noise series with the reference autocovariance and dispersion. We run each resulting series through the DNS chronology detector, and measure the average recovery rate statistics as above. We calculate the means of the statistics $\hat{\mathbb{R}}$ and $\hat{\mathbb{S}}$, and compare them to the underlying values in the reference chronology. We compute the standard errors of those statistics to demonstrate that errors arising from the simulation itself are completely negligible.

We interpret the deviation of the central tendency of $\hat{\mathbb{R}}$ from the known value $\mathbb{R}$ as the bias caused by the noise in the data. And we interpret the central tendency of $\mathbb{S}$ as the bias in the dispersion of the recovery rate across business cycles caused by the noise.

Table 1 shows the results of these calculations. The left column reports statistics for the reference case, taken here to be the true values. The right column reports the means of the same statistics across the 10,000 simulated values. It also reports the standard errors—the sampling error associated with the fact that the simulation is a sample and not an exact value. The standard errors are trivially small. The discrepancies between the left and right column reflect the noise introduced in each simulated value in the form of the disturbance ϵ.

The upper line of the table shows that the effect of the noise is to slightly overstate the recovery rate, which is actually 0.123 log points per year but is 0.149 on average in the sample with noise. The difference, 0.027, Is a number that could be subtracted to adjust for the upward bias in the recovery rate, provided that the conditions presumed in our experiment are similar to actual conditions.

Similarly, the lower line of the table shows that the effect of the noise is to slightly overstate the dispersion of the recovery rate across cycles. The true value of the standard deviation is 0.071, but the effect of the noise is to boost it to 0.078. Again, the difference, 0.007, could be subtracted from the observed value of the standard deviation.

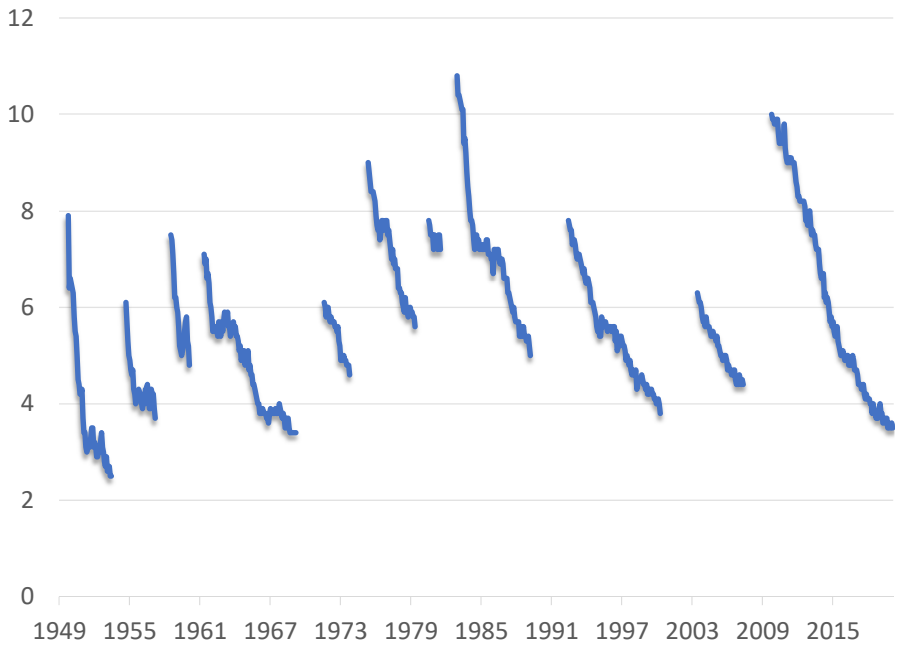


Figure 2: The Paths of Unemployment During Recoveries

We conclude that the blurring effect of the added noise creates errors in the chronology that bias measured recovery rates and cross-cycle standard deviation upward, but the biases are not large.

3 Uniform Unemployment Recovery across Recessions

In this section we show that the average rate at which unemployment recovers is similar across recoveries, regardless of the cause of the prior recessions.

Based on our finding that results based on the DNS chronology are very similar to results for the true chronology, we study the recovery rate in an econometric setup where we use the DNS recovery dates as if they were the true dates.

Figure 2 shows the unemployment rate during the recoveries in the DNS cycle chronology, with the recession spells left blank. The key fact about recoveries is apparent in the data: unemployment declines smoothly but slowly throughout most recoveries most of the time, at about the same rate.

Table 2 shows the results of estimation of the specification

$$\log u_t = \alpha_r - \beta_r(t - t_r) + \epsilon_t \tag{2}$$

Recovery		*α, intercept*	*β, rate of decline*
First month	*Last month*		
1949:10	1953:5	1.813	0.267
		(0.041)	(0.019)
1954:9	1957:3	1.621	0.112
		(0.029)	(0.019)
1958:7	1960:2	1.949	0.222
		(0.037)	(0.039)
1961:5	1968:9	1.887	0.090
		(0.012)	(0.003)
1970:12	1973:10	1.846	0.093
		(0.011)	(0.007)
1975:5	1979:5	2.181	0.112
		(0.009)	(0.004)
1982:11	1989:3	2.278	0.100
		(0.014)	(0.004)
1992:6	2000:4	2.010	0.082
		(0.007)	(0.001)
2003:6	2006:10	1.821	0.093
		(0.005)	(0.003)
2009:10	2019:9	2.377	0.115
		(0.007)	(0.001)

Table 2: Recovery Rates by Recovery

separately over the observations for recovery r. Here α_r is the initial fitted value for $\log u$, t_r is the value of t at the beginning of the recovery, and β_r is the rate of decline in log units during the recovery.

With only two exceptions, the recovery rates β_r cluster very close to 0.1. The exceptions are shorter recoveries that occurred in the 1940s and 1950s. Uniformity set in with the recession that began in 1961.

Why is there a widespread impression that the recovery from the 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 normal.

4 The Hidden Markov Approach

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

$$\log u_t = x_t + \epsilon_t, \tag{3}$$

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

$$\Delta \log u_t = \beta_i + \eta_t. \tag{4}$$

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. We have studied the success of the maximum-likelihood estimator, using the same simulation setup as in our work described above for the use of the DNS chronology extractor.

We confirm that the estimator is consistent—with a long enough time series with enough business cycles, it converges to the value assumed in constructing the data. But we find that the estimator has high dispersion when applied to the 866 observations in our sample. Several thousand months of data are needed to drive the standard error down to a usable level.

5 Recessions: Substantial but Short-Lived Spikes in Job Loss

In this section we document the substantial but short-lived spikes of job loss in recessions. We do not currently include the huge volume of job loss that occurred from the pandemic in the spring of 2020. We consider four measures of job loss:

- Layoffs, the flow of workers whose jobs ended at the initiative of employers
- Job destruction, the amount of employment decline among establishments with shrinking employment.
- Worker displacement, job loss among workers with at least three years of tenure at the lost job.
- Unemployment insurance claims

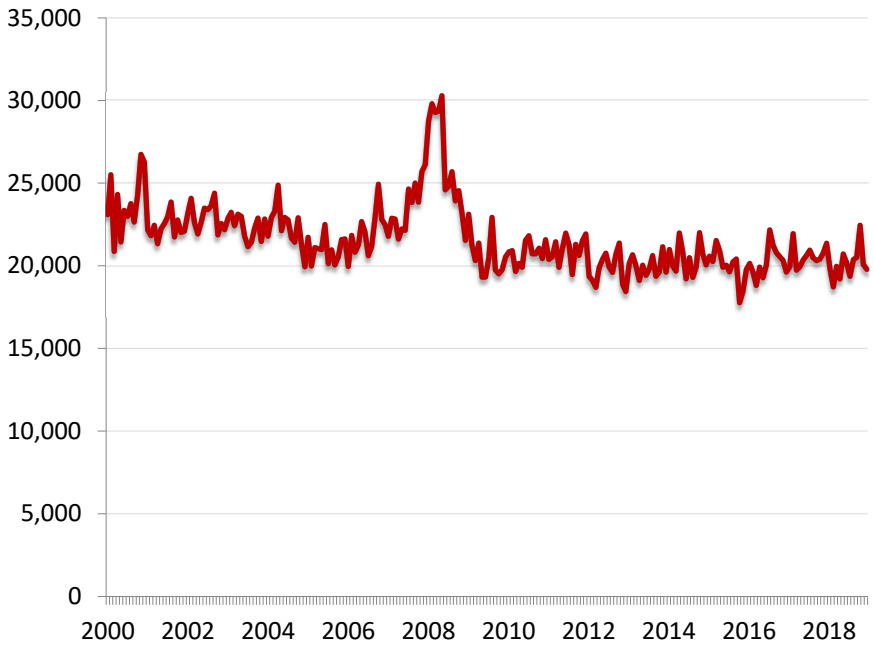


Figure 3: Layoffs Recorded in JOLTS, Monthly at Annual Rate, in Thousands of Workers

5.1 Layoffs

5.1.1 Layoffs in JOLTS

Figure 3 shows data on layoffs from the Job Openings and Labor Turnover Survey. A layoff occurs when an employer terminates a worker without prejudice, typically because continuing employment has become unprofitable. Most layoffs occur without any definite promise to rehire, but explicitly temporary layoffs are an important minority of layoffs. On average, 20 million workers lose their jobs each year in normal times. A substantial but short-lived burst of above-normal layoffs occurred soon after the financial crisis in the fall of 2008.

5.1.2 Mass layoffs

Mass layoff occurs when a relatively large number of firm's employees lose jobs. Such events often involve high-tenured workers who tend to suffer prolonged periods of joblessness following job loss (Jacobson, LaLonde and Sullivan (1993), Davis and von Wachter (2011)).

The Mass Layoffs Statistics program from the BLS tracks the effects of major job cutbacks using data from the states' unemployment insurance databases. A mass layoff is defined by the occurrence of fifty or more initial claims for unemployment insurance benefits being filed against an employer during a 5-week period. These employers are contacted by the state agency to determine whether the separations last more than 30 days. Such events are termed extended mass layoffs. The BLS obtains information on the total number of

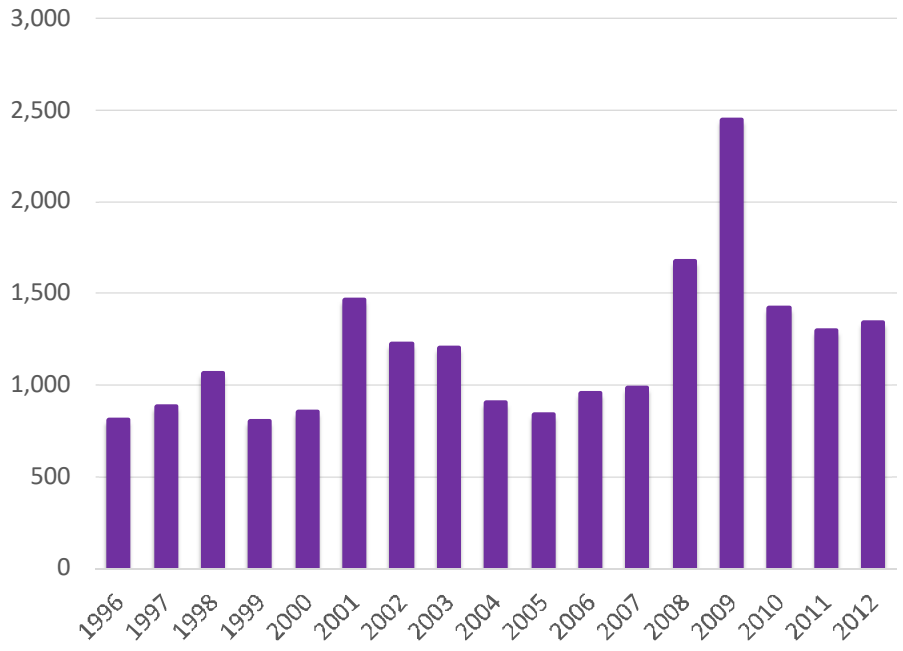
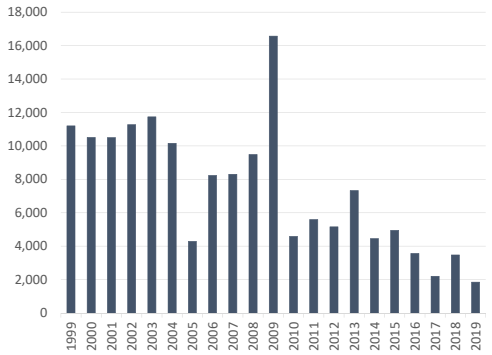


Figure 4: Extended Mass Layoffs, in Thousands of Initial Claimants per Year

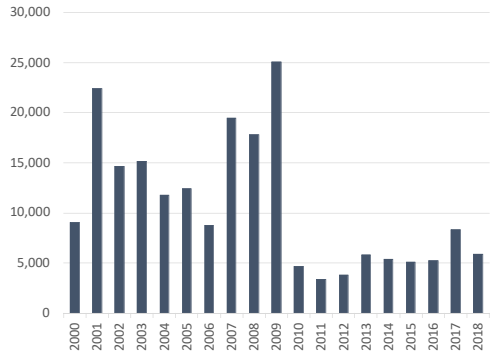
workers separated during the extended mass layoffs, including the workers who do not file for unemployment insurance, and the reasons for these separations according to the employer. These layoffs involve both people subject to recall and those who are terminated. The program ran from 1995 to the first quarter of 2013.

Figure 4 shows the number of initial claimants from extended mass layoffs. The number hovers around a million in normal times but spikes during recessions. A decline in business demand and financial difficulties are the main reasons cited behind the spikes. In 2009, extended mass layoffs spiked to 2.4 million.

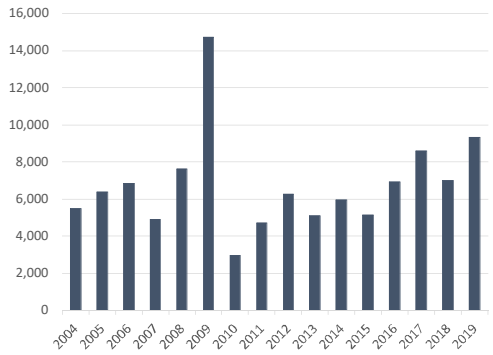
Another source of data on mass layoffs arises from the Worker Adjustment and Retraining Notification Act (WARN), which requires employers to provide notice 60 days in advance of covered plant closings, covered mass layoffs, or sale of business that result in an employment loss. Employers are covered by WARN if they have 100 or more employees, not counting employees who have worked less than 6 months in the last 12 months and not counting employees who work an average of less than 20 hours a week. The term *employment loss* means (1) an employment termination, other than a discharge for cause, voluntary departure, or retirement; (2) a layoff exceeding 6 months; or (3) a reduction in an employee's hours of work of more than 50 percent in each month of any 6-month period. A plant closing occurs if an employment site will be shut down, and the shutdown will result in an employment loss for 50 or more employees during any 30-day period. A mass layoff occurs without a plant closing if the layoff results in an employment loss at the employment site during any 30-day



(a) Alabama



(b) Michigan

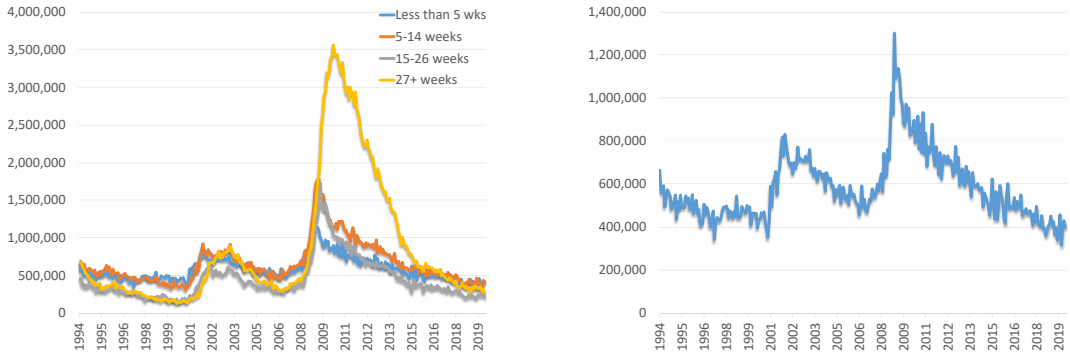


(c) Washington

Figure 5: Mass Layoffs, by State

period for 500 or more employees, or for 50-499 employees if they make up at least 33 percent of the employer's active workforce. Under certain circumstances, smaller employment losses also trigger notification requirements.

The WARN data over an extended period of time are publicly available for many states, but not all. Figure 5 shows the number of layoffs for Alabama, Michigan, and Washington, as examples. The data show clear spikes in layoffs in 2009. For Alabama and Washington the figure shows layoffs sorted by the effective date. For Michigan, we have information about the date of the WARN notice but not about the effective date of the layoff, so we sort the layoffs by the expected effective date, which is the notice date plus two months.



(a) All durations

(b) Less than 5 weeks

Figure 6: Unemployment due to Permanent Job Loss, by Duration

5.1.3 Permanent layoffs by duration

Figure 6 shows unemployment involving permanent job loss by duration, from the CPS. Layoffs with duration of 5 weeks or less is a flow of new layoffs. The flow spikes at the onset of recessions and declines only slowly afterwards.

5.2 Job destruction

The Business Dynamics Statistics data report job destruction. This measure is defined as the sum of all establishment-level reductions in employment. Davis and Haltiwanger (1992) and Davis, Faberman and Haltiwanger (2013) proposed job destruction as a measure of separations and validated the definition through study of the microdata from JOLTS. Although an employer could accomplish a reduction in employment by cutting back hiring and relying on normal attrition, in fact, almost all employment reductions take the form of separations. When an adverse shock hits the economy, separations jump even though quits fall. Layoffs account for more than all of the observed reduction in employment. Figure 7 shows data from the BDS on job destruction. It shows a considerable bulge of job destruction immediately after the financial crisis.

5.3 Displaced workers

Displaced workers are defined as those 20 years old and over who have worked for their employers for 3 or more years at the time of displacement, who lost or left jobs because their plants or companies closed or moved, because there was insufficient work for them to do, or

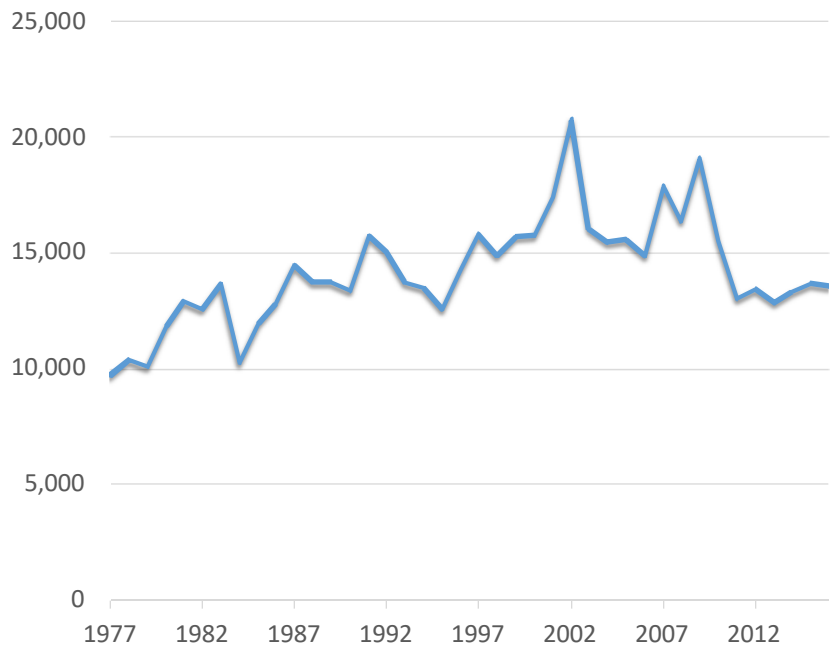


Figure 7: Job Destruction, in Thousands of Workers per Year

because their positions or shifts were abolished. These are job losses among workers with substantial tenure, in contrast to layoffs measured in JOLTS.

Table 3 shows the findings of the displaced workers supplement to the CPS taken in January of even-numbered years from 2002 through 2018. The survey inquires about current unemployment and displacement in the year ended the month before, and one and two years earlier. As far as we know, this is the only panel-type data reaching this far back between current unemployment and earlier displacement, and contains enough respondents who are unemployed for reasonable analysis. Administrative data rarely report unemployment and panel surveys do not contain many workers who are unemployed.

The design of the displaced workers supplement to the CPS poses an interesting challenge to inference about the time path of long-term displacements and the path of unemployment following displacement. Figure 8 shows an initial attempt. The annual estimates satisfy the overlapping three-year sums and are informed by the timing of layoffs and job destruction within each three-year span. The figure also shows a counterfactual path of displacements, which eliminates the two recession spikes present in the actual data.

Survey in January of	*Displacement occurring in calendar years*			*Number of displaced workers*	*Unemployed at time of survey*
2002	1999	2000	2001	3,969	841
2004	2001	2002	2003	5,329	1,076
2006	2003	2004	2005	3,815	511
2008	2005	2006	2007	3,641	655
2010	2007	2008	2009	6,938	2,505
2012	2009	2010	2011	6,121	1,634
2014	2011	2012	2013	4,292	893
2016	2013	2014	2015	3,191	507
2018	2015	2016	2017	2,981	429

Table 3: Total Long-Tenured Displaced Workers and the Number of Unemployed at the Time of the Survey, in Thousands

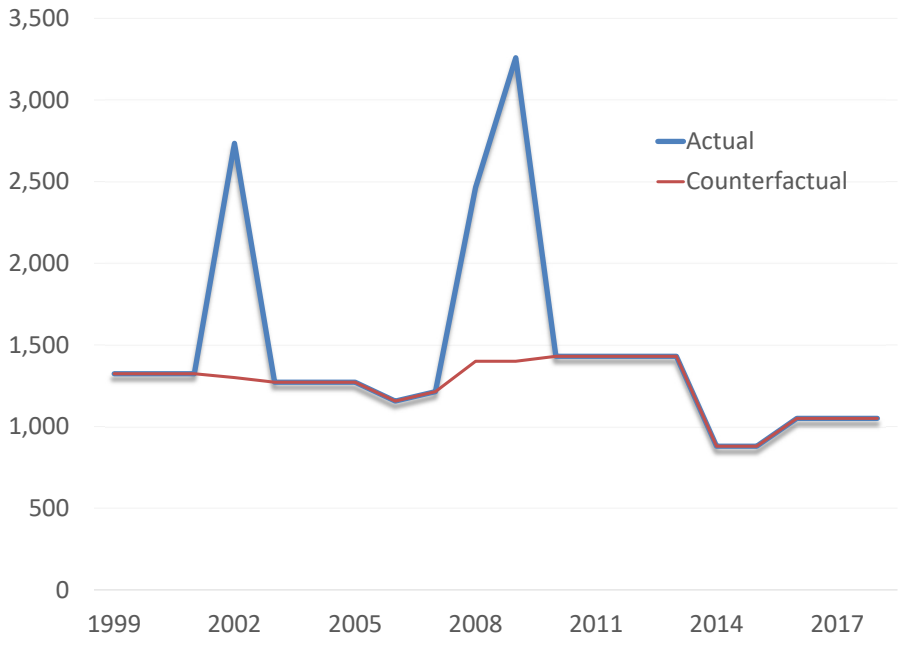


Figure 8: Estimated Annual Displacements, Actual and Counterfactual

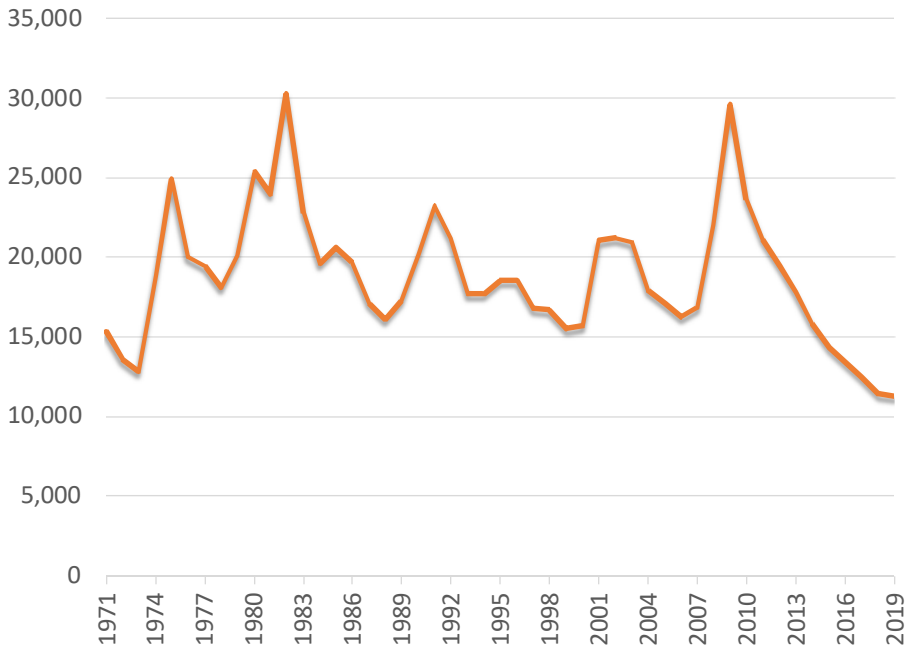


Figure 9: Initial Unemployment Insurance Claims, in Thousands

5.4 Initial unemployment insurance claims

Figure 9 shows the initial unemployment insurance claims. In contrast to layoffs (Figure 3) but similarly to unemployment, during recessions the initial UI claims go up like a rock and go down as a feather.

Why is there a discrepancy between the number of layoffs and the initial UI claims? One factor is that not all eligible unemployed claim the benefits. Building on Blank and Card (1991), Auray, Fuller and Lkhagvasuren (2019) find that from 1989 through 2012, the take-up rate averaged 77 percent. Research shows that the number of those who are eligible but do not claim benefits increases in recession and declines in recoveries (see Fuller, Ravikumar and Zhang (2012) and Auray et al. (2019)). Thus, fluctuations in take-up rates goes in the wrong direction as an explanation of the discrepancy between layoffs and the initial UI claims.

5.5 Comparison of measures of the spike of job loss in a recession

Figure 10 compares the estimated displacement counts to the tabulations of layoffs, job destruction, and extended mass layoffs. Although the normal level of displacement is far below the levels of layoffs or job destruction, the increase in displacements at the outset of the two recessions is an important fraction of the increases for layoffs and job destruction.

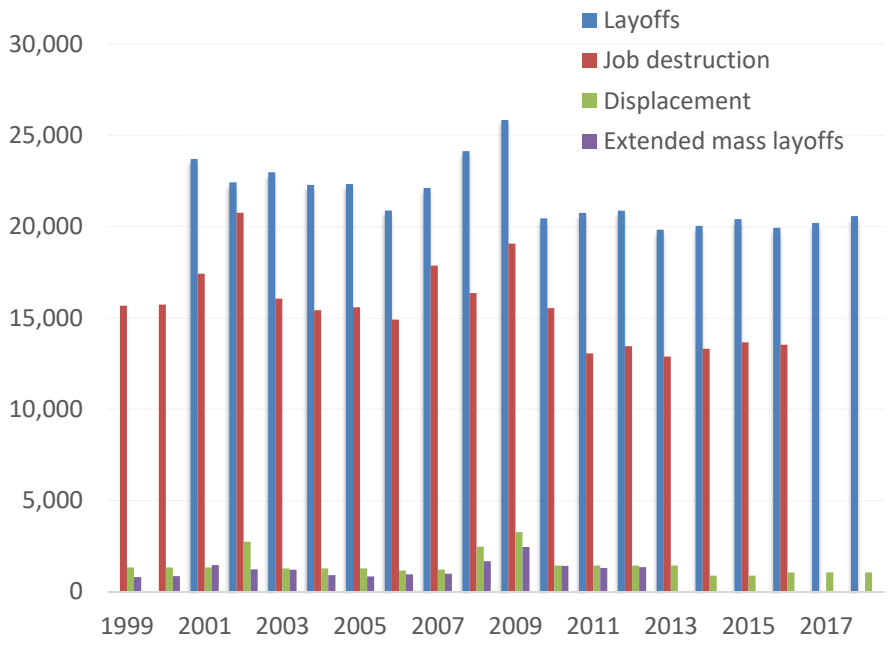


Figure 10: Annual Levels of Layoffs, Job Destruction, Displacements, and Extended Mass Layoffs, in Thousands

Figure 11 shows excess job loss associated with the 2009 recession by four measures of job loss, together with excess unemployment. All four job loss measures show a substantial but short-lived spike. Unemployment shows a substantial increase and slow return to its pre-recession level.

6 The Direct Channel from Job Loss to Subsequent Lingering Unemployment

We consider the hypothesis that excess job loss directly accounts for the spike and the subsequent long slow decline of excess unemployment. We call it the *direct-channel* hypothesis. According to this hypothesis, the extra individuals who become unemployed because of the recession shock follow a path similar to those found in research such as Jacobson et al. (1993) and Davis and von Wachter (2011) that tracks the post-displacement paths of workers who lose their jobs from layoffs. These paths often include multiple spells of unemployment.

6.1 Information about the subsequent role in unemployment from job displacement

The CPS survey supplement measuring job displacement is the only one with the crucial information about lingering unemployment among job-losers in the years following job loss.

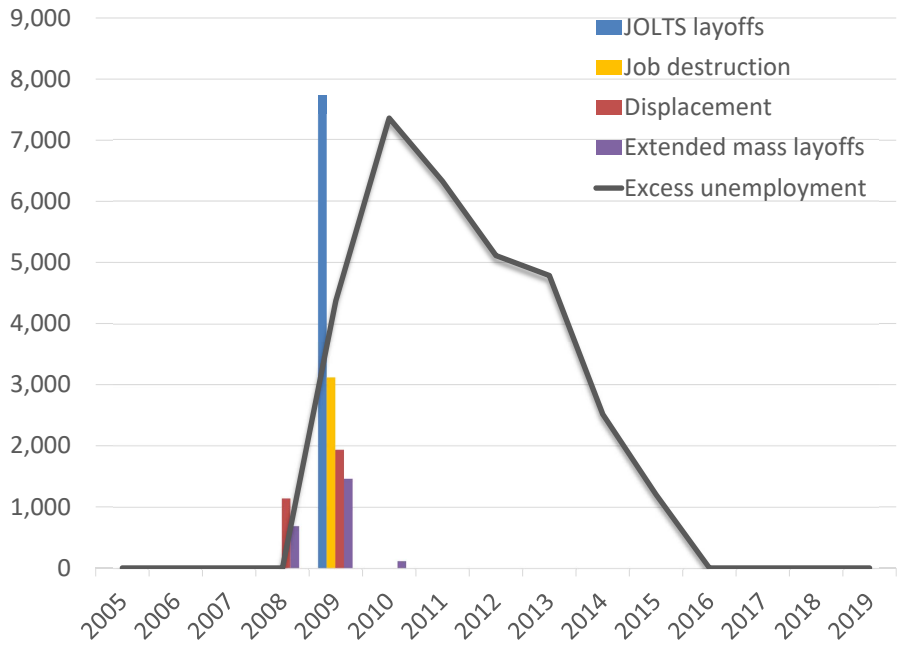


Figure 11: Excess Job Loss in 2009 and Excess Unemployment, in Thousands of Workers

We fit a simple time-series regression with the biennial data for unemployment in January of even-numbered years of workers suffering displacements in the previous three years as the left-hand variable and three lagged values of the estimated displacement counts as right-hand variables, along with a constant. The relation takes the form

$$u_t = f_1(D_{t-1}) + f_2(D_{t-2}) + f_3(D_{t-3}) \tag{5}$$

We linearize as

$$u_t \doteq \alpha + \beta_1 D_{t-1} + \beta_2 D_{t-2} + \beta_3 D_{t-3} \tag{6}$$

If market tightness were constant over time, β_1 would be the unemployment rate among workers who suffered displacement within the past year, β_2 one to two years ago, and β_3 two to three years ago. The design of the survey prevents learning about unemployment among people displaced more than 3 years ago. However, job-finding rates are lower in the same years that displacements are high, so $f(D)$ is a convex function of D. This property implies that the intercept α should be negative and the coefficients should be greater than the unemployment rates.

Table 4 shows the regression results. The good fit suggests that the imputation of annual timing for the displacements is reasonably successful. The fact that the 3rd-year coefficient is somewhat larger than the 2nd-year one is within sampling variation, but may also reflect the fact that a worker with displacement 3 years ago also suffered an earlier displacement as well. In addition, there may be a stronger convexity effect for the 3rd-year displacement.

Parameter	*Interpretation*	*Coefficient*	*Standard error*
α	Intercept	-991	(144)
β_1	Effect of prior year's displacements	0.76	(0.07)
β_2	Effect of displacement 1 to 2 years ago	0.21	(0.08)
β_3	Effect of displacement 2 to 3 years ago	0.37	(0.06)
R^2	0.991		
σ	116		

Table 4: Regression Results for the Relation between Lagged Displacements and Current Unemployment of Workers Suffering those Displacements

The negative intercept confirms the expectation of a convex relation between displacements and later unemployment.

Our first use of the regression is to impute unemployment of workers suffering displacements in the previous three years in January of the odd-numbered years when the supplement to the CPS was not performed. Figure 12 shows the fitted values from the regression for the years 2002 through 2018, in red, along with the actual unemployment counts for the even-numbered years when the supplement to the CPS occurs, in blue.

Our second use of the regression results is to calculate how much lower displacement-related unemployment would have been absent the spikes of displacement in the two recessions. Figure 13 shows the results of this counterfactual and compares displacement-related unemployment from the two recessions to overall unemployment in January in the years since 2001. The rise after the recession is material relative to the overall increase in unemployment following the recession that began in 2001, but is a small part of the large increase in unemployment following the financial crisis.

6.2 Application to other measures of job loss

We use the estimates from Table 4 to calculate excess unemployment from excess job loss by the four measures shown in Figure 11. Figure 14 shows the unemployment resulting from excess job loss in 2009 and total excess unemployment.

Figure 15 shows the contribution of unemployment from excess job loss in 2009 to the cumulative excess unemployment during the 2007-09 recession.

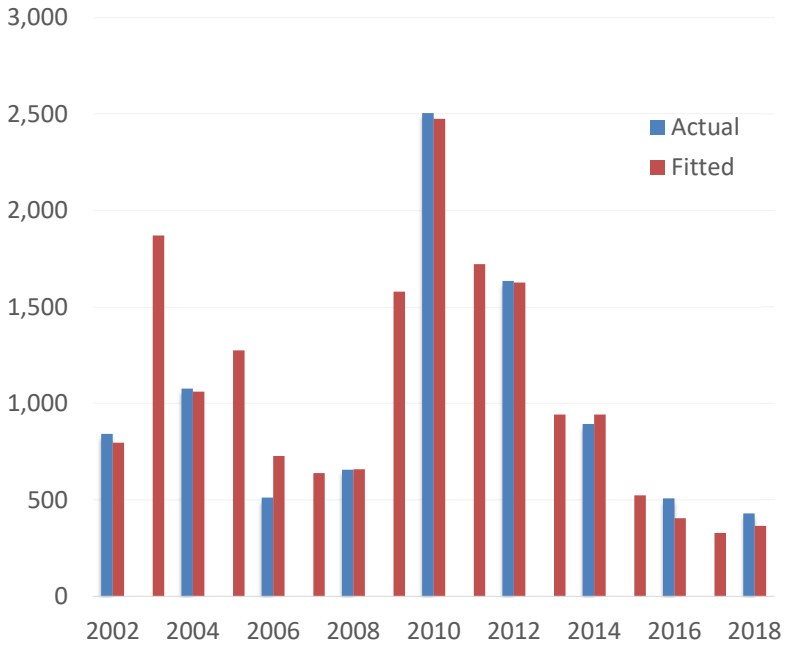


Figure 12: Actual and Fitted Unemployment Counts of Workers Suffering Displacements in the Previous Three Years, in Thousands

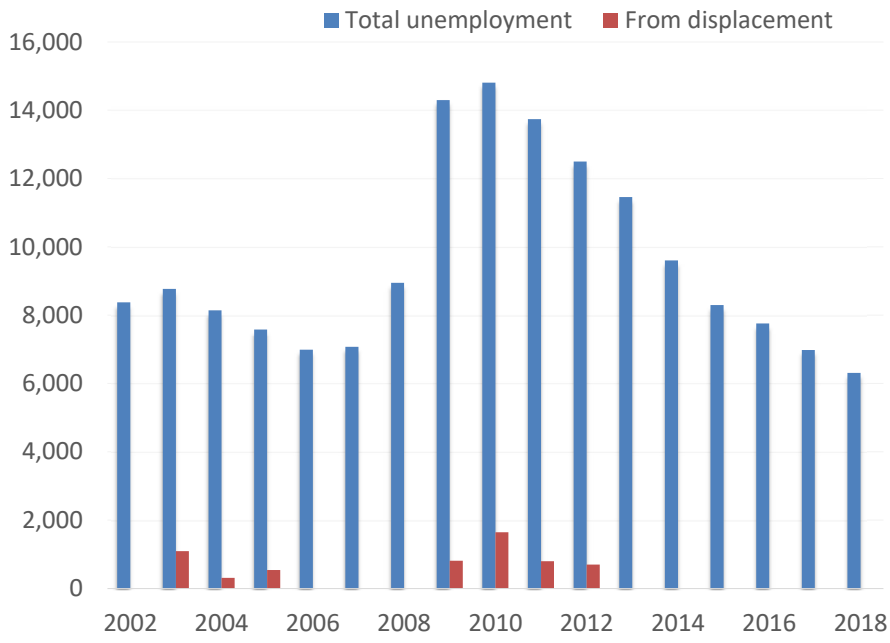


Figure 13: Total Unemployment and the Component Related to Bursts of Displacement Associated with Two Recessions, in Thousands

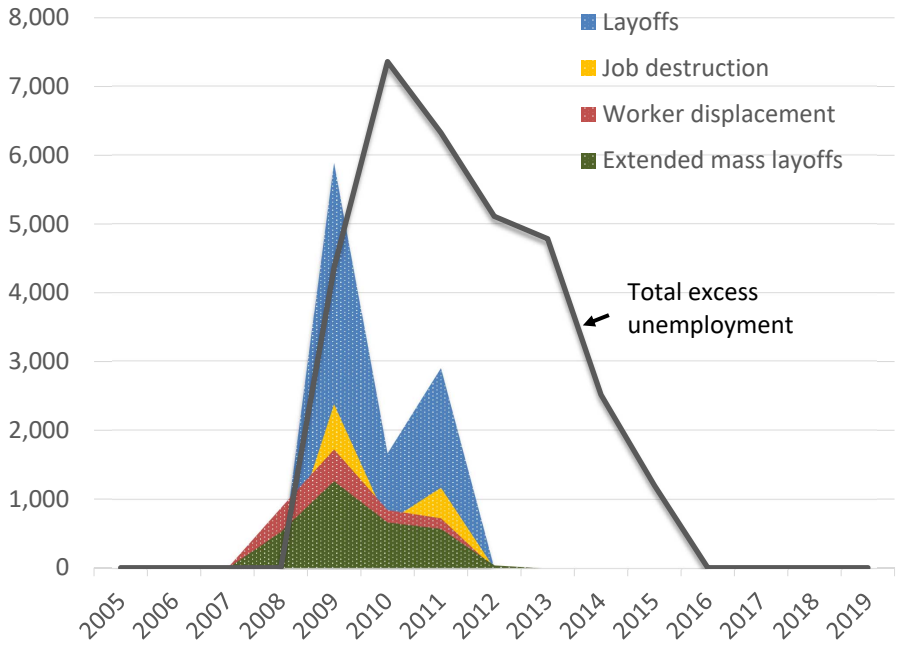


Figure 14: Unemployment from Excess Job Loss in 2009 and Total Excess Unemployment, by Four Measures of Job Loss, in Thousands of Workers

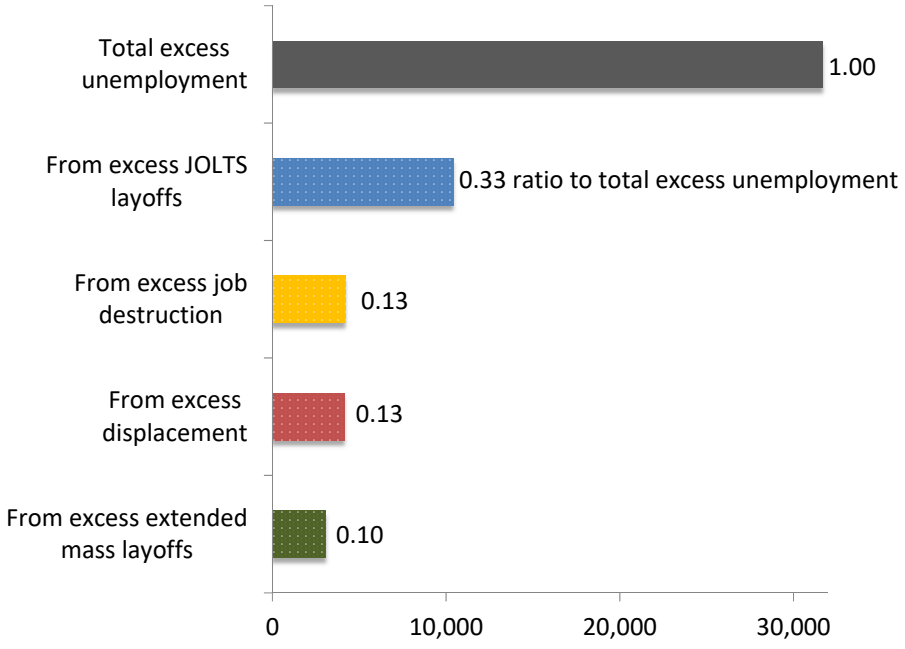


Figure 15: Contribution of Excess Job Loss to Cumulative Excess Unemployment, by Four Measures of Job Loss, in Thousands of Worker-Years

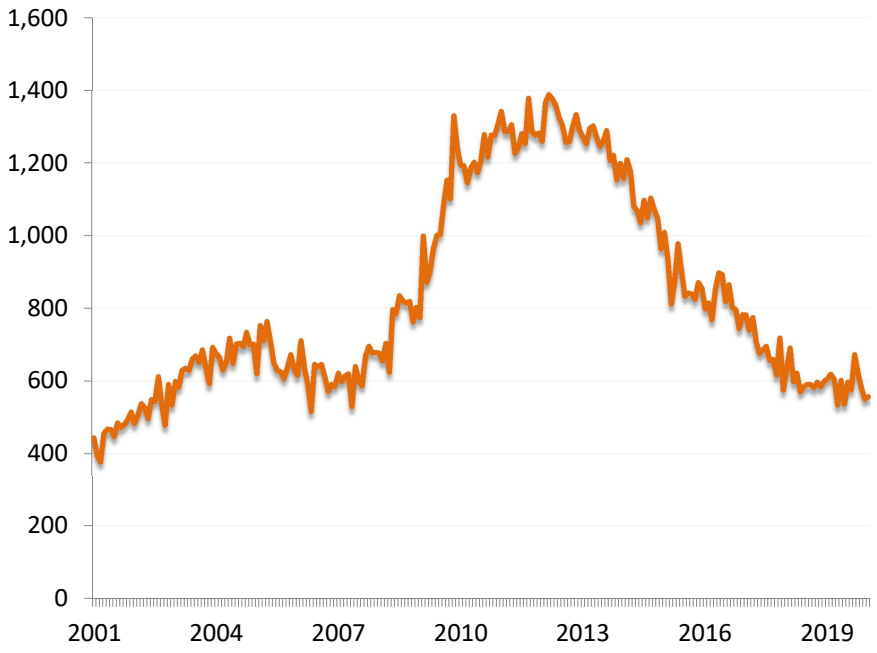


Figure 16: Number of Unemployed New Entrants to the Labor Force, in Thousands

6.3 Excess unemployment of new entrants

By definition, new entrants to the labor force are not the victims of job loss events. A bulge of unemployment of new entrants following an adverse shock indicates that either (1) unemployment is infectious, or (2) the bulge of unemployment arises from a decline in the incentives to create jobs. Figure 16 shows that new-entrant unemployment nearly doubled after the financial crisis. This finding rules out the hypothesis that the sole cause of lingering unemployment following the crisis was the slow absorption of workers who suffered job loss from the crisis. The direct channel cannot be the only link between a crisis and its subsequent gradual recovery.

6.4 Conclusions about the relation between the magnitude of the increase in unemployment following a recession shock and the measures of job loss

An unambiguous spike in regular and mass layoffs, job destruction, and displacement, and mass layoffs accompanies the shock that marks a recession. Figure 14 shows that the spike in layoffs more than fully accounts for the spike in unemployment in 2009 but cannot account for all of the excess unemployment afterwards. That is, the excess job loss accounts for the magnitude of the initial increase in unemployment, but not its persistence. The persistence is too large to be explained as reflecting only the personal experiences of the extra job-

losers dating from the spike. The direct channel is only part of the story of persistent high unemployment after the crisis. This conclusion is reinforced by the rise in unemployment among new entrants to the labor force.

7 Effective Exit Rate from Unemployment

7.1 Defining and measuring the effective exit rate

From the data on unemployment in the displaced workers supplement, we can estimate what we call the *effective exit rate from unemployment*, denoted f_t. We know the number of people in the survey who were displaced in the prior three years and who are currently unemployed. We also have our estimates of the number of people displaced in each of those years. The effective exit rate is based on the assumption that the probability of being unemployed ℓ months following a displacement is the product of the monthly exit rates from the time of displacement up to the survey. Here we adopt the perspective originated in Krueger, Cramer and Cho (2014), and expanded in Hall and Schulhofer-Wohl (2018) and Hall and Kudlyak (2019), that the typical path from initial unemployment to current labor-market activity often involves a mixture of spells of short jobs, time out of the labor force, and unemployment. As Krueger and co-authors showed, the probability of being unemployed a year later conditional on starting unemployed is much higher than would be expected from the monthly probability of unemployment ending raised to the 12th power. Our calculation here extends the calculation by two additional years, as we exploit the 3-year look-back in the displaced workers supplement of the CPS.

The implied relation between the observed number of people unemployed in the January survey of month t is

$$U_t = \sum_{\ell} \prod_{i} (1 - f_{t-i}) N_{\ell}. \tag{7}$$

We parameterize as

$$f_\tau = a - b\, u_\tau. \tag{8}$$

The parameter b is the negative sensitivity of the effective exit rate to the standard national unemployment rate u. Not surprisingly, it turn out to be essentially 1. We estimate a and b by minimizing the sum of squared residuals of the actual values to the implied values of U_t. The estimated value is $b = 1.00$ and the monthly effective exit rates range from 0.042 in 2010 to 0.099 in 2018. By contrast, the monthly exit rate is around 0.5.

7.2 Implications of low effective job-finding rates

This subsection provides evidence that recessions involve unemployment with long durations looking for work. During the long re-employment process, the unemployed often circle among unemployment, out of the labor force, and short-term jobs.

Figure 17 shows unemployment by reason as a share of the labor force, except for labor force new entrants. The figure shows that recession involve not only an increase in unemployment due to permanent and temporary layoffs but also due to completion of temporary jobs and labor force re-entry. This points towards an elevated number of individuals taking temporary jobs and circling between unemployment and OLF.

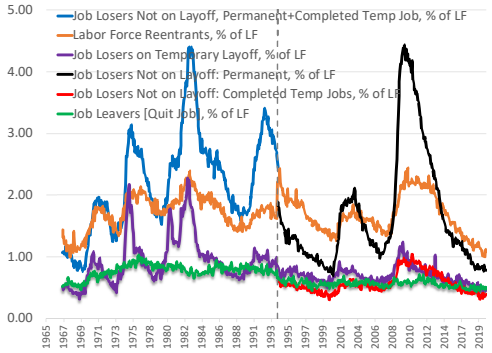
When a crisis causes a spike in unemployment, there is a shift away from stable jobs and toward brief jobs in the working-age population. This shift gradually subsides during the recovery. To demonstrate this phenomenon, we study the 8-month CPS activity paths, as in Hall and Kudlyak (2019). We define short employment spells as those lasting one or two months. These are the spells that are preceded and succeeded by unemployment or out of the labor force. We define an individual to have stable employment if employed in all 8 reported months. We calculate the average number of short employment spells among the of CPS respondents of working age. We also calculate the average number of respondents in stable employment. We create an index of the shift toward short jobs as the difference between first and second of these calculations. Figure 18 shows the short-spell index starting in 1976 for four demographic groups. The indexes jump upward in recessions and gradually declines during the ensuing recovery for all four groups.

8 Models and Evidence

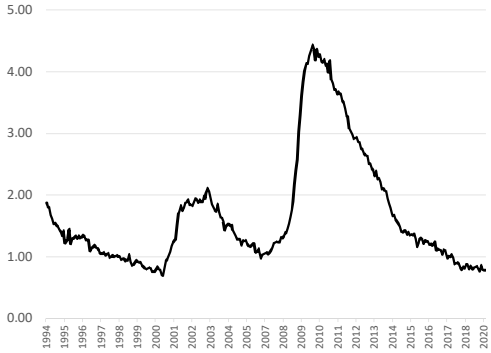
This section describes a sequence of versions of a model that may help understanding of the mechanisms lying behind our findings about what happens after a crisis. They may aid an understanding of the slow rate of recovery of unemployment as it continues to decline up to a decade after a spike in job loss.

We build on the DMP model of the labor market. Aspects of labor-market dynamics in recoveries that we consider in our investigation:

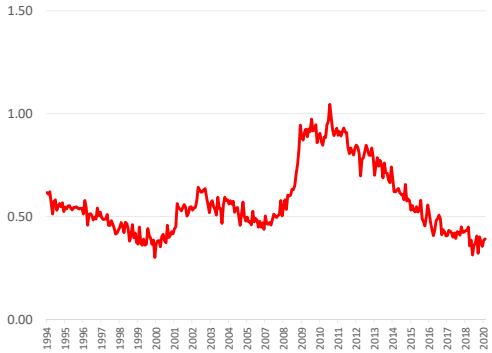
- Low rates of effective job finding despite high monthly exit rates out of unemployment
- A congestion effect arising from an external effect of unemployment on recruiting cost
- Credit frictions that decline during a recovery, raising workers' profitability to their employers



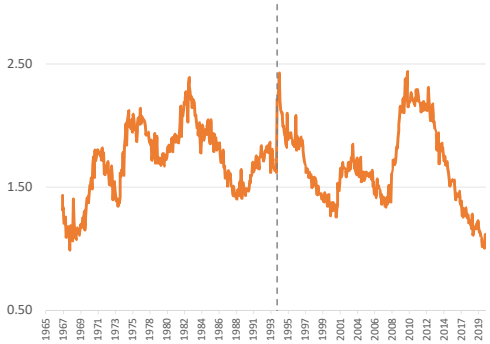
(a) Unemployment by reason



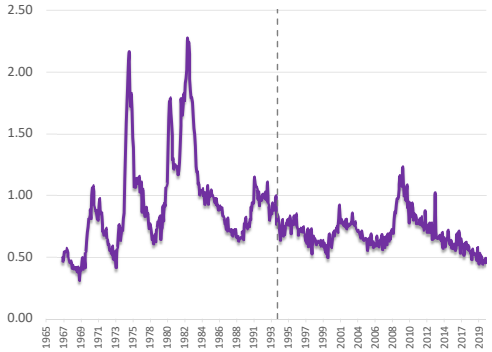
(b) Permanent job losers



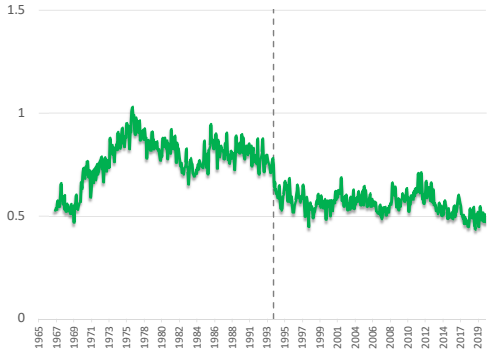
(c) Completed temporary job



(d) Labor force re-entrants

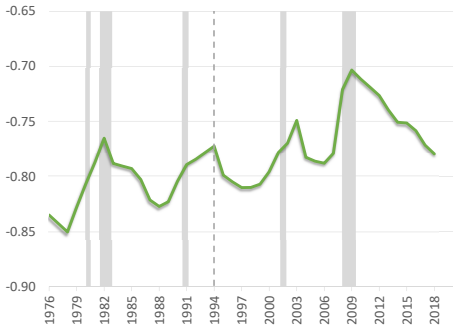


(e) Temporary layoffs

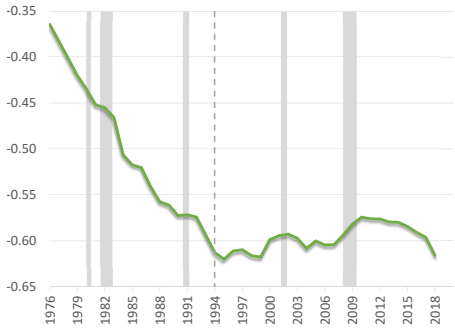


(f) Quits

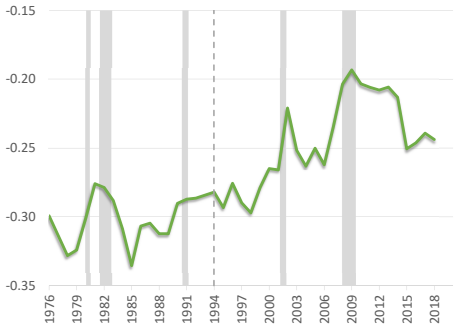
Figure 17: Unemployment by Reason, as Share of Labor Force



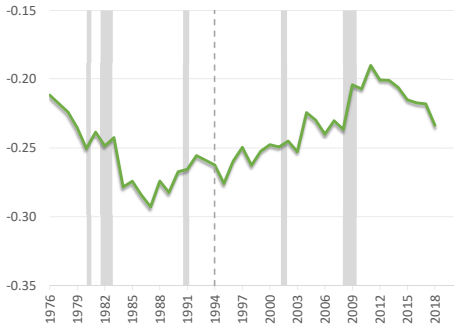
(a) Men aged 25-54



(b) Women aged 25-54



(c) Men aged 16-24



(d) Women aged 16-24

Figure 18: Indexes of the Duration of Employment Spells

Our model treats levels of unemployment and employment as state variables. At the outset, these state variables have values left behind by an event we call the crisis. The model traces the movements of the state variables for the next decade or so, trying to emulate the slow but steady return of the variables to their normal values.

8.1 The basic DMP model

The simple DMP model has three components: (1) wage determination, (2) tightness determination, and (3) the law of motion of unemployment. We make the simplifying assumption that productivity and the wage are constant. This assumption is relatively innocuous and greatly simplifies the model.

The labor market operates on the principles of random search. We measure the tightness of the market by the duration, T, of the typical vacancy. We let P be the present value of a newly hired worker's productivity and W be the present value of the wage, and $J = P - W$ be the net benefit to the firm from hiring a new worker. The flow cost of recruiting is a constant, κ. Recruiting satisfies the zero-profit condition,

$$\kappa T = J. \tag{9}$$

This condition pins down tightness:

$$T = \frac{J}{\kappa}. \tag{10}$$

A standard principle of the DMP model is that the job-finding rate increases with tightness:

$$f = \phi(T). \tag{11}$$

Our discussion of the model will follow the traditional principle that the only flow in the labor market that is sensitive to tightness is from unemployment to employment. To keep the exposition compact for now, we neglect variations in flows into and out of the labor market and from job to job without intervening unemployment. Thus the law of motion of unemployment is

$$\dot{u} = s \times (1 - u) - f\, u, \tag{12}$$

where s is the separation rate into unemployment. The job-finding rate f is also the exit rate from unemployment under the standard simplifying assumptions. The ergodic unemployment rate is $u^* = s/(s + f)$.

Figure 19 describes the behavior of the model in a phase diagram. The stationary value of unemployment is

$$u^* = \frac{s}{s + f} = \frac{s}{s + \phi\left(\frac{J}{\kappa}\right)}. \tag{13}$$

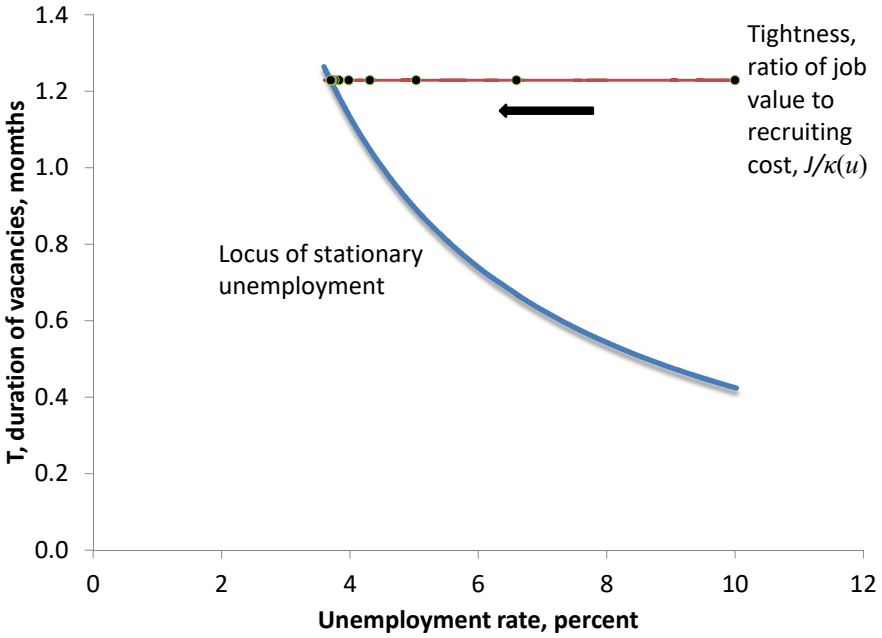


Figure 19: Phase Diagram for the Standard DMP Case

The movements of two variables, u and T, are along the horizontal line labeled J/κ. For example, the economy might start right after the crisis created high unemployment coupled with low tightness, at the right end of that line. As the recovery proceeds, the economy moves horizontally to the left. The dots on the horizontal path show the progress by month. In the first month, unemployment falls by more than two percentage points. As the unemployment rate falls during the recovery, the steps become smaller, but the calculations show that the economy closes most of the gap in just three or four months.

The curve in Figure 19 is the locus of stationary unemployment in the u, T space. At high unemployment rates, a given job-finding rate generates a higher outflow from unemployment simply because the rate applies to more workers. With lower unemployment, constancy of unemployment requires a higher jobfinding rate and thus a tighter labor market. The amount of the downward slope of the stationary locus is derived from an estimated matching function.

Figure 20 compares the path of unemployment starting from 10 percent implied by this model with the actual path starting from the peak in 2009. We use the standard job-finding rate of 0.3 per month together with a separation rate that yields an ergodic unemployment rate of 3.5 percent. The model's recovery is far speedier than actuality, the point made emphatically by Cole and Rogerson (1999). And the path is extremely convex, far away from the path of actual unemployment.

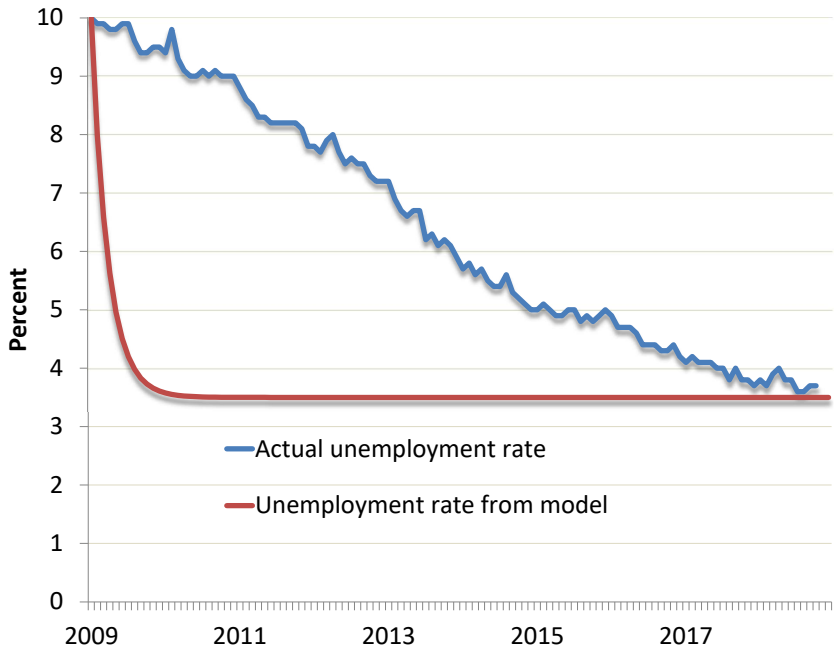


Figure 20: Recovery Path of Unemployment with the Standard DMP Model

8.2 Unemployment path with low effective unemployment exit rate

We know that the 30 percent per month transition rate from unemployment to employment in the basic DMP model considerably overstates the actual exit rate from unemployment. Figure 21 shows the model's unemployment path in the most recent recovery, together with the actual path, with the lower effective exit rate from unemployment of 0.042 per month estimated in a previous section. This alteration substantially delays the recovery but not nearly enough to match the actual path of unemployment.

We conclude that using the estimated effective unemployment exit rate of 0.042 makes an important contribution to matching the model's unemployment path but cannot be a full resolution of the low slope puzzle.

8.3 Congestion externality

Here we consider adding the property that the flow cost of maintaining a vacancy is increasing in unemployment, to capture a congestion effect on recruiting cost. We introduce the function

$$\kappa(u) = \kappa^* \left(\frac{u}{u^*}\right)^{\gamma} \tag{14}$$

to make the recruiting cost rise with unemployment. We show how congestion depresses the recovery rate in the simple DMP model with conventional parameter values. We continue

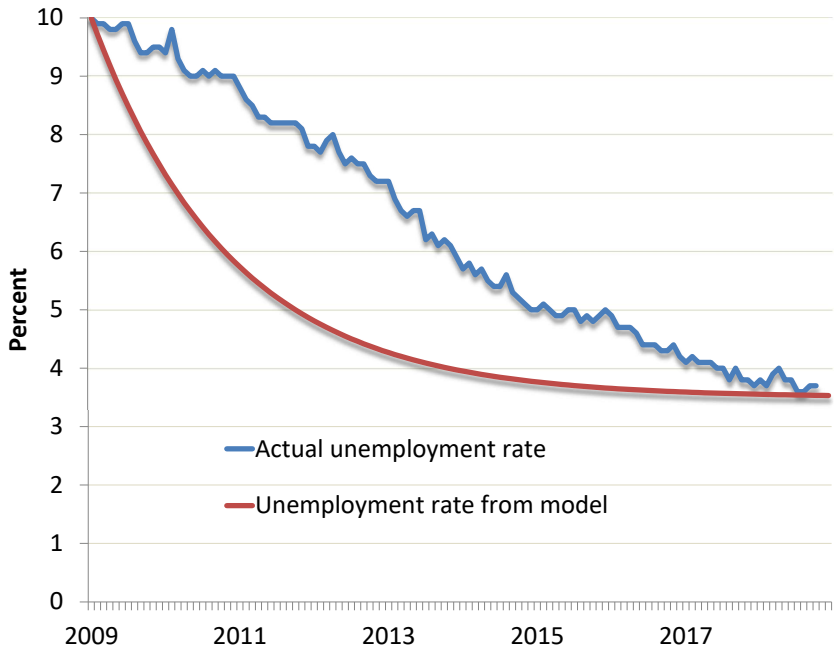


Figure 21: Recovery Path of Unemployment with Estimated Effective Exit Rate from Unemployment, and Actual Unemployment, 2010 to 2020

to take the job value J to be a constant—the downward slope of path of unemployment during the recovery is entirely the endogenous result of slow adjustment and not the result of improved incentives for job creation. If the crisis involved diminished incentives, those incentives snap back to normal at the beginning of the recovery.

With endogenous congestion, the recovery can be much slower. Figure 22 shows the effect of setting γ, the elasticity of the recruiting cost with respect to unemployment, to one, The locus of stationary unemployment remains the same, but the congestion cost bears more when unemployment is high, so the path from high initial unemployment to, ultimately, full employment, is vastly slower. The figure shows a multitude of tiny steps, each corresponding to a month.

It is clear that the story of the phase diagram only works if the congestion effect embodied in $\kappa(u)$ is reasonably strong—enough to twist the $J/\kappa(u)$ curve clockwise from flat to downward sloping, and lying close to the locus of stationary unemployment.

Figure 23 shows the model's unemployment path for the most recent recovery along with the actual unemployment path. Parameter values for the model are $\gamma = 0.95$ and $\psi = 0.3$. The model's path has unemployment falling faster than the actual path.

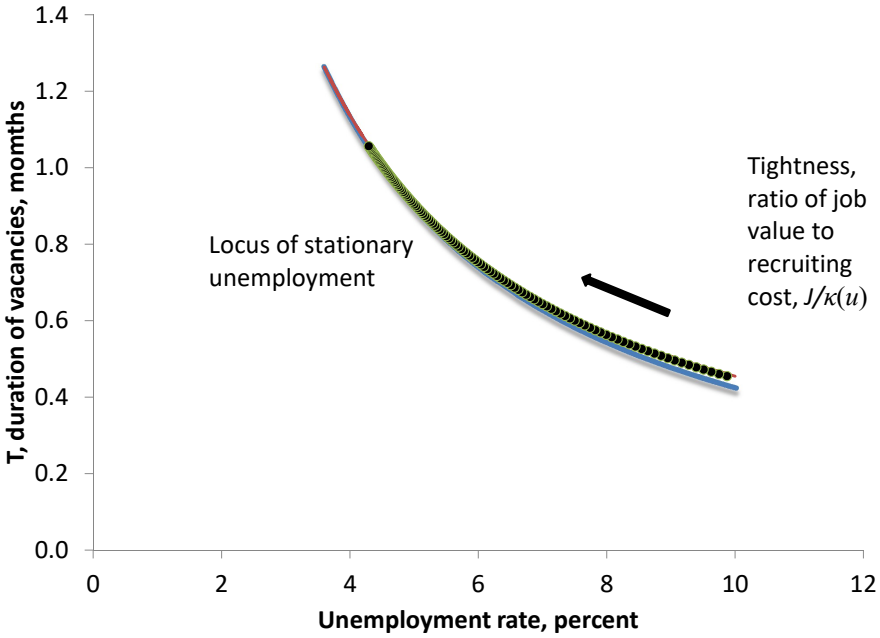


Figure 22: Phase Diagram for the DMP Case with Congestion in Recruiting

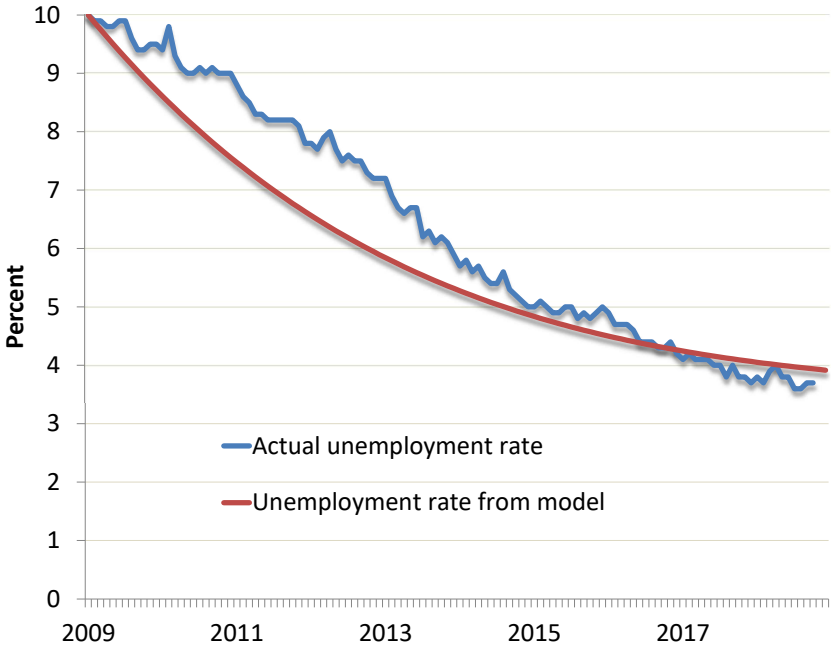


Figure 23: Recovery Path of Unemployment in the Congestion Model Compared to Actual Path

8.3.1 Ideas about the congestion effect

Engbom (2018) builds a model containing a congestion effect from unemployment. Job-seekers expand the scope of their search when jobs become harder to find in a recession. Recruiting firms' screening costs rise because of a higher volume of applications from job-seekers who are applying for less suitable jobs, on account of their widening scope. He cites recent evidence that screening costs, which rise in proportion to the number of applications, are a substantial fraction of total hiring costs. Advertising costs, which do not scale with applications, are relatively unimportant.

Molavi (2018) states "I show that a deterioration in the quality of the pool [of applicants] leads firms to post fewer vacancies and to demand a more positive signal of a worker's ability before hiring the worker. When the pool becomes more adversely selected, these two effects conspire to depress the job-finding rates of all workers irrespective of their ability."

Moscarini (2001) builds a model in which slack labor markets involve elevated levels of what he calls "excess worker reallocation." In the model, workers with specialized skills search selectively and contact few vacancies, where they are likely to be hired; while workers with weak comparative advantage apply to any vacancy, driven by the low anticipated acceptance rate. The latter workers produce movements across job types, both job-to-job and through unemployment–excess worker reallocation. He finds that in a tight labor market, comparative advantages dominate waiting costs and excess worker reallocation is lower and matches are more successful.

Lockwood (1991) considers a setup where employers may administer a test. When this happens in equilibrium, employers also consider unemployment duration as informative about how many times the job-seeker has flunked previous tests. When unemployment is higher, this problem may worsen, creating a congestion externality.

Blanchard and Diamond (1994) develop a "ranking" model of unemployment in which when firms receive multiple acceptable applications, they hire the worker who has been unemployed for the least amount of time. In the model, the exit rate from unemployment depends on duration, and the effect of duration is stronger the higher the rate of unemployment. "In a tight labour market, a long-term unemployed worker may be the only applicant at a given vacancy. In a depressed labour market, most vacancies receive many applications, so that the probability of being hired decreases quickly with duration."

Eeckhout and Lindenlaub (2019) observe that, in a recession, the composition of job-seekers shifts toward the unemployed and away from on-the-job searchers. Incentives for job creation are diminished because the productivity of job-seekers is lower.

8.3.2 Other externalities

The channel behind the spillovers from a large sudden increase in the number of jobseekers depends on whether the increase is accompanied by a decline in the marginal product of labor. Typically, the literature considers a scenario in which the marginal product of labor declines. This depresses the incentives to create jobs and slows down the absorption of jobseekers via the decline in labor demand. We consider this possibility in the next subsection. Here we briefly consider how studies of local spillovers might identify a congestion externality along the lines of the model in this subsection.

We look to the literature on local labor markets for evidence. An ideal natural experiment is a situation when there is a large influx of jobseekers into a local labor that is exogenous to the local labor market conditions and leaves the marginal product of labor unchanged. Such an increase in jobseekers can occur as a result of job loss, that is, from the reduction in local labor supply, or as a result of inflows from OLF or a population increase, that is, from the increase in local labor supply. An example of a reduction in local labor supply is a plant closing because of outsourcing of production overseas. An example of an increase in local labor supply is an influx of immigrants such as the Mariel boatlift.

Gathmann, Helm and Schonberg (2020) quantify the spillover effects of large local mass layoffs using administrative data on firms and workers in Germany. They find sizable and persistent negative effects on the regional economy: The regions generally, and especially firms producing in the same broad industry as the layoff plant, lose many additional jobs. Employment in these indirectly affected firms starts to decline one year after the mass layoff takes place. Four years after the initial layoffs, 35 percent of local employment losses stem from spillover effects in plants not directly affected by the mass layoff. The authors argue that the negative effect on other firms is the result of the loss of agglomeration externalities. In contrast, negative employment effects on workers employed in the region at the time of the mass layoff are considerably smaller because these workers are able to relocate to other regions. Vom Berge and Schmillen (2015) and Jofre-Monseny, Sanchez-Vidal and Viladecans-Marsal (2017) analyze the potential indirect effects of plant closures and do not find evidence for spillover effects. Gathmann and co-authors argue that this is because the size of the mass layoffs in these papers is much smaller.

Our reading of these papers is that they study the effect of local layoffs in terms of an overall external effect on other firms, without considering the separate effects on the marginal product of labor and the cost of recruiting workers. That distinction is relevant to identifying the source of the slow decline of unemployment during recoveries.

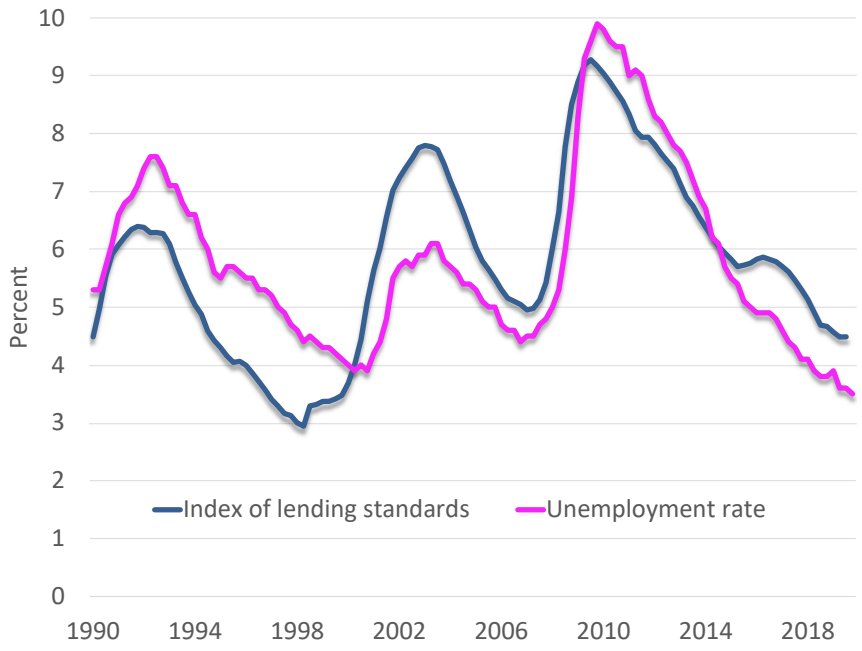


Figure 24: Scaled Index of Loan Availability Compared to Unemployment Rate

8.4 Impaired profitability of employment

Recall that $T = J/\kappa$ and that $J = P - W$, the net benefit of the employment relationship to the employer. If some force depresses P by more than it depresses W, J falls and T falls as well, assuming that the source of the decline in J does not also affect recruiting cost κ. Tightness falls, and unemployment rises, along the path set by its law of motion. One potential source of this effect is a crisis-induced cut in the availability of credit, which raises discount rates and thereby cuts J. As the availability of credit gradually returns back to normal, unemployment also returns to normal.

To measure the availability of credit, we use data from the Federal Reserve Board's survey of Senior Loan Officers. Respondents in the survey answer in terms of tightening and easing of commercial loan standards. We cumulate these answers using the statistical model in Hall (2011) to form an index of loan availability. The scale of the index is arbitrary. In Figure 24, we scale it to have the same standard deviation as unemployment and compare the scaled index to actual unemployment. The two variables move closely together.

To account for the effect of loan availability L in the model, the equation for the job-finding rate becomes

$$f = \psi \cdot (1 - \lambda L). \tag{15}$$

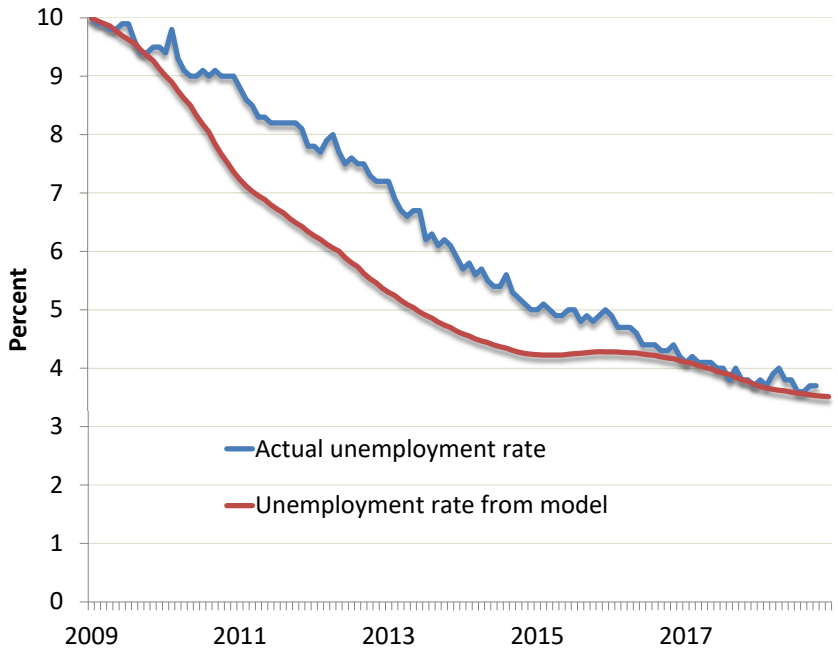


Figure 25: Comparison of Model's Recovery Path of Unemployment to Actual Unemployment

We pick the coefficient λ to give the best fit of the model's implied unemployment path following the financial crisis. Figure 25 shows the unemployment recovery path from the model in comparison to the actual path.

Among the models in this section, this one comes closest to matching the actual declining path of unemployment.

8.5 Conclusions about the variants of the DMP model

In general, the models we reviewed in this section do not fully resolve the puzzle of low recovery speeds of unemployment. The simple DMP model with a job-finding rate calibrated to the high observed rate of monthly unemployment-to-job transitions is a disaster, as Cole and Rogerson (1999) observed early in the development of the model. Using an effective unemployment exit rate inferred from the 3-years histories of displaced workers improves the fit but still does not match the line over the 10 year span in our example. The congestion model does somewhat better but still overstates the job-finding rate in early years. . The model based on credit availability also falls in that way. We have spent some time checking combinations of these forces, but because none of them match the data, no combination is fully successful.

9 Other Forces Operating during a Recovery

9.1 Other variables

Next we take a look at a variety of macro variables that may be involved in recoveries. These are policy instruments—government spending and monetary policy—and influences that might be considered exogenous determinants—productivity, labor-force growth, and the stock market.

Government purchases. Figure 26 displays consolidated government purchases of goods and services divided by the CBO's potential GDP series. The dates of peaks in the employment rate appear along the bottom—not the peaks in the purchases series itself. Essentially all macroeconomic models agree that an increase in government purchases stimulates output. The figure shows that purchases in the first recovery,1949 through 1953, grew rapidly because of the Korean War. The Reagan military buildup in the 1980s also accounted for rising purchases relative to potential GDP in that recovery—in all other recoveries, even the one in the 1960s containing the Viet Nam war, purchases failed to keep up with potential GDP. The conclusion with respect to those, notably including the most recent recovery, is that fiscal policy taking the form of deliberate expansion of purchases—such as the American Recovery and Reinvestment Act—provided stimulus when the economy was weak. As the economy recovered, the stimulus was withdrawn. This influence was much greater, relative to potential GDP, in the current recovery, compared to earlier recoveries.

Government transfers. The US has large and effective countercyclical government transfer programs and practices. Figure 27 shows the history of dollar benefits in terms of our recovery chronology. We standardize the data by dividing by nominal disposable income. Some of the countercyclical pattern arises from automatic stabilizers—programs that enroll more dependents in bad times—and some from discretionary expansion of programs and creation of new ones—such as extending unemployment insurance benefits to cover more weeks.

The figure shows that there is a good deal of heterogeneity across the recoveries. Transfers declined remarkably in the first recovery, starting in 1949. In the next four recoveries, transfers grew relative to disposable income. In four of the recent five recoveries, transfers declined.

Monetary policy. The central instrument of monetary policy in the US is the Federal Reserve's policy interest rate. The standard way to state its effect as an instrument is to define it as the margin of the economy's natural or equilibrium short interest rate over the policy rate. To expand, the Fed depresses the policy rate and increases the margin. And to

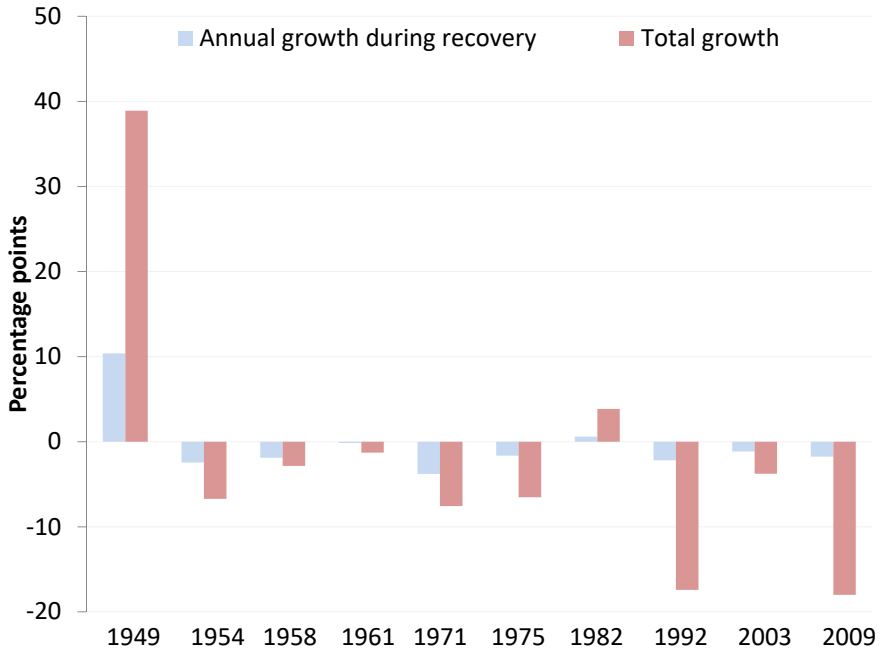


Figure 26: Annual Growth in Real Government Purchases of Goods and Services, as a Ratio to Potential GDP, by Recovery, and Total Growth during Each Recovery

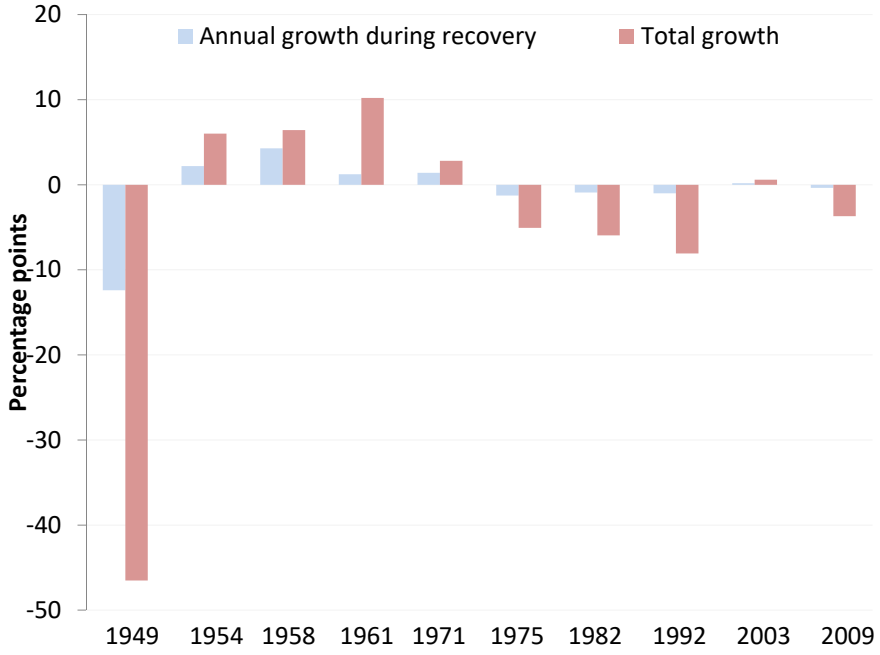


Figure 27: Government Transfer Payments as a Ratio to Disposable Income, by Recovery, and Total Growth during the Recovery

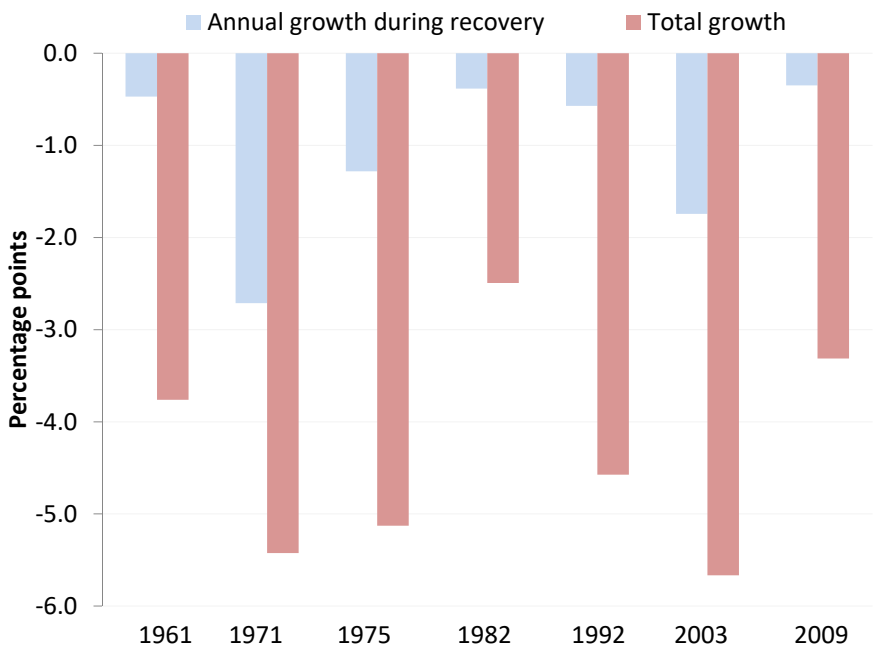


Figure 28: The Expansionary Margin of Interest-Rate Policy

contract, the Fed raises the policy rate above the natural rate to drive the margin negative. Laubach and Williams (2003) is a widely used estimate of the natural short rate.

Figure 28 shows the average value of the expansionary margin of interest-rate policy, according to Laubach and Williams. The Fed has chosen net expansion in four expansions and net contraction in two. In the current recovery, the Fed has chosen substantial expansion, almost as much as in the recovery of second half of the 1970s. Oddly, the late 1970s were a period of high and rising inflation, so the Fed was failing in its duty to lean against the wind. The choice of high expansion in the current recovery is in accord with optimal policy, given slightly substandard inflation during the period.

As with the other policy instruments, we find heterogeneity in the setting of the Fed's interest-rate margin during the recoveries of the past 70 years.

9.2 Other driving forces during recoveries

Financial discounts. Forces other than macroeconomic policy may influence unemployment declines during recoveries. For example, a recent literature has described a relation between financial discounts and unemployment. See Hall (2017) in the context of the aggregate labor market and Kilic and Wachter (2018) and Kehoe, Lopez, Midrigan and Pastorino (n.d.) in general equilibrium. These papers consider DMP-type models of unemployment and events that alter economy-wide discount rates, thus changing the job-value, which is

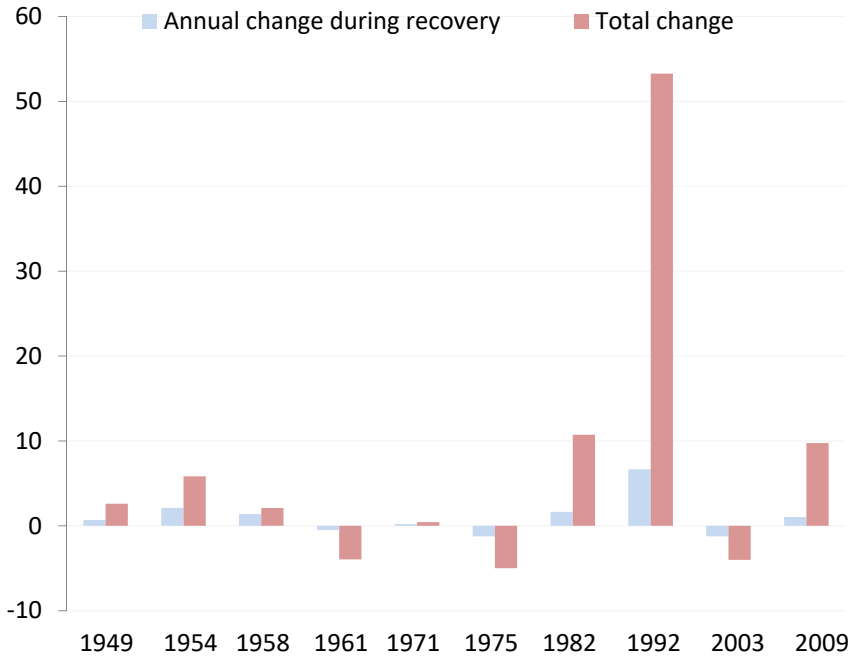


Figure 29: Price/Dividend Ratio of the S&P 500, by Recovery, and Total Growth during the Recovery

the present value of the contribution of a newly hired worker net of the wage paid to the worker. Discounts sometimes jump upward almost discontinuously, as they did immediately after the Lehman bankruptcy in 2008. The job value represents the incentive to recruiting. When it declines, the labor market slackens and unemployment rises. In the recovery phase, falling discounts raise the job value and unemployment falls.

According to principles of modern finance elucidated in Campbell and Shiller (1988), discount rates for risky future cash payouts are equal to the expected rates of returns associated with those payouts. In a recovery, the stock market rises, the price/dividend ratio rises, and expected rates of return decline. According to the literature linking financial events to the labor market, unemployment declines back to normal. Figure 29 shows the history of the ratio for recoveries since 1949. The ratio rose dramatically during the recovery of the 1990s. It fell substantially during the financial crisis in 2008 and 2009, but recovered during 2010, when unemployment was still rising. Its relation to the chronology of the ratio of GDP to potential in earlier years is less apparent. A rising price/dividend ratio is sometimes important for a recovery, but does not explain the reliability of US business-cycle recoveries.

Productivity growth. Another aggregate influence of unquestioned importance for GDP growth is productivity growth. If the topic of this paper were real GDP growth in recoveries, productivity would receive top billing. But the relation of productivity growth to the gradual reduction of unemployment in recoveries is ambiguous and may well be small.

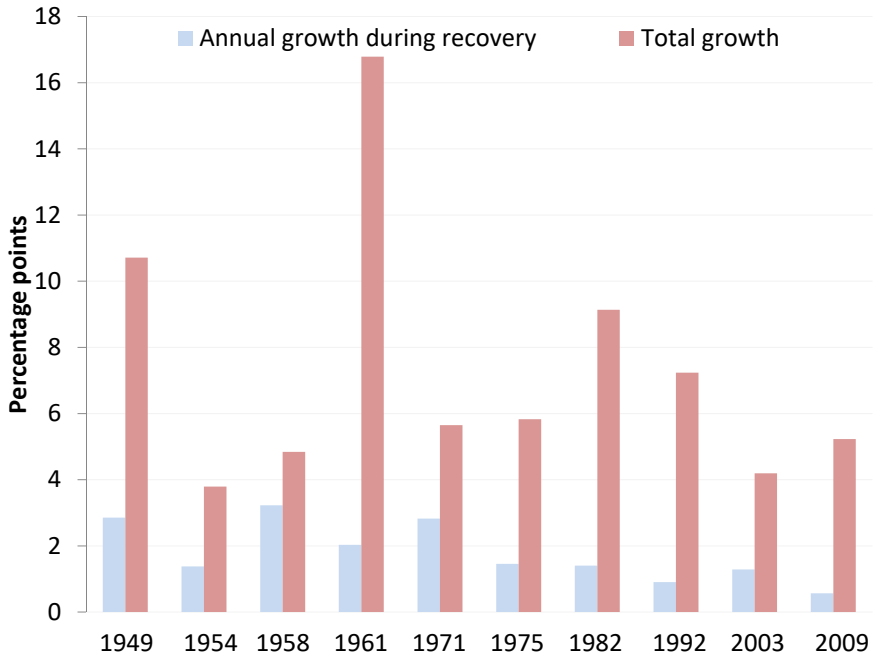


Figure 30: Annual Growth in Total Factor Productivity, by Recovery, and Total Growth during Each the Recovery

Figure 30 shows that productivity growth tended to be high in recoveries through the 1980s, had a small comeback in the recovery starting in 2003, and had a spectacular shortfall in the current recovery. Overall, productivity growth tended to be irregular in recoveries.

Variations in labor-force growth. The DMP model of Mortensen and Pissarides (1994) has a constant labor force. Extensions to endogenous participation may involve positive or negative co-movements of participation and unemployment. Figure 31 shows that the participation rate grew during the years up to 1990 when the rising rate for women was a key factor for overall participation (to achieve a basic adjustment for demographic influences, the data refer to ages 25 through 54). In the three recoveries since 1990, participation was essentially unchanged during recoveries.

10 Relation to the Natural Rate of Unemployment

During a recovery, unemployment gradually declines to converge to a minimum unemployment rate, which we take to be 3.5 percent. This minimal rate is well under estimates of the natural unemployment rate or NAIRU. The view we develop in this paper does not involve the concept of a natural rate. Instead, we believe that unemployment gradually declines with equal proportional decrements over the duration of a recovery, until it reaches the minimal rate of 3.5 percent. Only in the most favorable case of a long recovery is it possible for the

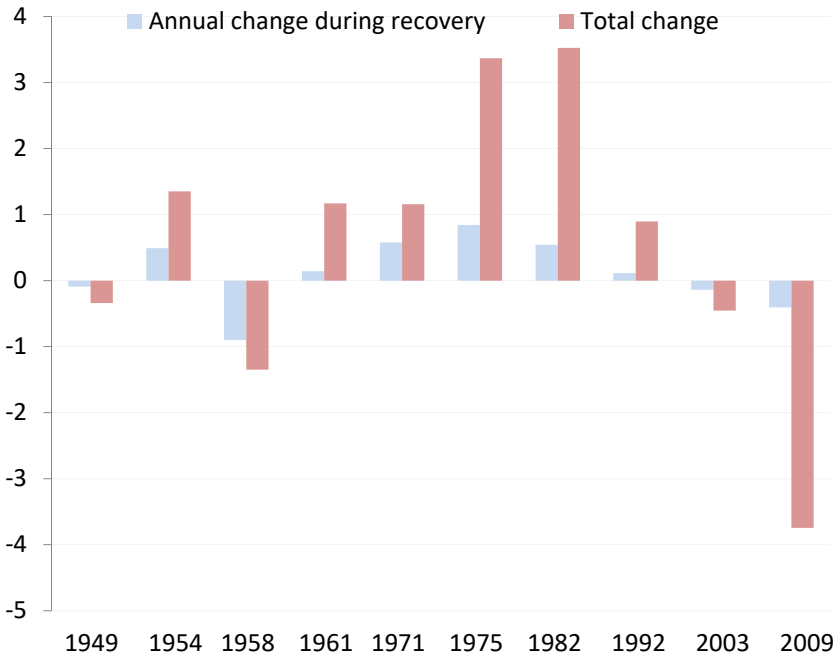


Figure 31: Percentage-Point Annual Change in Labor-Force Participation, by Recovery, and Total Percentage Change during Each Recovery

rate to be that low. As the recovery progresses, log unemployment declines along a straight line with constant slope.

The natural rate is defined as the rate that would prevail absent recent dislocations from shocks. Another definition within New Keynesian thinking is the rate that would prevail in the absence of sticky wages and prices. These definitions spring from a view that the economy vibrates around a stable underlying growth path. Our take is somewhat different. At widely spaced intervals—in recent decades, around 10 years—a significant adverse shock causes a substantial jump in the unemployment rate. There follows, during the relative calm period before the next bad shock, a predictable decline of about 0.10 log points points per year in the unemployment rate. We believe not in a constant natural rate, but in a natural declining path of the rate.

11 Related Literature

Pries (2004) builds a DMP model that explains the high persistence of unemployment as the result of recurrent spells of unemployment following a shock. Once a job match is made, the parties are at risk of an adverse productivity realization that reveals that the match should end and the worker should return to the labor market. In normal times, most matches will have become known to be reliable and no longer at risk of being found unproductive. At

random, a cloud may form over the labor market that calls into question the earlier belief that a match is good—the parties need to receive a new signal of reliability for a fraction of the existing matches. Some of the matches end immediately and the others are exposed to the possibility that they will be found to be unproductive from a later draw of productivity.

If the aggregate shock simply knocks out some of the existing matches, the model generates little persistence—the victims of the shock regain reliable employment almost as quickly as they would without the learning-by-experience feature of the model (see Pries's Figure 3). The broader version of the shock, which induces the parties to wait to determine who are the job losers, makes the effect of the shock realistically persistent. Thus Pries has found a way to make the shock effectively persistent.

11.1 Cyclicality of the consequences of job loss

Davis and von Wachter (2011) find that the present-value losses for displacements that occur in recessions are nearly twice as large as for displacements in expansions. They conclude that that the present-value earnings losses associated with job displacement are very large, and they are highly sensitive to labor market conditions at the time of displacement.

Mueller (2017), using micro-data from the Current Population Survey, documents that in recessions the pool of unemployed shifts toward workers with high wages in their previous job and that these shifts are driven by the high cyclicality of separations for high-wage workers.

Fujita and Moscarini (2017) document that a large share of workers return to their previous employer after a jobless spell, and experience different unemployment and employment outcomes from job switchers. They add a recall option to a canonical search and matching model and find that the recall option amplifies the cyclical volatility of new job-finding and separation probabilities. Fujita, Moscarini and Postel-Vinay (2020) argue that the labor market policy response to the pandemic must save aggregate matching capital.

11.2 Cyclicality of recruiting intensity

Modestino, Shoag and Balance (2019), Modestino, Shoag and Balance (2016) and Hershbein and Kahn (2018) use data on the education and skill requirements of online job postings to study the relationship between the unemployment rate and job requirements. They argue that changes in the skill requirements of vacancies are cyclical. When job seekers are plentiful, firms become choosier and increase skill and educational requirements of vacancies. And, conversely, when job seekers are scarce, firms relax skill and educational requirements of vacancies. Modestino et al. (2019), Modestino et al. (2016) calculate that approximately 20 percent of the aggregate increase in upskilling during 2007-10 can be accounted for by their

cyclical labor supply mechanism. The findings are consistent with evidence of pro-cyclical recruitment intensity provided by Davis et al. (2013) who find that in booms firms obtain more hires per vacancy than in recessions.

Hershbein and Kahn (2018) argue that not all upskilling during 2007-09 was cyclical and that some firms laid off workers during the downturn, invested in labor-saving technologies during the recovery and subsequently required higher skills from new hires than in 2007.

11.3 Heterogeneity in rates of growth of employment across recoveries

We noted early in this paper that the homogeneity of the rate of recovery of unemployment does not carry over to the rate of growth of employment, which is sensitive to variations in the growth of labor-force participation. Leduc and Liu (2020) modify the DMP model by introducing a cost of creating vacancies, which results in variations in the growth of employment even if participation is constant.

12 Concluding Remarks

We propose a rather different notion of the recovery of the US economy from a recessionary shock. The shock leaves a legacy of high unemployment. In the standard view, unemployment is high because there is a shortfall of demand. As time passes, demand recovers and unemployment returns to normal. In our view, unemployment remains high and declines only gradually because of frictions in rebuilding employment. We believe that the frictions comprise low effective rates of movement back into jobs and congestion externalities that raise the effective cost of job creation. elevated frictions in lending standards after a crisis are another friction that dissipates slowly during a recovery.

clearpage

References

Auray, Stephane, David L. Fuller, and Damba Lkhagvasuren, "Unemployment Insurance Take-up Rates in an Equilibrium Search Model," *European Economic Review*, 2019, *112*, 1–31.

Berge, Philipp Vom and Achim Schmillen, "Direct and Indirect Effects of Mass Layoffs: Evidence from Geo-Referenced Data," 2015. IAB-Discussion Paper, No. 11/2015.

Blanchard, Olivier Jean and Peter Diamond, "Ranking, Unemployment Duration, and Wages," *The Review of Economic Studies*, 1994, *61* (3), 417–434.

Blank, Rebecca M. and David Card, "Recent Trends in Insured and Uninsured Unemployment: Is There an Explanation?," *The Quarterly Journal of Economics*, 1991, *106* (4), 1157–1189.

Campbell, John Y. and Robert J. Shiller, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1988, *1* (3), 195–228.

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.

Davis, Steven J. and John Haltiwanger, "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics*, 1992, *107* (3), 819–863.

___ and Till von Wachter, "Recessions and the Costs of Job Loss," *Brookings Papers on Economic Activity*, Fall 2011, (2), 1–55.

___ , R. Jason Faberman, and John C. Haltiwanger, "The Establishment-Level Behavior of Vacancies and Hiring," *Quarterly Journal of Economics*, 2013, p. 581–622.

Dupraz, Stephane, Emi Nakamura, and Jon Steinsson, "A Plucking Model of Business Cycles," Working Paper 26351, National Bureau of Economic Research October 2019.

Eeckhout, Jan and Ilse Lindenlaub, "Unemployment cycles," *American Economic Journal: Macroeconomics*, 2019, *11* (4), 175–234.

Engbom, Niklas, "Application Cycles," July 2018. Federal Reserve Bank of Minneapolis.

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.

___ , ___ , and Fabien Postel-Vinay, "The Labour Market Policy Response to COVID-19 Must Save Aggregate Matching Capital," 2020. CEPR VOX Policy portal, https://voxeu.org/article/labour-market-policy-response-covid-19-must-save-aggregate-matching-capital.

Fuller, David L., B. Ravikumar, and Yuzhe Zhang, "Unemployment Insurance: Payments, Overpayments and Unclaimed Benefits," *FRB St. Louis Regional Economist*, 2012, *October*, 12–13.

Gathmann, Christina, Ines Helm, and Uta Schonberg, "Spillover Effects of Mass Layoffs," *Journal of the European Economic Association*, 2020, *18* (1), 427–468.

Gregory, Victoria, Guido Menzio, and David G Wiczer, "Pandemic Recession: L or V-Shaped?," Working Paper 27105, National Bureau of Economic Research May 2020.

Hall, Robert E., "The Long Slump," *American Economic Review*, April 2011, *101* (2), 431–69. 2011 AEA Presidential Address.

——— , "High Discounts and High Unemployment," *American Economic Review*, February 2017, *107* (2), 305–330.

——— and Marianna Kudlyak, "Job-Finding and Job-Losing: A Comprehensive Model of Heterogeneous Individual Labor-Market Dynamics," 2019. Hoover Instiution.

——— and Sam Schulhofer-Wohl, "Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Jobseekers," *American Economic Journal: Macroeconomics*, January 2018, *10* (1), 1–32.

Hershbein, Brad and Lisa B. Kahn, "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings," *American Economic Review*, 2018, *108* (7), 1737–1772.

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan, "Earnings Losses of Displaced Workers," *American Economic Review*, 1993, *83* (4), 685–709.

Jofre-Monseny, Jordi, Maria Sanchez-Vidal, and Elisabet Viladecans-Marsal, "Big Plant Closures and Local Employment," *Journal of Economic Geography*, 2017, p. 163–186.

Kehoe, Patrick, Pierlauro Lopez, Virgiliu Midrigan, and Elena Pastorino, "Asset Prices and Unemployment Fluctuations," November. Federal Reserve Bank of Minneapolis.

Kilic, Mete and Jessica A. Wachter, "Risk, unemployment, and the stock market: A rare-event-based explanation of labor market volatility," *The Review of Financial Studies*, 2018, *31* (12), 4762–4814.

Krueger, Alan B., Judd Cramer, and David Cho, "Are the Long-Term Unemployed on the Margins of the Labor Market?," *Brookings Papers on Economic Activity*, Spring 2014.

Laubach, Thomas and John C Williams, "Measuring the natural rate of interest," *Review of Economics and Statistics*, 2003, *85* (4), 1063–1070.

Leduc, Sylvain and Zheng Liu, "The weak job recovery in a macro model of search and recruiting intensity," *American Economic Journal: Macroeconomics*, 2020, *12* (1), 310–43.

Lockwood, Ben, "Information Externalities in the Labour Market and the Duration of Unemployment," *The Review of Economic Studies*, 06 1991, *58* (4), 733–753.

Modestino, Alicia Sasser, Daniel Shoag, and Joshua Balance, "Downskilling: Changes in Employer Skill Requirements Over the Business Cycle," *Labour Economics*, 2016, *41*, 333–347.

——— , ——— , and ——— , "Upskilling: Do Employers Demand Greater Skill When Workers are Plentiful?," 2019. Working Paper, Northeastern University.

Molavi, Pooya, "A Theory of Dynamic Selection in the Labor Market," December 2018. Department of Economics, MIT.

Mortensen, Dale T. and Christopher Pissarides, "Job Creation and Job Destruction in the Theory of Unemployment," 1994, *61*, 397–415.

Moscarini, Giuseppe, "Excess Worker Reallocation," *The Review of Economic Studies*, 2001, *68* (3), 593–612.

Mueller, Andreas I., "Separations, Sorting, and Cyclical Unemployment," *American Economic Review*, July 2017, *107* (7), 2081–2107.

Pries, Michael J., "Persistence of Employment Fluctuations: A Model of Recurring Job Loss," *Review of Economic Studies*, 01 2004, *71* (1), 193–215.

Romer, Christina D. and David H. Romer, "NBER Business Cycle Dating: History and Prospect," December 2019. University of California, Berkeley.