

OLIVIER COIBION

University of Texas at Austin

YURIY GORODNICHENKO

University of California at Berkeley

DMITRI KOUSTAS

University of California at Berkeley

Amerisclerosis?

The Puzzle of Rising U.S. Unemployment Persistence

ABSTRACT The persistence of U.S. unemployment has risen with each of the last three recessions, raising the specter that future U.S. recessions might look more like the “Eurosclerosis” experience of the 1980s than like the traditional V-shaped recoveries of the past. We revisit several explanations for this rising persistence, decomposing them into three possible sources: business cycle fluctuations, changing policy responses, and propagation mechanisms. First, we find that financial shocks do not systematically lead to more persistent unemployment than monetary policy shocks, casting doubt on the hypothesis that different drivers of business cycles are the primary explanation. Second, we find that changing monetary and fiscal policy responses account for approximately one-third of the rise in unemployment persistence. Third, after examining three propagation mechanisms we find that jointly they cannot account for any rising persistence of unemployment. The three propagation mechanisms we focus on—declining labor mobility, changing age structures, and the decline in trust among Americans—are consistent with four other cyclical patterns that have evolved since the early 1980s: a rising cyclicity in long-term unemployment, lower regional convergence after downturns, rising cyclicity in disability claims, and missing disinflation. We exploit regional variation in labor market outcomes across Western Europe and North America during 1970–91 to assess the predictive capacity of each propagation mechanism for unemployment persistence. In summary, two-thirds of the rise in unemployment persistence is unexplained.

When the U.S. unemployment rate surged by four percentage points between 1979 and 1983 in the midst of what was then the most severe slump since the Great Depression, Western European countries experienced, on average, an almost identical rise in unemployment. But whereas by 1987 rapid job growth in the United States had offset all of this rise in joblessness, the average unemployment rate among West European countries declined only half a percentage point over the same time period. Thirty years later, and nearly six years after the start of the 2007–09 Great Recession, the recovery in the U.S. labor market is a pale shadow of the U.S. experience in the early 1980s. From its peak of 10 percent in October 2009, the United States has seen its unemployment rate fall only halfway back to its prerecession levels in the four years since, placing it midway between the pace of recovery of the United States and Western Europe in the early 1980s.

But while the pace of the U.S. job market recovery looks downright anemic relative to the rapid rebound experienced after the Volcker recessions of the early 1980s or prior recessions in the post–World War II era, the degree of persistence in unemployment since the Great Recession only modestly exceeds that following the 2001 recession, which in turn modestly exceeded that of the 1990 recession. From this perspective, we observe a gradual trend of increasingly weak recoveries over the last three recessions that contrasts sharply with the previous U.S. experience. If this trend reflects more than a historical coincidence and is to continue, future U.S. recessions are likely to display extended periods of depressed labor markets that will increasingly resemble the experience of many Western European countries in the 1980s.

To support the notion that there may have been common factors at work in the 1990, 2001, and 2007 recessions, we document four related properties of the Great Recession that are puzzling in comparison to historical (pre-1990) recessions but are not unusual when compared with the 1990 and 2001 recessions. They are (i) missing disinflation—given the historical link between inflation and unemployment, one would have expected inflation to fall much more in the Great Recession than it actually did; (ii) the unusually large share of long-term unemployed; (iii) the slow rate of convergence in regional labor markets; and (iv) the rise in disability claims during the Great Recession. We show that each of these properties is indeed puzzling relative to pre-1990 recessions, but that the experience in the Great Recession is either in line with or only modestly exceeding what would have been expected given the patterns of the 1990 and 2001 recessions and the severity of the Great Recession itself. Furthermore, each

was a feature of the Western European experience of the early 1980s. These four features of the Great Recession, when unemployment persistence was high, are puzzling when compared with pre-1990 recessions, when unemployment persistence was low. But these features are not puzzling when compared with the 1990 and 2001 recessions, when unemployment persistence was high, and the same features were also present in West European economies in the early 1980s, when unemployment persistence was also high. These facts at least suggest that common forces are at work, underlying the persistence of unemployment.

What then are the common forces that lie behind this rising persistence of unemployment in the United States? We consider three classes of explanation. First, the composition of shocks driving business cycles may have changed. If the channels through which shocks affect the economy vary across different types of shocks, then one might expect some shocks to have more persistent effects on the economy than others. This view is frequently advocated in the context of financial shocks, based on the evidence of Reinhart and Rogoff (2009) and Jorda, Schularick, and Taylor (2011) that financial crises have historically been associated with longer-lived downturns than typical recessions. Second, policy responses to business cycles may have changed. In the context of the Great Recession, one can point to a monetary policy response that has been severely constrained by the zero bound on nominal interest rates and by fiscal policy turning increasingly contractionary since 2010. Third, the economic mechanisms by which shocks propagate through the economy (propagation mechanisms) may have changed. This last view would imply that the same shocks that generated little unemployment persistence in the economy of the 1980s may now have much longer-lived effects. We address each of these potential explanations in turn.

With respect to a changing composition of shocks, a common interpretation is that most post-World War II recessions were driven by the Federal Reserve's desire to clamp down on inflation, leading to rapid recoveries once interest rates were loosened, whereas recent recessions have been driven by financial factors that could inherently have longer-lived effects. To assess this explanation, we compare the persistence of unemployment after financial shocks, identified as in Gilchrist and Zakrajsek (2012), to the unemployment persistence after monetary policy shocks, identified as in Romer and Romer (2004), and find no meaningful difference between the two. Thus, our empirical evidence does not support the argument that financial shocks, as a source of recent business cycles, can explain the rising unemployment persistence.

The second explanation is that changing policy responses to business cycles are responsible for the rising unemployment persistence, with the zero bound on interest rates and a turn toward fiscal austerity being the sources of contractionary monetary and fiscal policies respectively in the Great Recession. We construct for each recession the monetary policy deviations from an average response function conditioning on the Fed's real-time beliefs about economic conditions following Romer and Romer (2004), which allows us to quantify the extent to which monetary policy was unusually expansionary or contractionary in recent recessions (including during the zero bound period) relative to the pre-1990 average. We perform a similar exercise for fiscal policy, using changes in the cyclically adjusted federal budget balance as a share of potential GDP to quantify the extent to which fiscal policy was unusually expansionary or contractionary in each of the last three recessions relative to pre-1990 recessions. From these, we construct counterfactual paths of unemployment for each of the last three recessions under the assumption that monetary and fiscal policies had followed their pre-1990 behaviors. We find that monetary and fiscal policies have contributed significantly to the persistence of unemployment in each of the last three recessions, accounting for approximately one-third of the excess persistence observed relative to pre-1990 recessions.

We then turn to the third class of explanations: changes in the propagation mechanisms of the economy. Because the range of factors that can affect the propagation of shocks is vast, we use the four stylized facts from the Great Recession—the missing disinflation, the rise in long-term unemployment, the declining convergence rate in regional labor markets, and the changing cyclicity of disability claims—as guideposts in selecting possible candidate explanations. For example, much of the recent discussion about the missing disinflation has centered on downward wage rigidity, which in a low-inflation environment can hinder the downward adjustment of real wages needed to facilitate the adjustment of labor markets during economic downturns. Proponents of this view point to the fact that since the early 1980s a rising share of workers experienced no annual change in wages. But if downward wage rigidity is to hinder the downward adjustment of wages during a downturn, one would expect to find a larger increase in the incidence of zero wage changes during recent recessions than in the past, yet this is a feature that, as we document, is absent in the data. More broadly, for wage rigidities to be the source of the missing price disinflation, one would expect to see a missing wage disinflation as well, whereas no such pattern can be found in the data. Hence, downward wage

rigidities appear to be an unlikely source of the missing disinflation or the rising persistence in unemployment.

We consider instead three alternative potential sources of changing propagation mechanisms. Inspired by the declining rate of convergence in regional labor markets, the first possibility we consider is the falling rate of labor mobility in the United States. Low mobility can delay the adjustment of regional labor markets (Blanchard and Katz 1992) and could lead to lower quality job matches, thereby potentially inducing firms to defer hiring after recessions in the presence of hiring and firing costs. The second possible explanation is the aging of the workforce. Older workers who lose their jobs tend to be unemployed for longer periods of time than younger workers, in part because older workers may be more resistant to wage cuts, occupational changes, or geographic relocations (GAO 2012). Another mechanism is that older unemployed workers, particularly high-tenured displaced workers, are more likely to have obsolete skills that require retraining, which is particularly costly given their shorter remaining working years. Hence, an aging workforce could be a factor behind the rising persistence of unemployment. The third explanation we consider is a cultural one, inspired by the rising cyclicity of disability claims and larger share of long-term unemployment. Surveys over the last 30 years reveal, for example, that Americans increasingly find it justifiable to claim government benefits for which they do not qualify. These surveys also reveal increasingly cynical interpretations of others' motives and, more broadly, a decline in social trust. These changing social mores could naturally explain rising shares of long-term unemployment and increases in disability claims during downturns as well as more persistence in overall unemployment rates.

These explanations are potentially consistent not only with the time series variation in unemployment persistence but also the earlier cross-country differences in labor market outcomes from the 1980s. For example, in 1981 the U.S. population was unusually mobile and its demographic composition was much more heavily tilted toward the young than in any of the countries in Western Europe. The United States also stood out among developed countries in terms of many of its cultural mores, including its high levels of social trust. The falling U.S. labor mobility rate, the aging of the U.S. population, and the decline in Americans' social trust over the last 30 years therefore all represent a movement toward the characteristics of Western Europe in the early 1980s in the same way that the rising persistence of U.S. unemployment recalls that European experience.

Can the evolution of these characteristics account for the rise in U.S. unemployment persistence? We exploit the cross-sectional experience of the 1980s to assess the strength of correlations between each of these factors and unemployment persistence. Specifically, we build a new data set of regional labor market outcomes across 12 Western European countries, the United States, and Canada from 1970 to 1990, from which we construct regional measures of unemployment persistence. Regional data provide much wider cross-sectional variation than either time series or cross-country analysis. Furthermore, we integrate into our data a wide range of predetermined regional controls circa 1970 designed to capture in a reduced form the many other factors that could be related to regional persistence of unemployment. These controls include, among others, income per capita levels, years of schooling, female employment shares, urbanization rates, and the sectoral composition of employment. Finally, we construct regional measures of our key explanatory variables: labor mobility, age structure of the population, and levels of trust. The combination of detailed regional controls with country fixed effects that capture the effects of aggregate policies and institutions therefore allows us to isolate the partial correlations of mobility, demographics, and trust with regional unemployment persistence. While it is difficult to establish a direct causality from these factors to unemployment persistence, the strength of the correlation between them, combined with the evolution of these factors over time, can be informative about their predictive power for unemployment persistence.

We find no evidence that higher mobility is associated with lower unemployment persistence once controls (regional or aggregate) for other factors are included. This result is robust to a variety of checks and suggests that the rising persistence in unemployment cannot be explained through declining labor mobility. Second, we find that, all else being equal, a higher share of older workers relative to younger workers is generally associated with less persistence, the opposite of the effect needed to explain the rising persistence of unemployment through demographic effects. Third, we document a systematic negative correlation between regional levels of trust and unemployment persistence, even after regional controls and country fixed effects are included. Only this last result moves in the direction needed to account for the rise in unemployment persistence. We can then quantify how much of the rise in unemployment persistence could be accounted for through these propagation mechanisms. We find that while the decline in trust could account for all of the rise in

unemployment persistence observed since the 1980s, once we also incorporate the aging of the population the latter yields even larger predicted declines in persistence, so that the net effect from all three explanations combined predicts a *reduction* in unemployment persistence since the 1980s rather than an increase.

We interpret our results as supporting the idea that weak recoveries are likely to remain a prominent feature of future U.S. business cycles, since the high persistence in unemployment after the Great Recession appears to be a continuation of previous trends rather than something accounted for by special factors such as the financial crisis. While we cannot pin down a unique factor that explains the rising propagation of shocks in the United States, this feature of the data may require some rethinking of optimal countercyclical policies. In particular, if we can expect future business cycles to be much more long-lived events than the V-shaped recessions that characterized much of the post–World War II period, then perhaps discretionary fiscal policy responses should focus more on longer-lived investment projects than the transitory transfer payments that have become common in recent stimulus packages.

The structure of the paper is as follows. Section I documents the rising persistence of U.S. unemployment. Section II investigates whether this rising persistence can be explained—in whole or in part—by financial shocks or policy responses. Section III considers possible propagation mechanisms, while section IV quantifies the effects of our suggested mechanisms. Section V concludes.

I. The Rising Persistence of U.S. Unemployment

A useful starting point for characterizing the changing persistence of U.S. unemployment is to compare the Great Recession with the twin recessions of the early 1980s. For the United States, the twin recessions—induced by the Volcker disinflation policy—generated the largest increases in unemployment since the Great Depression, led to a persistent decline in inflation, and were ultimately followed by the Great Moderation. In contrast, many European countries experienced a similar increase in unemployment over the course of 1980s, but whereas the rise in unemployment in the United States was completely reversed by the mid-1980s, unemployment in many European countries persisted at high levels beyond the 1980s. This high persistence in unemployment, referred to as *hysteresis* among economists since Blanchard and Summers (1986) or as “Eurosclerosis” among

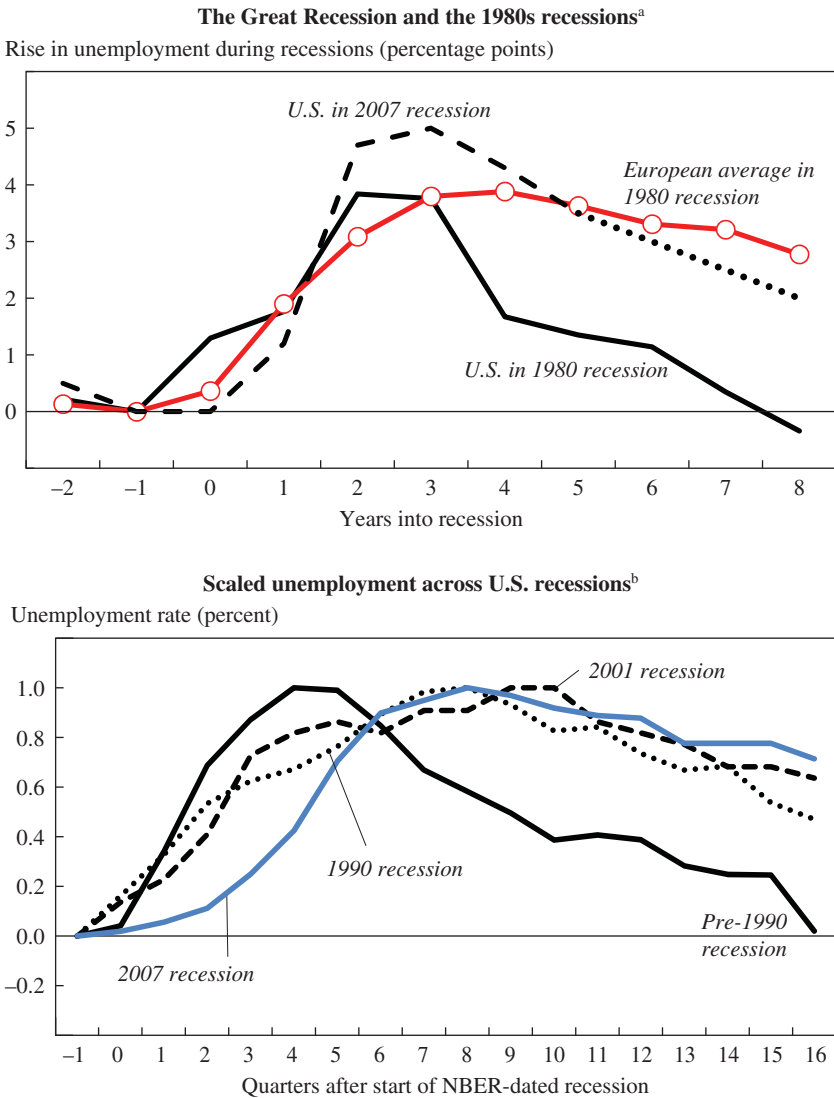
a broader audience, has been a scourge for many European policymakers ever since.¹

In figure 1, the top panel plots unemployment rates in the United States during the Volcker recession and thereafter relative to 1979. The unemployment rate rose nearly four percentage points by 1982 and 1983 but then declined sharply thereafter. The figure also plots the average rise in unemployment rates for Western European countries. While the rise in unemployment there was of approximately the same size and speed as in the United States, unemployment in Europe stayed persistently high thereafter. The same panel also plots the unemployment rate in the United States during the Great Recession, showing that the initial rise in unemployment exceeded that of the 1980s recession in both the United States and Western Europe by about one percentage point. But the more striking difference is the absence of a strong decline in unemployment in the United States since the peak in unemployment in 2010. This asymmetry in the response of U.S. unemployment during the Great Recession is in sharp contrast with the experience of the 1980s. From 2013 onwards, the figure also shows the projected path of unemployment from the Survey of Professional Forecasters, which implies only a very gradual decline in unemployment with much more persistence in unemployment than occurred in the early 1980s. Furthermore, some of the recent decline in unemployment has been driven not by rising employment but rather by declining labor force participation. Broader measures of employment in the United States, like employment-to-population ratios, suggest even less improvement in labor markets since the height of the recession (for example, see Erceg and Levin 2013). The U.S. recovery from the Great Recession is therefore following a path approximately midway between the recovery path of the United States and that of Western Europe over the 1980s.

The labor market recovery following the Great Recession appears anemic not only compared to the recessions of the early 1980s but also compared to most recessions of the post–World War II era. The bottom panel of figure 1 plots the average path of the unemployment gap (the unemployment rate minus the CBO’s estimate of the long-run natural rate of unemployment) across U.S. recessions from the 1948 recession up to and including the 1981–82 recession, relative to the unemployment gap prior to the start of each recession and normalized by the maximum rise in the gap during

1. While the term *hysteresis* is sometimes used narrowly to refer only to cases in which transitory shocks lead to permanent changes in unemployment rates, we adopt the broader interpretation of the word as characterizing long-lived changes in unemployment that need not necessarily be permanent.

Figure 1. The Changing Persistence of U.S. Unemployment



Source: Authors' calculations.

a. The top panel plots annual unemployment rates for Europe (average for 12 West European countries in our sample) and the United States relative to their levels in 1979 and annual unemployment rates for the United States in the Great Recession relative to the 2007 levels. Values from 2013–15 are forecasts from the Survey of Professional Forecasters.

b. The bottom panel plots unemployment rates across U.S. recessions relative to their level in the quarter prior to the start of the recession and normalized by the peak increase in unemployment in each recession. The pre-1990 average is the average value across all recessions since 1948 and prior to 1990, dropping all periods following one recession that become part of the subsequent recession.

each recession (to control for the size of recessions).² In recessions until the 1980s, on average, unemployment peaked approximately one year after the start of the recession and had returned to the natural rate within four years. In contrast, the peak in the unemployment rate during the Great Recession did not occur until two years after the start. Moreover, in the two years following the peak rise in unemployment during the Great Recession, only about one-fourth of its rise was undone. Thus, the persistence in unemployment following the Great Recession, measured both as the length of time before the peak and as the pace of the decline in unemployment after the peak, contrasts sharply with the historical pattern of recessions.

However, as the bottom panel of figure 1 also illustrates, the labor market recovery following the Great Recession looks much less exceptional when compared to the 1990 and 2001 recessions. In each of those recessions, unemployment gaps rose for around two years, and the subsequent declines were much slower than the average declines following the pre-1990 recession. There appears to be a trend toward increasingly inertial labor market recoveries over time, with the 2001 recovery being more inertial than the 1990 recovery and the Great Recession recovery in turn being more inertial than the 2001 recovery. Thus, while the labor market recovery following the Great Recession stands out relative to those recessions before 1990, it is not much slower than the “jobless recoveries” of 1990 and 2001. This suggests that there is likely a common cause to the high persistence in unemployment in all three recessions.³

II. Can Financial Shocks and Policy Responses Account for Rising Unemployment Persistence?

In this section, we consider two prominent explanations for the rise in U.S. unemployment persistence since the 1980s. The first is that recent business cycles have been driven by financial factors and that financial crises are, in

2. When recovery from one recession is interrupted by a subsequent recession, we drop all values as of the start of the new recession in constructing the average path of pre-1990 recessions. But almost identical results obtain if we drop the interrupted recessions.

3. In online appendix figure C.1, we document that rolling time-varying estimates of the persistence of the unemployment rate confirm a rise in its persistence over time. Specifically, starting in 1980Q1 and iterating forward until 2013Q2, we estimate an AR(2) model for the quarterly unemployment rate over the previous 30 years and report the estimated sum of the autoregressive coefficients in the online appendix figure C.1. The point estimates range from a low of approximately 0.92 in the sample ending in 1980Q1 to a high of 0.97 in the sample ending in 2011Q3. (Online appendixes for papers in this volume may be found at The Brookings Papers website, www.brookings.edu/bpea, under “Past Editions.”)

general, associated with more persistent declines in economic activity. The second explanation focuses on the response of policymakers to business cycles, such as the inability of monetary policymakers to reduce interest rates in the face of the zero bound.

II.A. The Role of Financial Shocks

A commonly suggested reason for the delayed recovery during the Great Recession is that financial crises such as that of 2007–09 have historically been associated with more persistent economic downturns than recessions not associated with financial crises (Reinhart and Rogoff 2009; Jorda, Schularick, and Taylor 2011). While the persistence of economic downturns following financial crises could be higher than normal because financial shocks themselves lead to more persistent downturns than other shocks, it could also reflect other underlying factors responsible for both the appearance of financial shocks and more persistent economic downturns. To assess the causal nature of this claim, we compare the persistence of unemployment after identified financial shocks to the unemployment persistence following monetary policy shocks. We focus on monetary policy shocks in particular because these are one of the leading candidates for explaining post–World War II recessions up to and including the Volcker recessions (Romer and Romer 2004; Coibion 2012).

To identify monetary policy shocks, we replicate Romer and Romer (2004) with the same sample and data and regress changes in the target federal funds rate decided upon at Federal Open Market Committee (FOMC) meetings on real-time information available to FOMC members through the Greenbook forecasts. The estimated residuals are defined as monetary policy shocks. Impulse responses of unemployment to monetary policy shocks are constructed by regressing the unemployment rate on 24 monthly lags of itself and 36 monthly lags of the monetary policy shock, which we then convert to quarterly frequency.

For financial shocks, we focus on the credit spread shocks identified in Gilchrist and Zakrajsek (2012).⁴ They construct a measure of excess corporate bond premiums from corporate spreads, which has significant predictive power for economic activity, and then identify exogenous inno-

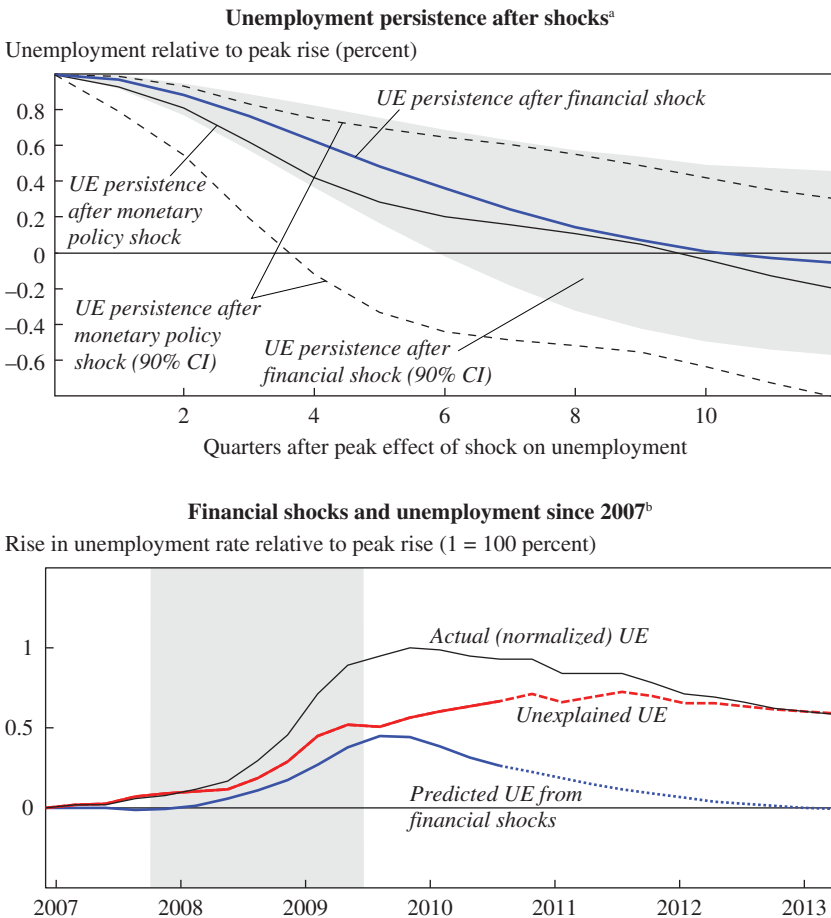
4. We focus on this measure of financial shocks for several reasons. First, it captures rising bond spreads associated with credit crunches. Second, changes in the excess bond premium have strong predictive power for macroeconomic aggregates. Third, responses to these shocks are consistent with their being a potential source of recessions, with a rise in the bond premium being followed by declines in consumption, investment, output, inflation, and interest rates.

vations to excess bond premiums from a vector autoregression (VAR). The specific VAR they use includes nonfinancial variables, their measure of the bond premium, and other financial variables. The identifying assumption is that nonfinancial variables do not respond on impact to a shock to the excess bond premium, whereas other financial variables do. We replace real GDP in the VAR with the unemployment rate, but otherwise we follow Gilchrist and Zakrajsek in running a VAR (2) with data from 1973Q1 until 2010Q3.

As a first step to measuring the conditional persistence of unemployment, we estimate the half-life of unemployment after each shock, defined as the number of quarters until unemployment reaches half its maximum level. Specifically, we estimate the impulse response of unemployment to each kind of shock along with standard errors of coefficients. We then repeatedly draw from the distribution of parameters and construct impulse responses for each draw. For each set of impulse responses, we find the number of quarters after the shock that it takes for the unemployment rate to fall halfway below its peak level. Monetary policy shocks, on average, require 13 quarters before unemployment falls below half of its peak level, while financial shocks take 12 quarters on average. The 90 percent confidence intervals are [11,20] and [8,20] quarters for monetary and financial shocks, respectively. Thus, the two sets of shocks yield very similar results for the overall persistence of unemployment.

We also construct a narrower measure of persistence that focuses on the rate of decline of unemployment after its peak. Specifically, we estimate the impulse response of unemployment to each kind of shock along with standard errors of coefficients. We then repeatedly draw from the distribution of parameters and construct impulse responses for each draw. For each set of impulse responses, we normalize the dynamic response of unemployment by the peak rise in unemployment in that response and store the resulting normalized path of unemployment starting with the period in which unemployment peaks. This yields a distribution of responses after the peak rise in unemployment for each of the two shocks. The top panel in figure 2 plots the median response of unemployment after the peak for each of the two shocks as well as the associated 90 percent confidence intervals. The rate at which unemployment declines after the two shocks is broadly similar: both median responses point to unemployment returning to normal approximately 10 quarters after the period in which unemployment peaks. The confidence intervals largely overlap, so there is little statistical evidence that an average financial shock leads to more persistence in unemployment rates than monetary policy shocks.

Figure 2. Financial Shocks and Unemployment Persistence



Source: Authors' calculations.

a. The top panel plots the estimated median and 90 percent confidence interval (CI) of unemployment responses to financial and monetary policy shocks in periods after the peak rise in unemployment.

b. The bottom panel plots the estimated contributions of financial shocks to unemployment in the 2007 recession as well as the component of unemployment due to other factors.

Even if a typical financial shock does not lead to more persistence in unemployment than monetary policy shocks, one could still explain high unemployment persistence over recent recessions through financial shocks if the latter were themselves unusually pervasive during each period. To assess this possibility, we construct the predicted path of unemployment during the 2007 recession from financial shocks starting in the year prior to each recession. We then plot in the bottom panel of figure 2 the path of

unemployment (relative to the level in the quarter prior to the start of the recession and normalized by the size of the peak increase in unemployment), the predicted path from financial shocks, and the unexplained component of unemployment. The credit bust accounts for half of the initial rise in unemployment, but its effects should have dissipated rapidly, so the persistence of unemployment is again unexplained.⁵

In short, our results suggest that the effects of financial shocks on unemployment are no more persistent than the effects of monetary policy shocks, and therefore it is not sufficient to appeal to financial factors to account for the excess unemployment persistence of recent recessions.

II.B. The Role of Monetary and Fiscal Policies

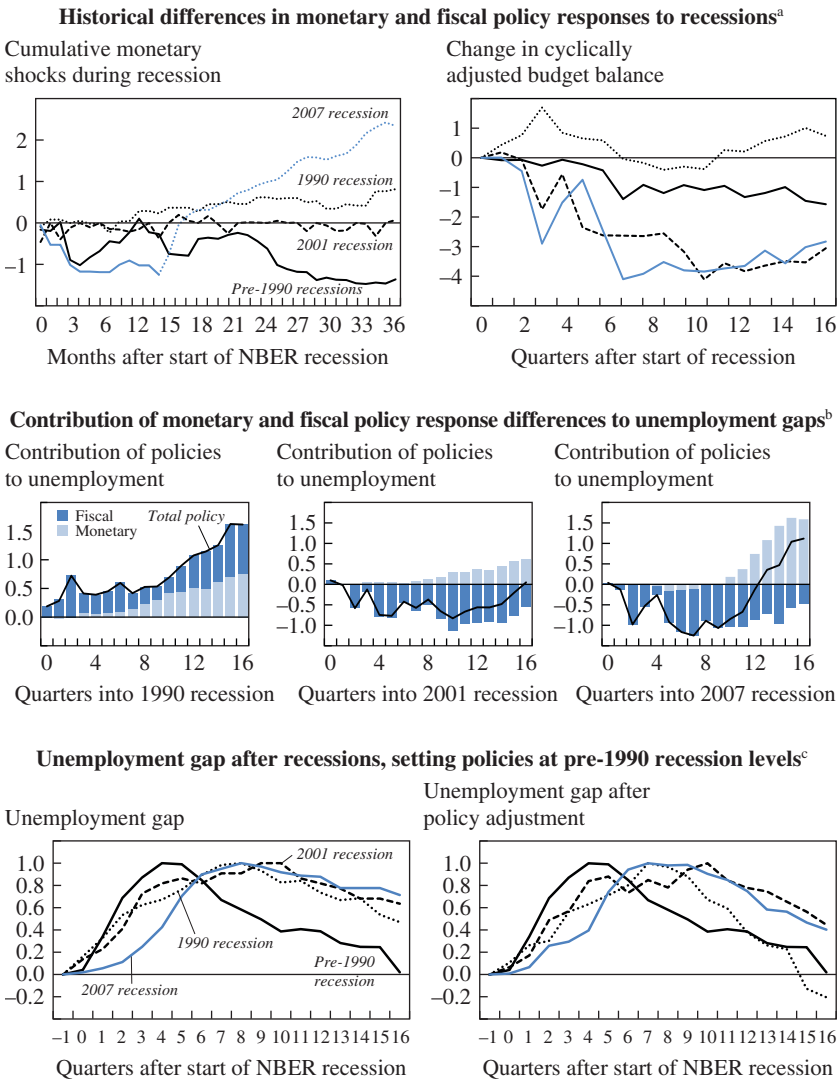
Since financial shocks cannot account for the rising persistence of unemployment experienced in the United States since the 1980s, we turn to a second factor often suggested in academic and policy discussions: differences in policy responses. We focus on both monetary and fiscal policies. For monetary policy, we again follow Romer and Romer (2004) after extending their data set to December 2012.⁶

We estimate the same reaction function as Romer and Romer estimated until December 2008 when the zero bound on interest rates became binding. We interpret the reaction function as capturing the average or systematic response of the Federal Reserve to real-time economic conditions, and we interpret the residuals as the innovations to monetary policy. We then construct the cumulative sum of residuals for each recession since 1969 as a measure of unusual monetary policy actions taken in each recession. We plot the resulting series for the 1990, 2001, and 2007 recessions as well as the average across pre-1990 recessions in the left-hand graph of the top panel in figure 3.

5. In our working paper (Coibion, Gorodnichenko, and Koustas 2013), we documented similar results for the 1990 and 2001 recessions. While one concern with focusing on Gilchrist-Zakrajsek shocks is that we could be missing other financial shocks that have longer-lived effects, we also find no excess unemployment persistence after shocks to excess stock market returns. Another possibility is the deleveraging hypothesis, but the 2001 recession was followed by a leveraging up of households of the same order as the deleveraging after the Great Recession. Hence, deleveraging could at best account for dynamics after the Great Recession but would make the unemployment persistence in the 2001 recession even more puzzling. Nonetheless, we cannot rule out that other unspecified financial shocks could yield more persistent effects on unemployment.

6. Forecasts from January 1997 to December 2007 are from the Greenbooks. However, because more recent Greenbooks have not yet been released by the Federal Reserve Board of Governors, we use forecasts from the Blue Chip Survey of Forecasters starting in 2008.

Figure 3. The Contribution of Policy Differences to Unemployment Outcomes



Source: Authors' calculations.

a. Upper panel shows cumulative shocks to monetary policy (identified as in Romer and Romer [2004]) and to fiscal policy (identified as changes in cyclically adjusted fiscal balance in percent of potential output; source: CBO).

b. Middle panel shows the contribution of monetary and fiscal policy to the unemployment rate. The contribution of monetary policy is calculated from the sequences of shocks to monetary policy and estimated impulse response function of unemployment to monetary policy shocks. The contribution of fiscal policy is calculated with the estimated Okun's law (sample period 1949:2013) and output multiplier of government spending assumed to be equal to 1.5.

c. Lower panel shows actual (left) and counterfactual (right) dynamics of unemployment gaps. The counterfactual shows the dynamics when the contribution of monetary and fiscal policies is eliminated. In both panels, y axis is scaled by the maximum change in the unemployment gap, with 1 = 100 percent. Online appendix figure C.2 shows that this result is qualitatively similar using measures of total government employment rather than federal budget balances to compare historical responses of fiscal policy.

Monetary policy shocks averaged out in the first 12 months of the pre-1990 recessions, but thereafter interest rates progressively fell by two percentage points more than would have been expected given economic conditions. While we observe a similar pattern of shocks averaging out to zero over the first 12 months of the 1990 and 2001 recessions, differences arise thereafter. In the 2001 recession, monetary policy shocks continue to average around zero, such that monetary policy ultimately is two percentage points more contractionary than during the average pre-1990 recession. The 1990 recession is similar, except that monetary policy averages an extra half-percentage point above and beyond what would have been expected from economic conditions. Hence, monetary policy in both the 1990 and 2001 recessions was more contractionary than in previous recessions.

For the Great Recession, we observe a more rapid decline in interest rates over the first 12 months than would be expected from economic conditions, leading to a cumulative decline in interest rates of one percentage point below the experience in other recessions. Thus, through December 2008, when the zero bound became binding, monetary policy was unusually expansionary. To construct monetary policy shocks after 2008, we construct a predicted path of interest rates for FOMC meetings starting in January 2009 under the assumption that the Federal Reserve would have followed the estimated reaction function in subsequent periods absent the zero lower bound (ZLB). From this sequence of predicted rates, we define the shocks each month as the negative of the federal funds rate changes that would have been decided upon had monetary policymakers been free to lower interest rates further. This leads to the sequence of shocks plotted in the figure, which point toward a dramatic and continuing sequence of contractionary monetary policy shocks as the economy continued to deteriorate but monetary policymakers were unable to implement traditional monetary policy responses. Our estimates imply that the inability to lower interest rates starting in 2009 led to a sequence of monetary policy shocks summing to over 3 percentage points over a two-year period.

To assess differences in the stance of fiscal policy, we focus on changes in the cyclically adjusted federal budget balance normalized by potential GDP. Cyclically adjusted balances are needed to control for automatic changes in spending and revenues as business cycle conditions change. We normalize these measures by potential GDP (rather than actual GDP) because this normalization is largely insensitive to business cycle conditions. We plot the changes in cyclically adjusted balances relative to their levels in the quarter prior to the start of each recession in the right-hand

graph of the top panel in figure 3. As with monetary policy measures, we present these responses for the 1990, 2001, and 2007 recessions as well as an average across pre-1990 recessions. The average pre-1990 recession is characterized by little change in the cyclically adjusted budget balance over the first year of the recession, after which a 1-to-1.5-percentage-point decline in the budget balance occurs and is maintained over the subsequent two years. During the 1990 recession, there is a sharp increase in the budget balance over the first few quarters, which is sustained relative to the pre-1990 recessions for the next two years. Ten quarters after the start of the recession, there is a second and sustained increase in the budget balance, leading to a 3-percentage-point gap between the budget balances after the 1990 recession and pre-1990 recessions. Hence, fiscal policy was distinctly more contractionary in the 1990 recession than the average over previous recessions. In contrast, both the 2001 and 2007 recessions were followed by significant increases in cyclically adjusted deficits, approximately 4 percentage points of potential GDP after two years. In both cases, these deficits were gradually reduced over the subsequent two years.

To quantify how policy responses may have contributed to rising unemployment persistence, we next construct counterfactual paths of unemployment in which monetary and fiscal policies are set equal to the pre-1990 averages. For monetary policy, we estimate the impulse response of unemployment to monetary policy shocks (using the extended sample of shocks from 1969–2008) as in Romer and Romer (2004). For each of the 1990, 2001, and 2007 recessions, we use these estimated impulse responses to predict the path of unemployment coming from the difference in monetary policy shocks between each of these recessions and the pre-1990 recessions.

These contributions are plotted in the middle panel of figure 3 for each of the three recent recessions. For example, the fact that monetary policy was more contractionary during the 1990 recession than during pre-1990 recessions points to a cumulative increase in unemployment over the course of the recession and beyond, pushing the unemployment rate 0.70 percentage points higher four years after the start of the recession than it would have been had policy followed the same path as in pre-1990 recessions. Monetary policy similarly contributed to higher unemployment after the 2001 recession. The zero lower bound is estimated to have contributed much more to unemployment in the Great Recession, adding approximately 1.5 percentage points to the unemployment rate after four years. This increase in unemployment is estimated not to have begun until two years after the recession started, because the zero bound was not

reached until one year into the recession and monetary policy shocks have delayed effects on economic activity. Thus, for each recession, monetary policy contributed to the persistence of unemployment relative to pre-1990 recessions by raising unemployment rates two years after the start of each recession.

To quantify the contribution of fiscal policy, we assume instantaneous effects of changes in the cyclically adjusted fiscal balance on output growth, with a dollar multiplier of 1.5. This follows Auerbach and Gorodnichenko (2012), who document that government spending multipliers are significantly higher in downturns than during expansions. A multiplier of 1.5 is also in line with other studies, such as Romer and Bernstein (2009), Blanchard and Leigh (2013), and Nakamura and Steinsson (2011). The results are presented in the middle panel of figure 3. For the 1990 recession, the fact that the cyclically adjusted balance was approximately 1 percentage point higher (in terms of potential GDP) than in pre-1990 recessions leads to a prediction of a 0.5-percentage-point higher unemployment rate over the first two years, before rising 1 percentage point higher as the cyclically adjusted balance increases further relative to the pre-1990 recession. For the 2001 recession, the fact that the cyclically adjusted balance is significantly more expansionary than the pre-1990 average yields an unemployment rate lower by almost a full percentage point over much of the four-year period following the start of the recession. In the case of the Great Recession, the large decline in the cyclically adjusted balance over the first two years of the recession is predicted to have lowered unemployment by nearly a full percentage point, but this is almost completely reversed within the next two years as the cyclically adjusted balance converges back toward the pre-1990 average. Thus, as with monetary policy, fiscal policy likely contributed to higher persistence in unemployment by raising unemployment rates in later years of each recession episode relative to pre-1990 fiscal policies.

To quantify the contribution of these effects on unemployment persistence, we construct counterfactual paths of unemployment for each recession, assuming that both monetary and fiscal policies had followed their pre-1990 averages. We present in the lower panel of figure 3 the original dynamics of unemployment (normalized by predicted peak levels) as well as the counterfactual dynamics of unemployment (normalized by counterfactual peak levels) in which monetary and fiscal policies are held at their pre-1990 average responses. For the 1990 recession, monetary and fiscal policies can account for much of the excess persistence in unemployment. Four years after the recession, for example, unemployment with pre-1990

fiscal and monetary policies is predicted to have been even lower than in pre-1990 recessions. However, adjusting for fiscal and monetary policies cannot account for the delayed peak in the unemployment rate. The latter also obtains for both the 2001 and 2007 recessions. At the same time, the rate of decline in unemployment after the peak is significantly faster once one accounts for the differences in monetary and fiscal policies: in each case, unemployment four years after the start of the recession is at 40 percent of the peak level (after controlling for policy differences), whereas actual unemployment rates only fell to approximately 60 percent and 70 percent of peak levels in the 2001 and 2007 recessions, respectively. Thus, unusually contractionary monetary and fiscal policies in the years following the official ends of each recession appear to have played an important role in generating lackluster recoveries in employment, explaining approximately one-third of the excess unemployment persistence.

III. Have Propagation Mechanisms Changed over Time?

Given that neither financial factors nor policy responses can fully account for the rise in unemployment persistence, we turn to the factors that determine how long-lived the effects of an economic shock of a given size are, that is, propagation mechanisms. We present four additional facts about how the U.S. economy has changed in recent decades to discipline our choice of potential propagation mechanisms: the declining sensitivity of inflation to real economic conditions, the rising importance of long-term unemployment, the decline in regional labor market convergence, and the changing cyclical nature of disability claims. One potential explanation for these facts is downward wage rigidity, but we provide new evidence that this explanation is unlikely to lie at the root of these facts. So we suggest and document other mechanisms that might account for these facts and the rising persistence of unemployment.

III.A. Other Changing Features of the U.S. Economy

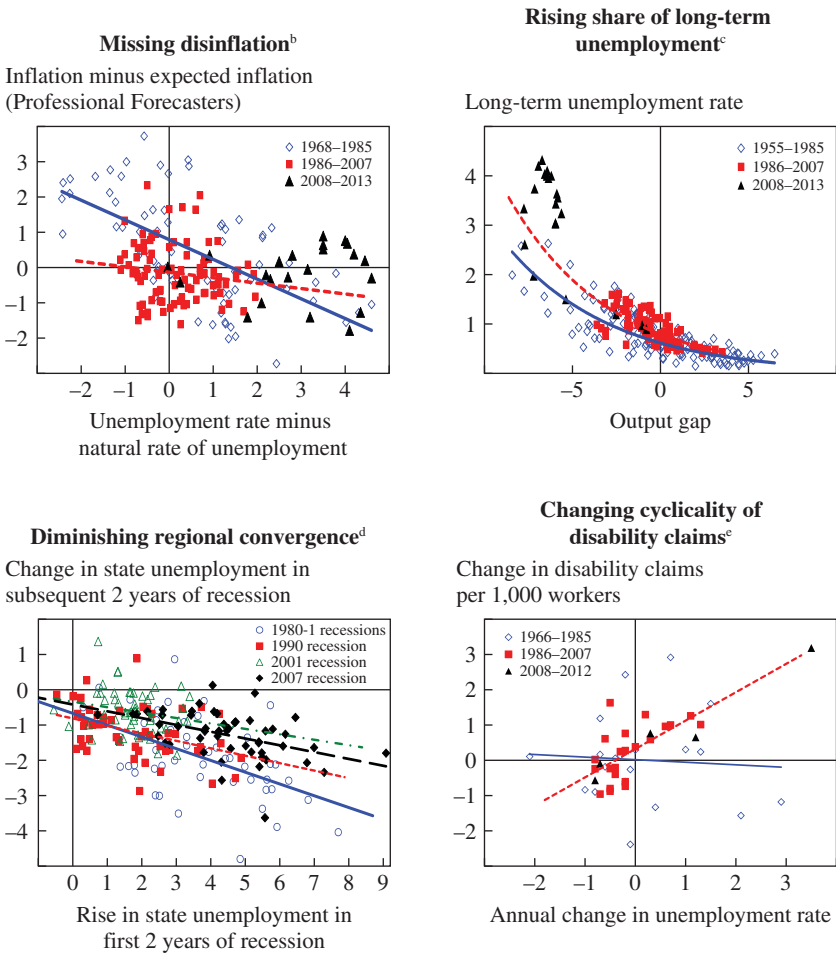
We consider four characteristics of the Great Recession that have received a lot of attention. In each case, we show that these features are not unique to the Great Recession but rather are a continuation of ongoing trends visible in the 1990 and 2001 recessions as well. These specific characteristics of the Great Recession are also notable because they recall stylized features of the European unemployment experience of the early 1980s, at least qualitatively, and are significant departures from U.S. behavior over the same period.

First we consider “missing disinflation,” that is, the fact that the large output gaps of the Great Recession should have led to much larger declines in inflation than what was actually observed (Ball and Mazumder 2011). We illustrate this in the upper-left panel in figure 4 using an expectations-augmented Phillips curve. We present a scatter plot of unemployment gaps (deviations from a Congressional Budget Office [CBO] estimate of long-term unemployment) against inflation net of inflation expectations (median forecast of inflation in the subsequent quarter from the Survey of Professional Forecasters). Many of the observations since 2008 have displayed unusually high levels of inflation relative to expectations given the high levels of unemployment. We also plot the average slope of the relationship between inflation net of expectations and unemployment for two sub-periods: up to the mid-1980s and the mid-1980s until immediately before the Great Recession. This suggests a decline in the slope of the Phillips curve over time, with little missing disinflation left during the Great Recession using the estimated slope of the Phillips curve from the mid-1980s on. Therefore, the missing disinflation need not be interpreted as a puzzle unique to the Great Recession, but rather should be interpreted as part of a broader trend toward a declining sensitivity of U.S. inflation toward real economic activity, a trend that began in the early 1980s. Furthermore, missing disinflation was also noted as a puzzling feature of the Western European experience of the early 1980s (Blanchard and Summers 1986), which suggests that there may be a systematic link between unemployment persistence and how strongly inflation responds to changes in economic activity.

A second unusual feature of the Great Recession is the rise in long-term unemployment, defined as 27 weeks of unemployment or more. The share of long-term unemployed among all unemployed was less than 20 percent in the first quarter of 2007 but climbed as high as 45 percent in the second quarter of 2010. This feature recalls the rise of long-term unemployment in many Western European countries in the early 1980s (Machin and Manning 1999). In contrast, the share of the long-term unemployed in the United States during the early 1980s peaked at only 26 percent in the second quarter of 1983. While the relationship between output and total unemployment has remained remarkably stable over time (Ball, Leigh, and Loungani 2012), this has not been the case with long-term unemployment.

In the upper-right panel in figure 4, we plot the relationships between the output gap and the long-term rate of unemployment, which illustrates how the Great Recession stands out in terms of how high long-term unemployment rates have been. But this pattern was already evident in the

Figure 4. Other Changing Features of the U.S. Economy^a



Source: Authors' calculations.

a. In all four panels here, the long-term unemployment rate and the output gap are measured in percent, the change in disability rate is the difference in share of applicants per 1,000 workers, and all other axes are in percentage points.

b. The upper-left panel plots quarterly unemployment rates against the difference between quarterly changes in GDP deflator inflation and expected inflation from the Survey of Professional Forecasters. See online appendix figure C.3 for rolling regression estimates of the slope of the Phillips curve.

c. Upper-right panel plots quarterly deviations of output from the CBO measure of potential against long-term unemployment rates.

d. Lower-left panel plots, for each state and each recession, the rise in unemployment rate from the year prior to the start of the recession to two years into the recession (x-axis) versus the change in unemployment rate over subsequent two years (y-axis). Lines show average relationships between the two for each recession.

e. Lower-right panel plots annual changes in unemployment rates against annual changes in the number of disability claims per 1,000 insured workers from 1968 to 2012. See section III.A for details.

1990 and 2001 recessions, with the slope of the relationship between long-term unemployment and output gaps having become much steeper even prior to the Great Recession. As with the missing disinflation, the high long-term unemployment during the Great Recession is therefore much less of a puzzle if one compares it to the experience since the early 1980s. This finding again suggests that the unusual features of the Great Recession appear to be a continuation of longer-term trends in U.S. business cycle behavior but made more visible by the sheer size of the recent recession.

A third unusual property of the Great Recession lies in the behavior of regional U.S. labor markets. As noted in Blanchard and Katz (1992), changes in unemployment across states in the years after 1982 were negatively correlated with the magnitude of the rise in unemployment over the prior years. Blanchard and Katz argue that the high degree of labor mobility in the U.S. led workers to move from high unemployment to low unemployment regions, thereby facilitating the adjustment of the U.S. labor market. In contrast, Decressin and Fatas (1995) document that the adjustment of European regional labor markets to regional shocks was not characterized by the same degree of regional mobility, leading to more persistence in unemployment rates. The lower-left panel of figure 4 illustrates the relationship between the initial rise in unemployment rates from 1979 to 1982 and the subsequent decline in rates between 1984 and 1982 for U.S. states emphasized by Blanchard and Katz (1992).⁷ Like them, we find a strong negative correlation between the initial rise in unemployment and the subsequent change in unemployment. But when we replicate this analysis for the period during the Great Recession, using the change in unemployment from 2007 to 2010 relative to the change between 2012 and 2010, we find a much weaker relationship between the two (the difference in slopes across samples is significant at the 10 percent level). The slope of the relationship is cut in half after the mid-1980s: regional convergence in labor markets during the Great Recession is significantly slower than what was experienced in traditional recessions. But this characteristic is not unique to the Great Recession and was already visible in previous recessions. To show this, we also plot the equivalent state-level increases in the first two years of the 1990 and 2001 recessions against the subsequent declines in unemployment. In both recessions, the slope of the relationship is indistinguishable from that of the Great Recession.

7. Wisconsin is a larger outlier because its unemployment rate rose sharply between 1982 and 1983 rather than peaking in 1982, so it is not included in the figure for the 1982 recession.

Hence, the slow convergence in regional labor markets during the Great Recession appears to be in line with that experienced in the previous two recessions.

Finally, we consider the behavior of disability claims over the cycle. Social Security disability claims spiked during the Great Recession, with the ratio of claimants to insured workers rising by 30 percent from 2007 to 2010. This is particularly striking because, during the early 1980s, disability claims also went up dramatically in many Western European countries, whereas no such rise occurred in the United States (Emerson and Dramais 1988). Although the absence of a rise in U.S. disability claims in the early 1980s partly reflected policy changes designed to reduce the incidence of disability claims (Autor and Duggan 2003), there was no strong relationship between disability claims per 1,000 insured workers and unemployment prior to these policy reforms either.

For example, between 1972 and 1973 disability claims per 1,000 insured workers rose 10 percent while the unemployment rate fell from 5.6 to 4.9 percent. Then, during the subsequent recession, as the unemployment rate spiked from 5.6 percent in 1974 to 8.5 percent in 1975, disability applications per 1,000 insured workers fell 7 percent. The lower-right panel of figure 4 plots all the annual changes in disability claims per 1,000 insured workers against annual changes in unemployment rates from 1965 to 1985 and confirms that there was little cyclical in disability claims during this early period, in sharp contrast to the European experience. However, the rate of disability claims per insured workers has been strongly counter-cyclical since then. The lower-right panel of figure 4 illustrates a strong positive correlation between annual changes in disability claims and annual changes in unemployment rates from 1986 to 2007, as noted in Duggan and Imberman (2009). Annual changes in disability since 2008 have been almost completely predictable, given annual changes in unemployment over this period, when using the relationship between the two from 1986 to 2007.⁸ Hence, the experience of the Great Recession appears to have very closely followed that of the 1990 and 2001 recessions, but for the scale of the crisis.

8. The Social Security Disability Benefits Reform Act of 1984 expanded the types of ailments which allowed one to qualify for disability and shifted the criteria for eligibility to broader measures of a person's ability to work, thereby making it easier to qualify for disability (CBO 2012). One possibility is that workers who experience long-term unemployment have been increasingly filing claims for disability since 1984 upon the expiration of their unemployment benefits. Mueller, Rothstein, and von Wachter (2013), however, find no link between the expiration of extended unemployment benefits and disability claims.

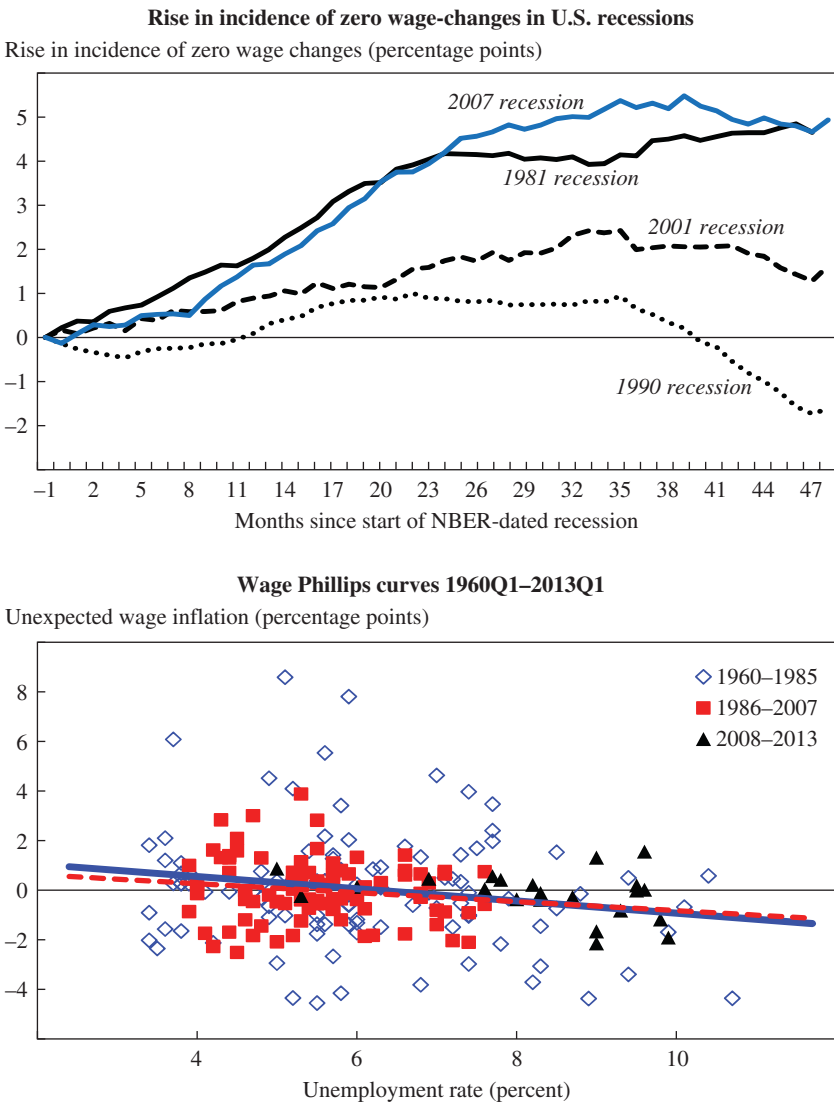
Jointly, these four changing properties of the U.S. economy conform in timing to the rising unemployment persistence documented in section I. Each also represents a growing qualitative similarity with the West European experience of the 1980s, in which very high levels of unemployment persistence were a defining characteristic. This gradual evolution in some of the properties of the U.S. business cycle suggests that the underlying explanation is likely to be rooted in slow-moving fundamental factors driving up the propagation of economic shocks rather than in unique forces at work during the Great Recession. To explore this possibility, we turn next to review some of the fundamental factors that might account for this evolution.

III.B. The Roles of Downward Wage Rigidity and Low Inflation

One potential source of additional propagation has been extensively discussed in conjunction with the missing disinflation, namely a combination of downward wage rigidities and low inflation rates. At least since the work of Tobin (1972), economists have recognized that wages tend to be downwardly rigid, as employees chafe at nominal wage declines, and that in a low inflation environment downward wage rigidity can prevent declines in real wages from occurring during downturns. Given that the share of workers receiving no change in wages reached higher than 16 percent in 2011—more than twice the share in 1981—and that inflation rates have come down significantly since the early 1980s, this interpretation has not surprisingly become a leading explanation for the missing disinflation (e.g., Daly, Hobijn, and Lucking 2012).

However, one can make at least two counterarguments. First, if an economic downturn calls for a decline in real wages, the severity of downward wage rigidity should be captured by the change in the incidence of zero wage changes across the population, capturing the extensive margin of workers whose nominal wages remain uncut. Using the zero wage change dataset constructed by Daly, Hobijn, and Lucking (2012), we quantify the changes in this incidence for the 1990, 2001, and 2007 recessions relative to their initial levels and compare them to the changes in the incidence of zero wage changes in the recession starting in July 1981 (the only pre-1990 recession covered in the aforementioned data set). As illustrated in the upper panel of figure 5, the rise in the incidence of zero wage changes during the Great Recession exceeds that of the 2001 and 1990 recessions but is in line with what was observed during the 1981 recession. This pattern is inconsistent with the explanation that downward wage rigidity has been playing a significantly larger role in slowing the adjustment of wages in recent recessions.

Figure 5. Evidence on Wage Rigidities



Source: Authors' calculations.

Notes: Upper panel plots the change in incidence of zero wage changes during the last four U.S. recessions relative to their level in the month prior to the start of each recession. Lower panel plots a scatter of quarterly unemployment rates against quarterly unexpected wage inflation. The wage series is the average hourly earnings of production and nonsupervisory employees in manufacturing. Trend lines for 1960Q1–1985Q4 and 1986Q1–2007Q4 are shown as the solid and dotted lines, respectively.

A related issue is that if downward wage rigidity were preventing wages from falling as rapidly as in the past given economic conditions in the Great Recession, then one would expect to encounter the puzzle of missing wage disinflation just as one has encountered the puzzle of missing price disinflation. But while the latter has already been well documented, there has been much less emphasis placed on the former. To determine whether wage dynamics have been unusual in recent years, we present in the lower panel of figure 5 a scatter plot of unemployment rates against annualized wage changes net of expectations. We model expectations as backward-looking and equal to the average over annualized quarterly wage changes over the previous four quarters. Unlike what one sees in the slope of price Phillips curves, in this wage Phillips curve there is no evidence of a change since the early 1980s. Also, wage changes during the Great Recession line up very closely with what historical experience would have predicted from the rise in unemployment since 2007. Coibion and Gorodnichenko (2013) document similar results for alternative wage measures and expectations of wage changes from professional forecasters. Hence, there is no missing wage disinflation puzzle to match the missing price disinflation puzzle. This suggests that downward wage rigidity is unlikely to be the key factor underlying the missing disinflation of the Great Recession.

III.C. Role of Changes in Key Structural Characteristics

Given that financial shocks, monetary and fiscal policies, and downward wage rigidity all seem insufficient to account for the rising persistence in U.S. unemployment, we turn to long-term trends in key structural characteristics which could explain this phenomenon through propagation mechanisms. We specifically consider three possibilities: (i) declining labor mobility, (ii) changing age composition of the population, and (iii) the declining culture of “trust” in the United States.

The decline in regional convergence across states points toward changing mobility patterns as a potential explanation. Kaplan and Schulhofer-Wohl (2012) document that mobility in the United States has declined since the early 1990s, with the fall in mobility rates being as high as one-half for some measures of mobility. This decline is visible for immigrants, for individuals of all ages, for both single and dual earner households, for both low and high levels of education, and for different occupational groups. Using data from the Decennial Census, the American Community Survey, and the Current Population Survey (CPS), we construct the interstate turnover mobility rate since 1970, with turnover mobility defined as the average of each state’s immigration and outmigration rates. This rate was stable

between 1970 and 1980 but has been falling since. Mobility as reported in the Census has declined by about 20 percent since 1980 and by about 40 percent in the CPS since 1982 (the first available year).

While the sources of this decline in mobility are not fully understood and could be an endogenous response to changing economic patterns, one could also consider causal effects of exogenous variation in mobility rates, due perhaps to the types of declining information costs suggested by Kaplan and Schulhofer-Wohl (2012). An exogenous decline in mobility over time could naturally account for declining regional convergence after recessions and could explain other features of the data. For example, if falling mobility reduces the expected quality of job matches, then firms may choose to delay hiring more during recovery periods if there are costs to subsequent separations. One might also expect to see more long-term unemployment if workers were less willing to move away from the most seriously affected regions. This increase in expected duration of unemployment would also tend to increase the appeal of receiving disability payments, potentially accounting for the changing cyclical nature of disability claims over time. Another factor consistent with a mobility explanation is that, in the early 1970s, the United States stood out relative to most Western European countries in terms of how high its mobility rate was (see figure C.4 in online appendix C to this paper). A decline in mobility therefore represents a growing similarity with one of the defining characteristics of Western European economies during this early period. The mobility rates were all significantly lower there than in the United States at the time, but since 1980 the gap between the rates has closed by 52 to 97 percent.

A second changing characteristic of the U.S. economy is the age composition of its workforce. In 1970, the United States stood out relative to Western European countries in its age distribution: it had a larger share of young people (0 to 24 years old) and a smaller share of seniors (over 65) than any Western European country other than Spain. But this age composition has changed significantly over time, with the share of people ages 0 to 24 falling from 45 percent in 1970 to 34 percent in 2010 and the share of those ages 55 to 65 rising from 19 to 25 percent. The age structure of the labor force can be important because the experience of losing a job affects people differently depending on their age groups. For example, Johnson and Butrica (2012) document how, during the Great Recession, unemployment durations have been significantly longer for older workers than for younger workers and have been associated with larger declines in earnings upon reemployment. Since older workers are also less mobile on average

(Plane 1992), their lower mobility could account for declining regional convergence. The aging of the workforce could also be consistent with the rising cyclical of disability claims, since the higher expected duration of unemployment for older workers raises the benefit of receiving disability, and older workers may find it easier to justify claiming disability by appealing to a wider range of health impairments than younger workers.

A third changing feature of the United States lies in its cultural characteristics. The United States has long been an outlier among developed economies along a number of cultural dimensions, such as religiosity, views on private enterprise, and views on the role of the state. Another metric along which the country has traditionally differed from other developed economies is in people's perspective on claiming government benefits along with broader measures of trust toward others. For example, in the early 1980s Americans were much less likely than citizens of major Western European countries to report in the World Values Survey that they ever found it acceptable to claim government benefits for which they did not qualify. A greater unwillingness to claim benefits can have direct effects on labor markets: workers are more likely to accept job offers when unemployed to minimize the time spent receiving unemployment benefits and they are less likely to claim disability benefits when they do not necessarily qualify for them. Thus, this cultural difference could have played a role in explaining why the U.S. experience was so different from that of major Western European economies in the early 1980s. This cultural feature of the U.S. has changed over time (see figure C.4 in the online appendix). By 2006, the share of Americans reporting that they sometimes consider it acceptable to claim benefits for which they do not qualify had risen to the same level as reported, on average, by Western Europeans in 1981. Thus, like the mobility and age composition of the population, this cultural more is another dimension along which the United States was unusual in the 1970s and 1980s relative to Western Europe but on which it has gradually converged toward the levels reported there.

These three factors—declining mobility, age composition, and cultural values—are not meant to be an exhaustive list of the potential factors underlying the changing propagation of shocks in the United States. But we view these three patterns as consistent with the other stylized facts about the American economy that we have documented, so they constitute a natural starting point. Because of the gradually evolving nature of these factors, there is little hope of measuring their effects through time series methods. Instead, our strategy is to exploit the variation in labor market outcomes not just across countries but also across regions within countries over the

course of the 1970s and 1980s, a period characterized by large common shocks, to shed light on the potential quantitative impact of each factor on unemployment outcomes.

IV. Propagation Mechanisms in the 1970s and 1980s

In this section, we use the cross-sectional variation in labor market outcomes from the 1970s through the 1980s to quantify the extent to which labor mobility, age composition of the population, and cultural values might affect unemployment persistence. A key novelty in our analysis is that we focus on the variation across regions within Western European countries, Canada, and the United States. Whereas most previous work has focused on variation across countries, our approach significantly increases the size of the cross-section, which allows us to more precisely isolate the contributions of each factor to unemployment persistence. To do so, we construct a novel data set of regional labor market variables across 14 advanced economies, as well as a wide range of control variables designed to help isolate the marginal effects of labor mobility, age composition, and cultural factors.

IV.A. Regional Labor Markets

A central component of our analysis is the construction of a data set of regional unemployment rates going back to 1970 for 14 countries: Austria, Belgium, Canada, Denmark, Finland, France, (West) Germany, Italy, Netherlands, Spain, Sweden, Switzerland, the United Kingdom, and the United States. For most European countries, the smallest regional unit for which the unemployment series and other variables of interest are typically available is at the regional level of the Nomenclature of Territorial Units for Statistics (NUTS)-2. NUTS regions are the modern statistical areas of Europe. Modern NUTS-2 regions, which are smaller than NUTS-1 and bigger than NUTS-3 regions, range in size between 800,000 and 3 million people.⁹ (A detailed description of the sources of the data is in online appendix A.) For most countries, regional unemployment rates are available at the

9. Examples of NUTS-2 regions include the 39 government regions (or equivalents) of Germany, the 19 autonomous communities and cities of Spain, the 26 administrative regions of France plus the *Départements d'Outre Mer*, etc. Our study also includes the 48 states of the continental United States and the 10 provinces of Canada, yielding a total of 232 regions. We exclude, in addition to Alaska and Hawaii, islands (e.g., the Azores), West Berlin, and Washington, D.C. Because some regions were split or merged over time, we combine regions in a handful of instances where available data did not allow us to create consistent series for a region.

annual frequency, although for some countries, such as Switzerland and the United Kingdom, comparable unemployment rates are only available in Census years or other select years before 1996.

Early research on the differences in unemployment patterns over the course of the 1980s focused on the contrast between the United States and Western Europe as a whole, but subsequent work emphasized that differences within Western Europe were also large (Nickell 1997, Blanchard and Wolfers 2003). Differences in labor market outcomes within countries are even more challenging to explain in terms of aggregate policies and labor market institutions. In Belgium, for example, the North (Flemish) and South (Walloon) regions experienced diverging unemployment paths over the course of the 1980s despite a common set of aggregate institutions and policies. While both regions had low rates of unemployment in the early 1970s, the South saw a larger increase in unemployment in the late 1970s and the regional gap continued to widen from then on. By 1990, the North of Belgium had seen its unemployment rate decline to less than 6 percent from over 10 percent in 1980, whereas the unemployment rate in the South was still over 10 percent (see figure C.9 in online appendix C). Within a single country, therefore, one can find the same range of labor market outcomes as across countries. Other countries also experienced widening regional disparities over the 1980s, including France, Germany, Italy, and Spain. By contrast, a few countries saw only transitory regional disparities arise during this decade, among them Sweden.¹⁰

To measure the persistence of unemployment at the regional level, we focus on multiyear averages of unemployment rates for each region because of concerns about measurement error in regional rates at the annual frequency (this follows Blanchard and Wolfers 2003). Specifically, we calculate

$$\bar{u}_{1990}^{i,c} \equiv \frac{1}{3}(u_{1989}^{i,c} + u_{1990}^{i,c} + u_{1991}^{i,c}),$$

$$\bar{u}_{1980}^{i,c} \equiv \frac{1}{3}(u_{1980}^{i,c} + u_{1981}^{i,c} + u_{1982}^{i,c}),$$

and

$$\bar{u}_{1970}^{i,c} \equiv \frac{1}{4}(u_{1970}^{i,c} + u_{1971}^{i,c} + u_{1972}^{i,c} + u_{1973}^{i,c}),$$

10. The rise in within-country dispersion of unemployment rates tracked the rise in cross-country dispersion rates throughout the 1970s and 1980s, as documented in figure C.5 in the online appendix.

where $u_t^{i,c}$ is the annual rate of unemployment in year t for region i in country c .¹¹ For some regions, data are not available for all years, so we use available data to construct $\bar{u}_t^{i,c}$.

From these multiyear averages, we construct our primary measure of unemployment persistence (or hysteresis) $h_{i,c}$ for region i in country c as

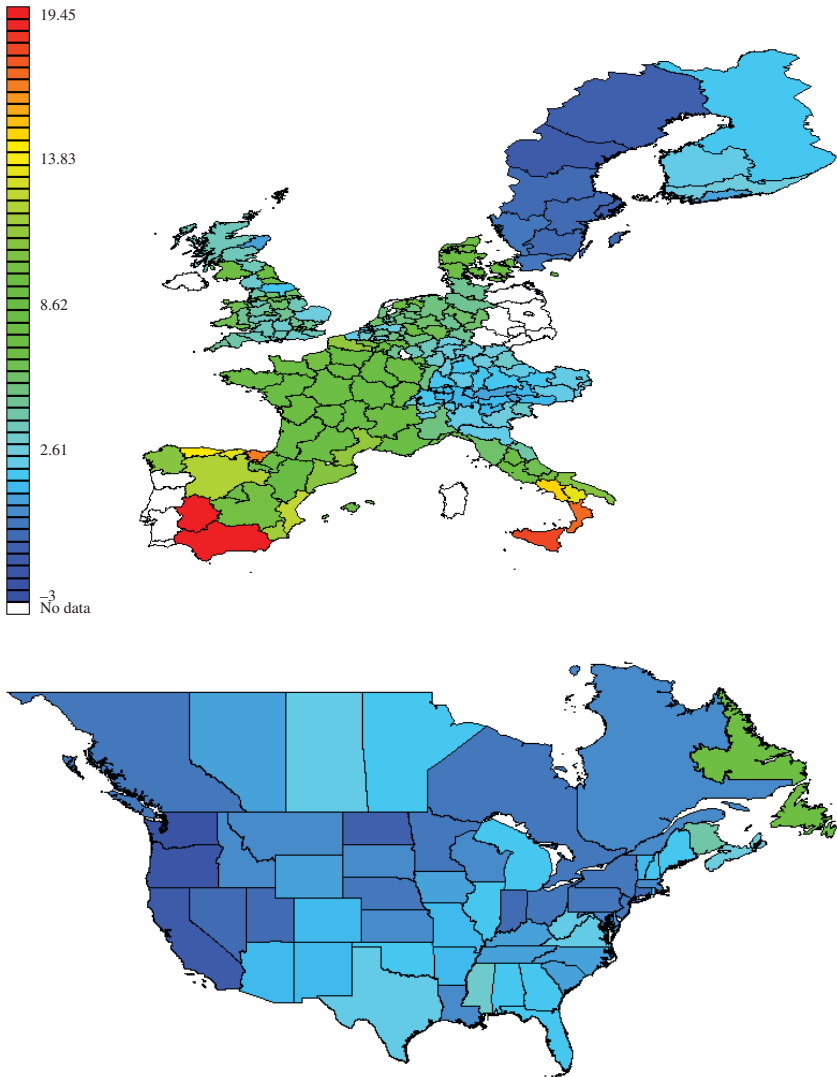
$$h_{i,c} = \ln(1 + \bar{u}_{1990}^{i,c} - \bar{u}_{1980}^{i,c}) + \ln(1 + \bar{u}_{1980}^{i,c} - \bar{u}_{1970}^{i,c}).$$

The first term measures the change in unemployment over the 1980s, and the second measures the change in unemployment over the course of the 1970s. High values of $h_{i,c}$ denote regions with more persistence, that is, in which unemployment fell only gradually (or even rose) relative to the initial rise in unemployment during the 1970s. We take logs to ensure a balanced distribution and minimize outliers. In online appendix B, we show how $h_{i,c}$ can be related to the persistence of unemployment measured by the magnitude of the AR(1) coefficient. We prefer our measure of unemployment persistence over alternatives based on time series regressions, because it minimizes the influence of short-term fluctuations, measurement error, and missing observations and thus provides a robust metric of persistence. One limitation of our measure is that it does not exploit the full time series variation in unemployment rates.

Figure 6 displays the resulting values of unemployment persistence for each region on a map. There are clear differences in persistence both across and within countries. For example, the United States and Canada both had relatively low levels of unemployment persistence on average, although there was some variation across regions. In the United States, southern states experienced more persistence than others, whereas northwestern states and New England states had the lowest. In Canada, Newfoundland and New Brunswick, two sparsely populated provinces, experienced significantly more unemployment persistence than the rest of the country. In Europe, Sweden experienced by far the least persistence, with very little interregional variation to be found, and Finland had a similar experience

11. The timing of the multiyear averages was determined by data availability as well as macroeconomic developments in Europe. Specifically, we include an additional year in calculating the average unemployment rate before the mid-1970s to improve the precision, since earlier years have more missing observations. We shift the timing of the average for 1990 by one year because later years were marked with crises and recessions in Europe (including the ERM crises, the collapse of trade with the Soviet Union, and German reunification).

Figure 6. Geographic Distribution of Unemployment Hysteresis



Source: Authors' calculations; see online appendix for data sources.

Notes: The figure shows the degree of hysteresis (as defined in the text) for each region. Higher values denote more hysteresis.

albeit with higher persistence on average. Unemployment persistence was relatively mild in two other general regions of Europe: Central Europe (including Austria, Switzerland, northern Italy, and southern Germany) and the southern region of the United Kingdom. The Atlantic coast had significantly higher persistence on average, which grew gradually as one descended further south. Denmark was an exception, with much higher persistence levels than neighboring regions. Finally, the most extreme regions of unemployment persistence were in the far south of Spain and Italy, with some high-persistence outliers along the Atlantic coast of Spain (including the Basque region).¹²

IV.B. Regional and Aggregate Determinants

To estimate the partial correlations of labor mobility, age composition, and cultural factors with employment persistence, we need to disentangle them from other potential sources of variation. While aggregate differences can be captured through country fixed effects, one would expect many factors to be potentially related to regional variation in unemployment rates. Thus, a second key element of our analysis is the construction of a set of “reduced form” predetermined regional controls. We refer to these as reduced form controls because they are not meant to identify precise causal mechanisms. Rather, we expect these controls to be correlated with many of the channels which have been suggested in the hysteresis literature, so that controlling for them should allow us to evaluate the marginal correlations of labor mobility, age composition, and cultural values.

We use 10 regional control variables (each is plotted against regional levels of unemployment persistence in figures C.7 and C.8 in the online appendix). Each of these variables is for 1970 or the nearest available year, as described in the data appendix. The first three include average education (years of schooling) in 1970, the log of GDP per capita (in PPP U.S. dollars from Penn World Tables) in 1970, and the share of female employment in total employment in 1970. All three measures are extremely highly correlated with subsequent regional labor market outcomes, with higher persistence regions being initially poorer, less educated, and with fewer women participating in the labor force. A fourth control is the net migration rate, defined as the annual net inflow (that is, in-migration minus out-migration)

12. There is a systematic relationship between the average level of persistent unemployment in a country and the cross-sectional dispersion in unemployment outcomes. Online appendix figure C.6 shows that from 1980 on, there is a strong positive correlation between the rise in within-country dispersion in unemployment rates and subsequent persistence, with the relationship continuing throughout the 1980s and into the early 1990s.

of people normalized by total population. The next set of controls consists of the industry composition of employment in the region in 1970, that is, the share of the primary sector in employment (agriculture, fishing, mining, and energy) and the share of the secondary sector in employment (manufacturing and construction), with the tertiary sector consisting of all other sectors. Regions that experienced more unemployment persistence had, on average, larger shares of employment in the primary sector and smaller shares in the tertiary sector. We also include as controls the density of highways, household size, home ownership rates, and the share of the population living in urban environments. None of these display strong unconditional correlations with unemployment persistence. Finally, we also include dummy variables for regions that consist of a larger metropolitan area (such as Hamburg) and another dummy variable for regions in which the capital city is located.

A striking feature of these regional controls is that they can account for much of the regional variation in unemployment persistence. For example, we regress regional unemployment persistence on these predetermined regional controls with no aggregate controls or fixed effects with equal weight on each country while weighting regions within countries by their population. We apply equal weights on countries to avoid having the results dominated by the United States. Weighting regions by population also serves to avoid placing undue weight on smaller regions, where measurement error is more likely to be an issue. This regression yields an R-squared of 0.70, implying that much of the variation in labor market outcomes is very highly correlated with the set of regional controls included here. If we also include country fixed effects as a simple way to capture cross-country differences in policies, shocks, and institutions, the R-squared of the regression rises to 0.84. Nearly all of the estimated coefficients on country dummies are not statistically significantly different from zero (Denmark is the main exception). This last outcome confirms that regional control variables are successful in accounting for both the within- and between-country variation in unemployment outcomes, leaving little room for improvement in terms of controlling for aggregate variables.

IV.C. The Role of Mobility, Demographics, and Culture

We similarly construct regional measures of mobility, age structure, and social trust. For mobility, our primary measure is the 1970 turnover migration rate, the average of inflows and outflows of a region normalized by the population of the region. The first column of table 1 presents results from regressing regional persistence on regional labor mobility. The coefficient

Table 1. The Effects of Mobility, Age Structure, and Culture on Unemployment Persistence

<i>Variables</i>	<i>OLS (1)</i>	<i>OLS (2)</i>	<i>FE (3)</i>	<i>OLS (4)</i>	<i>OLS (5)</i>	<i>FE (6)</i>	<i>OLS (7)</i>	<i>OLS (8)</i>	<i>FE (9)</i>	<i>OLS (10)</i>	<i>FE (11)</i>
Turnover migration rate	-0.34 (0.37)	0.43* (0.26)	0.41* (0.24)							0.37 (0.27)	0.29 (0.22)
Population share 0-24 years				0.16 (0.15)	0.18 (0.13)	0.60*** (0.13)				0.14 (0.13)	0.58*** (0.12)
Population share 55+ years				0.18 (0.18)	0.17 (0.14)	0.42*** (0.13)				0.17 (0.14)	0.42*** (0.12)
Trust							-1.34*** (0.45)	0.08 (0.30)	-0.97*** (0.34)	0.16 (0.33)	-1.00*** (0.32)
Controls											
Years of schooling		-0.87*** (0.15)	-1.67*** (0.66)	-0.94*** (0.15)		-1.60*** (0.64)		-0.88*** (0.17)	-1.55*** (0.66)	-0.98*** (0.17)	-1.57*** (0.63)
Share of employment in the primary sector		0.05 (0.04)	0.02 (0.03)	0.03 (0.04)		-0.01 (0.03)		0.02 (0.04)	-0.00 (0.03)	0.04 (0.04)	-0.01 (0.03)
Share of employment in the secondary sector		-0.02 (0.03)	-0.05** (0.03)	-0.03 (0.03)		-0.07*** (0.02)		-0.05* (0.03)	-0.07*** (0.02)	-0.02 (0.03)	-0.06** (0.02)
Women's share in employment		-0.20*** (0.04)	-0.26*** (0.09)	-0.19*** (0.04)		-0.21*** (0.08)		-0.18*** (0.04)	-0.28*** (0.09)	-0.18*** (0.04)	-0.23*** (0.08)
Household size		-0.23 (0.74)	-0.67 (0.72)	-0.60 (0.91)		-2.48** (1.17)		-0.32 (0.72)	-0.65 (0.69)	-0.08 (0.76)	-2.13* (1.18)
Home ownership rate		-0.05*** (0.02)	-0.04* (0.02)	-0.05*** (0.02)		-0.04* (0.02)		-0.05*** (0.02)	-0.02 (0.02)	-0.05*** (0.02)	-0.03 (0.02)
Net migration rate		-0.39 (0.27)	-0.26 (0.21)	-0.35 (0.28)		-0.11 (0.20)		-0.40 (0.28)	-0.26 (0.21)	-0.34 (0.27)	-0.10 (0.18)
Share of urban population		0.01 (0.01)	0.03** (0.01)	0.01 (0.02)		0.03** (0.01)		0.01 (0.01)	0.03** (0.01)	0.01 (0.01)	0.04*** (0.01)

(continued)

Table 1. The Effects of Mobility, Age Structure, and Culture on Unemployment Persistence (Continued)

Variables	OLS (1)	OLS (2)	FE (3)	OLS (4)	OLS (5)	FE (6)	OLS (7)	OLS (8)	FE (9)	OLS (10)	FE (11)
Log(Income p.c.) × 100		-0.05*** (0.01)	-0.04** (0.02)		-0.04*** (0.01)	-0.02 (0.02)		-0.05*** (0.01)	-0.03* (0.02)	-0.04*** (0.01)	-0.01 (0.02)
Road density		0.46*** (0.16)	0.49*** (0.12)		0.35** (0.17)	0.48*** (0.13)		0.42** (0.20)	0.47*** (0.12)	0.44*** (0.21)	0.44*** (0.14)
City region		3.50*** (0.56)	0.40 (0.82)		3.52*** (0.69)	0.57 (0.91)		3.52*** (0.59)	0.88 (0.81)	3.33*** (0.63)	0.47 (0.92)
Capital		-0.05 (0.51)	0.50 (0.50)		-0.12 (0.60)	0.42 (0.46)		-0.31 (0.54)	0.28 (0.49)	-0.01 (0.59)	0.35 (0.46)
Observations	232	232	232	232	232	232	232	232	232	232	232
R ²	0.01	0.71	0.85	0.01	0.71	0.86	0.12	0.70	0.85	0.71	0.87
BIC	5.690	4.741	4.415	5.712	4.772	4.342	5.570	4.759	4.355	4.803	4.286

Notes: Regression is estimated with population weights, constructed so that regions within a country receive weights proportional to their population but all countries have the same weight. Country fixed effects are included (but not reported) in columns "FE". Robust standard errors are reported in parentheses. *, **, and *** denote significance at 10 percent, 5 percent, and 1 percent levels.

is negative but insignificantly different from zero. To estimate the partial correlation of mobility and unemployment persistence after conditioning on other factors, we estimate the following regressions:

$$h_{i,c} = \theta_{mob} mobility_{i,c} + \beta' \mathbf{X}_{i,c} + \gamma_c + v_{i,c},$$

where $mobility_{i,c}$ is the turnover migration rate of region i in country c . We estimate this specification with and without country-fixed effects (γ_c) and report results in columns 2 and 3 of table 1. In both cases, we find a positive estimate of θ_{mob} that is statistically different from zero only at the 10 percent level. This implies that higher regional labor mobility is associated with more persistence, if anything. The absence of a negative relationship between labor mobility and persistence is robust to alternative measures of mobility (such as aggregate mobility, inflow rates vs. outflow rates) and is not driven by outliers (see figures C.7 and C.8 in the online appendix).

To measure the demographic makeup of the population, we decompose regional populations in 1970 into three groups: the share of people ages 0 to 24, the share of people ages 25 to 54 (prime ages), and the share of people age 55 and older. Since these shares are measured in 1970, the first category can be interpreted as capturing the share of young workers in the labor force in the mid-1980s, the second category corresponds to the share of older workers (that is, approaching or reaching retirement age) over the 1980s, while the third group captures the share of retirees in the 1980s. The fourth column of table 1 presents regression results for the share of young people and seniors in 1970 with no regional controls or fixed effects. Both age groups are associated with higher persistence on average, but the explanatory power of these two variables is limited. To assess how age composition correlates with unemployment persistence after controlling for other factors, we estimate

$$h_{i,c} = \theta_{<25} share_{i,c}^{<25} + \theta_{55+} share_{i,c}^{55+} + \beta' \mathbf{X}_{i,c} + \gamma_c + v_{i,c},$$

where $s_{i,c}^{<25}$ are the shares of people under age 25 and $s_{i,c}^{55+}$ are the shares of people over age 54 for region i in country c . We estimate this specification with and without country-fixed effects and report the results in columns 5 and 6 of table 1. For both variables, we find positive coefficients, although the magnitudes are much stronger when we include country-fixed effects. Both coefficients imply that larger shares of people ages 25 to 54 in 1970 are therefore associated with lower rates of unemployment persistence.

To construct a measure of trust, we aggregate over three different questions from the World Values Survey directly related to civic attitudes about trust. The questions relate to whether it is ever acceptable to claim benefits for which one does not qualify, whether others can be trusted, and whether others will try to take advantage of one. From these three questions, we create a regional index of “trust” (as described in online appendix D).¹³

One limitation of these cultural measures is that regional information on culture is available at the NUTS-level only in the 2008 wave of the World Values Survey, with previous surveys providing only coarser regional or country-level information.¹⁴ As a result, endogeneity of cultural variables with respect to prior economic conditions is a concern. For example, periods of high unemployment could tend to lower societal trust or reduce (or raise) the importance that people place on work. To investigate this possibility, we use the time-variation in aggregate values to each type of question and assess whether these are correlated with recent economic conditions, a test similar to the one used in Stevenson and Wolfers (2011).

For this purpose, we focus on questions that have been asked at least since the 1990s waves of the World Values Survey and for which we can therefore measure some time variation. For each question x_t^c at time t for country c , we regress it on the unemployment rate for that country in that year and pool across all countries for each question, allowing for country fixed effects. We run this regression for two questions in each class of cultural variables as well as the cultural index value for that country, constructed in the same way as the regional indexes (but using questions as they become available). In the results we fail to find consistent and robust evidence that survey responses or indexes of cultural dimensions are sensitive to unemployment fluctuations, which is consistent with cultural factors not responding strongly to short-term economic conditions (see table D.2 in the online appendix).

Given this lack of strong correlation between culture and the level of unemployment, we explore whether regional levels of trust are related to

13. In Coibion, Gorodnichenko, and Koustas (2013), we explore two additional cultural dimensions: “value of work” and “individualism.” We do not find that these measures change the main conclusions of the paper.

14. Data for the United States and Canada come from the 2006 wave of the WVS. Regional data begin in 1990 for these two countries. For the United States in all WVS waves, the data are only available at the aggregation of the nine census divisions (New England, Middle Atlantic, South Atlantic, East South Central, West South Central, East North Central, West North Central, Mountain, Pacific [excluding California], plus California). We tie the same responses from each Census division to each of the states within the division.

unemployment persistence over the 1970s and 1980s. As reported in column 7 of table 1, there is a strong unconditional correlation (significant at the 1 percent level) between trust and persistence, with trust accounting for 12 percent of the cross-sectional variation in persistence. We then report estimates allowing for regional controls and country fixed effects

$$h_{i,c} = \theta_i \text{trust}_{i,c} + \boldsymbol{\beta}' \mathbf{X}_{i,c} + \gamma_c + v_{i,c},$$

when $\text{trust}_{i,c}$ is the index of trust for region i in country c in columns 8 and 9 of table 1. The results are sensitive to the inclusion of fixed effects. In their absence, the coefficient on regional trust is close to zero and insignificant, whereas when fixed effects are included the estimated coefficient is negative and significant at the 1 percent level. The sensitivity to fixed effects reflects a single country, Denmark, for which our regional controls do not capture the average country fixed effect, with the Danish level of persistence being unaccountably high. Since Denmark also had the highest levels of trust of any country in our sample, this accounts in large part for the difference between estimates with and without fixed effects.

These baseline results are robust to a wide range of checks. For example, in columns 10 and 11 of table 1, we estimate the coefficients on mobility, demographics, and culture on unemployment persistence jointly with and without fixed effects and find little variation relative to the baseline results. In table 2, we reproduce baseline results with and without fixed effects, dropping, in turn, different subsets of countries. Dropping the United States and Canada, for example, has no qualitative effect on the results. This suggests that regional variation in the United States is accounted for by regional variables in a similar manner as in other countries, that is, there is little reason to believe that different coefficients on controls are needed to explain European and North American regional variation in persistence.

Dropping Italy and Spain generally reduces the estimated coefficients but does not otherwise affect our qualitative results. Since Spain and Italy are large outliers in terms of persistence outcomes, one might have expected them to drive the empirical results, but this is not the case. Dropping France and Germany also does not qualitatively alter the results, other than rendering the estimated coefficient on mobility statistically insignificant, as does dropping the Nordic countries. The Nordic countries have a more substantial effect on the estimated coefficients on trust: dropping these leads to a large negative and statistically significant coefficient on trust even when fixed controls are not included. This reflects the sensitivity to Denmark, which has already been noted. Using equal weights across

Table 2. Robustness to Country Selection

<i>Variables</i>	<i>Baseline</i>		<i>Exclude U.S. and Canada</i>		<i>Exclude Italy and Spain</i>		<i>Exclude France and Germany</i>		<i>Exclude Nordic countries</i>	
	<i>OLS (1)</i>	<i>FE (2)</i>	<i>OLS (3)</i>	<i>FE (4)</i>	<i>OLS (5)</i>	<i>FE (6)</i>	<i>OLS (7)</i>	<i>FE (8)</i>	<i>OLS (9)</i>	<i>FE (10)</i>
Turnover migration rate	0.37 (0.27)	0.29 (0.22)	0.17 (0.28)	0.43* (0.25)	0.39 (0.24)	0.07 (0.16)	0.35 (0.31)	0.29 (0.23)	0.36 (0.22)	0.24 (0.27)
Population share 0–24 years	0.14 (0.13)	0.58*** (0.12)	0.26* (0.15)	0.59*** (0.14)	0.25** (0.13)	0.29*** (0.11)	-0.14 (0.18)	0.68*** (0.13)	0.20* (0.11)	0.58*** (0.12)
Population share 55+ years	0.17 (0.14)	0.42*** (0.12)	0.12 (0.16)	0.39*** (0.14)	0.32** (0.13)	0.25** (0.10)	-0.14 (0.20)	0.47*** (0.13)	0.31** (0.13)	0.38*** (0.12)
Trust	0.16 (0.33)	-1.00*** (0.32)	-0.09 (0.32)	-1.05*** (0.32)	0.36 (0.31)	-0.61** (0.24)	0.15 (0.39)	-1.55*** (0.44)	-1.14*** (0.30)	-0.91*** (0.32)
Controls										
Years of schooling	-0.98*** (0.17)	-1.57** (0.63)	-0.85*** (0.20)	-1.62* (0.85)	-0.95*** (0.17)	-1.64*** (0.45)	-0.56** (0.27)	-1.55** (0.67)	-0.63*** (0.17)	-1.22* (0.63)
Share of employment in the primary sector	0.04 (0.04)	-0.01 (0.03)	0.07 (0.05)	0.01 (0.04)	0.08** (0.04)	0.03 (0.03)	0.08* (0.05)	-0.02 (0.03)	0.05 (0.03)	0.01 (0.04)
Share of employment in the secondary sector	-0.02 (0.03)	-0.06** (0.02)	-0.05 (0.04)	-0.08*** (0.03)	-0.00 (0.03)	-0.06*** (0.02)	-0.00 (0.04)	-0.06** (0.03)	-0.04 (0.03)	-0.04 (0.03)

Women's share in employment	-0.18*** (0.04)	-0.23*** (0.08)	-0.14*** (0.04)	-0.22*** (0.08)	-0.07** (0.04)	-0.15*** (0.04)	-0.24*** (0.05)	-0.24** (0.11)	-0.24*** (0.05)	-0.23*** (0.09)
Household size	-0.08 (0.76)	-2.13* (1.18)	-0.96 (0.90)	-2.43* (1.32)	-1.03 (0.71)	-1.74** (0.87)	-0.83 (0.99)	-3.27** (1.50)	2.21*** (0.85)	-2.21 (1.34)
Home ownership rate	-0.05*** (0.02)	-0.03 (0.02)	-0.01 (0.03)	-0.03 (0.02)	-0.08*** (0.02)	-0.03** (0.02)	-0.04* (0.02)	-0.04 (0.02)	-0.07*** (0.02)	-0.07*** (0.02)
Net migration rate	-0.34 (0.27)	-0.10 (0.18)	-0.42 (0.33)	-0.23 (0.23)	-0.13 (0.22)	-0.07 (0.16)	-0.46 (0.30)	0.03 (0.21)	-0.13 (0.26)	0.05 (0.22)
Share of urban population	0.01 (0.01)	0.04*** (0.01)	0.01 (0.01)	0.05*** (0.01)	0.00 (0.01)	0.01 (0.01)	0.01 (0.01)	0.04*** (0.01)	0.06*** (0.01)	0.04*** (0.01)
Log(Income p.c.) × 100	-0.04*** (0.01)	-0.01 (0.02)	-0.03* (0.01)	-0.01 (0.02)	-0.03*** (0.01)	0.01 (0.01)	-0.07*** (0.01)	-0.02 (0.02)	-0.04*** (0.01)	-0.04** (0.02)
Road density	0.44** (0.21)	0.44*** (0.14)	0.78*** (0.25)	0.39*** (0.14)	0.67*** (0.19)	0.46*** (0.11)	0.45** (0.21)	0.42** (0.17)	-0.09 (0.18)	0.48** (0.19)
City region	3.33*** (0.63)	0.47 (0.92)	3.54*** (0.76)	0.55 (0.91)	2.21*** (0.65)	0.83 (0.63)	3.77*** (0.78)	-0.79 (0.95)	1.80* (0.93)	1.54 (1.38)
Capital	-0.01 (0.59)	0.35 (0.46)	-0.16 (0.78)	-0.25 (0.52)	-0.08 (0.47)	0.30 (0.28)	-0.11 (0.64)	0.75 (0.47)	0.17 (0.48)	0.49 (0.48)
Observations	232	232	175	175	197	197	183	183	215	215
R ²	0.71	0.87	0.71	0.87	0.64	0.87	0.72	0.88	0.82	0.87
BIC	4.803	4.286	4.887	4.436	4.386	3.642	4.954	4.376	4.432	4.366

Notes: Regression is estimated with population weights, constructed so that regions within a country receive weights proportional to their population but all countries have the same weight. Country fixed effects are included (but not reported) in columns "FE". Robust standard errors are reported in parentheses. *, **, and *** denote significance at 10 percent, 5 percent, and 1 percent levels, respectively.

all observations, applying population weighting to all regions without imposing equal country weights, or using Huber robust regressions does not qualitatively affect the results (see online appendix table C.2). Similarly, allowing for interactions with the country-level employment protection or unemployment replacement rates (from 1970) from Blanchard and Wolfers (2003) suggests that our baseline estimates are more pronounced in countries with high employment protections, but the results with replacement rates are extremely mixed and do not suggest any robust interaction effects (see online appendix table C.3). Finally, considering a much broader range of interactions, such as with the nature and size of the shocks hitting each economy or with original controls, also does not qualitatively affect our findings or contribute to the explanatory power of the predictive regressions (see online appendix table C.4).

IV.D. Implications for Rising U.S. Unemployment Persistence

Given these estimated partial correlations, we can quantify how much variation in unemployment persistence one should have expected in the United States over recent years given the time variation in each of these factors, assuming that the estimated coefficients on each factor have not evolved over time.

For labor mobility and demographics, we replace their 1970 values for each state with their 2000 values. Holding constant other controls, we can then construct a counterfactual level of persistence for each state and aggregate these with population weights into a counterfactual nationwide level of persistence. For culture, regional data for the United States in the World Values Survey do not begin until 1990, so we cannot track the regional evolution of culture since the early 1980s. Instead, we lower the level of trust in each state by one standard deviation (that is, moving from the U.S. level of trust to the lowest level of trust in our sample of countries). This corresponds to the historical variation in U.S. responses to the question of whether it is ever justifiable to claim benefits for which one does not qualify, but it is much larger than the historical variation in responses to the question of whether others can be trusted. Since the third question used in the trust index is not available going back to 1982, we cannot measure the historical change in the overall trust index for the United States. But the time variation in the first two questions suggests that assuming a one-standard-deviation decline in trust likely presents an upper bound on the possible historical variation in trust. We present the results from this exercise in table 3.

Changes in labor mobility have trivial effects on the predicted level of unemployment persistence in the United States. The changing demographic

Table 3. Counterfactual Levels of Unemployment Persistence for the United States

<i>Scenario</i>	<i>Predicted hysteresis</i>	<i>Specification used for prediction^a</i>	
		<i>OLS (1)</i>	<i>FE (2)</i>
(1) ^b	All variables have values from 1970 counterfactuals	0.37 (0.18)	0.60 (0.20)
(2) ^c	Migration rate is from 2000; all other variables have values from 1970	0.30 (0.19)	0.55 (0.20)
(3) ^d	Demographic shares are from 2000; all other variables have values from 1970	-0.83 (0.25)	-4.91 (0.27)
(4)	Reduced index of trust; all other variables have values from 1970	0.21 (0.19)	1.61 (0.21)
(5) ^e	Migration rate and demographic shares are from 2000; reduced index of trust, all other variables have values from 1970	-1.06 (0.26)	-3.97 (0.28)

a. The table reports predicted hysteresis for various scenarios. To predict hysteresis, for column 1 (OLS) we use the estimated specification from column 10 of table 1 and for column 2 (FE) we use the estimated specification from column 11 of table 1.

b. Scenario (1) uses actual values of the variables from the 1970s to predict hysteresis.

c. Scenario (2) uses the U.S. migration rate from 2000 while all other variables continue to have their 1970 values.

d. Scenario (3) uses U.S. demographic shares (ages 0–24 and 55+) from 2000 while all other variables continue to have their 1970 values. Since the trust index is not available for the 1970s, Scenario (3) assumes that the trust index fell by one standard deviation (the standard deviation is calculated from the cross-section of U.S., Canadian, and European regions in the 2000s; the standard deviation by construction is approximately one) while all other variables continue to have their 1970 values.

e. Scenario (5) combines changes in scenarios (2) through (4) where only control variables continue to have their 1970s values. Standard errors are reported in parentheses.

makeup of the nation's population predicts lower persistence, modestly if we use the estimates without country fixed effects and dramatically if we use the estimates with fixed effects. The former is equivalent to reducing the level of persistence from French levels to Belgian levels, whereas the latter is equivalent to reducing the level of persistence from French levels to levels of Austria or Switzerland. The effects of demographics are primarily driven by the changing share of people ages 0 to 24, reflecting both its large decline over time as well as the fact that its estimated coefficient is higher than that on the share of people age 55 and over. Both labor mobility and demographic changes predict declines in U.S. unemployment persistence, given the estimated correlations of each with persistence from table 1.

The decline in trust in the United States is the only variable among the three that predicts increased unemployment persistence when fixed effects

estimates are used. A large, one-standard-deviation change in trust raises persistence by 1 point, the equivalent of going from U.S. levels of persistence to those of Austria or Switzerland. Converted to AR (1) equivalents (using the formula in online appendix B), this corresponds to moving from a quarterly unemployment persistence of 0.92 (the estimate from 1950–1980 in figure 1) to a quarterly persistence of 0.98. This actually exceeds the increase in persistence found for the United States between the 1980s and 2010s, so by itself the historical change in trust could account for the observed rise in U.S. unemployment persistence. However, as documented in the last row of table 3, this predicted increase in persistence due to culture is dwarfed by the effect of demographics and mobility when we construct joint counterfactuals for all three. Thus, mobility, demographics, and trust cannot jointly account for the rising persistence in unemployment in the United States.

V. Conclusion

A clear pattern has emerged: the response of the U.S. economy to recessions has become increasingly anemic, starting with the 1990 recession, and the Great Recession appears to be continuing a trend toward rising unemployment persistence. Understanding the sources of this rising persistence is a key question for economists, since its continuation suggests a future in which post-recession unemployment outcomes increasingly resemble the Euroclerosis experience of the 1980s rather than the rapid recoveries that characterized the U.S. economy through most of the post–World War II era.

This paper has called into question many of the “usual suspects.” We document that commonly suggested explanations for the persistence of unemployment in the Great Recession cannot fully explain this unfortunate new characteristic of U.S. business cycles. Financial shocks do not mechanically lead to more unemployment persistence than monetary policy shocks, so it is doubtful that the underlying explanation stems from any changing nature of the shocks driving the cycle. And while monetary and fiscal policies can account for a significant component of the rising unemployment persistence, much remains unexplained. This suggests that much of the answer must lie in the economic mechanisms that underlie the propagation of shocks through the economy.

We tested three potential propagation factors that are commonly discussed in the literature—declining labor mobility, changing demographic composition of the population, and changing cultural mores—which could account for rising unemployment persistence as well as other changes in

U.S. business cycle properties. Although a changing attitude on claiming government benefits and trusting others can quantitatively account for the observed rise in persistence, its predicted effect should have been more than offset by the other factors.

Our interpretation of these results is that there must be additional, more powerful factors at work. Identifying these other forces should be a key priority for the research agenda of macroeconomists. There is no shortage of places to look. A striking number of features of the U.S. economy changed around 1980. We highlighted four such changes in cyclical patterns in section III.A. Garin, Pries, and Sims (2013) emphasize another change: the changing cyclicity of labor productivity. But there are also pronounced changes in trends for a range of variables all of which appeared around 1980. An incomplete list includes rising economic inequality (Piketty and Saez 2003), a declining share of labor income (Karabarbounis and Neiman 2013), and changes in the demand for skills (Jaimovich and Sui 2013).

These results also have implications for policy responses to business cycles. First, we find that the contribution of monetary and fiscal policies toward the higher unemployment persistence of recent recessions is non-trivial. This suggests that improvements in the implementation of fiscal and monetary responses could go some way in reversing the trends toward more depressed labor market outcomes after recessions. Reducing the speed at which initially expansionary fiscal policies are reversed (as was the case in the 2001 and 2007 recessions) would be a step in the right direction.

More fundamentally, if future U.S. downturns will be more long-lived than pre-1990 recessions, then the nature of fiscal policy responses should likely be revisited. In the pre-1990 environment, recessions were short-lived events and there was little need to implement discretionary counter-cyclical fiscal policies, other than highly transitory ones such as the rebate checks of 2001, because the long decision lags involved in the legislative process meant that any positive effects of stimulus spending would likely occur too late. But if business cycles have become systematically more protracted affairs, as seems to be the case, then discretionary fiscal policy responses should target longer-lived projects rather than transitory transfer payments. Investment projects, for example, might be especially desirable because they tend to have larger stimulative effects per dollar (see for example Auerbach and Gorodnichenko 2012, Leduc and Wilson 2012) and tend to have long-run social returns that significantly exceed those of transfer payments. They also do not require legislators to vote on multiple stimulus packages, a fact that could facilitate their prolonged usage, in contrast to the repeated discretionary changes required with transfer payments.

On the other hand, if “timely, targeted and temporary” remains the mantra of future stimulus measures, then Amerisclerosis may not be so far away.

ACKNOWLEDGMENTS We thank Olivier Blanchard, Romain Ranciere, and the editors, as well as the conference participants at BPEA meetings and our discussants, Ed Glaeser, and Paola Sapienza, for very helpful comments. We also thank Florence Jaumotte for sharing updated Blanchard-Wolfers series on labor market institutions, and Alex Dombrowski, Noa Dreyman, Chuck Fang, and especially Donna Kim and Yury Yatsynovich for research assistance. Yuriy Gorodnichenko thanks the NSF and the Sloan Foundation for financial support. This paper was written in part while Olivier Coibion was a visiting scholar at the International Monetary Fund, whose support was greatly appreciated. The views expressed in the paper are those of the authors and should not be interpreted as representing those of any institutions with which they are or have been affiliated.

References

- Auerbach, Alan J., and Yuriy Gorodnichenko. 2012. "Measuring the Output Responses to Fiscal Policy." *American Economic Journal: Economic Policy* 4, no. 2: 1–27.
- Autor, David H., and Mark G. Duggan. 2003. "The Rise in the Disability Rolls and the Decline in Unemployment." *Quarterly Journal of Economics* 118, no. 1: 157–206.
- Ball, Laurence, and Sandeep Mazumder. 2011. "Inflation Dynamics and the Great Recession." *Brookings Papers on Economic Activity* 42, no. 1: 337–405.
- Ball, Laurence, Daniel Leigh, and Prakash Loungani. 2012. "Okun's Law: Fit at 50?" Paper presented at the 13th Annual Jacques Polak Research Conference, hosted by the International Monetary Fund, Washington, November 8–9, 2012.
- Blanchard, Olivier J., and Lawrence F. Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity* 23, no. 1: 1–76.
- Blanchard, Olivier, and Daniel Leigh. 2013. "Growth Forecast Errors and Fiscal Multipliers." Working Paper no. 13/1. Washington: International Monetary Fund.
- Blanchard, Olivier, and Thomas Philippon. 2006. "The Quality of Labor Relations and Unemployment." Manuscript available at <http://economics.mit.edu/files/712>.
- Blanchard, Olivier J., and Lawrence H. Summers. 1986. "Hysteresis and the European Unemployment Problem." In *NBER Macroeconomics Annual*, vol. 1, ed. by Stanley Fischer. Cambridge, Mass.: MIT Press.
- Blanchard, Olivier J., and Justin Wolfers. 2003. "The Role of Shocks and Institutions in the Rise of European Unemployment: The Aggregate Evidence." *Economic Journal* 110, no. 462: 1–33.
- Coibion, Olivier. 2012. "Are the Effects of Monetary Policy Shocks Big or Small?" *American Economic Journal: Macroeconomics* 4, no. 2: 1–32.
- , and Yuriy Gorodnichenko. 2013. "Is the Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation." Working Paper no. 19598. Cambridge, Mass.: National Bureau of Economic Research.
- Coibion, Olivier, Yuriy Gorodnichenko, and Dmitri Koustas. 2013. "Americ sclerosis? The Puzzle of Rising U.S. Unemployment Persistence." Conference Draft presented at the Fall 2013 Brookings Panel on Economic Activity, Washington, D.C., September 20, 2013.
- Congressional Budget Office (CBO). 2012. "Policy Options for the Social Security Disability Insurance Program" (http://cbo.gov/sites/default/files/cbofiles/attachments/43421-DisabilityInsurance_print.pdf).
- Daly, Mary, Bart Hobijn, and Brian Lucking. 2012. "Why Has Wage Growth Stayed Strong?" *FRBSF Economic Letter* 2012–10 (April 2).
- Decressin, Jorg, and Antonio Fatas. 1995. "Regional Labor Market Dynamics in Europe." *European Economic Review* 39, no. 9: 1627–55.
- Doepke, Mathias, and Fabrizio Zilibotti. 2008. "Occupational Choice and the Spirit of Capitalism." *Quarterly Journal of Economics* 123, no. 2: 747–93.

- Duggan, Mark, and Scott A. Imberman. 2009. "Why Are the Disability Rolls Skyrocketing? The Contribution of Population Characteristics, Economic Conditions, and Program Generosity." In *Health at Older Ages: The Causes and Consequences of Declining Disability among the Elderly*, ed. by David M. Cutler and David A. Wise. Chicago: University of Chicago Press.
- Emerson, Michael, and Andre Dramais. 1988. *What Model for Europe?* Cambridge, Mass.: MIT Press.
- Erceg, Christopher J., and Andrew T. Levin. 2013. "Labor Force Participation and Monetary Policy in the Wake of the Great Recession." Boston Fed Working Paper, available at http://www.bostonfed.org/employment2013/papers/Erceg_Levin_Session1.pdf.
- Garin, Julio, Michael Pries, and Eric Sims. 2013. "Reallocation and the Changing Nature of Economic Fluctuations." Manuscript available at http://www3.nd.edu/~esims1/gps_june27_2013.pdf.
- Gilchrist, Simon, and Egon Zakrajsek. 2012. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review* 102, no. 4: 1692–1720.
- Gorodnichenko, Yuriy, and Gerard Roland. 2010. "Culture, Institutions and the Wealth of Nations." NBER Working Paper no. 16368. Cambridge, Mass.: National Bureau of Economic Research.
- Government Accountability Office. 2012. "Unemployed Older Workers: Many Experience Challenges Regaining Employment and Face Reduced Retirement Security." Report to the Chairman, Special Committee on Aging, U.S. Senate. GAO-12-445. 2012 (April).
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2011. "Civic Capital as the Missing Link." In *Social Economics Handbook*, vol. 1, ed. by Jess Benhabib, Alberto Bisin, and Matthew O. Jackson. Elsevier.
- Jaimovich, Nir, and Henry Sui. 2013. "The Trend Is the Cycle: Job Polarization and Jobless Recoveries." Manuscript available at <http://faculty.arts.ubc.ca/hsui/research/polar20130613.pdf>.
- Johnson, Richard W., and Barbara Butrica. 2012. "Age Disparities in Unemployment and Reemployment during the Great Recession and Recovery." Issue Brief no. 2012-03. Washington: Urban Institute, Program on Retirement Policy.
- Jorda, Oscar, Moritz Schularick, and Alan Taylor. 2011. "Financial Crises, Credit Booms, and External Imbalances: 140 Years of Lessons." *IMF Economic Review* 59: 340–78.
- Kaplan, Greg, and Sam Schulhofer-Wohl. 2012. "Understanding the Long-Run Decline in Interstate Migration." Working Paper no. 18507. Cambridge, Mass.: National Bureau of Economic Research.
- Karabarbounis, Loukas, and Brent Neiman. 2013. "The Global Decline of the Labor Share." Working Paper no. 19136. Cambridge, Mass.: National Bureau of Economic Research.
- Leduc, Sylvain, and Daniel Wilson. 2012. "Roads to Prosperity or Bridges to Nowhere? Theory and Evidence on the Impact of Public Infrastructure Investment." Working Paper no. 12-04. San Francisco: Federal Reserve Bank of San Francisco.

- Machin, Stephen, and Alan Manning. 1999. "The Causes and Consequences of Long-Term Unemployment in Europe." In *Handbook of Labor Economics*, vol. 3, 1st edition, ed. by O. Ashenfelter and D. Card. Elsevier.
- Mueller, Andreas, Jesse Rothstein, and Till von Wachter. 2013. "Unemployment Insurance and Disability Insurance in the Great Recession." Working Paper no. 19672. Cambridge, Mass.: National Bureau of Economic Research.
- Nakamura, Emi, and Jón Steinsson. 2011. "Fiscal Stimulus in a Monetary Union: Evidence from U.S. Regions." Working Paper no. 17391. Cambridge, Mass.: National Bureau of Economic Research.
- Nickell, Stephen. 1997. "Unemployment and Labor Market Rigidities: Europe versus North America." *Journal of Economic Perspectives* 11, no. 3: 55–74.
- Piketty, Thomas, and Emmanuel Saez. 2003. "Income Inequality in the United States: 1913–1998," *Quarterly Journal of Economics* 118, no. 1: 1–39.
- Plane, David A. 1992. "Age-Composition Change and the Geographical Dynamics of Interregional Migration in the U.S." *Annals of the Association of American Geographers* 82, no. 1: 64–85.
- Reinhart, Carmen M., and Kenneth S. Rogoff. 2009. "The Aftermath of Financial Crises." *American Economic Review* 99, no. 2: 466–72.
- Romer, Christina, and Jared Bernstein. 2009. "The Job Impact of the American Recovery and Reinvestment Plan." January 9, 2009. Available at http://otrans.3cdn.net/45593e8ecbd339d074_13m6bt1te.pdf.
- Romer, Christina D., and David H. Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review* 94, no. 4: 1055–84.
- Stevenson, Betsey, and Justin Wolfers. 2011. "Trust in Public Institutions over the Business Cycle." *American Economic Review* 101, no. 3: 281–87.
- Tobin, James. 1972. "Inflation and Unemployment." *American Economic Review* 62, no. 1: 1–18.

Comments and Discussion


COMMENT BY

EDWARD GLAESER In 1969, only 5.1 percent of American males between the ages of 25 and 64 were not employed, which was only slightly below the post-1948 average of 6.3 percent. In the third quarter of 2013, 17.3 percent of prime-aged American males were not employed, which is close to the average level of nonemployment since 2009. This remarkable change partially reflects the lingering effects of the Great Recession, but it also appears to represent something deeper: a worrisome trend in American nonemployment. My figure 1 shows the nonemployment rate for prime-aged males during the first quarter of each year between 1948 and 2013.

Olivier Coibion, Yuriy Gorodnichenko, and Dmitri Koustas name the rise in U.S. unemployment persistence “Amerisclerosis,” a reference to Herbert Giersch’s description of Europe’s post-1970 economic malaise. The unemployment puzzle that these authors identify and analyze is not the steady upward drift seen in the nonemployment rates, but rather an increasing tendency of high unemployment rates to persist after an adverse economic shock. They show that in the wake of recessions before 1990, unemployment reverted to normal four years after each shock, but after 1990 fewer than half of the recessionary unemployment bursts had abated after 16 quarters. The three post-1990 recessions appear to have had similarly sluggish recoveries, although the wildly different magnitudes of the downturns limit our ability to lump the three together. Certainly, the persistence of unemployment after the Great Recession resembles the slow European recovery after the recession of the early 1980s, not the rapid American resurgence of those years.

Yet I cannot view the rise in nonemployment and the sluggish recoveries as distinct phenomena. The upward drift in nonemployment takes the form of sporadic upward bursts in unemployment associated with recessions that

Figure 1. Nonemployment Rate for Males Ages 25–54, 1948–2013



Source: Author’s calculations.

fade slowly only to be overtaken by a subsequent upward burst, pushing toward a new historic high. The persistence of unemployment after recessions occurs in an economy in which an ever larger share of American men leaves the labor force.

The paper’s title suggests that America’s labor markets have become as troubled as many European labor markets. However, labor conditions in the United States seem much healthier than those in France or Italy today. The U.S. Bureau of Labor Statistics indicates that labor force participation rates for men in the United States are about 10 percentage points higher than in those countries. Unemployment rates are also higher in Europe. Yet if America is compared not with Europe today, but with the Europe of the mid-1980s, when “Eurosclerosis” was first being discussed, the comparison looks far more apt. The United States’ 70.2 percent male labor force participation rate in 2012 was exactly the same as France’s male labor force participation rate in 1983.

Coibion, Gorodnichenko, and Koustas divide their explanations for increasingly persistent U.S. unemployment shocks into three broad groups of sources: (i) business cycle fluctuations, (ii) the changing nature of policy

responses to those fluctuations, and (iii) “propagation mechanisms.” The first two groupings are straightforward. The third is more of a grab bag of potential explanations for the puzzle. As a microeconomist, I might have been slightly more comfortable with a division that grouped explanations into bins related to labor demand, labor supply, and the matching mechanism that connects the two, but their structure serves well in many ways.

They follow Carmen Reinhart and Kenneth Rogoff (2009) in dividing recessions into busts associated with financial market shocks and downturns associated with more prosaic causes, such as monetary contractions. Their hypothesis is that unemployment has persisted since 1990 because post-1990 downturns have been primarily financial in nature. They test this hypothesis by estimating impulse-response functions connecting “exogenous” events with unemployment rates. They identify monetary policy shocks using the instruments devised by Christina Romer and David Romer (2004), and they identify financial shocks using the corporate bond premium–based classification of Simon Gilchrist and Egon Zakrajsek (2012).

While the magnitude of financially induced recessions may be larger than the magnitude of monetarily induced downturns (they do not discuss this), the persistence of the two types of events seems broadly similar. In both cases, they estimate, the unemployment shock disappears by about 10 quarters after the shock, and the pattern of decline is broadly similar. Their estimates also indicate that the financial shocks that occurred in 2007 do not predict an unemployment increase as large as that experienced after 2009, that is, where unemployment would persist until 2012.

I found this persuasive, but I am not likely to be the most critical reader of this section. I had no strong prior belief that financial crises should lead to more persistent unemployment than monetarily induced crises.

For those who want to maintain the financial crisis hypothesis, however, their natural criticism of the approach taken in the paper would be that the measure of financial crises used in the empirics is not a particularly clean or comprehensive measure of financial shocks, and consequently the authors’ estimates of the true impact of financial shocks are wrong and the Great Recession was a far more severe financial crisis than their measure would indicate.

The Gilchrist and Zakrajsek (2012) measure is surely an imperfect instrument for financial shocks, and those authors’ original paper does not suggest that their measure would be appropriate for such a purpose. Excess corporate bond premiums are certainly not orthogonal to much of anything in the larger economy, and these premiums should also be responsive to

expectations about future economic activity. Bond premiums will certainly be influenced by financial crises, but that does not make them a valid tool for identifying the impact of those crises. Moreover, there can be many aspects of a financial crisis that will not be captured by excess corporate bond premiums. As a result, the size of these premiums may radically understate the degree of financial dislocation experienced in the recent downturn.

However, I do think that Coibion, Gorodnichenko, and Koustas have taken a reasonable first step, and their conclusion is further bolstered by the similarities in unemployment persistence between the three post-1990 recessions. While it is plausible that the 2007 recession experienced a particularly destructive financial aftermath, the world enjoyed a credit orgy in the years after the 2001 recession. I cannot see how lingering financial turmoil could have been the cause of lingering unemployment after 2001, although it is worth emphasizing that this lingering unemployment was still only about 5 percent.

The second major hypothesis investigated by the paper's authors is that the policy response to recent downturns has been unusually anemic, allowing unemployment to persist more than the more robust policy responses of the past. The weak monetary response is seen as the result of the zero lower bound on nominal interest rates, which effectively barred the Federal Reserve Board from taking more aggressive steps. The weak fiscal response is understood by the authors to be the result of political factors that reined in U.S. spending after the midterm election of 2010.

The authors create a counterfactual series for both monetary and fiscal shocks, assuming that monetary policy would have been able to create negative nominal interest rates and that government spending had followed the pre-1990 norms. They use their previous impulse-response function linking unemployment with monetary shocks to estimate the impact that negative interest rate policies would have had on unemployment. They use a combination of Okun's law and a standard fiscal multiplier to estimate the impact of a more persistently positive fiscal policy.

While each aspect of their procedure is debatable, I found the overall approach compelling and ultimately convincing. The unusually slow response after 1990 is entirely accounted for by relatively weak monetary and fiscal responses. Each policy instrument appears to be responsible for about one-half of the persistence of unemployment during those years. For the latest two recessions, weak policy responses can explain perhaps one-third of the persistently high unemployment.

These results might give succor to those who advocated a more persistent fiscal response to the downturn. If the zero lower bound on interest

rates is inescapable, then fiscal policy would seem to be the natural alternative. Yet assuming a multiplier of 1.5 and Okun's law essentially guarantees that more government spending will imply lower unemployment. The foes of fiscal stimulus can plausibly disagree with both coefficients, however, and emphasize the large-scale deficit spending that carries other costs.

But even if the weak fiscal and monetary response to recent recessions can explain one-third of the persistence of unemployment, two-thirds of the effect is left to explain. The authors therefore consider four potential alternatives: downward wage rigidity, declining geographic mobility, a shifting age structure, and declining levels of trust.

DOWNWARD WAGE RIGIDITY The downward wage rigidity hypothesis is that in previous, higher inflation environments, it was easier for firms to cut real wages without cutting nominal wages. If nominal wage cuts are forbidden, either because of implicit bargains or behavioral ticks (nominal wage cuts can make employees angry and destructive), then inflation can make it easier to adjust wages downward. In an implicit contracts view of nominal wage rigidity, inflation is essentially allowing firms to renege on their implicit contracts, but this may have the salutary effect of reducing unemployment.

The evidence in favor of this hypothesis is that 16 percent of workers received no wage changes in 2011, double the rate of 1981, which suggests that there were plenty of cases in which firms would have reduced real wages further if they had the freedom to do so. The authors downplay this fact by pointing out that the increase in zero wage changes over the recession was no larger in 2011 than it had been in 1981 and that there is no change in the shape of the Phillips curve.

I am not sure that I found the similar increase in the level of zero wage changes to be compelling evidence against the importance of downward wage rigidities for layoffs. Consider a simple model in which the wages of employees are downwardly rigid, all employees have a common outside option, and firms and employees Nash bargain, so that wages at time t , denoted w_{it} , equal $w_t^{Res} + \sigma(\theta_{it} - w_t^{Res})$, where w_t^{Res} is the reservation wage, θ_{it} represents productivity for person i within the firm, and σ represents the Nash bargaining rule. Workers who are less productive than their reservation wage are laid off.

Between time t and $t+1$, reservation wages change, and productivity changes from θ_{it} to θ_{it+1} and the reservation wage changes to w_{t+1}^{Res} . I will treat these as real, time t values, where Δ_t represents the economywide shift in productivity, and ϵ_{it+1} represents a personal shock to productivity. The downward wage rigidity means that wage at time $t+1$ cannot be

below $\frac{1}{1 + \pi}$ times the wage at time t . If $\theta_{it+1} < \text{Max}\left(\frac{w_{it}}{1 + \pi}, w_{it+1}^{Res}\right)$, then the worker is laid off.

If $w_{it+1}^{Res} + \sigma(\theta_{it+1} - w_{it+1}^{Res}) < \frac{w_{it}}{1 + \pi} < \theta_{it+1}$, then the worker receives no wage change but remains employed. Finally, if $w_{it+1}^{Res} + \sigma(\theta_{it+1} - w_{it+1}^{Res}) > \frac{w_{it}}{1 + \pi}$, then the worker receives a wage increase.

I assume that $\theta_{it} = \theta_i$, hence $w_{it} = w_i$, so workers are initially identical (this can be understood as focusing on a particular subgroup of the population), and $\theta_{it+1} = \varphi_i w_i$, where φ_i is a density function $f(\cdot)$ and a cumulative distributive $F(\cdot)$. Moreover, I use the notation $w_{it+1}^{Res} = \delta w_i$.

If $\frac{1}{1 + \pi} < \delta$, then the nominal wage rigidity has no effect, and I ignore that case. If $\frac{1}{1 + \pi} > \delta$, then a fraction $F\left(\frac{1}{1 + \pi}\right)$ of workers are laid off, and the fraction $\int_{\frac{1}{1+\pi}}^{\frac{1}{\delta}} f(\varphi_i) d\varphi_i$ represents the excess layoffs due to the wage downward rigidity. A fraction $\frac{1}{\int_{\frac{1}{1+\pi}}^{\frac{1}{\delta}} f(\varphi_i) d\varphi_i} - \frac{(1 - \sigma)\delta}{\sigma}$ receive zero wage changes.

The empirical fact is that the change in the share of employees receiving zero wage changes before and after a recession is the same in both a high- and a low-inflation environment. If the advent of a recession represents a shift to a distribution $f_R(\varphi_i)$ of productivities and a shift of δ to δ_R , then the difference in zero wage changes equals:

$$(1) \quad \int_{\frac{1}{1+\pi}}^{\frac{1}{\sigma(1+\pi)} \frac{(1-\sigma)\delta_R}{\sigma}} f_R(\varphi_i) d\varphi_i - \int_{\frac{1}{1+\pi}}^{\frac{1}{\sigma(1+\pi)} \frac{(1-\sigma)\delta}{\sigma}} f(\varphi_i) d\varphi_i.$$

This difference appears empirically to be independent of the inflation rate, which suggests that $\frac{1}{\sigma} \left[f_R \left(\frac{1}{\sigma(1+\pi)} - \frac{(1-\sigma)\delta_R}{\sigma} \right) - f \left(\frac{1}{\sigma(1+\pi)} - \frac{(1-\sigma)\delta}{\sigma} \right) \right] = f_R \left(\frac{1}{1+\pi} \right) - f \left(\frac{1}{1+\pi} \right)$. The share of the population that is laid off during a recession due to wage rigidity is $\int_{\frac{1}{1+\pi}}^{\frac{1}{\delta_R}} f_R(\varphi_i) d\varphi_i$, and the impact of the inflation on this share is $-\frac{1}{(1 + \pi)^2} f_R \left(\frac{1}{1 + \pi} \right)$. The extra wage rigidity related

to layoffs in a recession equals $\int_{\delta_r}^{\frac{1}{1+\pi}} f_R(\varphi_i) d\varphi_i - \int_{\delta}^{\frac{1}{1+\pi}} f(\varphi_i) d\varphi_i$ and the inflation impact on this quantity equals $-\frac{1}{(1+\pi)^2} \left(f_R\left(\frac{1}{1+\pi}\right) - f\left(\frac{1}{1+\pi}\right) \right)$.

As such, if $\frac{1}{\sigma} \left(f_R\left(\frac{1}{\sigma(1+\pi)} - \frac{(1-\sigma)\delta_r}{\sigma}\right) - f\left(\frac{1}{\sigma(1+\pi)} - \frac{(1-\sigma)\delta_r}{\sigma}\right) \right)$
 $= f_R\left(\frac{1}{1+\pi}\right) - f\left(\frac{1}{1+\pi}\right)$, because $f_R\left(\frac{1}{\sigma(1+\pi)} - \frac{(1-\sigma)\delta_r}{\sigma}\right) = f\left(\frac{1}{\sigma(1+\pi)} - \frac{(1-\sigma)\delta_r}{\sigma}\right)$ and $f_R\left(\frac{1}{1+\pi}\right) = f\left(\frac{1}{1+\pi}\right)$, then this evidence does indeed imply that inflation will not have caused relatively fewer layoffs during a recession. However, that need not be the case. The quantity $f_R\left(\frac{1}{1+\pi}\right) - f\left(\frac{1}{1+\pi}\right)$ might be large and might simply be offset by $\frac{1}{\sigma} \left(f_R\left(\frac{1}{\sigma(1+\pi)} - \frac{(1-\sigma)\delta_r}{\sigma}\right) - f\left(\frac{1}{\sigma(1+\pi)} - \frac{(1-\sigma)\delta_r}{\sigma}\right) \right)$.

Moreover, even if $f_R\left(\frac{1}{1+\pi}\right) - f\left(\frac{1}{1+\pi}\right)$ is zero, then this does not imply that lower inflation had no impact on the absolute number of layoffs during the downturn. Those layoffs could easily have translated into more unemployment, especially if they led to congestion in the matching function.

By these comments I am not suggesting that I believe downward wage rigidity explains the increasing persistence of unemployment or even, more importantly, that it explains the decreasing employment ratio among adult males. Since wages are far more flexible for new hires, it seems likely that these rigidities should primarily lead to a modest increase in layoffs, which should presumably be reabsorbed if the labor market is otherwise healthy. My point, instead, is that the paper's evidence that there was the same change in the number of zero wage changes during the 2007 and 1981 recessions just does not imply that downward wage rigidity is irrelevant.

Better evidence for the authors' claim is the fact that inflation was also low during the pre-1970 period, when recoveries were particularly swift. The 1970–85 period was essentially the only period when inflation could easily enable nominal wage cuts. The key test of a link between downward wage rigidity and unemployment persistence is whether recoveries were

particularly swift during those years, not whether recoveries were particularly sluggish after that period.

ALTERNATIVE EXPLANATIONS The final section of the paper uses fascinating regional data to test whether regional mobility, age composition, or “trust” can explain the rising persistence of unemployment. A major issue in this paper is the authors’ dependent variable, which takes the form:

$$(2) \quad \ln(1 + u_{1990}^{i,c} - u_{1980}^{i,c}) + \ln(1 + u_{1980}^{i,c} - u_{1970}^{i,c}),$$

where $u_{Year}^{i,c}$ refers to the unemployment rate in country c in region i during a particular year. The first-order Taylor series approximation for this quantity (around one) is $u_{1990}^{i,c} - u_{1970}^{i,c}$, so we can think of these as essentially 20-year unemployment rate change regressions. The authors choose to interpret this as “hysteresis” but probably the more natural interpretation is that they are really just measuring the growth in unemployment over the time period. I will use that more transparent phrasing when discussing their result.

It is a major step forward that they use both within- and across-country sources of variation, but there are substantial differences between national and subnational empirical data. At the subnational level, especially in high-mobility countries like the United States, regional variables and outcomes reflect migration as much as anything else. This mobility, which is typically operationalized by assuming a spatial equilibrium, makes the interpretation of local and national effects quite different.

For example, at the country level, rising productivity levels are typically expected to translate into rising income levels. At the subnational level, in-migration should respond to higher productivity and, at the extreme, productivity boosts lead to no income gains but only increases in wages. The authors’ variable—essentially the change in the unemployment rate—could increase as a region becomes productive, because that productivity attracts a large flow of potentially marginal workers. The classic Harris-Todaro model suggests that high unemployment offsets higher wages, making the interpretation of rising unemployment rates particularly problematic.

A simple example of the differences between national and subnational data is the authors’ estimate of age effects. Without country fixed effects, age is essentially irrelevant. With country effects, areas with older populations (and areas with very young populations) initially are estimated to have faster unemployment growth. A natural interpretation is that this latter result says nothing about the actual impact of age but rather that more

dynamic local economies attracted more middle-aged people and had less growth in unemployment over the 1970–90 period.

The turnover variable is particularly problematic at the subnational level. The effect of having an abundance of turnover in a region in 1970 on unemployment growth afterward is marginally significant and positive. I can understand comparing high mobility countries with low mobility countries, but at the regional level the correlation is hard to interpret. Imagine that the United States is generally quite mobile and that population flows from high unemployment regions to low unemployment regions. The high mobility regions are likely to be at the extremes—those that attract and those that repel populations. The measured correlation at the regional level between mobility and either the level of unemployment or change in unemployment will be zero, if unemployment rates converge because of this mobility. But these measured correlations will tell us nothing about the impact that mobility has on unemployment across the country as a whole.

I have somewhat more faith in the authors' most robust result: the powerful correlation between unemployment growth and initial education levels. A large number of papers have found that more highly educated areas have grown more quickly than less educated areas (for example, Glaeser, Scheinkman, and Shleifer 1995) using any number of different outcomes. The correlation between unemployment in the recent downturn and area-level education is particularly extreme. However, since education has typically been rising, this variable offers little chance of enabling us to better understand persistent underemployment in the United States and elsewhere.

A particularly interesting result is the correlation that the paper's authors find between self-reported "trust" and the growth in unemployment. In the ordinary least squares specifications, the correlation is not robust, but it does appear to be strong when one looks only at subnational data. This variable is calculated from the 2008 World Values Survey, which makes it possible, if not likely, that the chain of causality runs the opposite direction: from rising unemployment to less trust rather than from less trust to rising unemployment.

But even if these regressions do not offer compellingly causal estimates of the impact of any of these variables, they are still interesting. The changing patterns of unemployment are important, and the authors have made an important contribution by looking at subnational data across nations. I agree with them that it doesn't seem likely that these hypotheses can succeed at explaining the phenomenon of increasingly persistent unemployment.

My own limited understanding of Eurosclerosis has been strongly shaped by Olivier Blanchard and Justin Wolfers' (2000) argument that institutions and shocks came together to create longer-term problems. The increasing prevalence of institutions meant either to protect dislocated workers or to safeguard existing employees causes relatively little problem when these institutions are first established, often in relatively good times. When an adverse shock hits, however, the laid-off workers face fewer incentives to find new work because they are protected by the more robust safety net. Simultaneously, employers face larger fixed costs of hiring new workers, so new jobs do not take up the slack.

This story has always been appealing to me, both because of its empirical fit and because it has a compelling connection with the microeconomics of labor markets. The microeconomic view of rising nonemployment, and also of persistent unemployment, is that these outcomes can be understood as reflecting weak labor demand, weak labor supply, or potentially some breakdown in the matching process that connects workers to employers, such as a decline in regional mobility. Labor market institutions that protect displaced workers work against labor supply. Institutions that require more from potential employers work against labor demand.

The long-run effects of these factors may reflect skill depreciation among out-of-work employees. Indeed, one of the impacts of a more skill-oriented world may well be that unemployment persistence will become more severe. A laborer's ability to work on a dock may not depreciate much after being out of work for six months, whereas an ability to use the latest technology may diminish significantly with absence from the workforce. If learning-by-doing is important, then the rise in returns to skill makes unemployment even more costly and it becomes more likely that workers will fall into long-term unemployment traps.

In the United States, the Blanchard and Wolfers (2000) argument also seems relevant, albeit with a slight lag. The rising nonemployment of prime-age males and the persistence of unemployment after recessions would seem to reflect the same phenomenon. Downturns turn into permanent unemployment when institutions reduce reemployment and skills depreciate.

There are several natural candidates for the institutional change that may be making the United States more like Europe. There has been a vast increase in the number of disabled workers, despite the fact that Americans are on the whole getting healthier and jobs are typically becoming less dangerous. This change is surely best seen as an institutional change meant to better protect workers. The downside is declining employment levels.

The problem identified in this paper is enormously important. The authors have done the world a service by focusing our attention on rising unemployment persistence after recessions. They have also successfully eliminated some candidate explanations for the change. One hopes that future work will more fully identify the causes of the change, which I suspect are more likely to reflect the interaction between institutions and shocks. Even more importantly, I hope that future work will point the way toward reversing the trend.

REFERENCES FOR THE GLAESER COMMENT

- Blanchard, Olivier, and Justin Wolfers. 2000. "The Role of Shocks and Institutions in the Rise of European Unemployment: The Aggregate Evidence." *Economic Journal* 110, no. 462: 1–33.
- Gilchrist, Simon, and Egon Zakrajsek. 2012. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review* 102, no. 4: 1692–1720.
- Glaeser, Edward L., José A. Scheinkman, and Andrei Shleifer. 1995. "Economic Growth in a Cross-Section of Cities." *Journal of Monetary Economics* 36, no. 1: 117–43.
- Reinhart, Carmen M., and Kenneth S. Rogoff. 2009. "The Aftermath of Financial Crises." *American Economic Review* 99, no. 2: 466–72.
- Romer, Christina D., and David H. Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review* 94, no. 4: 1055–84.

COMMENT BY

PAOLA SAPIENZA During the Great Recession, the U.S. economy shrank by 4.14 percent while payroll employment experienced the largest drop of any U.S. recession. While many European countries experienced a similar or larger drop in output, their employment rates did not respond so dramatically, at least not in the direct aftermath of the crisis. In the past, it may have been felt that while such differences in lost employment were worrisome, the United States often wins back the losses during recoveries, experiencing much lower unemployment persistence and a faster rebound than Europe.

However, this time there may be reason to worry. Job recovery in recent years has been very slow in the United States, and the labor force participation rate is at its lowest level since 1978. These developments pose questions about whether the U.S. labor market's prolonged job losses following a recession are becoming dangerously similar to those of many

western European economies, and thus whether the lessons learned in the 1980s regarding differences between the U.S. and European labor markets (Blanchard and Summers 1986) still apply. In the 1970s, the slow job recovery in many European countries was often linked to frictions in their labor market institutions. Is it possible that the United States is becoming “more European”? If so, what are the driving forces?

This paper makes an attempt to answer these important questions by looking at regional data within the United States and Europe. This approach makes sense because wide cross-sectional variation within countries can inform us about general trends. To describe the problem, the paper’s coauthors, Olivier Coibion, Yuriy Gorodnichenko, and Dmitri Koustas, compare the missing job recovery after the Great Recession to the United States’ response to the Volker recession in the 1980s. They point out that rather than being a recent phenomenon, jobless recoveries are part of a longer trend that has affected the country for the last three recessions. Indeed, another alarming aspect of this phenomenon is the reduction in labor force participation that originates with a recession and persists over time. In fact, I would argue that a valid explanation for the persistence in unemployment should be able to account for the rising fraction of discouraged workers in the United States.

To make sense of the discrepancy in job recoveries between earlier recessions and the last three recessions, the authors consider differences in the propagation mechanisms and fiscal and monetary policy responses. While they find some evidence that timid fiscal and monetary policy responses can account for a large fraction (one third) of the differences over time, most of the differences remain unexplained. They then identify three alternative, non–mutually exclusive causes for the unexplained differences. In particular, they focus on: (i) declining labor mobility; (ii) age changes in the working population; and (iii) cultural changes, specifically a decline in generalized trust. The authors admit that these are not exhaustive explanations; nonetheless, they consider them a first step.

However, there is a problem with this list. With the exception of age changes in the workforce, cultural norms and labor mobility are likely to be affected by labor demand or supply characteristics and, thus, to be endogenous. Even if we find a correlation between the labor market trends and some of these variables, we will not be able to provide a complete explanation of the driving forces affecting unemployment persistence. While frictions to labor mobility can contribute to it, labor mobility by itself can simply be the result of labor market characteristics. For example, reductions in labor mobility can be explained by either supply or demand factors.

On one side, labor mobility can be reduced by some structural constraints on job losers. This factor has often been mentioned as one of the labor market frictions in Europe. Alternatively, deficient demand may give workers fewer reasons and opportunities to move than in previous times. Understanding the causes of reduced worker mobility is essential to interpreting the data.

As for the culture explanation, the analysis is even more problematic. There is a long-standing debate in economics about whether culture affects economics or vice versa. While several papers show, in specific settings, the existence of a causal link between cultural norms and economic outcomes (Guiso, Sapienza, and Zingales [2006] provide an overview), there is no doubt that economic relations have an independent impact on culture. The authors highlight some interesting trends in cultural norms and behavior in the United States. First, the percentage of workers claiming Social Security disability benefits spiked during the last recession, with the ratio of claimants to insured workers increasing by about 30 percent from 2007 to 2010. Second, the percentage of people responding in surveys that it is acceptable to claim benefits to which one is not entitled has increased over time, from 3.6 percent in 1994–99 to 9 percent in 2005–07, surpassing the 7.2 percent average across European countries. Third, trust toward others, a measure often connected to social capital and civicness, has declined. These trends raise a legitimate question about whether unemployment persistence is somewhat related to cultural shifts.

The mechanics are not entirely spelled out in the paper, but one may argue that exogenous cultural changes may affect people's desire to work hard and increase their willingness to claim disability benefits rather than search for jobs. An alternative story, which is more compelling, reverses the causality: disenfranchised workers who have lost their jobs and have exhausted their unemployment benefits are more likely to find it justifiable to claim benefits they are not entitled to and to mistrust others in general. The authors rule out this latter reverse causal relationship by running an aggregate regression of unemployment rates on cultural norms and values.

But I am not convinced. The fundamental question is whether unemployed individuals are more likely than others to change their perceptions and cultural norms due to the dire economic conditions they are facing. A more reasonable specification should rely on disaggregate data and investigate whether unemployed individuals are more likely to mistrust others and find it acceptable to claim benefits they are not entitled to (controlling for country fixed effect and other individual characteristics) than

Table 1. Regression on the Cultural Norms of Unemployed Workers

	(1)	(2)
	<i>Undeserved benefits</i>	<i>Trust</i>
Unemployed	0.203*** (0.0621)	-0.0499*** (0.0154)
Age	-0.0192*** (0.000864)	0.000748*** (0.000216)
Education	-0.0564*** (0.00736)	0.0342*** (0.00183)
Male	0.172*** (0.0284)	-0.0159** (0.00711)
Income	-0.0354*** (0.00636)	0.0158*** (0.00158)
Constant	3.115*** (0.0996)	0.329*** (0.0273)
Country fixed effects	YES	YES
Time fixed effects	YES	YES
Observations	17,213	17,291
R ²	0.060	0.111

Note: Standard errors are in parentheses. Asterisks denote statistical significance at the *1 percent, **5 percent, and ***10 percent level.

similar employed people. The regression shown in my table 1 investigates this relation using the same source of data, but exploiting individual differences in work conditions.¹

My regression shows that unemployed individuals are indeed more likely to have a culture of low trust, controlling for basic demographic characteristics, time, and a country fixed effect. Compared to similar individuals, the unemployed are 8 percent more likely than the sample

1. The data in my table are from the World Values Surveys for all the waves in which the questions have been asked (from wave 2, 1989–93, to wave 5, 2005–08). The countries included in the regressions are those included in the analysis, namely Canada, Finland, France, Germany, Italy, The Netherlands, Spain, Sweden, Switzerland, Great Britain, and the United States. *Claim benefits not entitled to* is the answer to the question, “Please tell me . . . whether you think it can always be justified, never be justified, or something in between, using this card: Claiming government benefits to which you are not entitled, on a scale between 1 and 10 (1 = always, 10 = never).” *Trust* is the answer to the question, “Generally speaking, would you say that most people can be trusted or that you need to be very careful in dealing with people?” (0 = can’t be too careful, 1 = most people can be trusted). Education is coded on a scale between 1 (= inadequately completed elementary education) and 8 (= university with degree/higher education-upper-level tertiary certificate), and income is coded on a scale between 1 (= lower step) to 10 (= 10th step).

mean to judge it acceptable to claim benefits one is not entitled to. Also, other things being equal, a change in employment status from employed to unemployed lowers a person's average trust by 15 percent of the sample mean.

These regressions suggest that reverse causality could explain why cultural norms are changing. As the percentage of unemployed individuals grows over time, average cultural norms may shift. While it is entirely possible that a larger unemployed population perceives cultural norms differently, we are left in the dark regarding the causes of unemployment hysteresis.

While the paper nicely puts together many long-term trend variables, it is ultimately unsuccessful in explaining the reasons why, after the last three recessions, U.S. recoveries have been characterized by rising unemployment persistence. Thus, as discussant, it is probably my job to point to some possible avenues for future research. One compelling story ignored in this paper could be related to the increased job polarization that occurred in the U.S. labor market (Autor, Katz, and Kearney 2006). Nir Jaimovich and Henry Siu (2012) show that unemployment persistence is concentrated among routine manual workers in the United States. These low-skill workers have seen their jobs disappear due to technology and outsourcing in recent years, but especially during recession times. Since these jobs have been substituted by technology or foreign workers, it is likely that they are not coming back, causing longer, possibly permanent, unemployment spells.

To what extent could this last hypothesis fit the neat data trends presented in the paper? The disappearance of manual routine jobs certainly could be related to the increase in long-term unemployment and the decrease in labor force participation. As specific occupations are wiped out, it takes much longer for some types of workers to find jobs, and these workers may become discouraged. While this theory could explain why the phenomenon has emerged in recent decades, namely due to the emergence of substitute technology and outsourcing, the interesting question is why a replacement of workers is more likely to happen during recessions. As far as I know, this question has not been explored much. There are at least three possible, non-mutually exclusive explanations that could be investigated. (i) It is possible that it is politically more feasible for companies to outsource jobs or replace them with technology during recession times due to an understanding that labor force reductions are inevitable. (ii) In good times, firms can afford to be less efficient, but in bad times they are forced to optimize and eliminate jobs, saving money. (iii) Substituting

technology and foreign workers for local workers is often not a smooth process: it is easier to implement it during times of lower demand.

The job polarization hypothesis is also consistent with a few other trends highlighted in the paper. If jobs with specific skill-sets are eliminated, it is reasonable to expect to observe lower worker mobility. If jobs are wiped out by technology or outsourcing, unemployed people have fewer incentives to move. The hypothesis may also be consistent with increased demand for disability benefits. Since unemployment benefits are capped, the unemployed, or those with manual routine skill-sets foreseeing their limited chance to reenter the labor market, may be more prone to apply for disability benefits.

This hypothesis poses an important question for policymakers. If jobless recoveries are driven by technology and global shifts in the labor market, how do we design job market institutions to deal with these new trends and increased job polarization? The old U.S.-Europe dichotomy regarding friction versus frictionless labor markets may no longer be valid. If so, we should then expect that even in a relatively frictionless labor market like that in the United States, some workers with limited education may experience very long unemployment spells, or may possibly be excluded from the labor market forever.

Thus, the simple recipe of keeping friction low in the labor market may no longer be enough to ensure a socially acceptable recovery for those who lose jobs during recessions. New thinking may be required to design labor market institutions that interact with structurally jobless recoveries, especially considering a new duality in the labor market. Ultimately, jobless recoveries and their corresponding policy responses may also have a long-term effect on cultural norms, though one may need to wait a bit longer before those can be measured.

REFERENCES FOR THE SAPIENZA COMMENT

- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2006. "The Polarization of the U.S. Labor Market." *American Economic Review* 96, no. 2: 189–94.
- Blanchard, Olivier J., and Lawrence H. Summers. 1986. "Hysteresis and the European Unemployment Problem." In *NBER Macroeconomics Annual 1986*, vol. 1, edited by Stanley Fischer. Cambridge, Mass.: MIT Press.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2006. "Does Culture Affect Economic Outcomes?" *Journal of Economic Perspectives* 20, no. 2: 23–48.
- Jaimovich, Nir, and Henry E. Siu. 2012. "The Trend Is the Cycle: Job Polarization and Jobless Recoveries." NBER Working Paper no. 18334. Cambridge, Mass.: National Bureau of Economic Research.

GENERAL DISCUSSION Robert Gordon recalled that whereas in 2007 potential GDP growth was estimated to be about 2.5 percent, recent estimates have been lower, around 2.1 percent. This presents a paradox, because if potential GDP were growing at the more recently estimated rates the unemployment rate should have been rising and not declining. Reminding everyone that GDP could be decomposed into productivity per hour and hours worked, he pointed out that the downward trend in output growth could be attributed to a decline in the trend growth of total hours worked. One must then ask why the number of work hours has declined. Looking to the relationship between GDP and labor demand, Gordon agreed with the point Edward Glaeser had made in his comment that low labor demand was an explanation for why output remained low and unemployment remained high.

Erik Hurst argued that while the paper focused almost exclusively on cyclical variations, long-term trends could also explain much of the current decline in labor force participation. Citing David Autor's work on the decline in manufacturing and Paola Sapienza's work on the decline of routine jobs, Hurst noted that labor force participation has been declining for a long time.

Alan Blinder mentioned earlier work done by Assar Lindbeck on European hysteresis. Lindbeck had argued that uptake for unemployment insurance was higher in Europe because everyone was either applying for it themselves or knew someone who was applying. Blinder then suggested that the difference between Europe and the United States was not actually a difference in cultural norms. Rather, until recently Europe simply had many more unemployed persons than the United States. This could be tested econometrically, he suggested, and doing so might eliminate the need to include culture as an explanatory variable.

Raquel Fernandez disagreed with the widespread notion that culture changes only slowly. Pointing to the examples of changing attitudes toward women's rights and gay marriage, she argued that culture could change quite rapidly. She agreed with Blinder that the authors should focus on labor market fundamentals, but also thought that in the United States attitudes toward unemployment had changed during the Great Recession. She urged the authors to review the micro data more closely when considering cultural variables.

James Stock urged the panel to experiment with an AR(6) regression model on the monthly unemployment rate between 1984 and 2009, as he has found that such a regression predicts the post-2009 unemployment rate very well. Based on that, he believes the employment rate dynamics

have not changed much in the post-1984 period. It was different, he added, for the earlier period, but the commonalities observed since 1984 suggest to him that the authors are right to consider the changes examined here to have happened decades ago rather than reflecting something unique about the post-2007 experience.

Caroline Hoxby suggested that for many American workers in the Great Recession there was a permanent substitute—a replacement either by technology or by an outsourced worker—and that the share of American workers who could be permanently replaced was growing over time. As a result, when a worker was displaced in one of the more recent recessions, especially the Great Recession, he was more likely to drop out of the labor force permanently, perhaps by going on disability insurance. This would account for the greater hysteresis in recent recessions. Hoxby also suggested that because the well-educated cannot be easily substituted, this might explain why the authors' regressions show a negative relationship between the measures of hysteresis and schooling.

Steven Davis disagreed with some of the previous discussants who claimed that the dynamics of the Great Recession looked similar to that in previous recessions. Only the unemployment rate looked similar to what occurred in previous recessions, he argued, but the employment rate and the labor force participation rate did not look similar. He urged the authors to make a compelling argument for why they would examine unemployment and not either of the other two labor market indicators.

Paul Krugman suggested that the dynamics of the business cycle have changed, not just the dynamics of the unemployment rate. Prior to the 1990s, he noted, the standard explanation for the cause of recessions was that the Fed intentionally induced them in order to control inflation. Those recessions were short-lived because all that was necessary to begin the recovery was to loosen monetary policy. By contrast, he argued, since 1990 recessions have been caused by private-sector overreach and, therefore, have been associated with much longer recoveries. He felt that while many interesting changes had occurred in the labor market, the starting point for analyzing them should be to diagnose their causes by looking at changes in the business cycle.

George Perry agreed with Krugman that prior to the most recent recessions, the primary cause of recessions was monetary policy. Perry also agreed with Davis that the Great Recession has differed from the previous two because of changes in labor force participation. But he noted that the retirement of the baby boom generation has made much

of the decline in labor force participation predictable, and proposed that what could not be explained structurally could be explained through the demand side.

Laurence Ball asked whether the authors had accounted for the speed with which the economy recovered from the shallow recessions of the 1990s and early 2000s in seeking to explain the persistence of the Great Recession. Like Krugman and Perry, he believed deep recessions were associated with sharper recoveries. He noted with irony that during the Great Recession the Fed had kept interest rates at the zero lower bound even though the Taylor rule suggested that they should be 500 basis points below it.

Christopher Carroll observed that sometime in the mid-1980s, economists and policymakers became convinced that fiscal policy was not an effective response to recessions. However, the authors' regressions found that fiscal policy had robust explanatory power. This led Carroll to wonder whether the last three recessions were longer than usual because fiscal policy interventions, which actually had been effective before, were not being deployed.

Wendy Edelberg thought that the main concern of the authors' paper was to determine how much of the persistence in the unemployment rate is cyclical and how much of it is structural. She added that this seemed equivalent to trying to define the mismatch in employment—whether it is caused by geographic factors or cultural ones. In either case, she thought that analysts should be looking for evidence of wage pressure somewhere, since that is a natural result of mismatch.

Yuri Gorodnichenko replied to Stock and others by noting that if the labor market dynamics were similar between 1984 and 2009 but different after 2009, a better explanation was needed to explain the current labor market hysteresis. Moreover, if individuals were quitting the labor force the degree of hysteresis should be declining, but instead it remains high. He added that if the unemployment rate and cultural factors such as trust had been endogenous to one another, the paper's results would have placed an upper bound on the impact that culture has on the persistence of unemployment.

Olivier Coibion agreed with many of the discussants who emphasized the importance of labor force participation rates, which he believed were related to skill bias and technological change. He noted that whereas the overall labor force participation rate had not changed much, the rate among certain subpopulations had changed—for instance, it had fallen significantly for men.