The Great Recession and its Aftermath: What Role for Structural Changes?

August 2015

Jesse Rothstein*

University of California, Berkeley and NBER

Abstract

The last eight years have been disastrous for many workers, and particularly so for those with low human capital or other forms of disadvantage. Although the Great Recession officially ended in 2009, the labor market has been very slow to recover. One explanation attributes the ongoing poor labor market outcomes of young and non-college workers to the combination of deficient aggregate labor demand and greater sensitivity of marginal workers to cyclical conditions. A second attributes the recent outcomes to structural changes in the labor market. These have importantly different policy implications: Cyclical explanations imply that the main challenge is to raise aggregate labor demand and that if this is done many of the patterns seen in the last several years will revert to their prior trends. Structural explanations, by contrast, suggest the recent experience is the "new normal," absent policy responses to encourage more (or different) labor supply. This paper reviews recent data for evidence on the two explanations, focusing on wage trends as an indicator of the relative importance of labor supply and demand. I find little evidence of wage pressure in any quantitatively important labor markets before 2015. The most recent data is more mixed, but still suggests substantial ongoing slack.

* rothstein@berkeley.edu. This paper extends and expands on the results in Rothstein (2012a; 2014). I thank Apurba Chakraborty, Darian Woods, and Peter Jones for excellent research assistance, and Russell Sage conference participants for helpful comments. I am grateful to the Smith-Richardson and Russell Sage Foundations for financial support. Copies of the data and computer programs used to generate the results presented in the paper are available from the author.

I. Introduction

In a 2004 speech titled "The Great Moderation", Ben Bernanke – then a member of the Board of Governors of the Federal Reserve but soon to become the chairman – discussed the apparently substantial decline in macroeconomic volatility over the last decades of the 20th century. He argued that this was in important part due to improved monetary policy, and he expressed optimism that the moderation would persist into the future (Bernanke 2004).

Within four years of that speech, the U.S. had fallen into the "Great Recession." Between the fourth quarter of 2007 and the second quarter of 2009, real GDP fell by over 5 percent. The unemployment rate rose from a low of 4.4 percent in May 2007 to a high of 10.0 percent in October 2009, for a 29-month increase of 5.6 percentage points. This far exceeded the largest previous post-war increase over a similar duration, 3.9 percentage points in 1973-75.

The National Bureau of Economic Research (NBER) dated the business cycle trough in June 2009. But neither the real economy nor, especially, the labor market recovered quickly after that point. Real output recovered its pre-recession peak in the third quarter of 2011 (though it has yet to return to the pre-recession trajectory). But payroll employment took nearly seven years, until May of 2014, to reach its December 2007 level. The unemployment rate remained above 8% for 43 consecutive months (until August of 2012) and above 7% for 60 months (until November 2013), each the longest such period since World War II.

As of this writing, the unemployment rate is down to 5.3%, only slightly above its level on the eve of the recession. But most of the decline came from reduced labor force participation rather than increased employment: The employment-to-population ratio, which fell by an unprecedented 4.9 percentage points between December 2006 and December 2009, hovered around 58.5% for more than four years thereafter and has recovered less than one point since then. Its July 2015 level, 59.3%, is four full points below the pre-recession level.

By essentially every measure, low-skilled workers have fared even worse than these already dismal statistics suggest. Consider, for example, the employment-to-population ratio. For workers with high school diplomas but no college (aged 25 and up),

this ratio fell not by 4.9 but by 6.1 points from the peak to December 2009, and has fallen nearly an additional point since then with no sign of recovery to date.

Early in the recovery, some observers diagnosed unseen structural problems as impediments to what would otherwise have been a quick cyclical recovery. In a 2010 speech, for example, Narayana Kocherlakota, President of the Federal Reserve Bank of Minneapolis, stated that "Firms have jobs, but can't find appropriate workers. The workers want to work, but can't find appropriate jobs. There are many possible sources of mismatch—geography, skills, demography—and they are probably all at work" (Kocherlakota 2010).[1] In this view, poor outcomes for low-skill workers are attributable to their failure to supply the skills demanded by the labor market, and there is little hope of improving them short of raising their human capital.

The mismatch theory was eventually discredited by the evidence, as it became clear that – at least at that time – labor market slack was high in nearly all sizeable labor markets.[2] But as the weak recovery has dragged on, now into its seventh year, it has become harder to resist the view that this is the "new normal," and that we are destined for a future of low employment rates and a substantial class of individuals – disproportionately low-skilled – who are more or less permanently detached from the labor market.

This paper reviews data on labor market outcomes over the period since the recession's official end, focusing on the experience of less skilled workers. I argue that there is no basis for concluding that the recent past represents "the new normal" or that labor demand has tilted more rapidly away from low-skilled workers than at other times in recent decades. Rather, the evidence indicates that labor demand remains very weak for nearly all workers.

Less skilled workers' outcomes have always been particularly sensitive to the business cycle, worsening by far more in downturns than do those of more skilled workers and then improving by more when the economy recovers. Thus, if the economy

---

[1] See also Sahin, Song, Topa, and Violante (2011).

[2] See Diamond (2010); Mishel, Shierholz, and Edwards (2010); Mishel (2011); Rothstein (2012a); and Lazear and Spletzer (2012). The Congressional Budget Office (2012) was more favorable toward structural hypotheses but nevertheless concluded that aggregate demand shortfalls were the primary source of the high unemployment rate.

remains in a cyclical trough, it is not surprising that less skilled workers have suffered disproportionately, and it is reasonable to expect that this suffering will ease substantially if and when aggregate labor demand recovers.

An important possibility is that cyclical labor demand shortfalls that extend for many years may *create* structural problems, as idle workers' human capital gradually depreciates and they become increasingly disconnected from the labor market such that they are unable or unwilling to take new jobs when they finally become available. This idea has gained currency as an explanation for our current situation. For example, Krueger, Cramer, and Cho (2014) argue that the long-term unemployed exert little or no pressure on the labor market, and conclude that extra-market measures such as expanded social welfare programs will be needed to support those who remain in this state.

This hypothesis has an important implication: If true, it means that even a labor market that appears to be quite slack from the perspective of workers can be tight from the perspective of employers, who see relatively few qualified, available workers to hire (see, e.g., Hall 2014). Employers facing tight labor markets should bid up the wage in order to attract workers. Labor demand shortfalls, by contrast, would have an opposite effect, as unemployed workers bid down equilibrium wages as they compete for the few available jobs.[3] I thus emphasize the examination of wage trends, both in the aggregate and in particular labor markets, for evidence about the appropriate diagnosis of the current situation.

The remainder of the paper proceeds as follows. In Section II, I review the overall state of the labor market and its performance since the onset of the Great Recession. In Section III, I explore the extent to which the downturn or subsequent slow recovery has proceeded evenly across skill groups, demography, or geography. Section IV presents estimates of the cyclical sensitivity of subgroup unemployment rates, and compares different groups' experiences to what one would have predicted given the heterogeneity in this sensitivity. Section V presents the analysis of wage trends. Section VI presents analyses of long-term unemployment and labor force participation. Section VII concludes.

---

[3] Of course, the failure of wages to fall quickly in response to labor demand shortfalls is a longstanding and still unresolved puzzle; see, e.g., Bewley (1999).

II. The state of the aggregate labor market

Figure 1 shows the time paths of aggregate employment, the unemployment rate, and the employment-population ratio from 2004 forward. The figure makes clear that the sharpest downturn was in late 2008 and early 2009, when over a six-month period the economy lost 4.5 million jobs. Job losses continued until February 2010, but employment has grown consistently since then – the only exception being a blip due to the end of temporary jobs associated with the 2010 census – at an average rate of about 191,000 new jobs per month. This is only a bit faster than is needed to keep up with population growth, however, and as a result the employment-population ratio, which fell from 62.9% in January 2008 to 58.2% in December 2009, remains well below its pre-recession level. There was no sign of improvement in this series for the first four years after the recession ended, and even in recent data the upward trend has been anemic. Thus, while the unemployment rate has fallen from its peak of 10.0% in October 2009 to 5.3% in July 2015, this decline is almost entirely attributable to falling labor force participation among the non-employed.

Figure 2 shows hires, involuntary displacements, layoffs, and quits, as measured in the Job Openings and Labor Turnover Survey (JOLTS). Layoffs, which have been extremely stable historically, rose by nearly 50% during the 2008-9 crisis. At the same time, the more volatile hires and quits series fell by about one-quarter and 40%, respectively.[4] Layoffs peaked in early 2009 and returned to their pre-recession levels by early 2010. Hires and quits bottomed out somewhat later, in the middle months of 2009, and have recovered only very slowly since then. Each remains well below its pre-recession peak. By contrast, job openings, which fell by fully 50% during the crisis, recovered more quickly and have exceeded their pre-recession peak since June 2014. In April 2015 they passed their historical record, set in the very first JOLTS survey in December 2000. They have continued to grow since.

Figure 3 shows the Beveridge Curve, relating job openings to the unemployment rate. One expects these measures to be inversely related: In tight labor markets with low unemployment, jobs are filled slowly and the job openings rate is therefore high, while when unemployment is high vacancies are filled quickly and there are few jobs open at

[4] These are counted from December 2007 to the respective series trough in 2009.

any given time. In search models of the labor market, shifts in the relationship between the two series can indicate changes in the efficiency of the labor market matching process (e.g., Blanchard and Diamond, 1989).

As Figure 3 illustrates, in 2008 and the first half of 2009 job openings fell steadily as the unemployment rate rose, tracing out a curve consistent with – though well beyond the support of – that seen in the prior business cycle. In the second half of 2009 and the first quarter of 2010, however, job openings rose substantially, with no change in unemployment; then, when unemployment began falling in mid-2010, the job openings rate continued to rise at a rate consistent with the pre-2009 slope. Thus, it appears that the Beveridge Curve has shifted upward, with a given unemployment rate supporting a job openings rate about 0.75 points higher than it would have before the crisis.

A number of commentators have interpreted this apparent shift in the Beveridge Curve as diagnostic of increases in structural unemployment. In this view, an increase in labor demand can be inferred from rising job openings, and the failure of the unemployment rate to fall faster than it has indicates that the currently unemployed are unable or unwilling to fill the newly created positions. This inference is supported by Krueger et al.'s (2014) recent analysis of the duration of unemployment, which argues that a Beveridge Curve that uses the short-term unemployment rate – the share of the labor force that has been unemployed for six months or less – does not appear to have shifted in the same way (see also Ghayad and Dickens 2012).

But while the shift in the Beveridge Curve is certainly consistent with a structural change, it is important to be cautious. There is at least some reason to think that part or all of the shift reflects changes in the meaning of a job opening rather than increases in the difficulty of finding qualified workers.

Job openings are well defined if hiring is a binary decision on the firm's part, as in many search models: Once a decision is made to hire another worker, a job opening is posted and the first applicant who arrives (perhaps subject to some well-defined, fixed minimum qualifications) is hired. This, of course, is an extreme oversimplification. In

reality, both the wage and the required qualifications are choice variables that can influence the rate at which posted openings are filled.[5]

Consider a firm with labor demand curve $L^D = f(w)$, with $f'<0$. So long as wages are set exogenously, job openings are well defined as the difference between $f(w^*)$ – where w^* is the externally determined wage – and the firm's current employment. But if wages are not fixed there is no unique number of openings.[6] A firm might decide to offer wage $w_{low} < w^*$ for an additional $f(w_{low})-f(w^*)$ positions, knowing that these jobs are likely to remain open for longer than would a position offering w^*. Similarly, the firm might hold out for better-qualified workers, extending its search, or might be less choosy in order to hire more quickly (Diamond 2010). Either decision affects the number of measured job openings and the job-filling rate, but neither reflects changes in labor market matching efficiency.

These definitional issues may have become more important since the Great Recession. In previous business cycles, employers seem to have been unwilling to take advantage of labor market weakness by offering lower wages to new hires than they have in the past, or by substantially increasing their required qualifications. The reasons for this are not well understood, but appear to include concerns about the morale of the newly-hired and the incumbent workers and worries that workers who accept low wages when business conditions are weak will be likely to leave the firm once conditions improve (Bewley, 1999). These concerns appear to be less salient today. Anecdotally, two-tier wage structures that distinguish between incumbent and newly hired workers have become increasingly common (Vlasic 2011). Moreover, at least some employers seem to have taken advantage of their strong bargaining positions to be much choosier

[5] Even when the offered wage is not posted with the job advertisement the employer must decide on a bargaining stance once an otherwise suitable candidate is identified. Similarly, the employer sets both minimum qualifications to list with the position and its choosiness among workers meeting those qualifications. Finally, a firm planning to hire may do so without ever posting an official opening (Diamond 2010).

[6] This is of course the exact analogue to the somewhat more common claim that unemployment is always voluntary: Unemployment simply means that one's reservation wage has been set above the market price. In search models, there can be frictional unemployment and frictional job openings. But even in these models one might observe a range of reservation wages and wage offers, with frictional unemployment rising in the former and frictional vacancies declining in the latter.

among job applicants, raising qualifications and drawing out the hiring process with multiple rounds of interviews (Davis, Faberman and Haltiwanger, 2010). All of this could be raising the measured rate of job openings relative to the strength of underlying labor demand. This would be consistent with the divergent behavior of job openings and quits since mid 2009 – where the former appears to indicate the tightest market since the JOLTS survey began in 2000, the latter has recovered much more slowly and remains well below its 2006 peak. It is thus important to look beyond the Beveridge Curve for evidence that could confirm or disprove the indication that there have been structural changes in the labor market.

III. Heterogeneity across industry, geography, and demography

Poor aggregate outcomes may disguise heterogeneity. If demand is recovering in some labor markets but not in others, employers in the former markets will experience the market as tight even as workers in the latter experience it as slack, and worker shortages can coexist with high unemployment. For example, skill-biased technical change might lead to stronger demand in the market for high-skill workers than in the market for low-skilled workers, creating simultaneous tightness in the high-skill market and underutilization in the low-skill market. This sort of mismatch could rationalize low quit rates with high job openings, as the former reflect the tightness of the markets where existing workers are located and the latter the tightness of markets where employers would like to expand.

To test this mismatch hypothesis, we need a theory of the boundaries of labor markets. Unfortunately, the available data do not permit precise characterizations. I consider three dimensions that represent important labor market divisions: Industry, geography, and demography.

Industry

Everyone knows that the financial services and real estate industries led us into the recession. However, these sectors did not see disproportionate job losses: the employment contractions in these industries in 2007-09 – 5.6% in finance and 9.1% in

real estate – were comparable to the economy-wide average.[7] In both absolute numbers and percentage terms, job losses were much larger in construction and durable goods manufacturing, which contracted by 25% and 20%, respectively.

Since the trough, employment has grown in every major private sector industrial category (though it has declined in both federal and state and local government). The mining and logging sector is a clear outlier. It surpassed its pre-recession peak by mid-2011 and over the longer run saw net growth of 60% between early 2003 and the end of 2014 (before shrinking by over 8% in 2015, in the wake of low energy prices) . This is clearly a structural change, and there is compelling evidence of important mismatch-based impediments to the growth of this sector (for example, in North Dakota during the extraction boom there). But mining and logging accounts for only about 0.5% of national employment. Most other sectors have not yet made up their earlier losses. There are only a few sectors where employment growth between 2007 and 2014 has kept up with the growth of the working-age population: mining and logging, professional and business services, education, health and social assistance, arts and recreation, and food and lodging. Notably, one sector where anecdotal stories about labor supply shortages have been common – information – has since 2009 recovered only one third of the jobs lost in the downturn.

Thus, insofar as there is mismatch across industrial groups, the tight markets would need to be comprised mostly of low- and middle-skill sectors like lodging and food services and arts and recreation. (There has also been job growth in professional and business services, education, and health, each of which has substantial numbers of higher skill jobs.) This is unlikely to account for the poor outcomes of low- and middle-skill workers.

Geography

A second important source of heterogeneity is geographic. The recession hit some areas – most famously, Sun Belt cities like Las Vegas where the housing boom was most pronounced – harder than others, and the recovery has also been uneven. Mobility

[7] This of course does not rule out the idea that a shock that began in the financial sector was an important source of the general collapse in demand.

rates have fallen in recent decades (Kaplan and Schulhofer-Wohl, 2015), so geographically uneven labor demand growth might lead to mismatch. But while the evidence suggests that demand growth has been somewhat uneven, it also indicates that there are few places where demand has been robust enough to make up for the recessionary collapse. Figure 4 shows unemployment rates by state in December 2007 and December 2014. Across the 51 states plus Washington DC, only six had an unemployment rate in 2014 that was below its 2007 level. Moreover, these six include Michigan and Ohio, where the unemployment rates remain reasonably high. Only Kansas, Minnesota, North Dakota, and Vermont (which total only 3.4% of national employment) also had 2014 unemployment rates lower than 5%. At a finer geographic level, unemployment remains higher than it was before the recession in 302 of 394 metropolitan statistical areas. There is not enough scope for geographic mismatch to have made a meaningful contribution to current woes.

Demography

A third dimension of heterogeneity is demographic, based on gender and education. Construction and manufacturing employment is heavily male and largely non-college-educated, so one might expect that low-skill men would have suffered disproportionately in the recession. Figure 5 shows the unemployment rate by gender and education in 2007, 2009, and the 12-month period from July 2014 through June 2015. Consistent with the industrial composition of the cyclical collapse, we see that unemployment rates of less educated men rose more from 2007 to 2009 than did those of more educated men or of women of any education level. However, it is notable that low-skill workers had much higher unemployment rates than high-skill workers even in 2007, and that unemployment rates rose by similar proportions across all eight gender-education groups.

Importantly, Figure 5 also shows that declines in unemployment across groups since 2009 have been roughly proportional to the recessionary increases. In all eight groups, unemployment rates in 2014-15 are lower than the 2009 levels but notably higher than in 2007. The relative recovery has been somewhat stronger for men than for women at each education level, and stronger for non-college workers than for those who are more

educated, but the general pattern is proportionality. This appears more consistent with broad-based swings down and then part of the way back up, rather than with tilts favoring some workers; insofar as there have been tilts, they evidently are *toward* low-skill men.

IV. Accounting for heterogeneity in cyclical sensitivity

As noted above, the construction and manufacturing industries suffered dramatic employment declines in the downturn. This is not surprising; these sectors have always been more cyclically sensitive than the economy as a whole. Because these sectors disproportionately employ low-skilled men, such workers suffer more than others in every downturn and see better improvements in every recovery. Similarly, youth unemployment has always been highly sensitive to economic conditions (Clark and Summers 1982). This differential sensitivity can account for most of the heterogeneity in outcomes across groups seen in Figure 5.

Figure 6 illustrates this. For each gender-education group shown in Figure 5, I plot the actual change in the unemployment rate between 2007 and 2014, as well as the change in unemployment that would have been predicted given the past cyclical sensitivity of that group's unemployment and the magnitude of the cycle. To form this prediction, for each group g I estimate a time series regression of the form:

$$u_{gt} = \alpha_g + u_{(-g)t}\beta_g + e_{gt}, \quad (1)$$

where u_{gt} is the unemployment rate for group g in month t and $u_{(-g)t}$ is the average unemployment rate in that month across all groups other than g.[8] The β_g coefficient is analogous to an equity price's beta, and measures the extent to which group g moves with the market: A value of β_g greater than 1 indicates that group g is more cyclically sensitive than others; less than 1 indicates relative insensitivity. I estimate α_g and β_g using monthly observations from 1978 through 2007, then use these coefficients and the observed path

[8] I compute $u_{(-g)t}$ using fixed weights for each group h≠g over time, proportional to the group's average labor force share over the 1978 to 2015 period.

of $u_{(-g)t}$ to forecast u_{gt} through 2014. The lighter bars in Figure 6 illustrate the change in average forecast values from 2007 to 2014.[9]

The figure shows that most of the across-group differences in unemployment growth between 2007 and 2014 are attributable to differences in cyclical sensitivity rather than to unique features of this business cycle. In particular, more educated workers are always less sensitive than less educated workers, and that has been true (particularly for women) in this period as well. However, the cyclical predictions are not perfect. The most important deviation from the prediction is men without high school diplomas: Based on past performance, we would expect their unemployment rate to have been nearly two percentage points higher in 2014 than was actually observed. By contrast, men with college degrees have an unemployment rate a full percentage point higher than before the recession, roughly double the increase that would have been expected based on past patterns. And women of all skill groups have also seen larger increases than past patterns would suggest. These comparisons suggest that any shifts in demand have been toward low-skill men, away from men with college degrees and from women.

Figure 7 provides another look at this. Here, I show the time series of the unemployment rate for men with high school diplomas but no college, along with predictions from equation (1), above. (Both series are smoothed; the dashed segment of the latter series indicates the out-of-sample portion.) The predicted and actual rates track each other extremely closely, through past business cycles and the current one. The actual experience for this group was worse than the prediction in the depths of the Great Recession (and even for about a year preceding the collapse), but the discrepancy was small relative to the predictable component and closed very quickly. While other work has documented long-run declines in demand for this group of low-skilled workers, this has evidently not shown up in the unemployment rate, which has been dominated by cyclical factors.

Figure 8 repeats the exercise from Figure 6, this time looking at age groups. Unemployment rates rose much more for young people between 2007 and 2014 than for

[9] Results are similar if I instead predict the change as the difference between the predicted 2014 rate and the *actual* 2007 rate. In either case, I focus on annual averages, though the prediction is conducted at the monthly level.

those aged 35 and over. But this is always the case in weak labor markets; the unemployment rate for those 25-34 actually rose slightly *less* than one would have expected, and in any case the difference is small. The only meaningful deviation from predictions is for those aged 65 and over: This group's unemployment rate is ordinarily not very cyclically sensitive, but between 2007 and 2014 it rose as much as those for prime-age workers.

One sees similar patterns when looking across industries (not shown). Insofar as there have been structural shifts, they have apparently been *toward* the goods-producing industries and *away* from the high-skill services. Unemployment rates remain higher than in 2007 for essentially all ages, education levels, genders, and industries. Sectors that have been more cyclically sensitive in the past saw larger increases, but there is remarkably little heterogeneity beyond this. This pattern appears consistent with a shortfall in aggregate labor demand, and less so with a gradual adjustment to a technological or demand-driven shock that changed the composition of labor demand.

V. Evidence from wages

The evidence presented thus far suggests that differential unemployment outcomes over the last several years across education, gender, or age have been largely consistent with what past cyclical patterns would have implied given the overall state of the labor market. But the unemployment rate is at long last now reasonably low, while other indicators – the employment-to-population ratio in particular – remain quite depressed. We thus face the important question of whether the economy has fully recovered, with recent outcomes "the new normal," or whether there remains substantial slack, either in the aggregate or in particular labor markets. There is some disagreement among indicators here, with the job openings rate indicating that (at least some) labor markets are tight and with other indicators – the employment-population ratio or quits – indicating that (possibly other) markets are slack.

To adjudicate this, we need to look beyond quantities data to prices. If employers are facing shortages of suitable, interested workers – either throughout the economy or in particular labor markets – they should be responding by bidding up the wages of those workers who can be found. By contrast, if the large number of people who have moved

out of the labor force in the last few years are still relevant to the labor market, then there are in fact a great many workers per job opening, and there is no reason for employers to increase wages. Thus, in this section I examine wage trends for evidence of tightness. I examine the aggregate labor market first, then turn to distinctions across sub-markets.

V.A. *Aggregate wages*

The solid line in Figure 9 graphs the 12-month change in real mean log hourly wages from 1999 through June 2015. These wages are calculated from the Current Population Survey Outgoing Rotation Groups, with imputed wages excluded; details are in the appendix.[10] The figure shows that the last sustained period of real mean wage growth ended in 2003. Average wages were largely stable between 2003 and late 2014, except for a period in late 2008 and 2009 when they rose at an annual rate of about 3% and a shorter period in 2011 when they fell at a -2% annual rate.[11] Outside of these periods, there was little movement. The most recent data, however, show some signs of strength, with growth at a 2-3% annual rate in the second quarter of 2015.

One concern about aggregate wage trends is composition changes: If the least skilled workers are the most likely to have lost their jobs in the Great Recession, changes in average wages will overstate what was experienced by individual workers. This could help to explain why wages appear to have risen quickly during 2008 and 2009, when employers were shedding workers quickly and there is little other evidence of labor market tightness, and could also confound trends in other periods.

To address this concern, I use the longitudinal structure of the CPS to match observations on the same individual from month m and month m+12, excluding

[10] The CPS is not the only source for information on wage trends, and observers more often rely on employer surveys. But the CPS is the best option for adjusting wages for worker observables or for focusing on newly hired workers. Other series show similar patterns for the "all worker" series depicted by the solid line in Figure 9.

[11] The price level was falling during much of the 2008-9 period of real wage growth; nominal wage growth actually slowed in the second half of 2008 and early 2009, from around 4% per year to under 2%. Similarly, the slowdown in real growth in late 2009 and early 2010 reflects stable nominal growth (at an annual rate of about 1.5%) and the return of mild inflation.

observations that cannot be matched or where the wage is unavailable in either.[12] The dashed line in Figure 9, labeled "composition-adjusted," shows the mean year-on-year change in mean wages for those who were employed in both periods. Note that this overstates the growth for workers with constant characteristics as in this sample the year-on-year change incorporates the effect of aging by one year as well as any calendar time effects. In part for this reason, this series is consistently about two points higher than the all-workers series, and shows average growth of about 3% per year between 1999 and 2009. The anomaly in 2008-2009 is much reduced here and plausibly consistent with sampling error. Average growth fell to near zero in 2009, however, not to reappear until late 2011; it remained around 2% for the next several years, only rising above that in 2015.

Workers rarely accept – or perhaps employers rarely demand – reductions in their nominal wages within existing jobs. This wage rigidity may have masked trends in the wages offered to new hires. To zero in on the latter, I take advantage of the fact that the CPS makes it possible to identify workers in the ORG sample who have started new jobs within the previous three months. The dotted line in Figure 9 shows the trend in mean wages for such workers. This series closely resembles the all workers series, with a similar pattern of rising real wages in 2008-2009 and falling real wages starting in early 2010. It began to outpace the all worker series in 2014, however, and shows real growth around 2-4% in 2015.

Thus, aggregate data suggest that the labor market remained weak, with too little demand to generate even modest wage growth, at least until late 2014. Wages appear to have begun growing modestly since then. This is perhaps consistent with the long-delayed arrival of labor market tightness that is forcing employers to offer higher wages in order to attract and retain workers. Even in the recent data, however, the growth rate is

[12] Roughly 40% of initial observations lack one-year-ahead wages, about two-thirds of the time because the individual cannot be matched to a year-ahead observation (due to having moved from the original home, to survey nonresponse, or to errors in the CPS identifiers) and the remainder because the person is surveyed in the follow-up but is no longer employed or lacks a valid wage. Attrition among the continuously employed may be correlated with wage growth. The reweighting exercise described in the text partially addresses this possibility.

relatively modest, and the recent trend would need to continue for some time to make up for the roughly zero average growth over the previous decade.

V.B. Individual labor markets

As discussed above, aggregate data can mask differences across individual labor markets. It is possible that particular labor markets are tighter and showing more dramatic wage growth than is apparent in Figure 9, or that some markets tightened earlier than 2015 even while others remained slack. To examine this, I examine changes in employment, hours, and wages of newly-hired workers by education, gender, age, and industry between 2007-8 and 2013-15. The latter window includes both the period of the recent runup in wages seen in Figure 9 and a preceding period of stagnant wages; unfortunately, sample sizes are insufficient for this sort of disaggregated analysis when limiting attention to the most recent data.

Columns 1-2 of Table 1 show estimated changes (and standard errors for these changes) in mean wages of newly hired workers in different sub-markets. To limit confounding due to changes in the composition of workers – as would occur, e.g., if manufacturing hires shifted from unskilled laborers to skilled machine operators – I use a regression to adjust for changes in workers' observed characteristics. Specifically, I regress log real hourly wages for new hires in 2004-2006 on a quadratic in age; indicators for education-by-gender, state, and industry-by-education; and separate linear age terms for each gender-education group. I then use the coefficients to form predicted log wages for new hires in 2007 and later, and compare these to the observed wages. Table 1 shows the change in the mean log wage residual in each cell and the standard error for this change.

Across education-by-gender and age cells, only over-65 workers (and, to a much lesser extent, those aged 55-64) saw nontrivial real wage increases over this period, amounting to a bit over 1% per year. Across industries, nontrivial wage changes are seen in mining and logging, finance/insurance, and education, though only the finance change is statistically significant. There is thus no sign of mean wages being bid up at the level of education groups, age, or broad industrial sectors.

Columns 3-6 of Table 1 present other indicators of labor market tightness: The change in total employment and the change in the average number of hours worked.

(Employment changes are measured in percentage points on the employment-population ratio for education-by-gender and age groups, and in percent for industry groups. Neither is adjusted for observables as in Columns 1-2.) While a few industries saw substantial employment growth over this period, this is not in general matched by substantial wage growth as well. The mining and logging sector is, as noted earlier, the exception that proves the rule; education also saw increases in both employment and wages, consistent with a demand increase unmatched by increases in labor supply. Across demographic groups, only the 65+ age category saw an increase in the employment rate. The story is similar for hours. While a few sectors – construction, finance – show small increases in average hours, these are not in general sectors that seem tight on the other margins, and taken as a whole there is little sign of tightness.

This analysis thus suggests that essentially none of the identified markets have tightened to the point where demand began to outstrip supply, at least by the end of 2014. It nevertheless remains possible that these market definitions are too coarse and that employers in more tightly defined submarkets have had trouble finding workers with suitable skills. Falling demand in other submarkets might make it impossible to detect rising wages for workers in short supply via examinations of highly aggregated averages.

One way to assess this is to examine points in the wage distribution other than the mean. If some markets are tight, it might be possible to see this manifested as a rightward shift in the relevant portion of the wage distribution. To assess this, I compared the 2007-8 and 2013:July-2015:June distributions of starting wages, adjusted for observable characteristics as in columns 1-2 of Table 1. The solid line in the upper left panel of Figure 10 shows the change in wages at different percentiles in the new-hires adjusted wage distribution between the two periods.[13] Thus, for example, it indicates that the 75th percentile of this distribution was roughly unchanged; that the upper quartile of the distribution shifted right by as much as three percentage points, and that everywhere below the 75th percentile the distribution shifted left, by a similar magnitude. This figure thus shows that any wage increases were concentrated in the upper quartile of workers, and that even here real wages increased by less than half a percentage point per year.

[13] To eliminate spurious changes due to changes in real value of the CPS topcodes, I censor weekly earnings at the real value of the 2015 topcode.

For comparison, Figure 10, Panel A also shows series for the change in the real wage distribution between 2000/1 and 2005/6 and between 1994/5 and 1999/2000. The former closely resembles the pattern in the recent data, consistent with the widespread view that the labor market was never very tight during the post-2001 expansion. The latter series is quite different, with real growth of about 8% throughout the distribution. Evidently, in a truly tight labor market, such as was seen during the late 1990s, real wages can grow substantially throughout the distribution.

The remaining panels of Figure 10 explore potential heterogeneity in the 2007/8 to 2013/5 distributional change. I do not track individual workers or labor markets over time, so the analysis here can only examine changes in the (conditional) distribution, not changes for workers at particular points in the distribution. This could lead me to miss tightness in particular markets. For example, increases in wages in a low-wage sub-market combined with decreases in a second, slightly higher-wage submarket, could offset each other with little effect on the overall distribution. (Note, however, that this is more plausible for *unconditional* wages; because I examine wages after adjusting for observables, heterogeneous shifts across different submarkets are unlikely to balance out in this way.)

The Current Population Survey sample is not large enough to permit detailed analysis of wage distribution changes within individual industries or geographic groups. As an alternative, I divide the sample into sub-groups based on non-wage indicators that might proxy for labor market tightness. In the upper right panel of Figure 10, I divide industrial sectors in half based on the increase in job openings between 2009 and 2014. The solid line shows the sectors with below average increases in openings and the dashed line those with above average increases. Wage declines at the bottom of the distribution are smaller in the latter industries, consistent with their labor markets being relatively tight. But even in these industries there is no sign of meaningful wage growth through the bulk of the distribution. Moreover, what growth there is at the top of the conditional wage distribution is coming from industries that have not seen large increases in job openings.

The lower left panel focuses on geography. Here, I divide metropolitan areas into two groups by their 2014 unemployment rates; the solid line shows the change in the

distribution of starting wages for areas with unemployment rates above 6.1%, the median in 2014, while the dashed line shows areas with rates at or below that point. There is no sign that wage growth was stronger in MSAs with lower unemployment rates.

The final panel of Figure 14 examines skill levels. Using the same flexible wage regression used to predict wages for Table 1, I predict a wage level for each newly hired worker in the CPS. I divide workers in half based on this predicted wage, plotting the change in the distribution of starting wages for the less skilled group with a solid line and the more skilled group with a dashed line. Insofar as the market for more-skilled workers has tightened more than that for less-skilled workers, we should see the dashed line systematically above the solid line. This is not at all apparent in the figure. Through the middle of the distribution the dashed line is somewhat above the solid line, though not by much and through most of this range still negative. In the tails, lower-skill workers seem to have done a bit better than have higher-skill workers.[14]

Across all the panels of Figure 10, there is no sign of any quantitatively important segment of the labor market where demand has meaningfully outstripped supply in recent years. Wage increases are modest throughout the distribution, and are no stronger in sub-markets that might plausibly be tight (e.g., industries with large increases in job openings, or MSAs with low unemployment rates) than in those that clearly are not. Wage pressures appear to have been modest, at least through the end of 2014, throughout the economy.

VI. Long-Term Unemployment and Labor Force Participation

The evidence of slack throughout the labor market appears to be somewhat at odds with the relatively quick decline of the unemployment rate, as seen in Figure 1. One hypothesis is that the unemployment rate is not adequately measuring the amount of slack in the labor market. This section investigates two aspects of this: Long-term unemployment, and labor market participation.

[14] At the lower end of the distribution, the increase for below-average predicted wage workers appears to be driven by increases in the real value of the minimum wage.

VI.A. Long-term unemployment

Much commentary about the labor market during and after the Great Recession has focused on the long-term unemployment share – the fraction of the unemployed who have been out of work for 26 weeks or more. This share rose from around 17-18 percent on the eve of the Great Recession to 45 percent in mid-2010, far higher than had ever been seen before. It has fallen slowly since then, but has only recently fallen below the pre-2007 series record, set in 1983.

Either reduced search effort among the unemployed or labor market mismatch would reduce the exit rate from unemployment and thus lengthen survival in unemployment. Of course, reductions in labor demand would have a similar effect. Nevertheless, many have interpreted the dramatic rise in employment durations in this cycle as indicative of reductions in job search effort (Barro, 2010). An alternative interpretation is that search is less effective for the long-term unemployed. Kroft, Lange, and Notowidigdo (2013) find that employer seem to discriminate against the long-term unemployed in hiring, and Krueger et al. (2014; see also Hall, 2014) conclude that these workers are largely disconnected from the labor market, exerting little supply pressure. If this is true, the labor market (for other workers) is much tighter than it appears based on the unemployment rate alone.

Table 2 shows unemployment rates and long-term unemployment shares across various demographic categories in 2006 and 2014. The long-term unemployment share does not vary nearly as much across gender and education as does the unemployment rate. More educated workers, whose unemployment rates are very low, tend to have higher long-term unemployment shares than do those with less education. Across ages, long-term unemployment is somewhat more prevalent among older workers, but the differences are relatively small except in the 16-24 year old group (many of whom have not been in the labor force for long enough to reach long-term unemployment).

Figure 11 shows the probability that a worker who is unemployed for one month is reemployed in the following month.[15] I show estimates for three categories of

[15] I compute reemployment rates following the procedure used by Rothstein (2012b) and Farber and Valletta (2013), counting individuals who transit from unemployment to employment, then immediately back to unemployment in the following month as having

unemployed workers – those who have been unemployed 13 weeks or less in the initial month; those who have been out of work for between 14 and 25 weeks; and the long-term unemployed whose durations exceed 26 weeks. The figure shows that reemployment rates decline with unemployment duration, primarily when comparing short-term to medium-term unemployed. This might reflect hysteresis, employer discrimination among the unemployed (Kroft, Lange, and Notowidigdo, 2013), declines in search effort with unemployment duration (Krueger and Mueller, 2011), or heterogeneity in reemployment probabilities. Interestingly, however, the gap in exit rates between the short- and long-term unemployed did not widen dramatically even as the level of the exit rate fell during the recession. Reemployment rates have recovered slowly since 2009 for all durations, though perhaps a bit more quickly for the short-term unemployed. Despite this, the between-group gap at the end of the period is similar to that seen before 2008.

Although there is no guarantee that this pattern of generally parallel movements will persist in the future, it appears likely based on recent history. If so, then we might expect that a robust labor market – if one ever arrives – will pull the long-term unemployed back from the margins of the market and into higher levels of attachment. Note, moreover, that there is no inconsistency between this optimistic view and the evidence (see, e.g., Kroft et al. 2013 and Krueger et al. 2014) that the long-term unemployed do not compete effectively with the short-term unemployed for work; under this story, those who have been out of work for many months are at the margins of the labor market, but if employer demand is robust enough to exhaust other sources of labor, firms will figure out ways to employ even them (Bernstein and Baker, 2003).

VI.B. Labor market participation and school enrollment

The evidence reviewed thus far suggests that there has been much more slack in the labor market than is indicated by the unemployment rate. A plausible explanation for this is that many workers may have given up their active searches for work, while nevertheless remaining available and ready to return when the market recovers. These

---

remained unemployed throughout. As discussed in the above-cited papers, many of these repeated transitions appear to derive from misclassification of labor force status in the middle month (see also Poterba and Summers, 1986).

"missing workers" (Shierholz, 2014) would represent additional slack in the labor market that is not captured by the unemployment rate.

As discussed above, one way to assess the potential importance of this group is to ignore the somewhat subjective unemployment concept, simply examining the fraction of the population that is employed. This fell roughly four percentage points between 2005/6 and 2013/4. The missing employment might represent remaining slack in the labor market, or might represent workers who have been left behind by changes in the composition of demand.[16]

The solid line in Figure 12 shows how the employment rate has changed for workers of different ages. To avoid spurious changes coming from vagaries of the academic calendar, I use only data from the academic year – the January-April and September-December surveys – in this figure. Note also that the figure does not compare the same worker over time – the employment rate of, say, 25-year-olds in 2005 is compared to that of 25-year-olds in 2013, born eight years later, not to that of the original workers in 2013, when they would have been 33 years old.

The figure indicates that the employment rate fell by much more than four percentage points at the youngest ages – as much as 10-12 percentage points for those under 20. The decline is fairly stable at 3-4 percentage points for those in their 30s, 40s, and early 50s. But employment rates *rose* substantially for workers over age 60. This last is almost certainly a labor supply effect, perhaps reflecting declining retirement balances that force people to work longer or generational differences in labor force attachment as baby boomers replace earlier cohorts among those in their 60s. (One plausible explanation turns out not to hold up, however: The increase does not reflect rising female labor force attachment in later-born cohorts, as it is similar for men as for women.)

But the dramatic decline in the participation rate of the youngest cohorts is a puzzle. It is unlikely that the non-participants in these age ranges are permanently out of the labor force. A more plausible explanation is that they have been unable to find a real toe-hold in a weak labor market that offers few entry points – many individuals in this

[16] Not all of the four percentage point decline corresponds to missing workers. The Council of Economic Advisers (2014) argues that roughly one-third to one-half of it derives from the aging of the population, with a fixed participation rate at each age.

age range in 2013-4 entered the labor market during or after the Great Recession and have likely never been able to find stable employment. Past evidence indicates that there is real reason for concern about their future earnings prospects (Kahn 2010; Oreopoulos et al. 2012). But it seems likely that they will be employable, at rates comparable to those seen for young people before the recession though likely at lower wages, if demand returns.

There is nevertheless the question of what these non-workers have been doing with themselves. To examine this, I consider an expanded concept of "non-idleness" that includes employment or enrollment in school but not unemployment or non-participation among non-students. The change in the non-idleness rate is shown as a dashed line in Figure 12. There is no decline here for young workers. Evidently, *all* of the young people who would have been employed in 2005/6 but were not in 2013/4 have instead decided to remain in school. Pooling all ages, the non-idleness rate fell by only 2.5 percentage points between 2005/6 and 2013/4, much less than the 3.9 p.p. decline in the employment rate.

It is not clear whether this is a labor supply effect – as would arise if it is taking longer for students to finish school or if full-time students are less interested in working while in school than in the past – or a demand effect, reflecting young people remaining in school as a way of waiting out the weak labor market. But in terms of the question of whether the lower employment rate reflects a permanent, structural decline in our economy's capacity, the evidence in Figure 12 is quite encouraging. It suggests, first, that the supply side of the market is responding to labor market weakness in a way likely to ameliorate any mismatch between changing skill demands and the fixed stock of skills supplied by existing workers and, second, that scarring effects from prolonged weakness are likely to be smaller than they would be if young people were truly remaining idle for years on end.

Additional analyses, not reported here, show that the educational attainment of recent cohorts is rising fairly dramatically, as the additional school enrollment is translating to additional college degrees. This is likely to translate into improved earnings capacity and prosperity in the years to come, though it remains an unresolved question whether the additional human capital earned in school will be enough to offset the

reduced experience and increased difficulty in getting a toehold on the job ladder experienced by cohorts coming of age in recent years.

VII. Discussion

The performance of the U.S. labor market since 2006 can fairly be described as catastrophic: The unemployment rate was above 8 percent for over three straight years; while it has since come down, the employment-population ratio, which fell by nearly 4 percentage points since 2007, remains extremely depressed. Many subgroups – particularly the young and less educated, along with racial minority groups – are still facing extremely high unemployment rates.

Many models that economists have used to understand business cycles have difficulty accounting for demand shortfalls that last for many years. In such models, sustained high levels of unemployment can arise only if there are structural impediments to labor market clearing – the unemployed are not looking very hard for work, have raised their reservation wages due to increased implicit taxes on work, or are in some sense unsuitable for the jobs that are available, perhaps because they lack the appropriate skills or are unwilling to move to where the jobs are.

Drawing in part on these models, many observers have concluded that structural impediments to recovery must be an important component of the 2010-2014 situation. The review of the evidence here offers no support for this diagnosis, however. The poor labor market outcomes for low-skilled workers are entirely consistent with cyclical explanations, as these workers have always been more sensitive to the business cycle. The most plausible sources of structural problems – labor supply disincentives due to conditional transfers like unemployment insurance or geographic immobility due to housing market frictions – do not appear to be quantitatively important.[17] And the Beveridge Curve provides at best weakly suggestive evidence regarding the state of the matching function.

[17] Unemployment insurance extensions can explain only about 0.3 percentage points of the 2011 unemployment rate (Rothstein 2012b). With regard to geographic mobility, declines are concentrated among renters who should not have been directly affected by the decline in home values (Farber 2012), and any "house lock" effect is quantitatively small (Schmitt and Warner 2011). The most likely explanations for the mobility decline are not consistent with mismatch stories (Kaplan and Schulhofer-Wohl 2015).

Indirect evidence also fails to support the claim. Structural explanations for inadequate recovery, whether due to supply reductions or to mismatch, imply that the labor market has actually been much tighter than it has appeared, at least as viewed from the perspective of potential employers. But there is no sign in the data that employers with jobs to fill have had trouble filling them, except perhaps in a few isolated and small submarkets such as resource extraction.

Finally, the unprecedented rise in long-term unemployment and decline in labor force participation, which some have pointed to in support of the structural unemployment hypothesis, is at a minimum more complex than this. The rise in long-term unemployment appears largely attributable to across-the-board reductions in reemployment rates, with no indication of particular declines for the long-term unemployed. And much of the decline in participation derives from young people who are remaining in school rather than from older workers taking early retirement, for whom hysteresis concerns may be most pronounced.

We can thus conclude that labor demand shortfalls continued to be an important feature of the labor market and the primary determinant of labor market performance, at least through the end of 2014. While there is some sign of wage growth in 2015, it remains modest and has only begun to make up the shortfalls from previous years. There is no support for the view that the anemic recovery to date has been driven by supply shortfalls due to changes in labor supply behavior or mismatch between employer needs and the available pool of labor.

With that said, several caveats are in order. First, while my results point to the importance of aggregate labor demand in understanding recent trends, this is not the place to address the question of how policy might stimulate additional demand. The results here speak to the importance of accomplishing that goal, but not to the best way to do so.

Second, I have not addressed longer-run structural changes, such as deindustrialization or skill-biased technical change, which may have proceeded smoothly previous to, during, and after the recession. Rather, I have focused exclusively on the very short run, looking for signs of structural explanations for changes between 2007 and the present. My analysis speaks to the question of whether increases in aggregate demand might return our labor market to something resembling its 2007 state, but not to

whether further increases could reverse longer-run trends toward reduced male employment-population ratios and higher inequality. Some policy responses – education and training programs and increased income support for low earners in particular – may make sense as a response to long-term trends, even if they cannot be expected to contribute meaningfully in the short run so long as the market is demand constrained.

Third, it is possible that structural changes did occur on the supply side of the market but that these were masked for many years by low aggregate demand. This could help to explain the real wage increases that we have seen at long last in 2015. If these persist, there will likely be room for policies aimed at improving job matching – e.g., search and mobility assistance – and thereby at expanding effective supply.

Finally, and most important: An extremely long downturn is likely to cast a long shadow over our future prosperity, even if this shadow falls more on wages than on employment rates. Productivity has been low in recent years, so even if wages had kept up with productivity the growth would have been unimpressive. Workers displaced in the early 1980s recession faced large declines in future earnings, amounting to 20% losses even 15 to 20 years after their initial displacement (von Wachter, Song, and Manchester 2011), and also saw substantial declines in their life expectancy (Sullivan and von Wachter 2009). Other research indicates that young people who enter the labor market during recessions see long-run negative earnings effects (Oreopoulos, von Wachter, and Heisz 2012; Kahn 2010) and that parental job loss hurts children's schooling and labor market outcomes (Oreopoulos, Page and Stevens 2008, Stevens and Schaller 2011, Ananat, Gassman-Pines, and Gibson-Davis 2011). This evidence implies that the extended period of weakness and slow recovery following the Great Recession will have negative repercussions that last for decades to come.

Data Appendix

This appendix describes the data used for the wage analyses in Section VI. The basis for these analyses is a sample constructed by pooling the CPS Outgoing Rotation Groups (ORGs) from May 2004 through March 2014.

For hourly workers who do not report that they usually receive overtime pay or who report that their weekly hours vary, I use the self-reported hourly wage. For other workers, I use weekly earnings divided by weekly hours. Hours are constructed as usual hours on the primary job if that is available. If not, I use actual hours in the previous week if the individual had only one job and if these hours are consistent with the self-reported part-time/full-time status. Otherwise, hours are set to missing (as are wages if the hourly wage is not reported directly).

CPS earnings are topcoded at $2,884 per week; I inflate topcoded earnings by 40%. For the distributional analyses in Figure 10, I impose a new topcode equal to the lowest real value of the topcode over the relevant period – 2007-2015 for most analyses, but 1994-2000 or 2000-2006 for others – but do not inflate this to avoid creating holes in the wage distribution. I adjust for inflation using the monthly CPI-U series, and trim at $1 and $200 (in January 2001 dollars). Observations with allocated hourly wages (or weekly earnings, if those are used) are excluded.

Many of the analyses focus on newly-started jobs. These are identified by merging the ORG observation to the regular CPS observations in each of the three previous months. This produces a panel of up to 4 months. An individual is coded as starting a new job if he/she reported in any but the first of these months that she was in a different job than the month before or that her duties or occupation had changed, or if she moved from non-employed (and not on layoff) to employed during the panel.

References

Ananat, Elizabeth O., Anna Gassman-Pines, and Christina M. Gibson-Davis. 2011. The effects of local employment losses on children's educational achievement. In Greg J. Duncan and Richard Murnane (Eds.), *Whither Opportunity? Rising Inequality and the Uncertain Life Chances of Low-Income Children*, pp. 299-314. New York: Russell Sage Publications.

Barro, Robert. 2010. The folly of subsidizing unemployment. *Wall Street Journal*, August 30.

Bernanke, Ben S. 2004. The Great Moderation. Speech delivered at the meetings of the Eastern Economic Association, Washington DC. February 20.

Bernstein, Jared, and Dean Baker. 2003. *The benefits of full employment: When markets work for people*. Economic Policy Institute.

Bewley, Truman F. 1999. *Why Wages Don't Fall During a Recession.* Cambridge: Harvard University Press.

Blanchard, Olivier Jean, and Peter Diamond. 1989. The Beveridge Curve. *Brookings Papers on Economic Activity*, 1: 1-76.

Clark, Kim B., and Lawrence H. Summers. 1982. The dynamics of youth unemployment." In Richard B. Freeman and David S. Wise (Eds.), *The Youth Labor Market Problem: Its Nature, Causes, and Consequences*, pp. 199-234. Chicago: University of Chicago Press.

Congressional Budget Office. 2012. Understanding and responding to persistently high unemployment." February. Accessed at http://www.cbo.gov/publication/42989.

Council of Economic Advisers. 2014. The labor force participation rate since 2007: Causes and policy implications. Executive Office of the President, July. Accessed on September 8, 2015, from https://www.whitehouse.gov/sites/default/files/docs/labor_force_participation_report.pdf.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2010. The establishment-level behavior of vacancies and hiring. NBER Working Paper no. 16265. Cambridge, MA: National Bureau of Economic Research.

Diamond, Peter. 2010. Unemployment, vacancies, wages. Nobel Prize lecture, December 8. Accessed at http://www.nobelprize.org/nobel_prizes/economics/laureates/2010/diamond-lecture.pdf.

Farber, Henry. 2012. Unemployment in the Great Recession: Did the housing market crisis prevent the unemployed from moving to take jobs? *American Economic Review*, 102(3).

Farber, Henry S., and Robert G. Valletta. 2013. "Do extended unemployment benefits lengthen unemployment spells? Evidence from recent cycles in the US labor market." Working paper no. w19048, National Bureau of Economic Research.

Ghayad, Rand and Dickens, William. 2012. "What can we learn by disaggregating the unemployment-vacancy relationship?" Federal Reserve Bank of Boston Public Policy Brief no. 12-3.

Hall, Robert E. 2014. Quantifying the lasting harm to the U.S. economy from the financial crisis. Working paper no. w20183, National Bureau of Economic Research.

Kahn, Lisa B. 2010. The long-term labor market consequences of graduating from college in a bad economy. *Labour Economics*, 17(2): 303-316.

Kaplan, Greg, and Sam Schulhofer-Wohl. 2015. Understanding the Long-Run Decline in Interstate Migration. Federal Reserve Bank of Minneapolis Working Paper 697, June.

Kocherlakota, Narayana. 2010. Inside the FOMC. Speech in Marquette, Michigan, August 17.

Kroft, Kory, Fabian Lange, and Matthew J. Notowidigdo. 2013. "Duration Dependence and Labor Market Conditions: Evidence from a Field Experiment" *The Quarterly Journal of Economics* 128.3: 1123-1167.

Krueger, Alan B., Judd Cramer, and David Cho. 2014. "Are the Long-Term Unemployed on the Margins of the Labor Market?" *Brookings Papers on Economic Activity*: 229-280.

Krueger, Alan B. and Andreas Mueller. 2011. "Job Search, Emotional Well-Being, and Job Finding in a Period of Mass Unemployment: Evidence from High Frequency Longitudinal Data" *Brookings Papers on Economic Activity*: 1-81.

Lazear, Edward P., and James R. Spletzer. 2012. "The United States labor market: Status quo or a new normal?" Working paper no. w18386, National Bureau of Economic Research.

Mishel, Lawrence, Heidi Shierholz, and Kathryn Edwards. 2010. Reasons for skepticism about structural unemployment: Examining the demand-side evidence. EPI Briefing Paper no. 279. Washington, DC: Economic Policy Institute.

Mishel, Lawrence. 2011. Education is not the cure for high unemployment or income inequality. EPI Briefing Paper no. 286. Washington, DC: Economic Policy Institute.

Oreopoulos, Philip, Marianne Page, and Ann Huff Stevens. 2008. Intergenerational effects of job displacement. *Journal of Labor Economics*, 26(3): 455-483.

Oreopoulos, Philip, Till von Wachter, and Andrew Heisz. 2012. The short- and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics*, 4(1): 1–29.

Poterba, James M., and Lawrence H. Summers. 1986. "Reporting errors and labor market dynamics." *Econometrica*: 1319-1338.

Rothstein, Jesse. 2012a. "The Labor Market Four Years into the Crisis: Assessing Structural Explanations." *Industrial and Labor Relations Review* 65.3: 437-500.

Rothstein, Jesse. 2012b. Unemployment insurance and job search in the Great Recession. *Brookings Papers on Economic Activity*, 1.

Rothstein, Jesse. 2014. "The Great Recession and its Aftermath: What Role for Structural Changes?" Paper prepared for *Building Human Capital and Economic Potential* conference, University of Wisconsin, July 2014.

Sahin, Aysegul, Joseph Song, Giorgio Topa, and Giovanni L. Violante. 2011. Measuring mismatch in the U.S. labor market." Federal Reserve Bank of New York Working Paper, October.

Schmitt, John, and Kris Warner. 2011. Deconstructing structural unemployment." Center for Economic Policy research report no. 2011-6.

Shierholz, Heidi. 2014. "Six Years from Its Beginning, the Great Recession's Shadow Looms Over the Labor Market." Report. Economic Policy Institute.

Stevens, Ann Huff, and Jessamyn Schaller. 2011. "Short-run effects of parental job loss on children's academic achievement. *Economics of Education Review*, 30(2): 289-299.

Sullivan, Daniel, and Till von Wachter. 2009. Job displacement and mortality: An analysis using administrative data. *Quarterly Journal of Economics*, 124(3): 1265-1306.

Vlasic, Bill. 2011. Detroit sets its future on a foundation of two-tier wages. *The New York Times*, September 12.

von Wachter, Till, Jae Song, and Joyce Manchester. 2011. Trends in employment and earnings of allowed and rejected applicants to the Social Security Disability Insurance program. *American Economic Review*, 101(7): 3308–29.

Figure 1. Employment, Employment–Population Ratio, and Unemployment Rate, 2004–2015

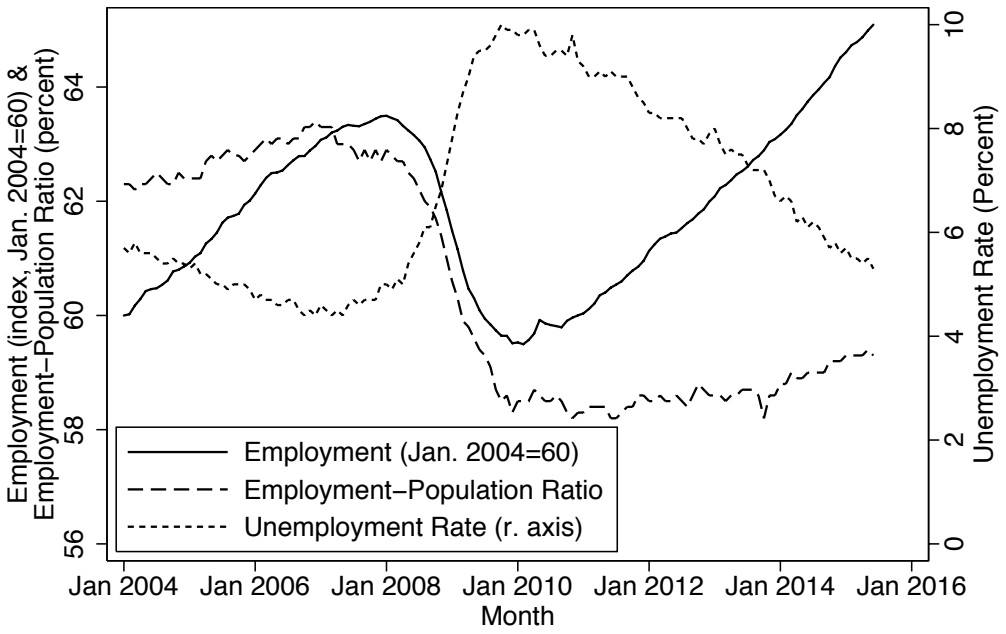

Source: Bureau of Labor Statistics, Current Population Survey and Current Employment Statistics. All series are seasonally adjusted.

Figure 2. Layoffs, Quits, Job Openings and Hires, 2001–2015

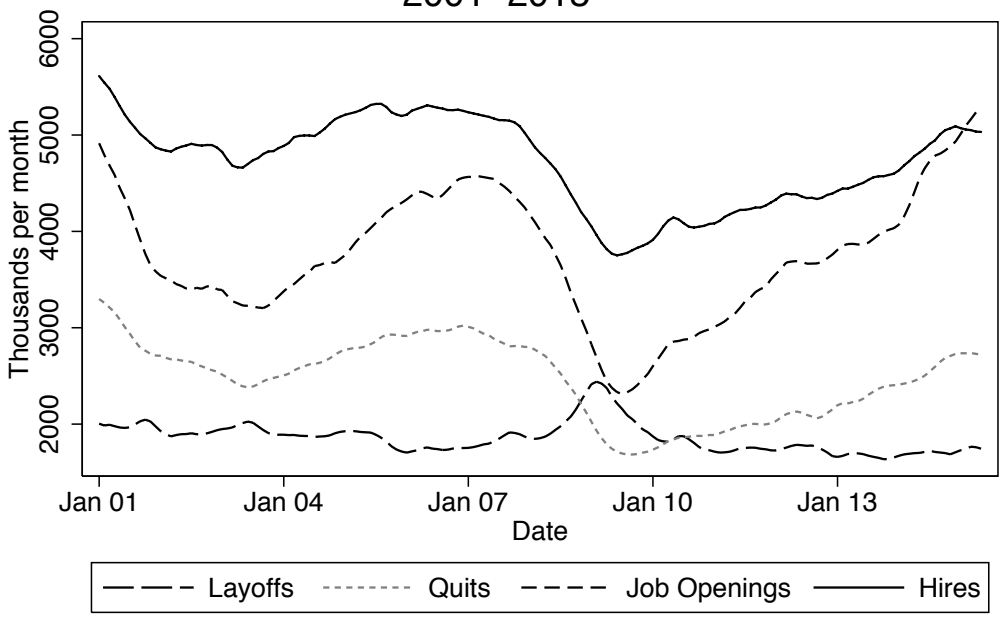

Source: JOLTS. Seasonally adjusted data, smoothed using a 3–month triangle smoother.

Figure 3. Beveridge Curve
(Job openings and Unemployment rates)

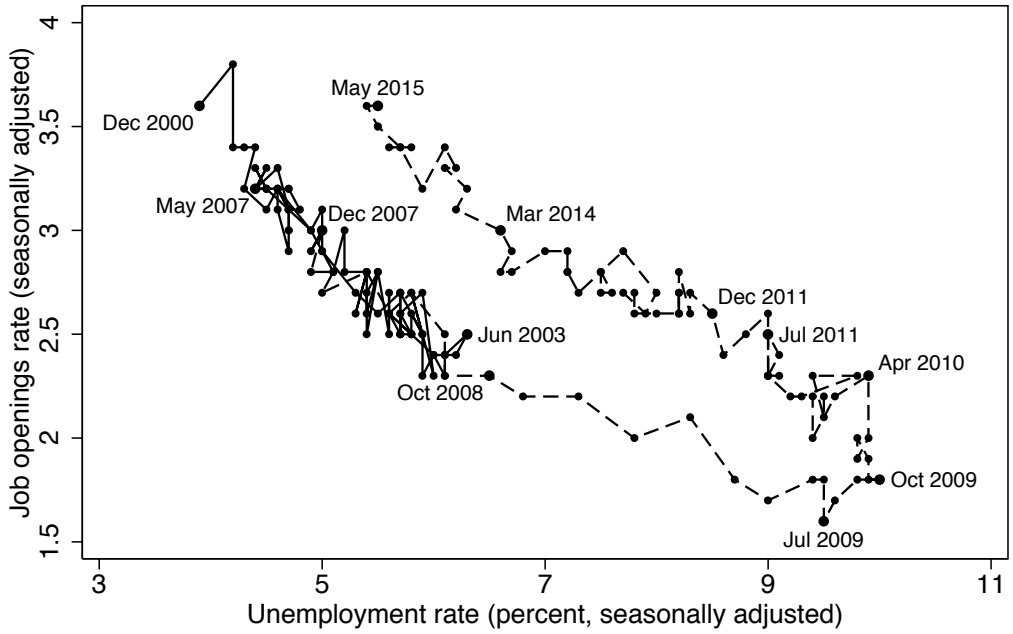

Sources: Bureau of Labor Statistics, Job Openings and Labor Turnover Survey (JOLTS) and Current Population Survey. The December 2007-May 2015 segment is indicated by a dashed line.

Figure 4. State unemployment rates, December 2007 and December 2014

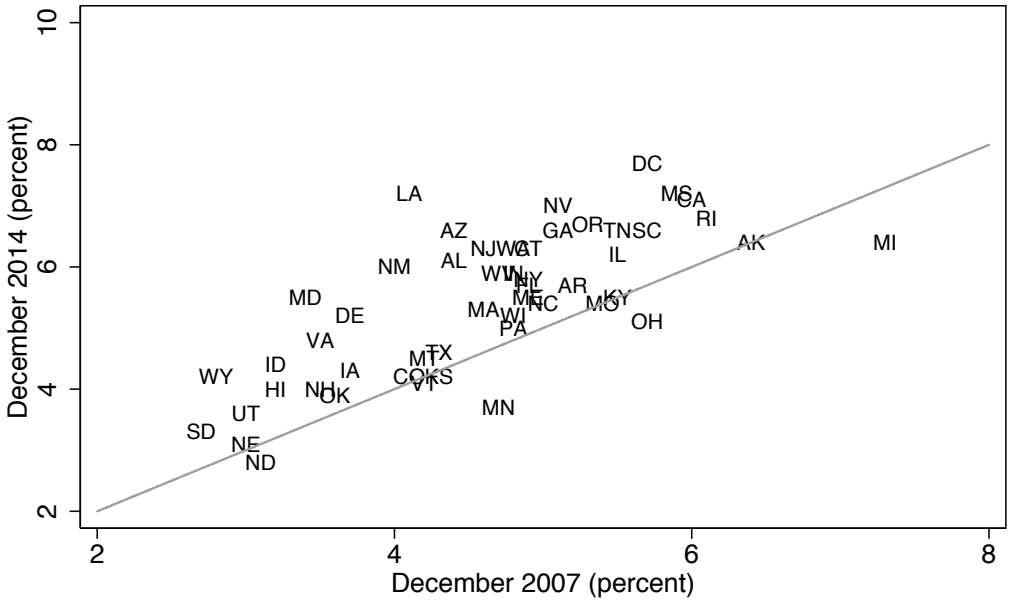

Source: Bureau of Labor Statistics, Local Area Unemployment Statistics. Seasonally adjusted unemployment rates are used.

Figure 5. Unemployment rates in 2007, 2009, and 2014–5, by gender and education

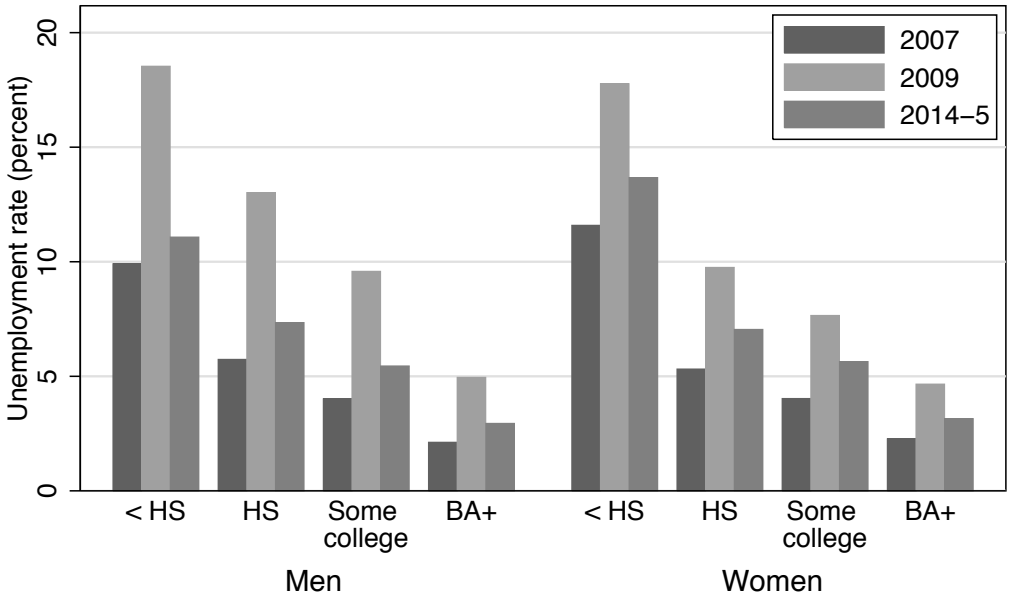

Source: Bureau of Labor Statistics, Current Population Survey. 2014-5 data span July 2014-June 2015. Annual unemployment rates are computed as equally weighted averages of non-seasonally-adjusted monthly estimates.

Figure 6. Actual and predicted change in unemployment rate, 2007–2014, by gender and education

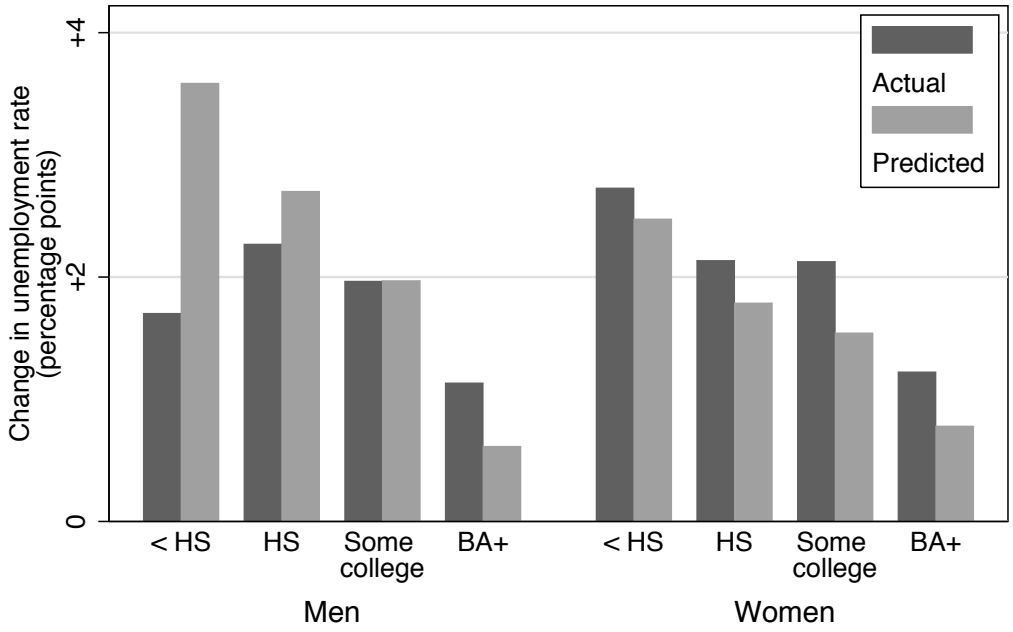

Source: Author's analysis of data from Bureau of Labor Statistics, Current Population Survey. Dark bars show the change between the 2007 and 2014 annual averages of non-seasonally-adjusted monthly unemployment rates. Lighter bars show differences in similar averages of predicted unemployment rates, obtained as the fitted values of a regression of the monthly unemployment rate in the gender-education group on calendar month dummies and the unemployment rate across the rest of the labor force, using data from 1978-2007.

Figure 7. Actual and predicted unemployment rate, HS Graduate Males, 1978–2014

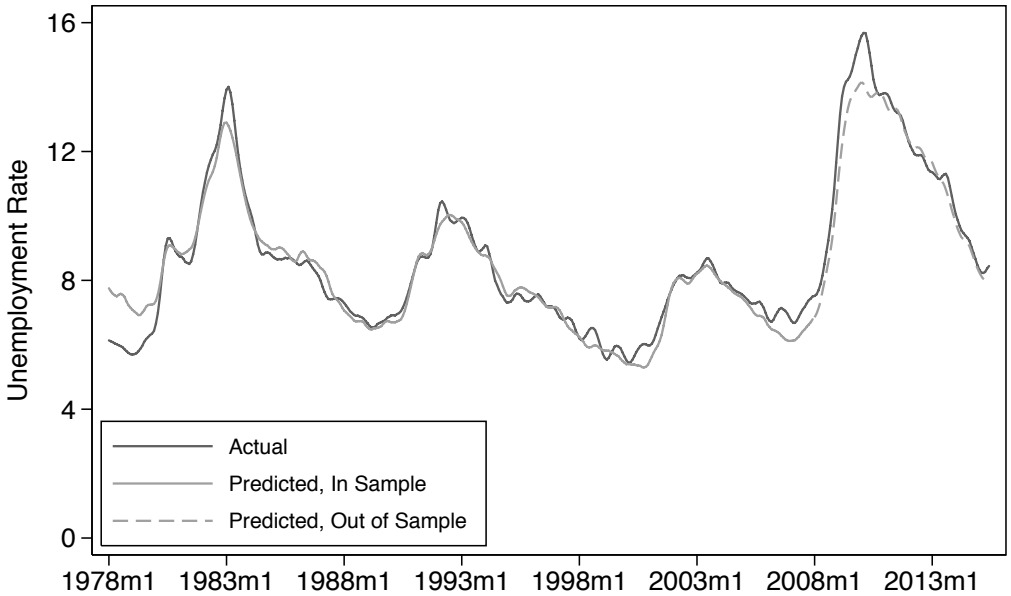

Note: Seasonally adjusted and smoothed using a 3–month triangle smoother.

Note: Both series are seasonally adjusted, and smoothed using a three-month triangle smoother. See notes to Figure 6 for explanation of predicted unemployment rate series.

Fig 8. Actual and predicted change in unemployment rate, 2007–2014, by age

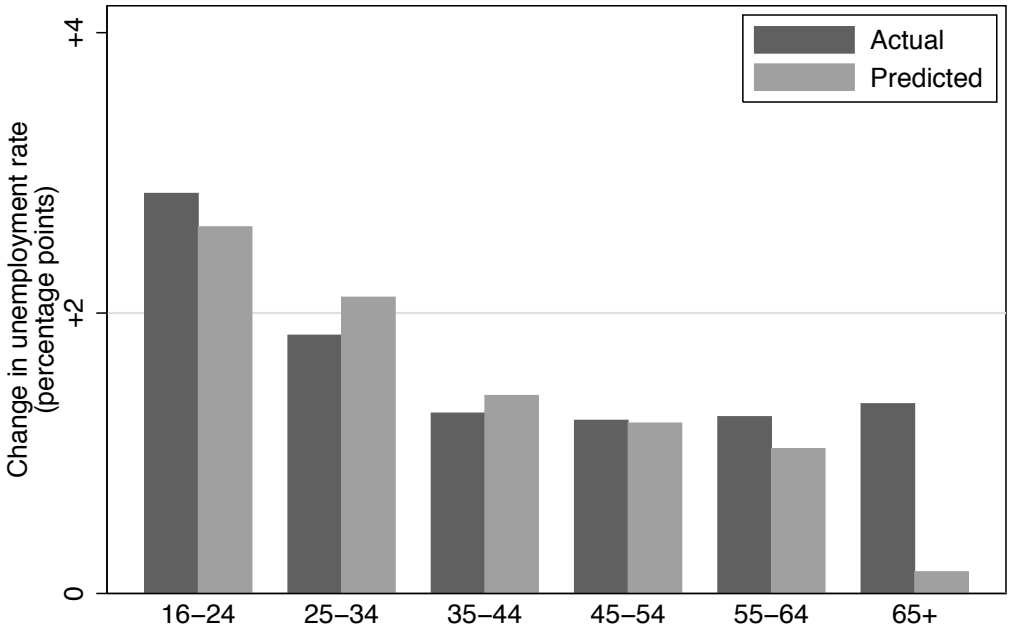

Note: See notes to Figure 6.

Figure 9. Twelve–month changes in mean wages, various subsamples

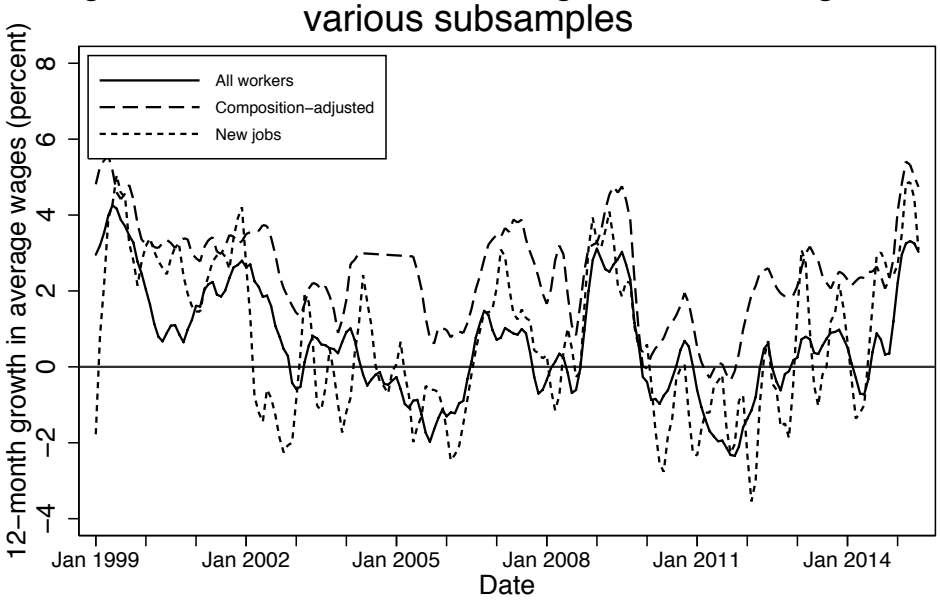

Notes: Composition-adjusted series compares wages within individuals across surveys 12 months apart. New jobs are those that started within the previous 3 months. All series are weighted by weekly hours and smoothed using a 3-month triangle smoother.

Figure 10. Change in distribution of starting wages for all workers, and by industry, MSA, and predicted wage

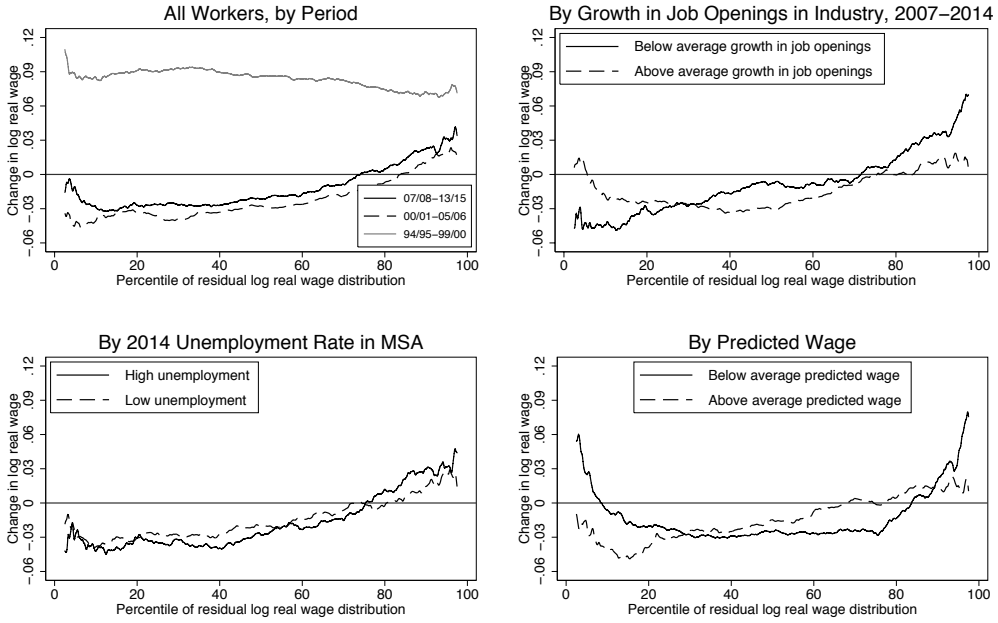

Note: Starting wages are those on jobs that started within the previous 3 months. Wage distributions are weighted by weekly hours. Percentile changes are computed for each 0.1 percentage point, then smoothed across three adjacent points using a triangle smoother.

Figure 11. Reemployment hazards for unemployed workers, by duration group

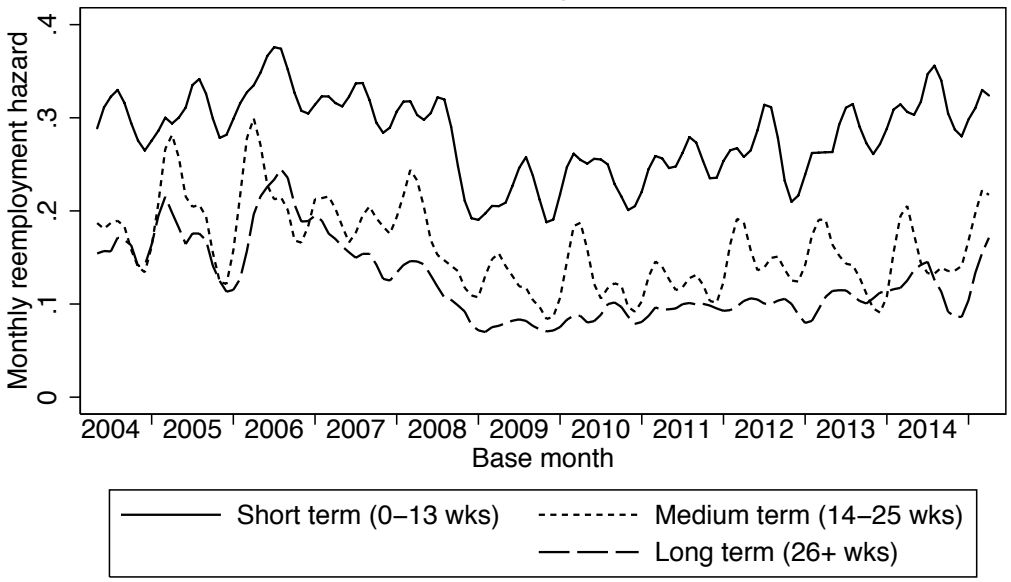

Notes: A worker who is unemployed in one month is counted as having been reemployed if he/she is employed in the next month and does not return to unemployment in the month after that. Series are smoothed using a five-month triangle smoother.

Figure 12. Change in employment rate and non–idleness rate by age, 2005–06 to 2013–14

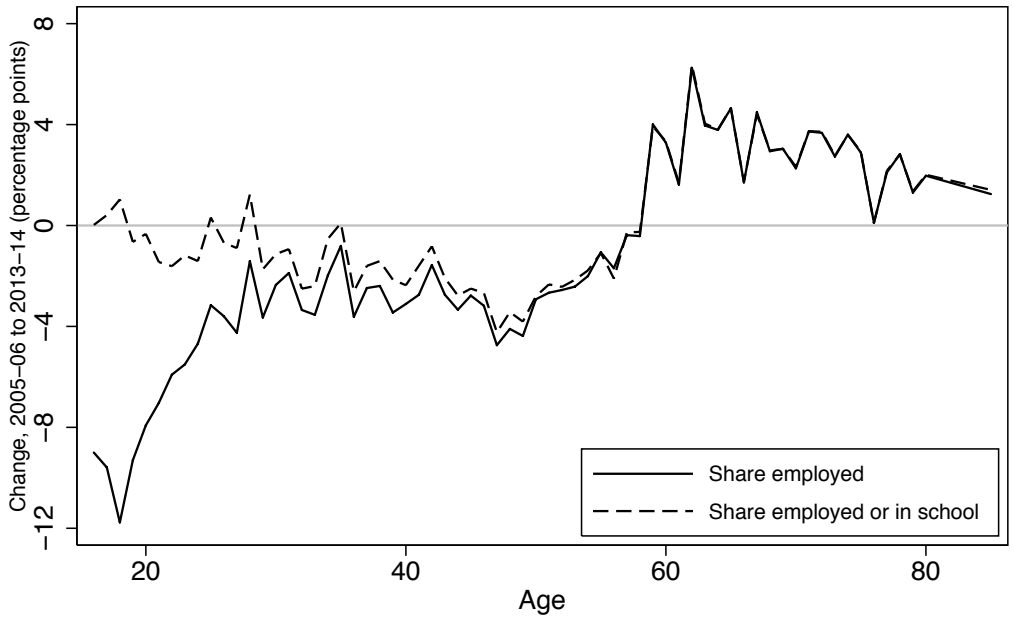

Notes: Computed from monthly CPS files from January-April and September-December of each year.

Table 1: Change in employment rates, working hours, and mean real wages adjusted for observables of new hires, 2007/8 to 2013/15

	Δ Mean Real Wages Adj. for observables (Pct.)		Δ Employment Rate (Pct. Points)		Δ Weekly Hours	
	Mean	SE	Mean	SE	Mean	SE
	(1)	(2)	(3)	(4)	(5)	(6)
Overall	-1.1	(0.3)	-3.6	(0.0)	-0.38	(0.02)
By education and gender						
Male, less than HS	-0.6	(1.1)	-4.8	(0.2)	0.02	(0.07)
Male, HS diploma	-1.6	(0.9)	-5.9	(0.1)	-0.68	(0.04)
Male, some college	-2.7	(1.0)	-5.8	(0.1)	-0.91	(0.04)
Male, BA+	-1.6	(1.0)	-3.7	(0.1)	-0.81	(0.04)
Female, less than HS	-0.9	(1.2)	-3.4	(0.2)	-0.36	(0.09)
Female, HS diploma	0.2	(0.9)	-5.1	(0.1)	-0.68	(0.04)
Female, some college	-1.3	(0.8)	-5.2	(0.1)	-0.51	(0.04)
Female, BA+	-0.5	(0.9)	-2.9	(0.1)	0.07	(0.04)
By age						
16-24	-3.8	(0.6)	-4.5	(0.1)	-0.65	(0.05)
25-34	-1.8	(0.7)	-3.0	(0.1)	-0.52	(0.03)
35-44	-1.7	(0.8)	-2.4	(0.1)	-0.35	(0.03)
45-54	-1.8	(0.8)	-3.0	(0.1)	-0.31	(0.03)
55-64	2.2	(1.0)	-0.7	(0.1)	-0.02	(0.04)
65+	7.4	(2.1)	2.0	(0.1)	1.16	(0.08)
By industry			Pct. Δ			
Agriculture	-3.4	(3.0)			-0.72	(0.17)
Mining and logging	3.3	(3.9)	23.8		-0.74	(0.22)
Construction	-2.7	(1.4)	-19.5		0.12	(0.05)
Durable goods mfg	-1.1	(1.3)	-12.7		-0.05	(0.05)
Nondurable goods mfg	0.6	(1.8)	-11.2		-0.32	(0.06)
Wholesale trade	-5.9	(2.3)	-3.1		-0.06	(0.08)
Retail trade	-3.6	(0.9)	-1.0		-0.84	(0.05)
Transport and utilities	-1.3	(1.8)	1.8		-0.38	(0.08)
Information	0.2	(2.7)	-9.6		-0.05	(0.10)
Finance and insurance	3.8	(1.6)	-4.0		0.40	(0.05)
Real estate	-2.8	(3.0)	-5.7		0.03	(0.12)
Prof and bus svcs	-2.3	(1.1)	6.4		-0.30	(0.05)
Education (private)	2.9	(2.0)	16.2		-0.01	(0.10)
Health and soc assistance	-2.0	(1.0)	15.2		-0.11	(0.04)
Arts and recreation	-1.2	(2.3)	6.7		-0.85	(0.13)
Lodging and food services	1.5	(1.0)	10.0		-0.92	(0.06)
Other services	-0.4	(1.8)	1.5		-0.84	(0.08)
Federal government	-2.8	(2.1)	-0.3		-0.10	(0.07)
State/local government	-1.0	(1.0)	-1.8		-0.04	(0.04)

Notes: 2013-5 data reflect July 2013 through June 2015. New hires are those who began their jobs within the previous three months. Adjusted estimates in columns 5-6 are the changes in mean residuals from a log wage regression, estimated on 2004-6 data, with controls for education-by-gender, state, and industry-by-education indicators, an age quadratic, and interactions of a linear age term with education-gender indicators. SEs in column 6 do not account for sampling error in the regression coefficients. Industry percent change in employment based on CES.

Table 2: Unemployment rates and long-term unemployment shares by demographic group, 2006 and 2014.

	Unemployment rate (%)		Long-term unemployment share	
	2006	2014	2006	2014
	(1)	(2)	(3)	(4)
Overall	4.6	6.2	21.2	37.3
By education and gender				
Male, less than HS	9.3	11.6	19.5	33.3
Male, HS diploma	5.7	8.0	22.9	39.4
Male, some college	3.9	6.0	21.6	37.3
Male, BA+	2.1	3.3	26.4	40.5
Female, less than HS	11.5	14.3	17.5	32.4
Female, HS diploma	5.4	7.5	21.1	38.5
Female, some college	4.2	6.2	20.1	35.8
Female, BA+	2.3	3.5	20.8	38.1
By age				
16-24	10.5	13.4	15.1	26.4
25-34	4.7	6.5	19.6	37.5
35-44	3.6	4.7	24.1	39.9
45-54	3.1	4.4	27.0	44.3
55-64	3.0	4.4	31.4	48.7
65+	2.9	4.6	24.4	47.4

Notes: The long-term unemployment share is the fraction of the unemployed who have been out of work 27 weeks or more.