

**The Labor Market Four Years Into the Crisis:
Assessing Structural Explanations**

March 2012

Jesse Rothstein *

University of California, Berkeley and NBER

Abstract

Four years after the beginning of the Great Recession, the labor market remains historically weak. Many observers have concluded that “structural” impediments to recovery bear some of the blame. This paper reviews such structural explanations. I find that there is little evidence supporting these hypotheses, and that the bulk of the evidence is more consistent with the hypothesis that continued poor performance is primarily attributable to shortfalls in the aggregate demand for labor.

* I thank Jared Bernstein, David Card, Larry Kahn, Michael Reich, and Larry Katz for helpful conversations. Tommy Drake and Sarena Goodman provided excellent research assistance. I am grateful for financial support to the Institute for Research on Labor and Employment and the Center for Equitable Growth, both at UC Berkeley, and to the Russell Sage Foundation. Copies of the data and computer programs used to generate the results presented in the paper are available from the author at rothstein@berkeley.edu.

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 attributable 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 the labor market has been extremely slow to recover, and all of the available metrics indicate substantial continuing weakness. Although real output recovered its pre-recession peak in the third quarter of 2011, as of this writing payroll employment remains 3.8% below its December 2007 level.¹ The unemployment rate has been above 8% for 37 consecutive months, the longest such period since World War II. (The prior record was 26 months, in 1981-83.) Moreover, while the unemployment rate has begun to decline, much of the recent decline reflects reduced labor force participation rather than increased employment. The male employment-to-population ratio, which fell by an unprecedented 7.1 percentage points between December 2006 and December 2009, hovered around 64 percent – nearly four percentage points lower than had ever been recorded before the current cycle – for the subsequent two years, and has only recently begun to recover slowly.

Economists will be debating the causes and interpretations of this cycle for decades to come. Current views about the state of the labor market can be divided into

¹ Current figures refer to the most recent available data at the time of writing: The February 2012 employment and GDP releases from the Bureau of Labor Statistics and the Bureau of Economic Analysis, respectively.

two rough camps, though of course there is heterogeneity of opinion within each one. One camp, of which Paul Krugman is perhaps the most prominent member (see also Romer 2011), argues that recent poor outcomes are primarily reflective of a shortfall of aggregate demand. This camp prescribes aggressively stimulative monetary policy – which would have to take unconventional forms, as the federal funds rate has been at or near its zero lower bound since late 2008 – and additional fiscal stimulus to raise effective demand.

A second camp points to “structural” factors as important impediments to labor market recovery. This diagnosis comes in several flavors. Some focus on mismatch between the types of labor supplied by workers and the types demanded by employers. As Narayana Kocherlakota, President of the Federal Reserve Bank of Minneapolis, described it in a 2010 speech, “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).

Others focus on workers’ labor supply decisions. For example, Mulligan (2009; see also Mulligan, 2011a) concludes that reductions in labor supply, due either to changing worker preferences or to increases in labor market distortions, explain much or all of the decline in employment in 2008. Mulligan (2011b) points to a particular source of such distortions, noting that safety net spending has increased dramatically over the last several years.

The common element of the various structural explanations is that they posit that, as expressed by participants at the January 2012 Federal Reserve Board Open Markets Committee meeting, “a substantial part of the increase in unemployment since the beginning of the recession reflected factors other than a shortfall in aggregate demand.”² These explanations thus generally militate against activist policies aimed at spurring labor demand. Kocherlakota (2010) concludes that “[m]ost of the existing unemployment represents mismatch that is not readily amenable to monetary policy,” while Mulligan (2011b) suggests that countercyclical fiscal policies like unemployment insurance extensions that reduce the private return to employment may slow rather than

² <http://www.federalreserve.gov/monetarypolicy/fomcminutes20120125.htm>.

hasten recovery. To the extent that the primary labor market problem is structural, policy solutions would have to fall on the labor supply side: Depending on the specific source of the structural problems, one might want to emphasize job training, mobility assistance, or reductions in effective labor income tax rates.

This paper reviews labor market data with an eye toward assessing the plausibility of structural explanations for recent performance.³ I focus on short-run labor market dynamics, distinguishing structural hypotheses about these dynamics from the related but distinct hypothesis that long-run, slow-moving structural trends in our economy have harmed some groups and helped others.⁴ (For example, Autor and Dorn, 2011, argue that technical change since 1980 has worked to the disadvantage of middle-skill workers in occupations requiring “routine” work that can be easily computerized; see also Reich, 1992.) This is not to say that long-run structural trends are not of interest for my analysis. They are relevant insofar as these trends can help to explain the collapse in labor market outcomes between 2007 and 2011 or the prospects for labor market recovery in response to aggregate demand increases. For example, long-run declines in labor market flexibility might have been expected to slow job losses during the downturn but also to prolong the labor market recovery after demand returns.

I examine the various forms of evidence that have been seen as indicating structural problems. I also propose a new source of information about the importance of structural impediments to labor market recovery. Importantly, most structural explanations for our current predicament (though not changes in adjustment costs) imply that the labor market appears tight from the perspective of some or all potential employers: Despite high measured unemployment, there are relatively few workers who are both interested in and qualified for the jobs on offer. Employers facing tight labor

³ Diamond (2010) examines different evidence than that considered here but comes to a similar conclusion, as do Mishel, Shierholz, and Edwards (2010) and Mishel (2011). The Congressional Budget Office (2012) is more favorable toward structural hypotheses but nevertheless concludes that aggregate demand shortfalls are the primary source of the high unemployment rate. Sahin, Song, Topa, and Violante (2011) conclude that across-industry mismatch is more important than the data reviewed here indicate.

⁴ I set aside until Section VIII the hypothesis that extended cyclical unemployment might eventually become structural, as idle workers’ human capital gradually depreciates, except insofar as this hysteresis has already taken place to a sufficient degree to represent an important impediment to rapid recovery.

markets should bid up the wage in order to attract workers. Labor demand shortfalls, by contrast, would have an opposite effect, as unemployed workers to bid down equilibrium wages as they compete for the few available jobs.⁵ I thus emphasize the examination of wage trends for evidence of structural impediments to growth.

Section II begins by reviewing the state of the labor market. Section III develops a specific definition of “structural” impediments to labor market recovery and argues for the value of price as well as quantity data in identifying such impediments. Sections IV through VII explore the various forms of available evidence: Aggregate data on GDP and the job openings rate (Section IV); estimates of the effects of Unemployment Insurance extensions (Section V); wages for newly-hired workers (Section VI); and the long-term unemployment share (Section VII). Section VIII concludes.

II. The state of the 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 temporary hiring associated with the 2010 census – at a rate of 144,000 new jobs per month. This is only a bit faster than is needed to keep up with population growth, and as a result the employment-population ratio, which fell from 62.9% in January 2008 to 58.2% in December 2009, has hovered in a very narrow range around 58.5% for over two years. Thus, while, the unemployment rate has fallen from its peak of 10.0% in October 2009 to 8.3% in February 2012, this decline is almost entirely attributable to falling labor force participation among the non-employed.

{{Place Figure 1 about here}}

Figure 2 shows the change in employment by broad industry group between December 2007 and December 2009 – roughly the period of employment contraction –

⁵ 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).

and between December 2009 and December 2011.⁶ Industry shares of December 2007 employment are shown in parentheses. The standard narrative holds 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 – 5.8% 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, each of which contracted by well over 15%. Only four of the sectors portrayed here saw rising employment over the first two years of the downturn. None of these grew by more than 4%, and among reasonably large sectors only health care grew by more than 1%.

{{Place Figure 2 about here}}

The light bars in Figure 2 show mostly small changes in the period since the end of 2009. An exception is the mining and logging industry, which has grown by more than 20% since 2009 and thus has more than made up for the 10% decline between 2007 and 2009. Beyond this small sector – which has only about 0.5% of national employment – the data show only a bit of unevenness around the general slow growth trend since 2009. Employment in education and health has continued to grow since 2009, while the professional and business services, hospitality, and “other services” industries have each recovered half or more of their initial losses. Trade, transportation, and durable goods manufacturing have also added jobs but remain well below their 2007 employment levels. Other industries, including construction, nondurable goods manufacturing, information, finance, and real estate, have continued to lose jobs, though much more slowly than in the earlier period. Government employment, which grew slightly in the first two years of the recession, has declined more recently, led by states and local governments.

Construction and manufacturing employment is heavily male and largely non-college-educated, so one might expect that low-skill men would have suffered

⁶ I focus on December to December comparisons to avoid seasonal adjustment concerns. All changes are expressed as shares of employment in the industry in December 2007.

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

disproportionately in the recession. Figure 3 shows the unemployment rate by gender and education in 2007 and in 2011. Consistent with the industrial composition of the cyclical collapse, we see that unemployment rates of less educated men rose more than did those of more educated men. 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 dramatically for low-skill women as well as for men. Indeed, across all eight gender-education groups, unemployment rates in 2011 were roughly double their 2007 levels.⁸ Once more, this appears more consistent with a cyclical decline than with structural shifts that favor particular subgroups.

{{Place Figure 3 about here}}

One deviation from the general doubling is that unemployment rates rose by somewhat more (in percentage terms) for men than for women at each education level. This pattern has led some commentators to refer to the Great Recession as a “man-cession.” Similarly, some have concluded from the extremely high unemployment rates among young people – 14.4% in December 2011 for those aged 20-24 – that dysfunctional labor market institutions are effectively reserving jobs for insiders.

However, it is not clear that the current cycle is unusual in these regards. The construction and manufacturing sectors have historically been more cyclical than the economy as a whole, and so one would expect low-skilled men to suffer more in any downturn. Similarly, youth unemployment has always been highly sensitive to economic conditions (Clark and Summers 1982).

Figures 4, 5, and 6 illustrate this. For each industry (Figure 4), gender-education group (Figure 5), and age group (Figure 6), I plot the actual change in the unemployment rate between 2007 and 2011, as well as the change in unemployment that would have been expected given the past cyclical sensitivity of that group’s unemployment and the

⁸ Figure 3 shows that the male unemployment rate was lower than the female rate in three of the four education groups in 2007 and in two out of four in 2011. Nevertheless, the overall male unemployment rate was higher in each year, as men have less education, on average, than do women.

magnitude of the cycle.⁹ 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 .¹⁰ A value of β_g greater than 1 indicates that group g is more cyclically sensitive than others. I estimate α_g and β_g using monthly observations from 1978 through 2007, then use these coefficients and the observed path of $u_{(-g)t}$ to forecast u_{gt} through 2011. The lighter bars in figures 4-6 illustrate $\hat{u}_{g,2011} - \hat{u}_{g,2007}$.¹¹

{{Place Figures 4, 5, and 6 about here}}

The figures show that the vast majority of the across-group differences in unemployment growth between 2007 and 2011 are attributable to differences in cyclical sensitivity rather than to unique features of this business cycle. However, there are some anomalies. First, contrary to many discussions of the housing bust, construction industry unemployment – which is extremely cyclical – has risen by 1.5 percentage points *less* than would be expected given the weak overall labor market, and the anomaly in durable goods manufacturing is more than twice as large. In contrast, unemployment has risen more than predicted in the usually less cyclical agriculture, information, finance, real estate, professional services, education, and health industries. Insofar as there have been structural shifts, they have apparently been *toward* the goods-producing industries and *away* from the high-skill services, though these shifts have been masked by the across-the-board cyclical decline.

Turning to Figure 5, the unemployment rate for men without high school diplomas – again, a group that is ordinarily very cyclically sensitive – has risen by 1.4 percentage points less than expected, while women’s unemployment rates have risen slightly more than expected. The term “man-cession” is thus wholly inappropriate, as at

⁹ Unemployed workers are assigned to the industry in which they last worked.

¹⁰ I compute $u_{(-g)t}$ using fixed weights for each group $h \neq g$ over time, proportional to the group’s average labor force share over the 1978 to 2011 period.

¹¹ Results are similar if I instead predict the change as the difference between the predicted 2011 rate and the *actual* 2007 rate, $\hat{u}_{g,2011} - u_{g,2007}$.

least as of 2011 the recession has hit *women* unusually hard.¹² Finally, Figure 6 indicates that over-65 workers, whose unemployment rate has risen less than have those of younger workers, have nevertheless been much more affected by this business cycle than by past cycles.¹³ These data suggest that many of the gender, industry, and age patterns that have been the focus of public discussions are simple characteristics of severe recessions and not unique to this one.

Another important source of heterogeneity is geographic. The recession has hit some areas – most famously, Sun Belt cities like Las Vegas where the housing boom was most pronounced – harder than others. Figure 7 plots state unemployment rates in December 2007 and December 2011. Unemployment rates rose over this period in all 50 states plus the District of Columbia, but by different amounts: Nevada, North Carolina, California, Florida, and Rhode Island saw notably large increases, while Alaska, Minnesota, Nebraska, Vermont, and the Dakotas have been relatively unaffected. Notably, the five states with December 2011 unemployment rates below 5.5% together account for less than 2% of national employment.

{{Place Figure 7 about here}}

Taken together, Figures 2-7 show that the rise in unemployment between 2007 and 2011 was quite broad based, affecting workers of all ages, education levels, genders, and industries. Sectors that have been more cyclically sensitive in the past saw larger increases, but there was little heterogeneity beyond this. Developments since 2009 have shown a bit more heterogeneity than did those in the initial collapse, but the labor market

¹² One possible explanation for the spread of the idea that the recession was uneven in its impacts is that outcomes as of 2009 were a bit different than those seen in 2011. But even in 2009 the unemployment rate for non-high-school men was a bit lower than would have been expected based on past patterns, while for more educated men the excess unemployment was quite small.

¹³ Of course, unemployment rate changes combine changes in employment and in labor force participation. The over-65 anomaly in Figure 6 is much reduced when I instead examine the employment-population ratio, suggesting that much of the anomaly reflects increased labor force participation in this group, perhaps related to collapses in retirement account values due to equity market declines. By contrast, the labor force participation of less-educated men appears to have fallen by more than expected given the severity of the recession.

is still quite depressed essentially across the board. 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.

III. A focus on structural explanations

Discussions of economic aggregates often distinguish between cyclical and structural components. The budget of a government with strong automatic stabilizers may swing from surplus to deficit as the economy weakens, then return to surplus during the recovery. Deficits observed during the downturn would generally be described as cyclical. By contrast, a government that ran a deficit even at the business cycle peak would generally be agreed to be running a structural deficit.

The distinction between cyclical and structural components becomes muddier, however, when one examines outcomes that are constituent of the business cycle itself. If we define the structural component of an outcome as the outcome when measured at a business cycle peak, then by definition there can be no structural component to the variables used to identify that peak.¹⁴ For example, the employment-population ratio was never above 63.4% during the 2001-2007 economic expansion, and thus was always well below the level seen throughout 1998, 1999, and 2000. Does this necessarily mean that structural non-employment rose between 1999 and 2006, due perhaps to changes in government policies or in the rate of technical or demographic change? One would like a definition that allows for the possibility that the economy never attained its full potential during the unusually weak 2001-2007 expansion.

In practice, discussions of structural components of labor market outcomes often use a different (though related) definition. Many commentators appear to describe as structural any labor market outcome that would not be improved by a balanced increase in labor demand to levels ordinarily seen in business cycle expansions. For example, if

¹⁴ If recessions are defined solely based on measured economic output, employment and unemployment are not mechanically related to cycle measurement. But it does not seem reasonable to define any unemployment observed while GDP is growing as *prima facie* structural. Many business cycle dating efforts recognize this by considering other measures. The NBER, for example, bases its decisions on “economic activity,” considering various measures – including “economy-wide employment” – to assess this (NBER 2010).

employment fell during the 2007-2010 period because the labor supply curve shifted inward, as Mulligan (2011a) argues, then even if labor demand were to return to its 2007 level employment would not. One might then refer to the component of the employment decline that derives from falling supply rather than from falling demand as structural.

As another example, unemployment increases coming from changes in the *composition* of labor demand that are unmatched by corresponding changes in labor supply might be considered structural. For example, consider Figure 8, which shows traditional supply and demand curves for two types of labor, labeled “A” and “B.” Imagine that the demand curve shifts inward for A labor and outward for B labor, as in the D' curves, with no change in the total amount of labor demanded at the pre-change equilibrium wages w_A and w_B . (That is, $D'_A(w_A)+D'_B(w_B)=D_A(w_A)+D_B(w_B)$, as in the figure.) If the supply of type-B labor is less elastic than is the supply of type-A labor, as shown in the Figure, total employment will fall: $L'_A+L'_B<L_A+L_B$. Moreover, equal outward shifts in both labor demand curves large enough to restore the original aggregate employment level – marked as D'' , with $L''_A+L''_B=L_A+L_B$ – would require average wages to rise above the pre-shock level. In this case, it seems reasonable to refer to at least a portion of the reduction in employment from $(L_A + L_B)$ to $(L'_A + L'_B)$ as structural.

{{Place Figure 8 about here}}

In this example, if the downward shift in relative demand for type-A workers is permanent¹⁵ then the only non-inflationary way to restore the pre-shock employment level would be for workers to shift from labor market A to market B, perhaps by retraining (if labor types are distinguished by occupational skills) or by moving (if the distinction is geographic). If these shifts are possible, one would expect the type-B supply curve to become more elastic over time, increasing employment in market B.

Both of these examples make clear that structural and cyclical unemployment can coincide, and that declines in total employment can be consistent with either type. The

¹⁵ It is possible for the sort of relative demand shift depicted in Figure 8 to be wholly cyclical. For example, if B is an inferior good, a negative shock to aggregate incomes could produce the shifts shown. In this case, one would expect the demand shift to reverse itself during the business cycle upturn.

distinguishing feature of structural unemployment is (the possibility of) tight labor markets even when employment remains below its potential. In other words, even if monetary or fiscal policy produced a balanced outward labor demand shift sufficient to restore the pre-shock wage level, employment would remain depressed. Larger demand shifts might in principle restore the employment level, but only with higher (and potentially inflationary) equilibrium wages.

Of course, these characteristics of structural unemployment might be impossible to observe so long as cyclical unemployment remains high. In the two-labor-markets example illustrated in Figure 8, for example, I posited that the increase in demand for type B labor was large enough to offset the decline in type-A demand, at pre-shock wages. A more realistic characterization of recent history might involve a small decline in type-B demand accompanied by a larger decline for type A. The resulting rise in unemployment would have both structural and cyclical components.

In the next four sections, I investigate the available data for patterns that might support or contradict the hypothesis that structural factors contributed importantly to the slow labor market recovery in 2009-2011. I consider first the aggregate quantities data that have been most often interpreted as indicative of a structural problem, then discuss the evidence regarding labor supply disincentives, evidence from wage trends, and finally the long-term unemployment share.

IV. Aggregate Quantities: Okun's Law and the Beveridge Curve

Two macroeconomic phenomena that appeared relatively early in the Great Recession suggested that cyclical explanations might not be sufficient to explain the weakness of the labor market.

The first of these was that Okun's Law appeared to have been violated (Elsby, Hobijn, and Sahin, 2010). A rule of thumb that has been useful in the past is that the unemployment rate rises by approximately 1% for every 2% that real output growth falls short of its potential. Early in the recession, there appeared to be a relatively large deviation from this historical pattern. Consider the period through the third quarter of 2009: The "advance" estimate released in October 2009 indicated that real GDP in

2009:Q3 was 2.8% below that in 2007:Q4.¹⁶ If we assume that potential GDP grew 2.5% per year over this period,¹⁷ the shortfall, expressed as a share of 2009:Q3 potential GDP, was 6.9%. However, the unemployment rate rose 4.8 percentage points between 2007:Q4 and 2009:Q3, 1.3 percentage points more than Okun's Law would have predicted.

However, data revisions and subsequent performance have entirely eliminated this anomaly. First, the GDP decline between 2007 and 2009 was revised to be much sharper. Current data (as of March 2012) indicate that the deviation from potential GDP growth over the 2007:Q4-2009:Q3 period was 8.8 percent, 1.8 percentage points more than it initially appeared. This eliminates about two-thirds of the anomaly over that period. Moreover, in the 9 quarters from 2009:Q3 through 2011:Q4, real GDP grew at an annual rate of exactly 2.5 percent. Thus, Okun's Law would predict that the unemployment rate should have been stable over this period; instead, it fell by almost a full percentage point. Thus, over the entire 2007:Q4-2011:Q4 period the unemployment rate increase of 3.9 percentage points is actually substantially smaller than the Okun's Law prediction of 4.4 percentage points. From the current vantage, the puzzle is that the unemployment rate is so *low*, not that it has not fallen enough. Of course, as the recent decline in the unemployment rate has overwhelmingly reflected falling labor force participation rather than rising employment, this is less of a puzzle than it first appears.

A second phenomenon that has been often interpreted as an indication of recent structural problems is an apparent shift in the Beveridge Curve that relates job openings to the unemployment rate. Figure 9 illustrates this curve, using job openings data from the Bureau of Labor Statistics' Job Openings and Labor Turnover Survey (JOLTS). One expects job openings and unemployment to be inversely related: In tight labor markets with low unemployment, jobs are filled slowly and the job openings rate is therefore

¹⁶ The advance estimate is the first preliminary estimate of quarterly GDP produced by Bureau of Economic Analysis; there are two scheduled updates that reflect late-arriving data, plus later revisions as needed. Historical estimates are available from the BEA website, at <http://www.bea.gov/newsreleases/relsarchivegdp.htm>.

¹⁷ In January 2007, the Congressional Budget Office (2007) forecast that potential GDP would grow at an average annual rate of 2.8% between 2007 and 2012. Its more recent estimates have been somewhat lower, in part reflecting the effects of the recession (via, e.g., the rate of capital investment).

high, while when unemployment is high vacancies are filled quickly and there are few jobs open at 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 (Blanchard and Diamond, 1989).

{{Place Figure 9 about here}}

As Figure 9 illustrates, in 2008 and the first half of 2009 job openings fell steadily as the unemployment rate rose. Starting in mid-2009, however, unemployment stagnated but the job openings rate began rising. Since early 2010, the job openings rate has remained a bit more than 0.5 point above what was seen with similar unemployment rates in early 2009.

A number of commentators have interpreted this apparent shift in the Beveridge Curve as diagnostic of increases in structural unemployment. In this view, the rise in job openings indicates that labor demand has shifted outward, while the stability of the unemployment rate suggests that the currently unemployed are unable or unwilling to fill the newly created positions.

A number of objections might be raised against this inference: Nonlinearity of the Beveridge Curve might make it difficult to identify outward shifts while unemployment remains high; the unemployment rate may be a poor proxy for labor market tightness when there are many discouraged workers who are temporarily out of the labor force; or the rise in job openings might simply be part of a normal counterclockwise rotation of the Beveridge Curve in an economic recovery.¹⁸ None of these, however, can explain more than a small share of the sustained apparent outward shift of the curve.

A more important concern is that measured job openings data and the openings-to-unemployment ratio are only loosely related to the efficiency of the economic matching process, particularly in an unprecedentedly long period of labor market weakness. This is because the definition of a job opening used by the JOLTS survey

¹⁸ Tasci and Lindner (2010) note that the Beveridge Curve has often rotated counterclockwise during the early part of recoveries, and hypothesize that this is because initial increases in labor demand draw discouraged workers back into the labor force and prevent the measured unemployment rate from falling.

does not closely correspond to any economically meaningful concept. Thus, the increase in job openings provides at best weak support for the view that labor demand has increased substantially since 2009.

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 not realistic. Both the wage and the required qualifications are choice variables that can influence the number of measured openings independent of actual labor demand.¹⁹

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.²⁰ 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.²¹

¹⁹ 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).

²⁰ 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.

²¹ Davis, Faberman and Haltiwanger (2010) examine the related idea that the employer's choice of "recruiting intensity" can influence the rate at which vacancies are filled. However, they do not have a direct measure of recruiting intensity, and must proxy for it using the ratio of hires to vacancies. This makes it impossible to evaluate whether labor supply changes have altered the job filling rate that would be seen for a given labor demand and a fixed recruiting effort level.

These definitional issues may be unusually important now. In the past, 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 may be quite likely to leave the firm once conditions improve (Bewley, 1999). The saliency of these concerns may be in decline: Anecdotally, two-tier wage structures that distinguish between incumbent and newly hired workers have become increasingly common (Vlasic 2011), and in a downturn that is expected to be prolonged an employer may not worry as much about retaining its workers after the economy recovers.

If indeed employers are taking advantage of the weak labor market to reduce offered wages or to hire more qualified workers, one would expect this to reduce the rate at which posted vacancies are filled and therefore to raise the job openings rate. This limits our ability to diagnose labor market tightness based solely on the aggregate Beveridge Curve.

{{Place Figure 10 about here}}

Figure 10 presents some evidence that rising job openings may not indicate labor market tightness. The horizontal axis shows the increase in job openings by industrial sector from July 2009 to July 2011, while the vertical axis shows employment growth in the sector over the same time period. The mining and logging sector is excluded from the figure as it is an enormous outlier: Its job openings rate rose by 3.5 percentage points and employment grew by 17.4%. Across all industries, the changes in the job openings rate and in employment are modestly correlated ($\rho=0.55$). But this is entirely driven by mining and logging; when this industry is excluded, the correlation falls to 0.04. Moreover, the other industries with the biggest increases in job openings are information, arts and recreation, durable goods manufacturing, and professional and business services. All four of these sectors had lower employment in December 2011 than in December 2007, and one (information) even lost employment between 2009 and 2011. Also, as figure 4 showed, in each of these sectors the unemployment rate grew by more than 3.5

percentage points between 2007 and 2011. It thus appears unlikely that any of them are suffering from shortages of appropriately-skilled workers.

Only in mining and logging does the combination of employment and job openings growth appear to indicate meaningful labor supply shortages (though even here the unemployment rate was 2.8 percentage points higher in 2011 than in 2007). Increasing energy prices have led to dramatic expansions of this sector, and both job openings and employment have risen substantially. There may be structural impediments that are preventing the sector from growing even more quickly than it has. However, mining and logging accounts for only about 0.5% of national employment, so even if its growth is being importantly hampered by supply shortages this cannot explain a large share of overall labor market weakness.

The data thus appear consistent with the view that the increase in job openings reflects reduced recruiting effort, lower offered wages, or higher minimum qualifications rather than labor supply shortages in fast-growing sectors. However, it is also possible that *intra*-industry shifts in labor demand have created shortages of some particular types of workers within individual sectors that are masked by weakness in other subsectors. This is perhaps most plausible for the information sector, where one can easily imagine shortages of workers with experience with particular technologies, or for the extremely heterogeneous professional and business services sector. Unfortunately, job openings data are not available for detailed industries. However, in Section VI I use wage data to explore the possibility of heterogeneity in labor market tightness within sectors.

V. Labor Supply Disincentives from Unemployment Insurance Extensions

We have seen little sign that shifts in the composition of labor demand across industrial sectors have been large enough to create meaningful mismatch. Another possible source of structural impediments to recovery is reductions in labor supply. Mulligan (2011b) points to expansions in safety net programs over the recession as a potential source of such reductions.

The most prominent such program is unemployment insurance (UI). Potential UI durations increased from 26 weeks before the recession to up to 99 weeks in 2010 and 2011. This far exceeds previous maximum durations in the United States, though is still

unremarkable in comparison to European UI systems (which also typically offer more generous weekly benefits than are available in the U.S.; see, e.g., Table 3.2 of Organization for Economic Co-operation and Development 2006). Extended UI durations reduce the incentive for recipients to find work quickly. With over 12 million UI recipients at the peak, these reduced incentives might have had quantitatively important impacts on the labor supply curve.

Estimates calibrated from past UI research suggest that reduced job search among UI recipients contributed about 1 percentage point to the unemployment rate (Mazumder 2011), and some commentators have argued for much larger effects (Grubb 2011, Barro 2010). However, these are at best extrapolations. At this point, four years after the recession began, we have enough data to estimate the effect of UI expansions directly. Using the uneven roll-out of the benefit extensions to generate plausibly exogenous variation in benefit durations, Rothstein (2012) finds that the total effect of UI extensions on the unemployment rate in early 2011 was only 0.3 percentage point, with more than half of this due to increased labor force participation among those who would in any case not be employed. Reductions in search effort due to the availability of extended benefits account for only 0.1—0.2 percentage point of the unemployment rate. This is far too small to create meaningful structural barriers to labor market recovery, particularly given the still-high ratio of job-seekers to job openings and the likely roll-back of the explicitly temporary UI extensions as the unemployment rate declines.²²

Mulligan (2011b) suggests that other work-discouraging benefits have increased over the recession as well, pointing in particular to implicit means-testing of mortgage modifications as a source of disincentives to work. However, there is no evidence regarding the effects of these incentives on labor supply. Rather, it appears that the reductions in homeowners' geographic mobility that many have hypothesized have been quite small (Farber 2012). Moreover, the roughly 11 million homeowners with negative

²² Although some commentators assume that any extension is permanent, this reflects neither the reality of past recessions, when extended benefits have always been allowed to expire, nor the political controversy that has accompanied each short-term extension of the Emergency Unemployment Compensation program during the current downturn.

equity represent only about 10% of households in the country.²³ Even large effects on these households' labor supply are unlikely to be quantitatively important to the macroeconomy.

VI. Evidence from Wages

I argued above that claims of an important structural component of unemployment imply that labor markets are tighter than they appear: Mismatches in the distribution of labor demand and labor supply across markets defined by skill or geography would produce tightness in at least some labor markets, while labor supply shifts would imply across-the-board tightness. This tightness should be directly observable in wages: If employers are facing shortages of suitable, interested workers, they should be responding by bidding up the wages of those workers who can be found. I examine the aggregate labor market first, then turn to distinctions across sub-markets.

VI.A. *Aggregate wages*

The solid line in Figure 11 graphs the 12-month change in real mean log hourly wages from 2005 through February 2012. These wages are calculated from the Current Population Survey Outgoing Rotation Groups, with imputed wages excluded. Details of the wage calculations are in the appendix. The figure shows that mean real wages have been largely stable since 2005, except for a period in late 2008 and 2009 when they rose at an annual rate of about 3%.²⁴ Since late 2009, real wages have been steady or falling.

{{Place Figure 11 about here}}

²³ See CoreLogic (2011). In a few states with much higher rates of negative equity – e.g., Nevada, where over half of homes with mortgages have loan-to-value ratios above one – it is plausible that mortgage modification-related incentives are reducing job search or slowing the out-migration process. But recall from Figure 7 that the labor market is extremely weak in nearly every state. In New York, where the CoreLogic data indicate that only 119,000 homeowners are underwater – less than 2% of all households – the unemployment rate is 8.0%.

²⁴ The price level was falling during much of this period; 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.

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 is experienced by individual workers. This could explain why wages appear to have risen quickly during 2008 and 2009, when employers were shedding workers and there is little other evidence of labor market tightness.

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 observations that cannot be matched or where the wage is unavailable in either month and holding constant the individual weight across the two periods.²⁵ The dashed line in Figure 11 shows the resulting year-on-year changes in composition-adjusted mean wages. This shows that individuals who remained employed saw rising real wages, at a rate of about 3% per year, throughout 2006-2009. The anomaly in 2008-2009 is much reduced here and plausibly consistent with sampling error. There is still no sign, however, that wage growth slowed before late 2009, when average changes fell to near zero. Weak growth reappeared for a period in mid 2010, then evaporated in late 2010 and early 2011, and has reappeared – still weak – in early 2012. With annual growth rates still below 2%, there is little evidence of meaningful tightness near the end of the sample.

A second way to adjust for composition changes is to reweight the data to offset changes in observables among those whose wages are observed. When I reestimate the “all workers” and “composition-adjusted” series using data reweighted to the 2007 all-worker age-education-race-gender distribution, results are quite similar to those in Figure 11. The most notable change is to reduce the 2008-09 anomaly in the “all workers” series, but the result of stable or falling wages since 2009 is unaltered.

Workers rarely accept – or perhaps employers rarely demand – reductions in their nominal wages within existing jobs. This wage rigidity may be masking trends in the wages offered to new hires. To zero in on the latter, I take advantage of the fact that the

²⁵ 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.

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 11 shows the trend in mean wages for such workers.²⁶ This series is somewhat more cyclical than the all-workers series, but shows a similar pattern of rising real wages in 2008-2009 and falling real wages starting in early 2010. This series thus shows even less sign of tightness in 2010 and 2011 than do the others. Moreover, it turns down sharply at the very end of the sample.

VI.B. Individual labor markets

Despite the aggregate slack in 2010 and 2011 that is evident in Figure 11, it is possible that particular labor markets were tighter. Table 1 shows the change in mean wages of newly-hired workers between 2007-08 and 2010-11, by education, gender, age, and industry. Across education-by-gender and age cells, only less educated women and over-65 workers saw nontrivial wage increases over this period. Across industries, nontrivial wage changes are seen in durable goods manufacturing, information, finance/insurance, and education, though these are at most marginally statistically significant.

{{Place Table 1 about here}}

Again, composition changes may confound changes in mean real wages. The strategy used above of limiting the sample to workers with two observed wages cannot be used when studying new hires. As an alternative, 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. Columns 3 and 4 of Table 1 show the change in the mean log wage residual in each cell and the standard error for these changes. Controlling for changes in observable characteristics substantially reduces the

²⁶ The “new jobs” series, like the others in Figure 11, is essentially unchanged when I reweight the data to a constant age-education-race-gender distribution over time.

apparent wage increases for less-educated women and for older workers, and eliminates most of the industry differences as well. Only the information sector appears to have meaningfully increased mean wages, adjusting for worker characteristics, since 2007-08, though even this change is insignificant.

One can also examine individual states, though samples are small and estimates are therefore noisy. There are only two states – Missouri and Oklahoma – in which mean residual real wages for new hires rose by more than 5% between 2007-08 and 2010-11. Changes smaller than this are generally not detectably different from zero, though there are ten additional states where point estimates indicate growth between 2% and 5%. This could indicate labor market tightness, though it could also be noise due to picking out the highest of 51 imprecise measures. One way to assess this is to focus on states where other indicators suggest labor market tightness. Recall from Figure 7 that the Dakotas, Nebraska, Alaska, Minnesota, Vermont, New Hampshire, and Iowa appeared to have relatively tight labor markets based on their 2011 unemployment rates compared with those seen in 2007. Of these, only in Nebraska did mean composition-adjusted wages rise (by 4.6%, with standard error 2.6%) between 2007-8 and 2010-11. The weighted average wage change across the eight states is -1.4% (SE 0.9%).

There is thus no sign of mean wages being bid up at the level of education groups, age, geography, or broad industrial sectors (with the possible exception of the information industries). It is possible, however, that these market definitions are too coarse and that employers in particular submarkets are having trouble finding workers with suitable skills. Falling wages in other submarkets might make it impossible to detect rising wages for workers in short supply via examinations of relatively aggregated averages.

The Current Population Survey sample is not large enough to permit detailed analysis of individual industries or occupations. However, it might be possible to see labor market tightness through rightward shifts in some portion of the overall wage distribution. To examine this, I examined the distribution of starting wages, adjusted for observable characteristics as in columns 3-4 of Table 1, in 2007-8 and 2010-11. The solid line in Figure 12 shows the change in wages at different percentiles in the new-hires wage distribution between the two periods. Not surprisingly, through most of the

distribution real starting wages fell by between 1% and 3% over this period. The upper tail – starting at around the 80th percentile – does seem to have shifted rightward somewhat, perhaps consistent with tightness in some individual labor markets. However, the changes are small – wage increases are below 3% everywhere and are below 2% at all but a couple of percentiles. Moreover, when I focus on the subset of industries where job openings increased the most between 2009 and 2011, in the dashed line in Figure 12, the wage increases are much more muted. There is no sign that wage growth in individual markets of any meaningful size is being masked by stagnation elsewhere.

{{Place Figure 12 about here}}

VII. The Long-Term Unemployment Share

A final source of evidence regarding structural impediments to recovery is the duration profile of unemployment. 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).

{{Place Figure 13 about here}}

Figure 13 shows the share of the unemployed who have been out of work for 27 weeks or more. The series increased dramatically in the later part of the Great Recession and the months immediately after its official end, starting from a level similar to that seen at the peak of past recessions and nearly doubling between December 2008 and December 2009. The long-term unemployment (LTU) share has shown no sign of decline during the labor market recovery, hovering near 45% for most of the last two years. However, Figure 13 also shows a long-run upward trend in the LTU share, as at each business cycle peak since the late 1960s the LTU share has been higher than at the previous cyclical peak.

One way to assess the potential contribution of structural factors to the recent rise in LTU is to decompose this rise into components attributable to other factors and a

residual that might reflect new structural challenges.²⁷ I focus on three factors that have been suggested as possible contributors to the rise in LTU. The first is demographic changes: Older workers have longer unemployment spells, so one might expect aging of the workforce to lead to higher LTU rates (Elsby et al., 2010). Second, the labor market has of course been weaker in the current recession than in most past recessions, and figure 13 shows that the LTU share is clearly countercyclical. Third, any additional long-run trend in the LTU share beyond the demographic changes identified above might plausibly have continued through the current recession. For example, Autor, Kerr, and Kugler (2007) argue that adjustment costs have risen over time. This would reduce employment flows, lengthen unemployment durations, and raise LTU shares.

I begin by computing LTU shares for each of 48 groups defined by the interaction of two gender cells, four education cells, and six age ranges.²⁸ Letting y_{gt} represent the share for group g in month t , the aggregate LTU share in month t can be written as

$$y_t = \sum_g y_{gt} w_{gt}, \quad (2)$$

where w_{gt} is the fraction of the unemployed in month t who belong to group g .

Using group-specific LTU shares from January 1990 through November 2007, I estimate a set of time series regressions of the form:

$$y_{gt} = \alpha_g + UR_t \beta_g + t \gamma_g + \varepsilon_{gt} \quad (3)$$

Here, UR_t is the overall unemployment rate in month t and $t \gamma_g$ represents a linear time trend that is allowed to differ for each demographic group. I use the coefficients from these regressions to compute residuals ε_{gt} in each month through 2011. Note that the residuals are by construction uncorrelated with t over the sample used for estimation of

²⁷ Aaronson, Mazumder, and Schechter (2010; see also Valletta and Kuang, 2012) carry out a similar exercise. Relative to their decomposition, my approach has two advantages: I allow for differences across demographic groups in the cyclical sensitivity of the LTU share as well as in the levels of that share, and I allow the lagged unemployment rate to have an effect independent of the current rate.

²⁸ The education groups are less than high school, high school diploma but no college, some college but no bachelor's degree, and a bachelor's degree or more. The age groups are 16-24, 25-34, 35-44, 45-54, 55-64, and 65 and over. Samples are small in some months for the over 65 age group; for this group, I collapse the education groups into two, one representing those with some college or more and the other those with a high school diploma or less. Thus, my actual analysis has only 44 groups.

(3), but that nothing constrains their values in the 2008-2011 period that is excluded from that estimation.

Equations (2) and (3) can be used to decompose the change in the LTU share between periods t' and t :

$$\begin{aligned} y_t - y_{t'} &= \sum_g y_{gt} w_{gt} - \sum_g y_{gt'} w_{gt'} \\ &= \sum_g y_{gt'} (w_{gt} - w_{gt'}) + \sum_g (y_{gt} - y_{gt'}) w_{gt} \\ &= \sum_g y_{gt'} (w_{gt} - w_{gt'}) + (UR_t - UR_{t'}) \sum_g \beta_g w_{gt} + (t - t') \sum_g \gamma_g w_{gt} + \sum_g (\varepsilon_{gt} - \varepsilon_{gt'}) w_{gt} \end{aligned} \quad (4)$$

The four terms in the final line of (4) represent, respectively, (i) the contribution of changes in the demographic composition of the unemployed to the LTU share, (ii) the contribution of changes in economic conditions, (iii) the impact of long-run time trends, and (iv) other changes not explained by the first three factors (including Unemployment Insurance extensions and other policies).

{{Place Table 2 about here}}

The first column of Table 2 shows the decomposition applied to the 26.8 percentage point growth of the LTU share between June 2007 and June 2011. Roughly half of this can be explained by the rise in the unemployment rate over this period, given the historical relationship between unemployment and LTU. Not surprisingly given the short four-year window under consideration, neither demographic shifts nor long-run trends have much explanatory power. Thus, 10.0 percentage points of the 2007-2011 runup is unexplained, potentially indicating a meaningful role for growing structural problems.

Columns 2 and 3 present the decomposition of the change over longer time periods, first from June 1992 and then from June 1983. 1992 was the peak year for the LTU share in the early-1990s recession, while 1983 represents the pre-2007 global peak. (I focus on June-June changes to avoid seasonality issues.) The 2011 LTU share was 21.3 percentage points higher than in 1992 and 17.5 percentage points higher than in 1983, each comparable to but a bit smaller than the 2007-2011 change. However, the unemployment rate was much higher in 1992 and 1983 than in 2007, reducing the potential explanatory power of this variable. Changing demographic composition has more of a role than in column 1, but still explains only 2.8 percentage points of the 1992-

2011 growth and 3.5 percentage points of the 1983-2011 growth, while the long-run time trend now explains 8.9 and 13.1 percentage points, respectively. The unexplained component of the increase in LTU is 5.6 percentage points for the 1992-2011 change and 3.4 percentage points for the 1983-2011 comparison, each much reduced from column 1. Nevertheless, each leaves room for a nontrivial component of the 2011 LTU share to be attributable to post-2007 structural changes.

An important limitation to the decompositions in columns 1 - 3 is that they force the contemporaneous unemployment rate to serve as a sufficient statistic for economic conditions affecting the LTU share. If unemployment exit hazards depend on current labor market conditions, the share of the unemployed who survive in that status to six months must depend on labor market conditions over the entire six month window. This suggests that the LTU share is likely to depend on conditions in the recent past conditional on current conditions.²⁹ This is potentially important given the extended duration of the current labor market weakness.

To address this, I augment equation (3) with controls for the average unemployment rate over the previous six and previous twelve months. Columns 4 – 6 of Table 2 present the resulting decompositions. This change increases the explanatory power of changes in economic conditions for the short- and medium-term increases, but does little to explain the long-run increase because unemployment was high for an extended period in 1982-83 as well. The contributions of demographic and time trends are largely unchanged.

The unexplained component of the rise in the long-term share over the 1983-2011 and 1992-2011 periods is reduced to nearly zero in columns 5 and 6, though 6.6 percentage points of the 2007-2011 increase remain after removing the components explained by demographics, economic conditions, or long-run time trends. This suggests a somewhat different gloss on the long-term unemployment share trend than was

²⁹ Indeed, one might even expect the current unemployment rate to have a negative partial effect on the LTU share as once past conditions are controlled the primary effect of current conditions may be on the denominator of the long-term share. This is what I find when I augment equation (3) with controls for past conditions as described in the text.

discussed above: Evidently, the unusual aspect of this series in the current cycle is that it was so low in 2007, not that it was so high in 2011.

Even the specification used in columns 4-6 is quite restrictive. Most importantly, it rules out any long-run trends in the sensitivity of the LTU share to overall economic conditions.³⁰ There is reason to expect such trends. For example, Katz (2010) argues that employers are less prone than in the past to institute temporary layoffs during cyclical downturns. Insofar as workers on temporary layoff are recalled before 27 weeks (Katz and Meyer, 1990) the declining importance of this institution might plausibly have raised the cyclical sensitivity of the LTU share.

To address this, I further augment specification (3) by interacting the unemployment rate controls with a linear time trend. Point estimates are noisy with only three business cycles in the sample, but the interaction coefficients are jointly statistically significant and indicate that the sensitivity of the LTU share to the unemployment rate grew by about 15% between 1983 and 2011. In other words, the specification indicates that the LTU share rises further in recent recessions (and falls further in recent recoveries) than in earlier cycles. This increased sensitivity helps to explain the positive residuals seen for the simpler regression in 2009-2011. Importantly, as with the earlier analyses the structural change in the cyclical sensitivity of the LTU share is identified from pre-2007 data, so is not confounded by structural problems that are new in this recession.

Decompositions based on the interacted specification are presented in columns 7-9 of Table 2.³¹ The inclusion of interactions clearly raises the combined role of the two factors in explaining the 2007-2011 run-up in the LTU share. Indeed, this specification indicates that the LTU share should have been expected to increase by 1.9 percentage

³⁰ The specification also imposes a linear relationship between the unemployment rate and the LTU share. I have explored specifications that loosen this restriction, with little effect on the results.

³¹ With the interaction, the decomposition of the components due to changes in the unemployment rate versus the time trend is not unique (though the total explanatory power of the two variables is). I compute the pure UR contribution using the estimated cyclical sensitivity for June 2011, and the pure trend contribution using the June 2011 values for the unemployment rate and its lags.

points *more* than it actually did between 2007 and 2011, with even larger negative residuals in the longer-run comparisons.

Given the limited variation available for identification of the interacted specification, it probably should not count as strong support for the view that *all* of the increase in long-term unemployment in the current recession is due to the combination of exceptionally weak labor demand and long-run trends that predate the current recession. However, it is clear that the vast majority of the increase can be so attributed. Recent policy changes and new structural impediments to adjustment (as distinct from those operating over the longer run) can be blamed for no more than a small share of the recent rise in long-term unemployment.

VIII. Discussion

The U.S. labor market is extremely weak by any historical standard: The unemployment rate has been above 8 percent for nearly three straight years, the employment-population ratio has fallen by over 4.2 percentage points since 2008, and many subgroups – particularly the young and less educated, along with racial minority groups – are facing unemployment rates well into the double digits.

Many models that macroeconomists 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 – either the unemployed are not looking very hard for work or they are in some sense unsuitable for the jobs that are available, either because they lack the appropriate skills or because they 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 our current situation. The review of the evidence here offers no support for this diagnosis, however. 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.³² And the Beveridge Curve provides at best weakly suggestive evidence regarding the state of the matching function.

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 is actually much tighter than it appears, at least as viewed from the perspective of potential employers. There is no sign in the data that employers with jobs to fill are having trouble filling them, except perhaps in a few isolated and small submarkets such as resource extraction.

Finally, the unprecedented rise in long-term unemployment, which some have pointed to in support of the structural unemployment hypothesis, turns out not to support that hypothesis after all. The extended period of labor market weakness that we have seen, combined with long-run demographic and labor market trends that predate the current recession, explains all or nearly all of the rise in the long-term unemployment share relative to past downturns, leaving no need to appeal to recent structural factors for an explanation. There is an unexplained long-run upward trend in the long-term unemployment share, which has risen by about 0.5 percentage point per year over the last two decades and also appears to have become more sensitive to economic conditions, but I find no indication that either trend has accelerated recently.

We can thus conclude that labor demand shortfalls continue to be an important feature of the labor market and the primary determinant of labor market performance, four years after the Great Recession began.

Three caveats are in order. First, I have not addressed longer-run structural changes, such as deindustrialization or skill-biased technical change. 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

³² As discussed above, unemployment insurance extensions can explain only about 0.3 percentage points of the 2011 unemployment rate. With regard to geographic mobility, careful analyses indicate that mobility rates have changed little in recent years (Kaplan and Schulhofer-Wohl 2011), that any declines are concentrated among renters who should not have been directly affected by the decline in home values (Farber 2012), and that any “house lock” effect is quantitatively small (Schmitt and Warner 2011).

2007 state, but not to whether further increases could reverse longer-run trends toward reduced male employment-population ratios and higher inequality.

Second, although I have found no sign to date of labor market tightness, it is possible that structural problems that are now being masked by low aggregate demand would become apparent in a strong economic recovery. This bears watching as the recovery proceeds. There will likely be room for policies aimed at improving job matching – such as job training and search and mobility assistance – though these should be seen as complements rather than substitutes for policies aimed at stimulating demand.

Finally, and most important: An extremely long downturn is likely to itself create structural problems that will impede growth going forward. Those who go months or years without being able to find work are likely to become less employable and less productive as their skills deteriorate and/or become obsolete. 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 every month that unemployment remains high is making us poorer for decades to come.

Unfortunately, there is little hope of avoiding these consequences. In every month of the last 3 years, the unemployment rate has been higher than at any point since 1984. And we cannot reasonably hope for the labor market to recover quickly. Even if employment growth in 2012 and thereafter matches the pace seen in 1994 – the period of fastest sustained growth in recent history, when employment grew by an average of 321,000 jobs per month – it will still be years before we reach anything that might be characterized as full employment (Greenstone and Looney 2012). And at a more moderate growth rate of 208,000 jobs per month – matching the best year to date of the current century – recovery will take a decade or more. Thus, while aggressive policies

aimed at increasing aggregate demand quickly might help to limit the damage, even optimistic scenarios imply large ongoing costs.

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 December 2011.

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).

Constructed wages are topcoded at \$2,884 per week, and topcoded wages are multiplied by 1.4. Wages are then adjusted for inflation using the monthly CPI-U series, and trimmed 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


- Aaronson, Daniel, Bhashkar Mazumder, and Shani Schechter. 2010. What is behind the rise in long-term unemployment? *Economic Perspectives*, 2Q/2010: 28-51. Federal Reserve Bank of Chicago. Accessed at http://www.chicagofed.org/digital_assets/publications/economic_perspectives/2010/2qtr2010_part1_aaronson_mazumder_schechter.pdf.
- 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.
- Autor, David, and David Dorn. 2011. The growth of low-skill service jobs and the polarization of the U.S. labor market. MIT Working Paper. June.
- Autor, David H., William R. Kerr, and Adriana D. Kugler. 2007. Does employment protection reduce productivity? Evidence from U.S. states. *The Economic Journal*, 117(June): F189-F217.

- 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.
- 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. 2007. The budget and economic outlook: Fiscal years 2008 to 2017. January. Accessed at <http://www.cbo.gov/sites/default/files/cbofiles/ftpdocs/77xx/doc7731/01-24-budgetoutlook.pdf>.
- Congressional Budget Office. 2012. Understanding and responding to persistently high unemployment." February. Accessed at <http://www.cbo.gov/publication/42989>.
- CoreLogic (2011). CoreLogic third quarter 2011 negative equity data shows slight decline but remains elevated. Press release. November 29.
- 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.
- Elsby, Michael W., Bart Hobijn, and Aysegul Sahin. 2010. The labor market in the Great Recession. *Brookings Papers on Economic Activity*, 1: 1-65.
- 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).
- Greenstone, Michael, and Adam Looney. 2012. Shrinking job opportunities: The challenge of putting Americans back to work." Blog post, The Hamilton Project, January 6. Accessed at http://www.brookings.edu/opinions/2012/0106_jobs_greenstone_looney.aspx.
- Grubb, David. 2011. Assessing the impact of recent unemployment insurance extensions in the United States. OECD working paper, June 14. Accessed at http://www.iza.org/conference_files/UnIRE2011/grubb_d6938.pdf.

- 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. 2011. Interstate migration has fallen less than you think: Consequences of hot deck imputation in the Current Population Survey. Federal Reserve Bank of Minneapolis Staff Report 458, June.
- Katz, Lawrence F. 2010. Long-term unemployment in the Great Recession. Testimony for the Joint Economic Committee, U.S. Congress, April 29.
- Katz, Lawrence F., and Bruce D. Meyer. 1990. The impact of the potential duration of unemployment benefits on the duration of unemployment. *Journal of Public Economics*, 41(1): 45-72.
- Kocherlakota, Narayana. 2010. Inside the FOMC. Speech in Marquette, Michigan, August 17.
- Mazumder, Bhashkar. 2011. How did unemployment insurance extensions affect the unemployment rate in 2008-10? Federal Reserve Bank of Chicago Fed Letter 285, April. Accessed at http://www.chicagofed.org/webpages/publications/chicago_fed_letter/2011/april_285.cfm.
- 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.
- Mulligan, Casey. 2009. What caused the recession of 2008? Hints from labor productivity. NBER Working Paper no. 14729. Cambridge, MA: National Bureau of Economic Research.
- Mulligan, Casey. 2011a. Rising labor productivity during the 2008-9 recession. NBER Working Paper no. 17584. Cambridge, MA: National Bureau of Economic Research.
- Mulligan, Casey. 2011b. The expanding social safety net. NBER Working Paper no. 17654. Cambridge, MA: National Bureau of Economic Research.
- National Bureau of Economic Research. 2010. The NBER's Business Cycle Dating Committee. Accessed at <http://www.nber.org/cycles/recessions.html>.
- 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.

- Organization for Economic Co-operation and Development. 2006. *OECD Employment Outlook – 2006 Edition*. Paris: Organization for Economic Co-operation and Development.
- Reich, Robert. 1992. *The Work of Nations*. New York: Vintage Books.
- Romer, Christina. 2011. The continuing unemployment crisis: Causes, cures, and questions for further study.” Speech presented at Livable Lives Initiative Symposium, Washington University, April 12.
- Rothstein, Jesse. 2012. Unemployment insurance and job search in the Great Recession. *Brookings Papers on Economic Activity*, 1.
- 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.
- 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.
- Tasci, Murat, and John Lindner. 2010. Has the Beveridge Curve shifted? *Economic Trends*. Cleveland: Federal Reserve Bank of Cleveland. August.
- Valletta, Rob, and Katherine Kuang. 2012. Why is unemployment duration so long? *FRBSF Economic Letters* no. 2012-03.
- 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.

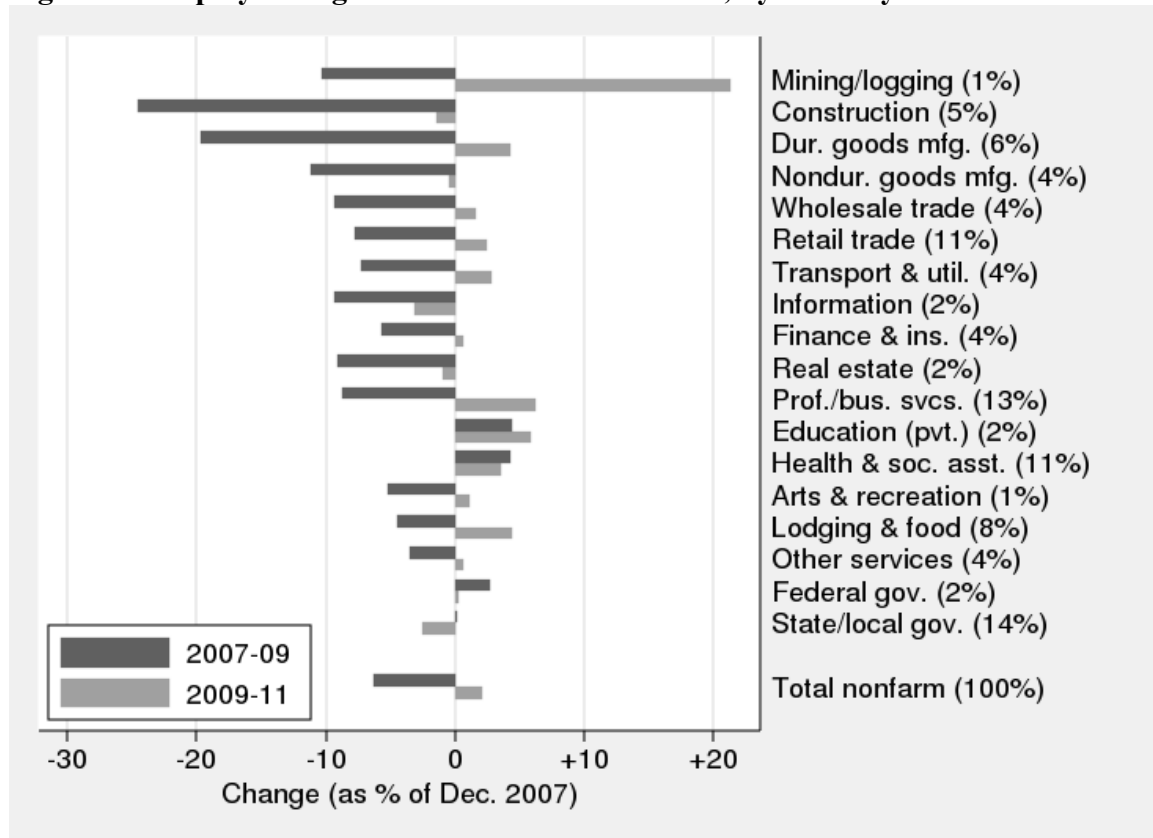
Figure 1. Employment, employment-population ratio, and unemployment rate, 2004 – 2011



Source: Bureau of Labor Statistics, Current Population Survey and Current Employment Statistics.

Notes: All figures are seasonally adjusted.


Figure 2. Employment growth since December 2007, by industry.



Source: Bureau of Labor Statistics, Current Employment Statistics.

Notes: December-December changes, based on seasonally adjusted employment counts, expressed as share of December 2007 employment level. Parentheses contain industry shares of December 2007 civilian non-farm employment.

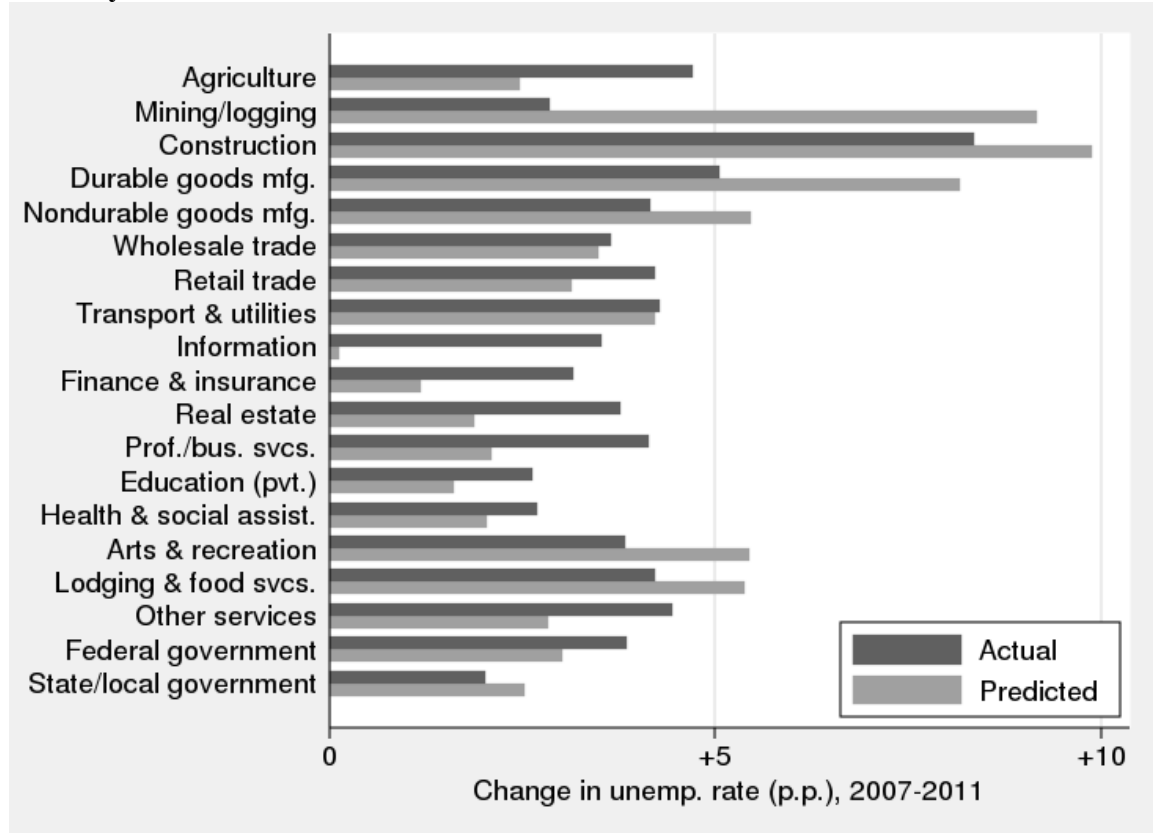
Figure 3. Unemployment rates in 2007 and 2011, by sex and education



Source: Bureau of Labor Statistics, Current Population Survey.

Notes: Equally-weighted averages of non-seasonally-adjusted monthly data.


Figure 4. Actual and predicted change in unemployment rate, 2007 to 2011, by industry



Source: Bureau of Labor Statistics Current Population Survey.

Notes: Change between the 2007 and the 2011 averages of non-seasonally-adjusted monthly unemployment rates. Predicted change is computed from the fitted values of a regression of the monthly unemployment rate in the industry on calendar month dummies and the unemployment rate across the rest of the labor force, using data from 1978-2007.


Figure 5. Actual and predicted change in unemployment rate, 2007 to 2011, by sex and education



Source: Bureau of Labor Statistics Current Population Survey.

Notes: Change between the 2007 and the 2011 averages of non-seasonally-adjusted monthly unemployment rates. Predicted change is computed from the fitted values of a regression of the monthly unemployment rate in the sex-education group on calendar month dummies and the unemployment rate across the rest of the labor force, using data from 1978-2007.

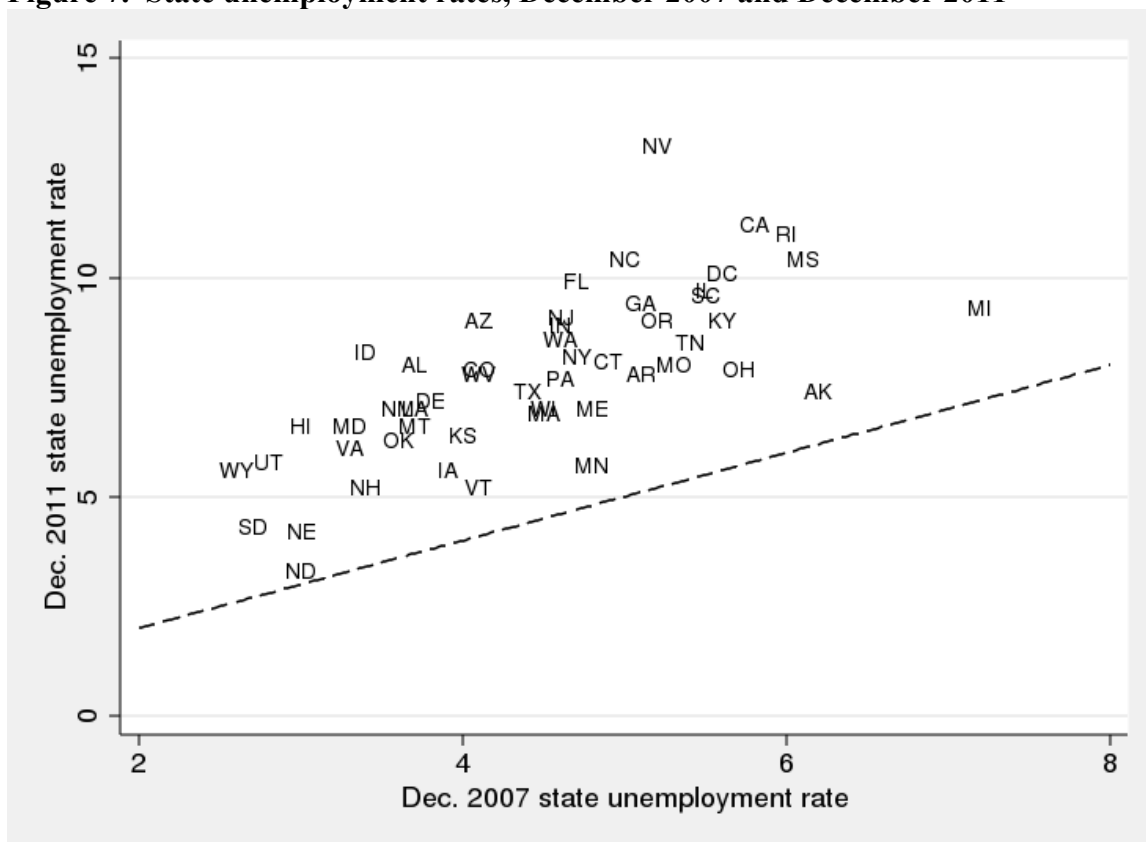
Figure 6. Actual and predicted change in unemployment rate, 2007 to 2011, by age



Source: Bureau of Labor Statistics Current Population Survey.

Notes: Change between the 2007 and the 2011 averages of non-seasonally-adjusted monthly unemployment rates. Predicted change is computed from the fitted values of a regression of the monthly unemployment rate in the age group on calendar month dummies and the unemployment rate across the rest of the labor force, using data from 1978-2007.

Figure 7. State unemployment rates, December 2007 and December 2011



Source: Bureau of Labor Statistics, Local Area Unemployment Statistics.

Notes: Rates are seasonally adjusted.

Figure 8. Structural unemployment due to sectoral mismatch


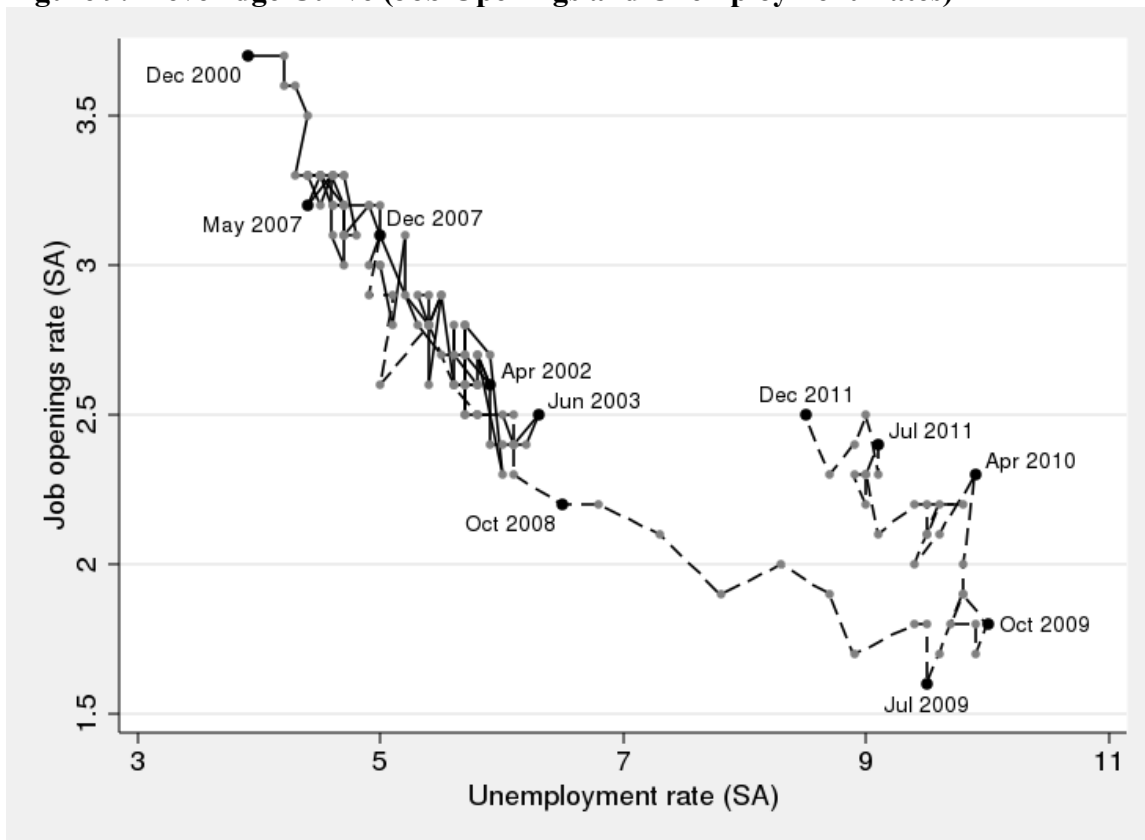


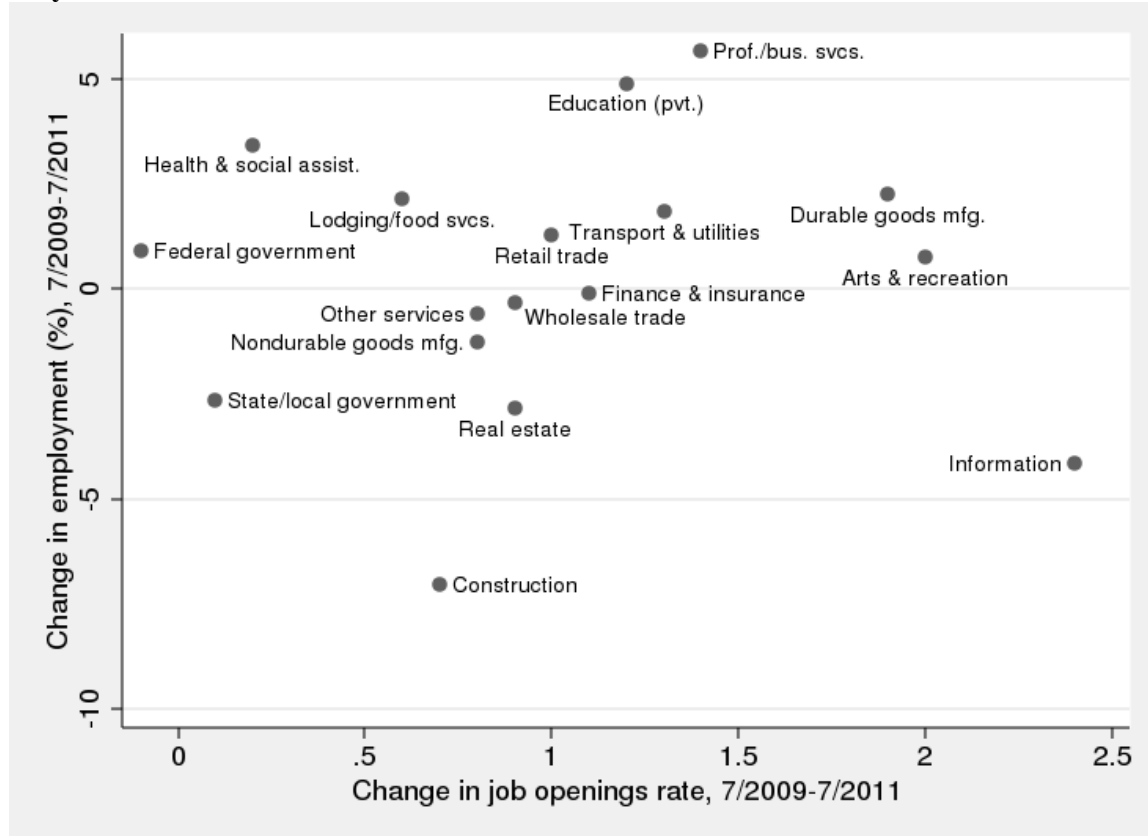
Figure 9. Beveridge Curve (Job Openings and Unemployment Rates)



Source: Bureau of Labor Statistics, Job Openings and Labor Turnover Survey and Current Population Survey.


Notes: Rates are seasonally adjusted. December 2007-December 2011 portion is indicated by a dashed line.

Figure 10. Job openings changes vs. employment changes by industry, July 2009-July 2011



Notes: Mining industry (change in job openings rate +3.5, change in employment +17.4%) not shown.

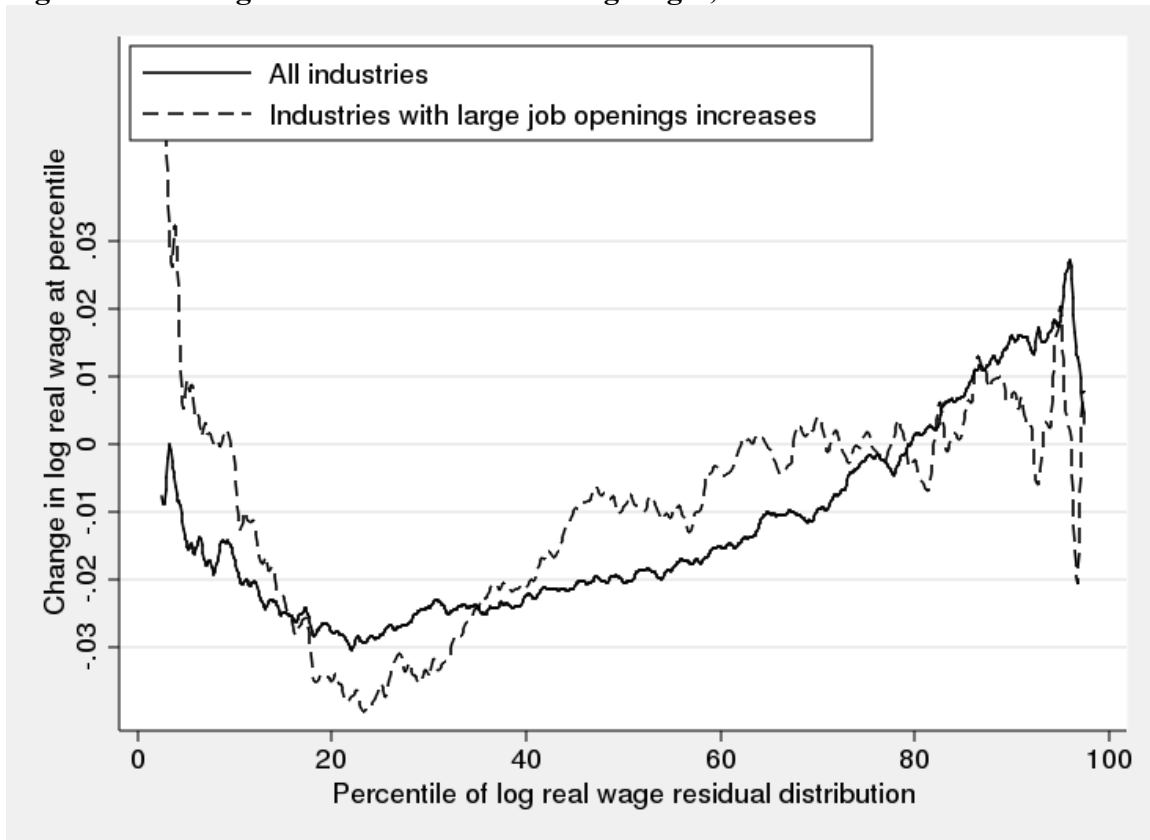
Figure 11. 12-month changes in mean wages, various subsamples



Source: Author's analysis of Current Population Survey data.


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

Figure 12. Change in distribution of starting wages, 2007-8 to 2010-11



Notes: Starting wages are those on jobs that started within the previous three months. Wage distributions are weighted by weekly hours. Percentile changes are computed for each 0.1 percentage points, then smoothed across three adjacent points using a triangle smoother. Changes are not shown below the 2.5th or above the 97.5th percentile.

Figure 13. Long-term unemployment share



Source: Bureau of Labor Statistics, Current Population Survey.

Notes: Series is seasonally adjusted.

Table 1. Change in mean real wages of new hires, 2007/8 - 2010/11

	Unadjusted		Adjusted for observables	
	Mean	SE	Mean	SE
	(1)	(2)	(3)	(4)
Overall	0.9%	(0.5%)	-1.2%	(0.4%)
By education and gender				
Male, less than HS	-0.2%	(1.3%)	-2.7%	(1.2%)
Male, HS diploma	-3.3%	(1.0%)	-3.0%	(0.9%)
Male, some college	-2.8%	(1.1%)	-2.6%	(1.0%)
Male, BA+	-1.9%	(1.2%)	-1.9%	(1.1%)
Female, less than HS	2.7%	(1.2%)	1.1%	(1.2%)
Female, HS diploma	1.4%	(1.0%)	1.3%	(0.9%)
Female, some college	-0.9%	(0.9%)	-0.5%	(0.8%)
Female, BA+	-0.9%	(1.0%)	-0.4%	(1.0%)
By age				
16-24	-1.8%	(0.6%)	-3.7%	(0.6%)
25-34	0.1%	(0.9%)	-1.6%	(0.7%)
35-44	0.5%	(1.1%)	-0.5%	(0.9%)
45-54	-1.3%	(1.1%)	-0.7%	(0.9%)
55-64	0.0%	(1.3%)	-0.5%	(1.1%)
65+	8.9%	(2.6%)	3.3%	(2.3%)
By industry				
Agriculture	-3.8%	(3.1%)	-5.7%	(2.9%)
Mining & logging	1.3%	(4.6%)	1.3%	(4.4%)
Construction	0.0%	(1.7%)	-4.0%	(1.5%)
Manufacturing - durable	3.5%	(1.9%)	-0.3%	(1.4%)
Manufacturing - nondurable	2.2%	(2.4%)	-1.7%	(1.8%)
Wholesale trade	-0.1%	(2.9%)	-3.8%	(2.4%)
Retail trade	-1.7%	(1.1%)	-3.9%	(0.9%)
Transportation & utilities	-1.8%	(2.1%)	-2.2%	(1.9%)
Information	5.4%	(3.2%)	3.2%	(2.6%)
Finance & insurance	4.1%	(2.2%)	0.3%	(1.8%)
Real Estate	-0.8%	(3.6%)	-3.8%	(3.2%)
Professional & business svcs.	-1.8%	(1.5%)	-2.1%	(1.2%)
Education	3.7%	(2.5%)	1.2%	(2.1%)
Health	1.5%	(1.4%)	0.1%	(1.1%)
Arts & entertainment	2.8%	(2.9%)	-2.0%	(2.4%)
Accommodation & food svcs.	1.8%	(1.2%)	0.0%	(1.1%)
Other services	-0.5%	(2.1%)	-0.4%	(1.9%)
Federal government	-0.8%	(2.7%)	-1.6%	(2.2%)
State & local government	1.4%	(1.2%)	0.1%	(1.0%)

Notes: New hires are those who began their jobs within the previous three months. Adjusted estimates in columns 3-4 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 4 do not account for sampling error in the regression coefficients.

Table 2. Decomposition of June-June long-term unemployment share changes

	2007- 2011	1992- 2011	1983- 2011	2007- 2011	1992- 2011	1983- 2011	2007- 2011	1992- 2011	1983- 2011
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Long-term unemployment share									
Year 1	15.2%	20.7%	24.5%	15.2%	20.7%	24.5%	15.2%	20.7%	24.5%
Year 2	42.0%	42.0%	42.0%	42.0%	42.0%	42.0%	42.0%	42.0%	42.0%
Change (p.p.)	+26.8	+21.3	+17.5	+26.8	+21.3	+17.5	+26.8	+21.3	+17.5
Decomposition of change (p.p.)									
Changing weights	+ 1.2	+ 2.8	+ 3.5	+ 1.2	+ 2.8	+ 3.5	+ 1.2	+ 2.8	+ 3.5
Change in current UR	+13.7	+ 4.0	- 2.5						
Change in current & lagged URs				+17.0	+ 7.9	- 4.1	+25.8	+12.8	- 5.5
Time trend	+ 1.9	+ 8.9	+13.1	+ 2.0	+ 9.7	+14.3	+ 3.2	+15.0	+22.1
UR * time trend interaction							- 1.4	- 3.8	+ 2.0
Unexplained	+10.0	+ 5.6	+ 3.4	+ 6.6	+ 1.0	+ 3.8	- 1.9	- 5.5	- 4.7

Notes: Long-term unemployment is defined as having been unemployed 27 weeks or more. See text for discussion of decomposition. In columns 7-9, the specification includes interactions between the time trend and the UR variables. The UR contribution is computed holding t at June 2011, and the t contribution holds UR and its lags at their June 2011 levels. The row labeled “interaction” reports the joint contribution of simultaneous changes in both, less effects of the two variables considered separately.