

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

**The Recent Evolution of
the Natural Rate of Unemployment**

Mary Daly
Federal Reserve Bank of San Francisco

Bart Hobijn
Federal Reserve Bank of San Francisco

Rob Valletta
Federal Reserve Bank of San Francisco

January 2011

Working Paper 2011-05

<http://www.frbsf.org/publications/economics/papers/2011/wp11-05bk.pdf>

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

The Recent Evolution of the Natural Rate of Unemployment

MARY DALY,* BART HOBIJN, AND ROB VALLETTA
Federal Reserve Bank of San Francisco
101 Market Street
San Francisco, CA 94105

January 17, 2011

ABSTRACT

The U.S. economy is recovering from the financial crisis and ensuing deep recession, but the unemployment rate has remained stubbornly high. Some have argued that the persistent elevation of unemployment relative to historical norms reflects the fact that the shocks that hit the economy were especially disruptive to labor markets and likely to have long lasting effects. If such structural factors are at work they would result in a higher underlying natural or non-accelerating inflation rate of unemployment, implying that conventional monetary and fiscal policy should not be used in an attempt to return unemployment to its pre-recession levels.

We investigate the hypothesis that the natural rate of unemployment has increased since the recession began, and if so, whether the underlying causes are transitory or persistent. We begin by reviewing a standard search and matching model of unemployment, which shows that two curves—the Beveridge curve (BC) and the Job Creation curve (JCC)—determine equilibrium unemployment. Using this framework, our joint theoretical and empirical exercise suggests that the natural rate of unemployment has in fact risen over the past several years, by an amount ranging from 0.6 to 1.9 percentage points. This increase implies a current natural rate in the range of 5.6 to 6.9 percent, with our preferred estimate at 6.25 percent. After examining evidence regarding the effects of labor market mismatch, extended unemployment benefits, and productivity growth, we conclude that only a small fraction of the recent increase in the natural rate is likely to persist beyond a five-year forecast horizon.

Keywords: equilibrium unemployment, Beveridge curve, structural unemployment, mismatch
JEL codes: E24, J3, J6.

* Corresponding author (email: mary.daly@sf.frb.org). The authors are grateful to Glenn Rudebusch and John Williams for their suggestions and comments. The views expressed in this paper are solely those of the authors and are not attributable to the Federal Reserve Bank of San Francisco or the Federal Reserve System.

1. Introduction

Over the past three years, the financial crisis and ensuing deep recession led to a sharp increase in the U.S. unemployment rate, which rose by an amount that is unprecedented during the postwar era. The overall unemployment rate peaked at 10.1 percent in October 2009; after adjusting for changes in the demographic composition of the labor force, this represents the highest unemployment rate reached since the Great Depression.¹ Moreover, in contrast to relatively rapid labor market recoveries following prior deep postwar recessions, the unemployment rate has edged down very slowly during the last 18 recovery months, to only 9.4 percent. Most economic forecasters expect very slow absorption of this high level of unemployment.²

Part of the continuing anemic labor market undoubtedly reflects the sluggish overall economic recovery, which is a common occurrence following financial crises (Reinhart and Rogoff, 2009). The resulting lackluster pace of job creation has barely kept up with trend labor force growth and therefore has not generated enough jobs to make a significant dent in the unemployment rate. Moreover, the unemployment rate has remained high relative to its historical relationship with other cyclical indicators. For example, the unemployment rate exceeds the level implied by Okun's Law, which relates it to output, and the level implied by the Beveridge curve, which relates it to the strength of labor demand as reflected in the number of vacancies posted.

This apparent disconnect between the unemployment rate and other aggregate indicators has raised the concern that rather than being purely cyclical, the recent elevated level of unemployment contains a significant structural component. Higher structural unemployment alters the role of short-run stabilization policies, including monetary policy, by increasing the possibility that expansionary policies will trigger inflation at higher rates of unemployment than otherwise. Economists have pointed to various reasons why the unemployment rate might substantially deviate from its normal cyclical pattern. Our analysis focuses on the three most

¹ Due to changes in the demographic composition of the labor force, the aggregate unemployment rate in the recent recession did not reach the peak from the 1981 recession (see Elsby, Hobijn, and Şahin, 2010)

² The participants of the 2010 Q4 Survey of Professional Forecasters predict an average unemployment rate of 9.3 percent in 2011, 8.7 percent in 2012, and 7.9 percent in 2013.

commonly discussed reasons: (1) a *mismatch* between the characteristics of job openings, such as skill requirements or location, and the characteristics of the unemployed; (2) the availability of *extended unemployment insurance* (UI) benefits, which may reduce the intensity of job search for UI recipients; and (3) *faster productivity growth*, which enables firms to raise output without extensive hiring of new employees.³ To the extent that these factors are persistent, they will raise the natural rate of unemployment rather than simply distort short-run cyclical movements in the unemployment rate.

In this paper, we assess the degree to which the natural rate of unemployment has changed and the reasons underlying such a change. We do so in two parts.

In the first part, we revisit Pissarides' (2000, Chapter 1) textbook model of frictional unemployment and use it as a framework to interpret recent patterns in the U.S. labor market that relate to the natural rate of unemployment. In this model, equilibrium unemployment is determined by the intersection of two curves: the Beveridge curve, which reflects the steady-state relationship between vacancies and the unemployment rate; and the Job Creation curve, which reflects employers' decisions to create job openings and can be loosely interpreted as an aggregate labor demand curve. Following Barnichon et al. (2010), we provide estimates of the empirical Beveridge curve and show that its recent movements are similar to those that occurred in previous severe recessions. We then introduce an estimate of the historical Job Creation curve based on the vacancy rate series by Barnichon (2010) and the historical estimates of the natural rate published by the Congressional Budget Office (CBO, 2010). This estimated Job Creation curve shows that the natural rate of unemployment tends to rise along with the average vacancy rate, which implies in turn that the recent shift in the empirical Beveridge curve exceeds the underlying increase in the natural rate.

In the second part of our analysis, we focus on the three primary factors noted above that may explain the increase in the natural rate: *labor market mismatch*, *extended UI*, and *productivity growth*.⁴ Because the mismatch argument has multiple components, we present an array of evidence, which on balance suggests that the increase in mismatch has been quite

³ Basu, Fernald, Kimball (2006) find that positive productivity shocks can be contractionary in terms of employment.

⁴ This evidence draws upon our previous work on the labor market during the recession and recovery: Daly and Hobijn (2010), Elsby, Hobijn, and Şahin (2010), Kwok, Daly, and Hobijn (2010), Valletta and Kuang (2010a,b), as well as Wilson (2010).

limited. Our analyses suggest a larger contribution arising from extended UI benefits, which are expected to lapse as the labor market recovers and therefore constitute a transitory influence on the natural rate. Unusually rapid productivity growth also may be contributing to elevated structural unemployment through the resulting suppression of hiring, although our evidence in this regard is largely speculative; like extended UI, the influence of this factor is likely to be transitory as well.

On balance, our analysis suggests that the natural rate of unemployment has risen by about 1.25 percentage points to 6.25 percent over the last few years. This higher natural rate nevertheless implies substantial labor market slack at the current unemployment rate and continuing over the next several years. Moreover, about half or more of this increase appears to be due to extended UI benefits, which are expected to lapse and therefore exert only a transitory influence on the natural rate. Thus, only a small portion of the elevated unemployment rate, about 0.5 percentage points, seems attributable to factors that are likely to raise the unemployment rate over a longer horizon of 5 years and beyond.

2. Equilibrium unemployment and the natural rate

As depicted in Figure 1, the recent increase in the quarterly U.S. unemployment rate is unprecedented during the postwar era. Between the first half of 2007 and the end of 2009, the unemployment rate rose 5.5 percentage points, from 4.5 percent to 10.0 percent; this exceeds the net increase of 4.8 percentage points during the two recessionary episodes between late 1979 and late 1982. Moreover, since peaking in late 2009, the unemployment rate has hardly budged, remaining above 9.5 percent in each of the subsequent four quarters.

Figure 1 also plots the estimate of the natural rate of unemployment from the CBO (2010). As defined in Brauer (2007), the natural rate “arises from all sources other than fluctuations in demand associated with business cycles. The natural rate is determined by the rate at which jobs are simultaneously created and destroyed, the rate of turnover in particular jobs, and how quickly unemployed workers are matched with vacant positions. Those factors in turn depend on the characteristics of jobs and of workers and on the efficiency of the labor market’s matching process.” The CBO estimate of the natural rate of unemployment has remained at 5.0 percent since the year 2000. If correct, it implies that the recession has generated a record high

unemployment gap, with the actual unemployment rate exceeding the natural rate by 5 percentage points in late 2009.⁵ Historically, shortfalls in labor demand of this magnitude were followed by relatively sharp declines in the unemployment rate—for example, following the 1981-82 recession. By contrast, in the recent downturn, the implied unemployment gap has exceeded 4.5 percentage points for 6 quarters. This pattern of elevated and persistent unemployment raises the possibility that in addition to the usual cyclical increase, an elevated structural component and corresponding higher natural rate may be contributing as well.

2.1 Frictional unemployment in equilibrium

To assess the factors affecting the unemployment rate in the short-term as well as its longer term level, we rely on the model of equilibrium frictional unemployment model from Pissarides (2000, Chapter 1). This model specifies two curves that determine equilibrium frictional unemployment: the Beveridge Curve (BC) and the Job Creation curve (JCC). We use this framework to analyze the potential increase in the natural rate of unemployment.

In models of unemployment due to search frictions, like the type we consider here, not every employer that is looking to hire a worker finds an employee and not every job searcher finds an employer. Therefore, the labor market does not fully clear in each period, and some job openings remain unfilled at the same time that some unemployed persons are unable to find a job.⁶ The rate at which job seekers are matched to job openings depends on the ratio of the vacancy rate, v , to the unemployment rate, u .

Because both the employer and employee benefit from a match, wages in this case are determined by the bargain between employers and employees over the surplus generated by the match, which occurs after the job match.⁷ As such, the equilibrium in this model is defined in terms of unemployment and vacancies—the intersection of the BC and JCC—rather than wages and the equilibrium level of employment. The BC is best understood by realizing that the change in the unemployment rate is determined by the difference between the number of people who

⁵ We use the CBO's natural rate measure as our benchmark throughout this paper but provide a discussion of relevant alternative measures below.

⁶ By contrast, in frictionless models of the labor market, wages adjust to equate labor demand to labor supply in a spot market, which excludes the existence of unemployment as an equilibrium outcome.

⁷ Assumptions about the type of wage bargaining are important for the cyclical properties of the model (Pissarides, 2009) but are not important for the equilibrium concept we focus on here.

flow into unemployment and the number who leave it. These worker flows are formalized in our expositional model.

For simplicity, we abstract from participation decisions and assume that the labor force is of size 1. The unemployment rate, u , is the fraction of the labor force that is unemployed. The rest of the labor force, $(1 - u)$, is employed.

At each instant, a fraction s of jobs is destroyed. We take the separation rate, s , as exogenously given. In the context of our simple framework it can be interpreted as the layoffs rate. Thus $s(1 - u)$ is the number of people flowing into unemployment. At the same time a fraction f of the unemployed get hired and, thus, fu persons flow out of unemployment. The change in the unemployment rate at each point in time is the difference between these flows:

$$\Delta u = s(1 - u) - fu. \quad (1)$$

Shimer (2005), among others, has argued that the U.S. aggregate labor market adjusts sufficiently fast that its dynamics are well approximated by its flow steady state:

$$\Delta u = 0, \text{ such that } u \approx s/(s + f). \quad (2)$$

This the approximation we use throughout the rest of our exposition.

Frictional unemployment in this model exists because not all unemployed find a job at each instant ($f < 1$). The rate of job matching f is assumed to depend on the ratio of the number of vacancies per unemployed, v/u . The relationship between f and v/u is known as the matching function. We use a Cobb-Douglas specification for the matching function, which has been found to fit the data well in various empirical settings (Petrongolo and Pissarides 2001):

$$f = \mu \left(\frac{v}{u} \right)^\alpha. \quad (3)$$

Here μ can be interpreted as the efficiency of the matching technology (e.g., the degree of mismatch in the labor market). Combining (2) and (3) yields an equilibrium relationship between the unemployment rate, u , and the vacancy rate, v :

$$u = s / \left(s + \mu \left(\frac{v}{u} \right)^\alpha \right). \quad (4)$$

This is the BC. For a given degree of labor market frictions, α and μ , and layoffs rate s , it defines the (flow-steady-state) unemployment rate as a function of the vacancy rate v . Note that if s increases or μ decreases then this curve implies a higher unemployment rate for a given

number of vacancies, i.e. the Beveridge curve shifts outward when the pace of layoffs rises or match efficiency declines.

The BC is one equation in two unknowns, u and v , and thus does not fully pin down the equilibrium unemployment rate. The missing equation is the JCC, which in its simplest form is determined by a free entry condition: if anyone can post a vacancy then, in equilibrium, the expected value of a job after a vacancy is filled, which we denote by J , should equal the expected search cost to fill the vacancy. The search cost is the cost of posting a vacancy, c , times the probability it is filled. This implies that the vacancy creation condition is

$$J = \frac{c}{\mu} \left(\frac{v}{u}\right)^{1-\alpha}, \text{ such that } \frac{v}{u} = \left(\frac{\mu J}{c}\right)^{\frac{1}{1-\alpha}}. \quad (5)$$

Equation (5) yields a closed-form solution to the vacancy-creation part of the model, or JCC. It implies that a decline in match efficiency μ (an increase in mismatch) affects vacancy creation, as do changes in the cost of posting a vacancy, c , and the value of jobs, J . More generally, it depends on the structure of the market in which firms operate and how they bargain over wages, as well as the separation rate, s , and other factors not explicitly incorporated here, such as the discount or interest rate.

The relevant literature on frictional unemployment has not attempted to pin down the quantitative determinants of the JCC, focusing instead on the shape of the BC and the types of frictions that underlie it. However, most models, including the simple model we discuss here as well as models with matching functions based on more elaborate micro-foundations (e.g., Shimer 2007), imply that the JCC is upward sloping. In our specific setting, the upward slope arises because as the pool of unemployed grows, employers can more easily fill open vacancies; this reduction in hiring costs causes more vacancies to be posted.

The key implication of this model for our analysis is that the equilibrium unemployment rate is determined jointly by the BC and JCC equations. Combining equations (4) and (5) yields:

$$u = s / \left(s + \mu^{\frac{1}{1-\alpha}} \left(\frac{J}{c}\right)^{\frac{\alpha}{1-\alpha}} \right). \quad (6)$$

This framework implies that the change in the equilibrium unemployment rate is determined not only by changes in match efficiency (e.g., mismatch) that shift the BC, but also by changes in vacancy creation, i.e. shifts in the JCC. Figure 2 illustrates equilibrium in our simple model of

frictional unemployment. It is based on Figure 1.2 in Pissarides (2000, p. 20). For a given BC and JCC, the equilibrium unemployment rate, or natural rate, at point a is determined by the intersection of the two curves.

2.2 Changes in equilibrium unemployment

The framework outlined in the preceding section enables us to determine what might lead to a rise in the equilibrium unemployment rate. This can occur due to an outward shift in the BC, a downward shift in the JCC, or a combination of both.

We first consider an outward shift in the BC, from BC to BC', shown in Figure 2. For a given JCC, this shifts the equilibrium from a to b . Because the JCC is upward sloping, equilibrium unemployment increases by less than the outward shift in the Beveridge curve, to a degree that depends on the slope of the JCC. In order for the unemployment rate to increase by the same amount as the rightward shift in the BC, the JCC either must be flat or must shift outward (or down) as well. This is illustrated by the shift in JC to JC' in Figure 2. This results in a shift in equilibrium from b to c .

This simple graphical analysis provides us with two key insights. First, to understand the driving forces of the rise in the unemployment rate, one must consider not only what is shifting the Beveridge curve and by how much, but also what is affecting job creation, i.e. labor demand. Second, to distinguish what part of the rise in the unemployment rate reflects purely cyclical fluctuations in labor demand and what parts are due to other factors, either transitory or permanent, that cause a rise in the natural rate, one has to consider what is driving the shifts in the BC and the JC curves and how long these effects are likely to last.

2.3 The recent shift in the Beveridge Curve

Figure 3 shows the empirical U.S. Beveridge Curve based on the monthly vacancy data from the Job Openings and Labor Turnover Survey (JOLTS), available beginning in December 2000. The data are divided into two groups, with the observations occurring prior to the recent recession in blue and the observations occurring since the start of the recession in orange. The last available data point (based on JOLTS release dates) is November 2010. In addition to the observed data, the figure contains a fitted Beveridge Curve, which is an update of the estimate provided in Barnichon, Elsby, Hobijn, and Şahin (2010). In terms of our simple framework the

fitted curve can be interpreted as the Beveridge Curve (4) solved using the pre-recession estimates of μ and α as well as the average pre-recession layoffs rate, s .⁸

In November 2010, the actual unemployment rate was 9.8 percent while the level of the unemployment rate on the fitted Beveridge Curve at the observed 2.4 percent vacancy rate is 6.3 percent. This implies the 3.5 percentage point unemployment gap, or implied outward BC shift, that is plotted in the figure. The size of the imputed current gap is not very sensitive to the estimation method applied. The non-linear OLS estimate in Valletta and Kuang (2010b) yields a similar size gap, as does the recalibrated version of Shimer's (2007) BC model presented by Kocherlakota (2010).⁹

The current estimate of a 3.5 percentage-point horizontal shift in the Beveridge curve is large. However, it is important to put this estimate into a longer-term historical perspective than that enabled by the JOLTS data. Construction of a historical Beveridge curve before 2000 involves merging data from JOLTS with the Help-Wanted Index published by the Conference Board. Two studies, Valletta (2005b) and Barnichon (2010), have independently constructed such merged historical series of the U.S. vacancy rate.¹⁰

Figure 4 displays the long-term BC based on this constructed historical time series for the vacancy rate. As can be seen from the figure, the BC shifted rightward about 4 percentage points between the 1960s and the early 1980s and then shifted back about 2.5 percentage points between 1984 and 1989. This suggests that the current outward shift in the BC is not very far outside the range of inward shifts that occurred during past labor market recoveries. Moreover, these outward shifts in previous recessions have coincided with increases in the natural rate of unemployment that are much smaller than the horizontal movement in the Beveridge curve. This is not only true for the CBO estimate of the natural rate in Figure 1 but also for the most

⁸ Barnichon, Elsby, Hobijn, and Şahin (2010) use data on vacancy yields (hires per vacancy) and quits rates from JOLTS rather than direct measures of the job finding rate, f . To do so, they use a slightly different flow-steady-state condition than the one we presented here in (2) to construct their fitted Beveridge curve.

⁹ Valletta and Kuang (2010b) point out that since we are on very flat part of the BC, small changes in the vacancy rate combined with a largely unchanged unemployment rate cause large monthly fluctuations in the estimated BC shift. For example, at the 2.1 percent vacancy rate and 9.5 unemployment rate that prevailed in June 2010, the implied unemployment gap with respect to the fitted BC was half the size observed five months later in November 2010.

¹⁰ The series from Valletta (2005b) and Barnichon (2010) are nearly identical. The main difference is their estimate of the relative importance of online job ads in the 1990's. Our substantive results are not sensitive to this difference. We use the series from Barnichon because it covers an additional decade.

commonly used alternative natural rate estimates (Orphanides and Williams, 2002). Credible estimates of the natural rate over these earlier periods suggest that it may have changed about half as much as the horizontal shift in the Beveridge curve.

An alternative way to consider the current shift in the Beveridge curve is presented in Barnichon, Elsby, Hobijn, and Şahin (2010). They observe that currently, the job-finding rate, f , is well below its level implied by the estimated matching function, (3), based on pre-recession data. This suggests a substantial decline in match efficiency, μ . They proceed by constructing the JOLTS-based equivalent of (4) for the currently observed percentage decline in match efficiency.

Figure 5 shows an updated estimate of their shifted BC, based on the average percentage decline in match efficiency in September 2010 through November 2010 relative to the pre-recession period. Two things are important to take away from this figure. First, in November 2010 the unemployment rate exceeded that implied by the shifted BC. This most likely indicates that unemployment is above its flow-steady-state level. It suggests that the 0.4 percentage point decline in the unemployment rate in December 2010, for which we do not yet have vacancy data, might be driven by a reversal of the unemployment rate to its flow-steady-state. Second, the figure shows that due to the non-linear nature of the BC, the horizontal shift is not uniform. At the 2.4 percent vacancy rate in November 2010 the estimated horizontal shift is 2.4 percentage points. At the 3.0 percent vacancy rate consistent with a pre-recession natural rate of unemployment equal to 5 percent, the horizontal shift is 1.9 percentage points.

Hence, our analysis of the shift in the BC highlights three difficulties with relying on simple plots of the BC to make inferences about changes in equilibrium unemployment. First, if the current unemployment rate exceeds its flow-steady-state level, such plots overestimate the shift in the BC. Second, our estimates of the shift in a non-linear BC suggest that the horizontal shift is not uniform but instead is larger at lower levels of the vacancy rate. Third, and most importantly, past horizontal shifts in the BC have coincided with much smaller movements in the natural rate of unemployment. This underscores the empirical relevance of the other curve in the model—the JCC—as discussed in the next section.

2.4 An estimate of the long-run JCC

Our estimate of the BC in Figure 5 suggests that, at that average pre-recession vacancy rate of 3 percent, the BC has shifted outwards by 1.9 percentage points. It is tempting to infer

from this that the natural rate of unemployment has thus increased from its pre-recession level of 5 percent to 6.9 percent now. However, as discussed in conjunction with Figure 2 earlier, this is equivalent to the claim that the JCC has shifted out by a corresponding amount or is flat, which is unlikely. The 6.9 percent value is more appropriately interpreted as an upper bound on the current natural rate. The actual change, i.e. the shift from a to b in Figure 2, depends on the shape of the JCC, in particular its upward slope.

To our knowledge, there are no existing estimates of the historical U.S. JCC. We therefore provide a rudimentary estimate of the long-run JCC. This estimate is based on the observation from Figure 4 that, historically, outward shifts of the BC have coincided with increases in the natural rate of unemployment as published by the CBO. These historical data also show that when the BC shifts outwards, the average vacancy rate generally increases. The average vacancy rate at a particular value of the natural rate of unemployment can be interpreted as an estimate of the *natural vacancy rate*. Hence, Figure 4 suggests that the natural rate of vacancies is increasing in the natural rate of unemployment.

We analyze this observation more formally based on Figure 6, which plots the historical vacancy rate series used earlier against the CBO's estimate of the natural rate of unemployment. The vertical lines created by sub-sets of the points in the figure depict cyclical fluctuations in labor demand for a given level of the natural rate. Each of these lines coincides with a different location of the BC in Figure 4.

A regression of the historical vacancy rate series on the natural rate of unemployment, using data points observed prior to the recent recession, yields the statistically significant upward sloping relationship depicted by the dashed line in Figure 6. This regression line can be interpreted as an estimate of the long-run JCC; it reflects typical vacancy creation, i.e. labor demand, as a function of the natural rate of unemployment.

The estimated JCC slope of 1.2 in Figure 6 probably is at the high end of the range of plausible estimates. The reason is that there is neither agreement on the right historical measure of the U.S. vacancy rate, i.e. the dependent variable in the regression, nor on the estimates of the natural rate of unemployment, the explanatory variable in the regression. Abraham (1987) pointed out that some of the variation in the Help-Wanted Index data used for the construction of the historical vacancy rate time series reflect a longer-run trend due to the occupational mix of job openings, the consolidation in the newspaper industry, and the increased requirements to post

job openings for Equal Employment Opportunity purposes. These factors likely drove up the index relative to the actual number of vacancies during the period of rising unemployment in the 1970s and 1980s. This might lead to an overestimate of the slope of the long-run JCC.¹¹

Other estimates of the natural rate of unemployment, including those computed by Orphanides and Williams (2002) that allow for more time variation in the natural rate, would produce different estimates of the slope of the JCC as well. In these cases the estimated regression line in Figure 6 would be flatter, i.e., the responsiveness of the average level of vacancies to the natural rate of unemployment would be weaker. That said, the empirical JCC would remain upward sloping.

2.5 Putting the BC and JC Curves together

To assign some magnitude to the potential increase in the natural rate of unemployment we combine the fitted and shifted JOLTS-based BC from Figure 5 with the estimated long-run JCC curve from Figure 6. As can be seen from Figure 7, the shifted BC and the empirical JCC curve intersect at an unemployment rate of 5.6 percent.¹² Note that this is substantially lower than the 6.9 percent upper bound on the natural rate of unemployment that we obtained by simply considering the horizontal shift in the BC.

Given the simplicity of our estimation strategy for the JCC, the estimated natural rate of 5.6 percent reflects considerable uncertainty. If one were to use alternative time varying estimates of the NAIRU to estimate the empirical JCC, it would flatten out and the estimate of the natural rate would increase. For this reason, we interpret the 5.6 percent estimate as a lower bound on the current natural rate of unemployment.

Thus, we find that if the currently estimated shift in the BC is permanent and the economy returns to its long-run JCC, then the long-run natural rate of unemployment has increased from its 5 percent level in 2007 to somewhere between 5.6 and 6.9 percent as of late 2010. In the absence of additional evidence to pin down its exact value, we regard 6.25 percent, as our preferred estimate of the current long-run natural rate of unemployment.

¹¹ Unreported results based on a vacancy rate measure that corrects for these issues still yield a significantly upward sloping long-run JCC.

¹² In Figure 7, the empirical JCC intersects the fitted pre-recession Beveridge curve at a 5 percent unemployment rate, which is the CBO estimate of the pre-recession level of the natural rate.

This conclusion, however, depends on two important assumptions. The first is that the labor market will return to its long-run JCC and that there are no factors that will permanently reduce job creation. The second is the assumption that the entire estimated shift of the Beveridge curve is permanent rather than attributable at least in part to temporary factors whose influence is likely to fade in the future. This raises the question of what factors affect the current positions of the BC and JCC curves and whether these factors reflect permanent or transitory influences. In the next part of this study we discuss the factors, beyond weak aggregate demand and layoffs, that are most commonly mentioned as moving the BC and JCC.

3. What is shifting the Beveridge and Job Creation Curves?

The previous section laid out a framework for distinguishing shifts in the BC from normal cyclical movements along it. The results suggested that in recent years both types of movement have occurred, implying that the aggregate labor market has been adjusting along an outward-shifted BC. In this section we turn to the second part of our analysis and ask whether the recent shifts in the BC and the JCC owe to structural or transitory factors.

Table 1 lists the five factors we consider, how they affect the BC and the JCC, and whether they are likely to be transitory or persistent. The factors are divided into two groups: normal cyclical shifters, and other factors. Our main focus in this section is on the other factors. We do not provide a detailed discussion of the shortfall in aggregate demand or the elevated level of layoffs.

As for weak aggregate demand, we consider this to be the determinant of the unemployment fluctuations beyond those captured by the natural rate. That is, it drives the shortfall in labor demand that depresses the JC and, for a given Beveridge curve, generates the cyclical move along the Beveridge curve. For this reason we offer no additional discussion of this topic.

The omission of layoffs relates to the fact that elevated layoffs rates tend to play an important role during the onset of recessions but they are of much less importance during

recoveries.¹³ Bowden (1980), Blanchard and Diamond (1990), and Mortensen (1994) all emphasize that counter-clockwise loops in (u, v) -space are commonly observed during U.S. business cycles and consistent with the simple theoretical framework we consider. Such loops are partly driven by an initial outward shift in the Beveridge curve driven by layoffs and a subsequent return inward when the economy recovers, layoffs dissipate, vacancies get posted, and the unemployment rate declines. These loops are evident in Figure 4. However, since measured layoffs rates from the JOLTS data have returned to their pre-recession levels, we do not consider them a concern for the labor market recovery going forward and therefore do not address them below.

We then proceed to provide more detailed evidence related to the three other factors in Table 1. The first is a mismatch between the characteristics of job openings, such as skill requirements or location, and the characteristics of the unemployed. This increases the matching frictions and moves the BC outwards. It also increases the costs of hiring since it reduces the probability of filling a vacancy and, through that, flattens the JCC. If mismatch is large enough to cause major disruptions in the labor market then it is unlikely to be resolved in a couple of years. As such, mismatch is generally regarded as the main potential cause of a long-run increase in the natural rate.

The second factor is the effect of extended availability of unemployment insurance (UI) payments. Increased availability of UI benefits is likely to increase the unemployment rate by increasing the net value of job search, thereby causing longer unemployment durations. While this effect is likely to be temporary, extended UI availability has been unprecedented during the recession and its aftermath, raising the possibility of large effects.

The final factor is the recent strong productivity growth during the latter part of the recession and the onset of the recovery. We focus on how these data are reconciled with the recent apparent outward shift of the Beveridge curve and the implications for job creation going forward. Since we have no reason to expect, nor are others forecasting, a structural upward shift in trend productivity growth, we consider this a transitory phenomenon.

¹³ This has been documented for the U.S. by, among others, Darby, Haltiwanger, and Plant (1985, 1986) and Fujita and Ramey (2009). Elsby, Hobijn, and Şahin (2008) show that this is also true across countries.

In the remainder of this section we provide recent empirical evidence on the potential importance of each of these three factors.

3.1 Mismatch

The mismatch argument for sustained increases in the unemployment rate and the natural rate of unemployment is predicated on imbalances in labor supply and demand across industry sectors, geographic areas, or skill groups. One way to assess the extent of such imbalances is by examining employment growth and unemployment rates across industry sectors, states, and occupations, as shown in Figures 8, 9, and 10. The underlying data used are the BLS payroll employment series and the Census Bureau's Current Population Survey microdata, which are available beginning in 1976. Industry and occupation unemployment rates are based on the most recent jobs held by unemployed individuals and exclude new labor market entrants. Only the unemployment series is shown for occupations because consistent occupational employment data are not available from BLS and cannot be reliably calculated from Census Bureau data.

Employment mismatches would be expected to increase if job growth were uneven, with some sectors gaining while others were shrinking. The actual extent of such job reallocation can be roughly calculated by examining differences in employment growth among sectors. Figures 8 and 9 show that dispersion in employment gains and losses spiked in the most recent recession as a result of severe and unevenly distributed job declines. For example, construction employment declined nearly 25 percent from the start of the recession through the end of 2009, while health and education jobs grew about 4 percent. Similar to past recessions, job losses were concentrated in cyclically sensitive sectors such as construction and manufacturing. The job losses also were distributed unevenly across states, causing the cross-state dispersion of employment growth to surge, although a downward trend over time caused it to remain well below the peak from previous recessions. As aggregate employment stabilized, the dispersion of employment gains and losses across industries and states returned to its pre-recession level, suggesting very little imbalance in the pace of employment growth in recent months.¹⁴

¹⁴ Some observers have argued that elevated unemployment rates in certain geographic areas reflect the geographic immobility of unemployed individuals whose home values have fallen below the amount owed on their mortgages ("house lock"). The earlier plot showing relatively balanced employment growth across states suggests that such geographic imbalances are minimal. More direct analysis by Valletta (2010), which relates unemployment durations to local housing market conditions, found no

While dispersion in employment growth has receded, Figures 8-10 show that dispersion in unemployment rates has remained high in recent months. Job growth has been too slow to significantly reabsorb workers idled by massive job losses in certain sectors, states, and occupations. The overhang of unemployed workers in the hardest-hit sectors suggests that some of them will need to look for work in other sectors as employment is reallocated in the economy. There is no straightforward way to assess the extent of required reallocation because it depends on unpredictable patterns of future demand. However, it is important to note that the recent peak in unemployment dispersion across industries differs little from the peak attained during the early 1980s recession, while unemployment dispersion across states and occupations is below that from the early 1980s recession. That earlier recession is generally acknowledged to have resulted from tight monetary policy in response to elevated inflation rates. It did not have a large mismatch or structural component, and the observed unemployment rate and the NAIRU dropped to low levels during the subsequent recovery. Given this historical precedent, current imbalances seem as likely to reflect cyclical as structural factors.

On the other hand, the sharp reduction in construction employment and the persistence of unemployment among workers in this sector probably reflects the aftermath and ensuing adjustments required by the busting of the housing bubble. The seasonally adjusted unemployment rate for construction workers has been hovering around 20 percent in recent months, compared with a more typical rate from 2003 to 2006 of about 7 to 8 percent. This represents about 1.25 million more unemployed construction workers in the current recovery than was typical during the preceding expansion. Assuming that half of them are unemployable, structural unemployment would increase about 0.4 percentage point. Although also affected by the housing bust, similar adjustments do not appear to be necessary in the financial sector. The increase in the unemployment rate for financial sector workers during the recent recession was below the average for all industries and of a relative magnitude similar to that posted in the early 1980s recession.

evidence of a meaningful house lock effect during 2009 and the first half of 2010. This conclusion is consistent with recent papers, most notably Schulhofer-Wohl (2010), that are based on alternative data sources. Moreover, concerns about a house lock effect were partly predicated on a measured decline in interstate mobility since 2006, which Kaplan and Schulhofer-Wohl (2010) show was largely a statistical artifact arising from a change in survey procedures for missing values. Finally, recent theoretical work by Sterk (2010) suggests that although house lock will lead to an outward shift in the BC, the likely shift is much smaller than the one depicted in our Figure 3.

Another way to address whether skill and sector mismatches are causing the recent shifts in the BC is to compare more detailed measures of labor market strength, namely vacancy yields (hires per vacancy) and quits rates across industries. Using data from JOLTS Barnichon, Elsby, Hobijn, and Şahin (2010) show that, consistent with the matching function (3), vacancy yields generally rise when unemployment is high and fall when labor markets tighten. Among private sector industries, labor market tightness, i.e. the ratio of the number of vacancies to the number of unemployed persons, explains more than two thirds of the variance in vacancy yields. Quits also tend to respond to the strength of the labor market; the strength of the labor market explains between one quarter and one half of the variance in industry quits rates. We exploit the responsiveness of vacancy yields and quits rates to labor market conditions to consider whether mismatch is present in particular industries of the economy.

Using the data and methods developed in Barnichon, Elsby, Hobijn, and Şahin (2010), Figure 11 plots the “shortfalls” in vacancy yields and quits rates by industry. These shortfalls are computed as the difference between actual average vacancy yields and quits rates from September 2010 through November 2010 and their predicted values based on a regression fit using data restricted to the period prior to the recent recession. The bars for the total economy in the figure show that both the aggregate vacancy yield and quits rate are well below where one would expect them to be at the current stage of the recovery.

If mismatch were present, we would expect to see considerable heterogeneity in the shortfalls, with growing industries (manufacturing, health care) facing the largest shortfalls in vacancy yield and struggling industries (construction, government) facing minimal shortfalls in vacancy yield. Similarly, a high degree of skill mismatch would suggest variation in quits rates across industries. In particular, one would expect firms that have a hard time finding persons to fill open positions to make job offers to employees at competitors, resulting in an elevated quits rate in the sector.

As the figure shows, the pattern by industry is not consistent with mismatch. The shortfall in vacancy yields and quits is broad-based, with little evidence of labor markets in growing sectors behaving differently than those in stagnant and shrinking sectors. Moreover, there is no evidence that sectors with especially low vacancy yields have especially high quits rates, as one would expect if sectoral mismatch was causing the shift in the BC. In contrast, the correlation of the shortfall in vacancy yields and quits across industries is large and positive

(0.77), rather than large and negative as would be expected if sector or skill mismatch were a significant problem (Table 2).

As a final check for evidence of sectoral mismatch, we compare the recent path of wages across the same industries for which we examined vacancy yields and quits, using the Employment Cost Index and Average Hourly Earnings series for the same industries used for the analysis of vacancy yields and quits rates above. The results are displayed in Figure 12. It reveals a broad-based weakness in wage gains relative to pre-recession rates. There is little variation across the growing and shrinking industries, providing scant support for the mismatch hypothesis. Most importantly, there is no consistent evidence of more rapid wage growth occurring in industries with higher quits rates or lower vacancy yields. Indeed, rather than being large and negative, the correlation between the wage series and the vacancy and quits series is small and positive (Table 2).

Taking all this information together we find little evidence to support the idea that structural mismatch in worker skills and available jobs across industries and geographic areas is driving recent shifts in the Beveridge curve.

3.2 Extended unemployment benefits

Another potentially important factor that may be influencing the search behavior of unemployed workers and boosting the underlying structural unemployment rate is the extended availability of unemployment insurance (UI) benefits. UI extensions are a standard policy response to elevated cyclical unemployment, and the sharp increase in the unemployment rate during the 2007-2009 recession resulted in an unprecedented increase in the potential duration of UI receipt. Beginning in late June, 2008, the maximum duration of UI benefits was extended multiple times, reaching 99 weeks for most UI-eligible job seekers as of late 2009.¹⁵ Congress has allowed the primary extension program to expire twice, most notably for nearly 2 months in

¹⁵ The joint Federal-State UI program provides up to 26 weeks of normal benefits. The recent benefit extensions reflect the impact of two federally funded programs: the permanently authorized Extended Benefits program, which provides up to 20 additional weeks of UI benefits, and the special Emergency Unemployment Compensation, which provides up to 53 weeks of benefits, depending on the unemployment rate in the recipient's state of prior employment (which causes the share of unemployed workers eligible for the 99-week maximum to change over time). The previous maximum eligibility was 65 weeks under the Federal Supplemental Benefits (FSB) program in the mid-1970s.

June-July of 2010, but in each case renewed the extensions, which currently are effective through January 3, 2012.

Increased availability of UI benefits is likely to increase unemployment duration through two primary behavioral channels. First, the extension of UI benefits, which represents an increase in their value, may reduce the intensity with which UI-eligible unemployed individuals search for work. This could occur because the additional UI benefits reduce the net gains from finding a job and also serve as an income cushion that helps households maintain acceptable consumption levels in the face of unemployment shocks (Chetty 2008). Alternatively, the measured unemployment rate may be artificially inflated because some individuals who are not actively searching for work are identifying themselves as active searchers in order to receive UI benefits (a “labeling effect”). These behavioral effects on job search will increase the non-cyclical component of the unemployment rate during the period over which extended benefits are available.

Given the unique nature of recent economic conditions and the UI extensions, assessing the magnitude of the extended UI effect is challenging. Based on existing empirical research using U.S. data, Chetty (2008) noted that a 10% increase in the overall value of UI benefits increases unemployment durations by 4–8%. Other estimates, particularly those that focus on extension periods rather than the dollar value of benefits, lie below this range (see for example, Card and Levine 2000). As such, there is a wide range of uncertainty around the implied estimates of the impact of the recent UI extensions on unemployment duration. Moreover, as noted by others (e.g., Katz 2010), the impact of UI benefits on job search likely was higher in the 1970s and 1980s than it is now, due to the earlier period's greater reliance on temporary layoffs and the corresponding sensitivity of recall dates to unemployment insurance benefits. As such, reliance on past estimates of the effects of UI generosity and duration on unemployment is likely to lead to overestimates of the effects of extended UI in the current economic environment. Hence, it is important to use recent labor market data to obtain estimates of extended UI effects.

A few studies conducted by other researchers within the Fed system have assessed the impact of extended UI on the unemployment rate using recent labor market data. The estimated increase in structural unemployment ranges from 0.7 percentage points (Aaronson, Mazumder, and Schecter 2010; AMS) to a maximum of 1.7 percentage points (Fujita 2010). AMS based their analysis on direct examination of recent labor market flows, focusing in particular on the

recent low level of flows between unemployment and labor force withdrawal, hence the “labeling effect” noted above. Fujita also based his estimates on labor market flows, focusing on movements out of unemployment (to jobs or labor force withdrawal) around the time that normal UI benefits expire (26 weeks). In particular, he noted that after the UI extensions were implemented, these flows fell well below their levels observed prior to the recession. He attributed the difference to the impact of extended UI benefits and used the observed equilibrium relationship between labor market flows and the unemployment rate to calculate an impact on the unemployment rate ranging from 0.9 to 1.7 percentage points.¹⁶

Our own empirical assessment using recent data focuses on direct calculation and comparison of the duration of unemployment for individuals who are eligible or not eligible for UI receipt, as reflected in their reported reason for unemployment (as initially implemented and reported in Valletta and Kuang 2010a). We note that not all unemployed individuals are eligible to receive benefits, whether regular or extended. In particular, UI receipt generally is restricted to individuals who are unemployed through “no fault of their own,” to quote U.S. Labor Department eligibility guidelines, and have recent employment history that allows them to meet a base earnings test. The unemployed are categorized into four main groups in the Current Population Survey (CPS) data: involuntary job losers; voluntary job leavers; new labor market entrants; and re-entrants, meaning people who have worked before and are now actively searching for a job following a period of labor market inactivity.¹⁷ Individuals eligible for UI are concentrated in the job losers group, with few eligible individuals likely to identify themselves as job leavers or entrants. Under these conditions, unemployed job leavers and labor force entrants can be used as a control group for assessing the impact of extended UI benefits on eligible individuals.

Because we do not have access to precise administrative data on UI reciprocity and unemployment experiences, this experimental design is by necessity imperfect. It is likely that

¹⁶ This approach may overstate the impact of extended UI due to the implicit assumption that a worker who has been unemployed for nearly 26 weeks in 2009-2010 and hence is on the verge of receiving extended benefits faces the same job opportunities as a worker unemployed for nearly 26 weeks prior to the recession.

¹⁷ The CPS is conducted by the U.S. Census Bureau for the U.S. Bureau of Labor Statistics and is the official source of monthly aggregate labor force data in the United States. We used the individual observations on unemployment from the CPS microdata files, which at this time are only available through October 2010.

some CPS respondents mislabel their reason for unemployment, which will cause the presumed eligible and ineligible groups to overlap and produce an underestimate of the impact of extended UI on unemployment duration. This problem aside, calculations indicate that the number of UI-eligible individuals in our data is slightly larger than the number of actual UI recipients in the economy. This is not surprising given that UI benefits are unclaimed by some eligible individuals and expire for others, and it suggests that respondent misrepresentation is likely to be limited. Another potential concern is that re-entrants, who are in the control group of ineligible individuals, may in fact be eligible for UI receipt, with some of them having maintained the job history and search behavior required for UI receipt even though they reported a period of inactivity in the CPS survey. However, in analysis not reported, results for the effect of the UI extensions when re-entrants are excluded from the control group are similar.

For our specific test, we examine the increase in unemployment duration observed as the UI extensions were introduced and renewed in 2008 and 2009. We use the “expected unemployment duration” concept from Valletta (2005), which yields a monthly measure of the typical completed duration of unemployment for an individual who becomes unemployed in a particular month, based on the distribution of individual unemployment spells for the current and prior months. This measure more accurately reflects the overall duration of unemployment spells and changes in duration over time than do the average and median duration series published by the BLS, which are tallied from incomplete spells measured at the time each survey is conducted.

Figure 13 displays the resulting unemployment duration series for job losers and leavers/entrants from 2005 through October 2010 (the latest month of availability for CPS micro-data). The vertical lines identify the dates when the four separate extensions of maximum UI duration were implemented. Unemployment duration rose slightly in the early phase of the recession and then increased sharply after extended UI benefits became available, reaching a high of about 35 weeks in mid-2009 before declining back to about 25 weeks by about mid-2010. Both groups saw a substantial increase in unemployment durations, although the increase appears larger for the job loser group, due to the lower pre-recession level and higher recessionary peak.

The differential increase in unemployment duration between the two groups, measured from a pre-recession baseline to a period after the full UI extensions had taken effect, provides an estimate of the impact of extended UI benefits on job search and labor force attachment: both

groups faced the same deterioration in labor market conditions but are distinguished by whether they became eligible for extended unemployment benefits.¹⁸ Because the final extensions became effective in November 2009, the original Valletta and Kuang estimates focused on the fourth quarter of 2009 as the comparison period. As of that quarter, the expected duration of unemployment had risen about 18.7 weeks for job losers and about 17.1 weeks for leavers and entrants, using the years 2006–2007 as the pre-recession baseline. The differential increase of 1.6 weeks for job losers is the presumed impact of extended UI benefits on unemployment duration. It is straightforward to translate this increase in unemployment duration into an effect on the unemployment rate, based on their proportional relationship and adjusted for the share of job losers in overall unemployment, which was about 67 percent in December 2009. The implied increase in the unemployment rate is quite small, slightly less than 0.4 percentage point, indicating that without UI extensions, the measured unemployment rate would have been 9.6 percent in December 2009 rather than the observed 10.0 percent.

This focus on the fourth quarter of 2009 may be too narrow, however, particularly because the pattern in unemployment durations differed across the two groups over time. The increase in duration was initially very pronounced for job losers, but the duration gap between them and quits/entrants declined over time. Because the increase in durations became pronounced in 2009 and extended UI was available with little interruption through the CPS survey week in mid-June 2010, we provide an alternative comparison based on the entire period of January 2009 through June 2010. Using this broader comparison period, we obtain a larger estimated increase in duration for job losers: about 3.5 weeks, which implies an increase in the overall unemployment rate of about 0.8 percentage points.

To summarize, the results of this analysis suggests that the availability of extended unemployment benefits has increased the overall unemployment rate by about 0.4 to 0.8 percentage points. Our preferred estimate is toward the upper end of this range, because that higher estimate encompasses effects of extended UI spread out over a broader, more appropriate effective period than the smaller estimate. In either case, the impact of extended UI on the unemployment rate is expected to disappear as labor market conditions improve and the extended UI provisions are allowed to expire.

¹⁸ The availability of normal UI benefits (first 26 weeks) does not affect this comparison because it was unchanged before and after extended UI was implemented.

3.3 Productivity growth

The final factor we discuss that may be contributing to the persistently elevated unemployment rate is productivity growth. As documented by Daly and Hobijn (2010), since the onset of the recovery in June 2009 the U.S. unemployment rate has been substantially higher than would be implied by its historical relationship with overall economic activity. Figure 14 shows the historical relationship between the deviation of the unemployment rate from its natural rate and GDP from its potential—Okun’s Law.¹⁹ The red points in the figure show the data since the start of the 2007 recession. As the figure shows, during the recession Okun’s Law was relatively stable but began to break down in the summer of 2009. This suggests that the growth in economic activity since that time has not been accompanied by the normal (historically predicted) decline in the unemployment rate.

To understand the causes of the observed breakdown in Okun’s Law over the last year it is useful to refer to the output identity that underlies Okun’s original empirical analysis. There are numerous ways to express this identity; we rely on the following method which fully disentangles the contributions of important labor supply variables that have been the subject of concern during the recent recession. The output identity is written as:

$$Y = \underbrace{\left(\frac{Y}{H}\right)}_{ALP} \underbrace{\left(\frac{H}{E^*}\right)}_{AVH} \underbrace{\left(\frac{E^*}{E}\right)}_{EDD} \underbrace{\left(\frac{E}{L}\right)}_{EMP} \underbrace{\left(\frac{L}{N}\right)}_{LFP} N, \quad (7)$$

where Y is output, ALP is GDP per nonfarm hour, AVH is hours per worker, EDD is the ratio of the employment measure used for productivity statistics and household survey employment (the employment definition discrepancy), EMP is the employment rate, LFP is the labor force participation rate, and N is the working age population. Note that for our purposes a change in the employment rate is approximately equal to the arithmetic inverse of the change in the unemployment rate. Therefore, Okun’s Law holds as a rough empirical regularity because the product of the other factors in the equation, namely, $\left(\frac{Y}{H}\right)$, $\left(\frac{H}{E^*}\right)$, $\left(\frac{E^*}{E}\right)$, $\left(\frac{L}{N}\right)$, and N have co-moved over the business cycle with output and the unemployment rate.

During this recovery, we have seen a large deviation from this co-movement. In particular, as Wilson (2010) shows, we have seen a historically high level of average labor

¹⁹ The estimated unemployment and output gaps in the figure are based on CBO (2010).

productivity relative to its trend. The importance of this productivity growth for recent GDP growth is illustrated in Figure 15, which decomposes cumulative GDP growth since the start of the recession into the components listed in equation (7).

Standard search models of equilibrium unemployment that nest the simple framework we use here imply that such strong productivity growth would result in a surge of vacancy postings and a resulting decline in the unemployment rate, reflected in a steeper Job Creation curve. Although labor demand and employment have grown since the end of the recession, the increase has not matched productivity growth. This increase in productivity growth without any substantial increase in hiring is reminiscent of the jobless recoveries that followed the previous two U.S. recessions.

Both van Rens (2004) and Koenders and Rogerson (2005) offer explanations for jobless recoveries. They emphasize that rather than interpreting productivity growth as an exogenously given process, its pattern following recessions is best understood as resulting from directed decisions made by firms. In van Rens, firms face the option of either hiring workers or temporarily boosting productivity growth by moving workers from the production of intangibles to the production of measured output. In Koenders and Rogerson, rather than hiring, firms might choose to adopt organizational changes that improve productivity but were temporarily shelved during the prior expansion. In either case, the reorientation of production activity reduces the rate of hiring but raises productivity growth, consistent with the recent pattern of jobless recovery. Moreover, with fixed hiring and firing costs and uncertainty about the future state of aggregate demand, the option value of hiring new workers declines, thereby intensifying the jobless recovery (Bentolila and Bertola 1990, Bloom 2009). Combining these two insights implies that in recovery periods characterized by substantial uncertainty, firms will be reluctant to hire workers and instead will rely on temporary measures that boost their productivity to satisfy increased demand. This would put downward pressure on the Job Creation Curve despite strong productivity growth.

Uncertainty might also contribute to the outward shift of the Beveridge curve, by causing firms that create vacancies to become more selective about filling them. Such a change in firms' hiring decisions would cause a decline in the vacancy yield, as depicted in Figure 11 and discussed above; it is also consistent with a reduction in recruiting intensity in the sense of Davis, Faberman, Haltiwanger (2010).

Though this is a possible explanation for the joint weakness in vacancy creation and vacancy yields we have observed relative to the strong productivity growth during the recovery, we would like to emphasize that there are no studies that have tried to quantify this effect. This is why, in Table 1, we list the effect of productivity growth on the Beveridge curve in parentheses, indicating ambiguity about its role.

3.4 Where are we now?

The goal of this analysis has been to reach a conclusion about whether and to what extent the natural rate of unemployment has risen in the aftermath of the severe recession. The final task we undertake then is to put the empirical results just described into the model of equilibrium unemployment described earlier. These results are displayed in Figure 16 which shows the fitted and shifted BC, the empirical JCC, and the data on vacancies and unemployment through November 2010. The first thing to note is that the most recent data indicate that the unemployment rate exceeded that implied by the shifted BC (point A on the figure). We attribute this difference to a deviation of the unemployment rate from its flow-steady-state level, and we argue that this gap will likely be closed in coming months as the unemployment rate returns to its flow steady-state. We estimate that the BC has shifted out over the last several years but find little evidence that mismatch has been an important cause. Moreover, the BC appears to have shifted out by more than can be explained by a small degree of mismatch and the larger impact of extended UI benefits. We attribute the remaining decline in the vacancy yield to a fall in recruiting intensity related to endogenous productivity growth and/or uncertainty. While there is admittedly little research in this area, this story is consistent with previous work related to jobless recoveries. Finally, while we provide no direct estimates of movements in the JCC, we find that factors affecting it—the cyclical shortfall in labor demand, extended UI, and endogenous productivity growth and uncertainty—are likely to have rotated it downward, further exacerbating persistent unemployment by reducing hiring activity.

4. Conclusion

The stubbornly high rate of unemployment in the face of ongoing GDP growth and rising job openings has raised concerns that the level of structural unemployment, hence the natural or

non-accelerating inflation rate of unemployment, has risen over the past few years in the United States. This is an important policy issue because short-run monetary and fiscal stabilization policies are not appropriate for alleviating structural unemployment.

In this paper we revisited a simple framework laid out in Pissarides (2001, Chapter 1) that is useful for understanding the movements of equilibrium unemployment. This framework models equilibrium unemployment as an outcome from the interaction of two curves—the Beveridge curve (BC) and the Job Creation curve (JCC)—and highlights the pitfalls of drawing conclusions about structural unemployment from movements in the BC alone. In addition, this framework supports our itemization and description of the factors that can change equilibrium unemployment on a transitory or more durable basis.

Using this framework we estimated the empirical BC and JCC and the implied values of equilibrium unemployment. The results suggest that the natural rate of unemployment has risen from its pre-recession level of 5.0 percent to a value between 5.6 and 6.9 percent, with our preferred estimate lying at 6.25 percent. This value implies an unemployment gap of over three percentage points, which remains quite high. Even at the top of the range of estimated natural rates, however, considerable slack remains in the labor market.

We then turn to more direct investigations of the factors potentially pushing up the estimated natural rate of unemployment: mismatch between the characteristics of available jobs and workers, the extension of UI benefits, and recent increases in productivity. Mismatch would imply that the increase in equilibrium unemployment would persist, while the latter two would be associated with a transitory increase that would be expected to dissipate as the economy improves. Looking at data on employment growth and unemployment dispersion, vacancy yields and quits rates by industry, and wage growth by industry we find evidence of at most a modest impact of labor market mismatch, on the order of 0.5 percentage points or less. In contrast, about half or more of the 1.25 percentage point increase in the natural rate appears to result from the current temporary extensions of unemployment compensation. Rapid productivity gains, which likely reflect transitory influences on firms' organizational decisions, also appear to be having a small impact. Thus, only a modest portion of the elevated unemployment rate, about 0.5 percentage points or less, seems attributable to factors that have raised the natural or non-accelerating inflation rate of unemployment on a durable basis.

References

- Aaronson, Daniel, Bhashkar Mazumder, and Shani Schechter (2010). “What is behind the rise in long-term unemployment?” Federal Reserve Bank of Chicago *Economic Perspectives* 2Q/2010: 28-51.
- Abraham, Katherine G. (1987) “Help Wanted Advertising, Job Vacancies and Unemployment,” *Brookings Papers on Economic Activity*, 207-243.
- Barnichon, Regis (2010), “Building a composite Help-Wanted Index,” *Economics Letters*, Forthcoming.
- Barnichon, Regis, Michael Elsby, Bart Hobijn and Ayşegül Şahin (2010) “Which Industries are Shifting the Beveridge Curve?” *FRB SF Working Paper 2010-32*.
- Basu, Susanto, John G. Fernald, and Miles S. Kimball (2006) “Are Technology Improvements Contractionary?” *American Economic Review*, December 96(5), 1418-48.
- Bentolila, Samuel, and Giuseppe Bertola (1990) “Firing Costs and Labour Demand: How Bad is Eurosclerosis?” *Review of Economic Studies*, 57, 381-402.
- Blanchard, Olivier, and Peter Diamond (1989) “The Beveridge Curve,” *Brookings Papers on Economic Activity*, 1-76.
- Bloom, Nicholas (2009) “The Impact of Uncertainty Shocks,” *Econometrica*, 77, 623–685.
- Bowden, Robert J. (1980) “On the Existence and Secular Stability of u-v Loci,” *Economica*, 47, 33-50.
- Brauer, David (2007) “The Natural Rate of Unemployment,” *CBO Working Paper 2007-06*.
- Card, David, and Phillip B. Levine. (2000). “Extended benefits and the duration of UI spells: evidence from the New Jersey extended benefit program.” *Journal of Public Economics* 78: 107-138.
- Chetty, Raj. (2008). “Moral Hazard versus Liquidity and Optimal Unemployment Insurance.” *Journal of Political Economy* 116(2, April): 173-234.
- Congressional Budget Office (2010), *The Budget and Economic Outlook: An Update*, August 2010.
- Daly, Mary C., and Bart Hobijn (2010) “Okun’s Law and the Unemployment Surprise of 2009,” *FRB SF Economic Letter 2010-07*.
- Darby, Michael R., John C. Haltiwanger, and Mark W. Plant (1985) “Unemployment Rate Dynamics and Persistent Unemployment under Rational Expectations,” *American Economic Review*, 75, 614-637.
- Darby, Michael R., John C. Haltiwanger, and Mark W. Plant (1986) “The Ins and Outs of Unemployment: The Ins Win,” *NBER Working Paper 1997*.
- Davis, Steven, Jason Faberman, and John Haltiwanger (2010) “The Establishment-Level Behavior of Vacancies and Hiring,” *NBER Working Paper No. 16265*.

- Elsby, Michael, Bart Hobijn and Ayşegül Şahin (2008) “Unemployment Dynamics in the OECD,” *NBER Working Paper 14617*.
- Elsby, Michael, Bart Hobijn and Ayşegül Şahin (2010) “The Labor Market in the Great Recession,” *Brookings Papers on Economic Activity*, Spring 2010, 1-48.
- Fujita, Shigeru (2010). “Effects of the UI Benefit Extensions: Evidence from the Monthly CPS.” Working Paper No. 10-35, Federal Reserve Bank of Philadelphia, November.
- Fujita, Shigeru and Garey Ramey (2009) “The Cyclicalities of Job Loss and Hiring,” *International Economic Review*, 50, 415-430.
- Kaplan, Greg, and Sam Schulhofer-Wohl (2010) “Interstate Migration Has Fallen Less Than You Think: Consequences of Hot Deck Imputation in the Current Population Survey,” *FRB Minneapolis Working Paper 681*.
- Katz, Lawrence. (2010). “Long-Term Unemployment in the Great Recession.” Testimony to the Joint Economic Committee, U.S. Congress. April 29, 2010.
- Kocherlakota, Narayana (2010) “Inside the FOMC,” *Speech at Marquette, Michigan*, August 17, 2010.
- Koenders, Kathryn, and Richard Rogerson (2005), “Organizational Dynamics Over the Business Cycle: A View on Jobless Recoveries,” *Federal Reserve Bank of St. Louis Review*, 87, 555-579.
- Kwok, Joyce, Mary C. Daly, and Bart Hobijn (2010) “Labor Force Participation and the Future Path of Unemployment,” *FRB SF Economic Letter 2010-27*.
- Mortensen, Dale T. (1994) “The Cyclical Behavior of Job and Worker Flows,” *Journal of Economics Dynamics and Control*, 18, 1121-1142.
- Okun, Arthur M. (1962), “Potential GNP: Its Measurement and Significance,” *Proceedings of the Business and Economics Statistics Section of the American Statistical Association*, 98-104.
- Orphanides, Athanasios, and John C. Williams (2002) “Robust Monetary Policy Rules with Unknown Natural Rates,” *Brookings Papers on Economic Activity*, 2002-2, 63-118.
- Petrongolo, Barbara, and Christopher A. Pissarides (2001) “Looking into the Black Box: A Survey of the Matching Function,” *Journal of Economic Literature*, 39, 390-431.
- Pissarides, Christopher A. (2000), *Equilibrium Unemployment Theory*, Cambridge, MA: MIT Press.
- Pissarides, Christopher A. (2009) “The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?” *Econometrica*, 77, 1339-1369.
- Reinhart, Carmen M., and Kenneth S. Rogoff (2009) “The Aftermath of Financial Crises,” *American Economic Review*, 99, 466-472.
- Romer, Christina D., and David H. Romer (1994), “What Ends Recessions?,” *NBER Macroeconomics Annual*, 9, 13-57.
- Schulhofer-Wohl, Sam (2010) “Negative Equity Does Not Reduce Homeowners' Mobility,” *FRB Minneapolis Working Paper 682*.

- Shimer, Robert (2005) “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *American Economic Review*, 95, 25-49.
- Shimer, Robert (2007), “Mismatch,” *American Economic Review*, 97, 1074-1101.
- Sterk, Vincent (2010), “Home Equity, Mobility, and Macroeconomic Fluctuations,” *DNB Working Paper No. 265*.
- Valletta, Robert G. (2005a) “Rising Unemployment Duration in the United States: Causes and Consequences,” *FRB SF mimeo*.
- Valletta, Robert G. (2005b) “Why Has the U.S. Beveridge Curve Shifted Back? New Evidence Using Regional Data,” *FRB SF Working Paper 2005-25*.
- Valletta, Robert G. (2010). “House Lock and Structural Unemployment.” FRBSF mimeo, November.
- Valletta, Robert G., and Katherine Kuang (2010a). “Extended Unemployment and UI Benefits,” *FRB SF Economic Letter 2010-12*.
- Valletta, Robert G., and Katherine Kuang (2010b). “Is Structural Unemployment on the Rise?” *FRB SF Economic Letter 2010-34*.
- Van Rens, Thijs (2004), “Organizational Capital and Employment Fluctuations,” *CREI mimeo*.
- Wilson, Daniel J. (2010) “Is the Recent Productivity Boom Over?” *FRB SF Economic Letter 2010-28*.

Table 1. Factors that move the BC and JC curves.

<i>Shifter</i>	<i>JC</i>	<i>BC</i>	<i>Transitory or Permanent</i>
<u><i>Normal cyclical shifters</i></u>			
Shortfall in aggregate demand	↓		Transitory
Elevated layoffs rate	↓	⇒	Transitory
<u><i>Other factors</i></u>			
Decrease in match efficiency (mismatch)	↓	⇒	Persistent/permanent
Increased generosity of unemployment insurance	↓	⇒	Depends on the duration of the shift
Productivity increase	↑	(⇒)	Transitory

Table 2. Cross-industry correlations.

	Shortfall in...		Deceleration of growth of...	
	Quits rate	Vacancy yield	ECI	AHE
Quits rate		0.77	0.36	0.23
Vacancy yield	0.77		0.06	0.19
ECI	0.36	0.05		0.15
AHE	0.23	0.19	0.15	

Note: These are correlations between shortfalls from Figure 11 and wage disinflation from Figure 12. ECI is Employment Cost Index (wages and salaries) and AHE is Average Hourly Earnings.

Figure 1. Actual and natural rate of unemployment.

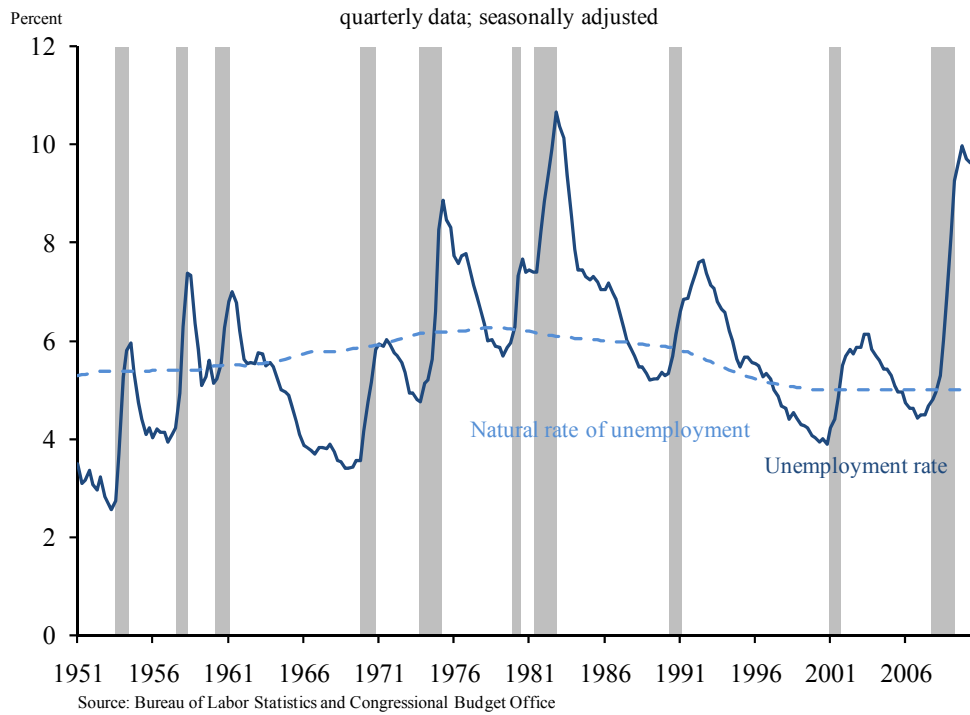


Figure 2. Determinants of shifts in equilibrium unemployment.

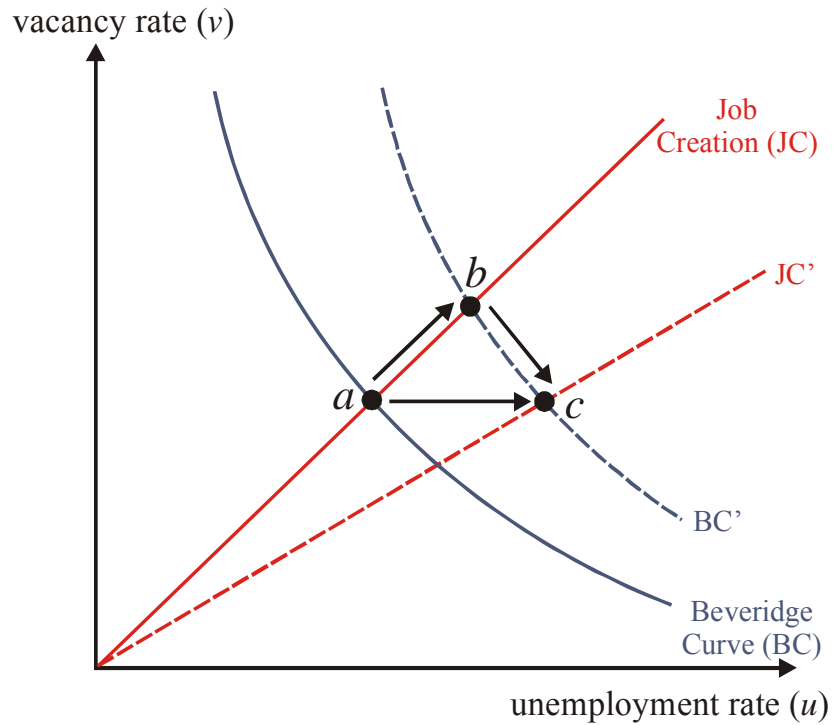


Figure 3. The U.S. Beveridge curve, January 2000–November 2010.

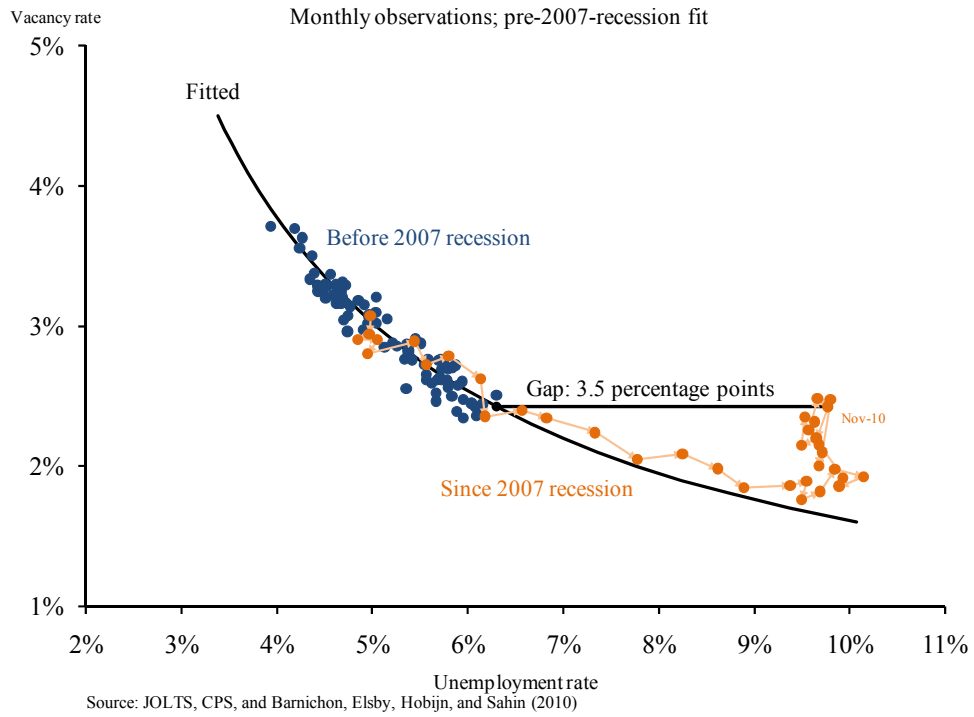


Figure 4. Historical shifts in the Beveridge curve.

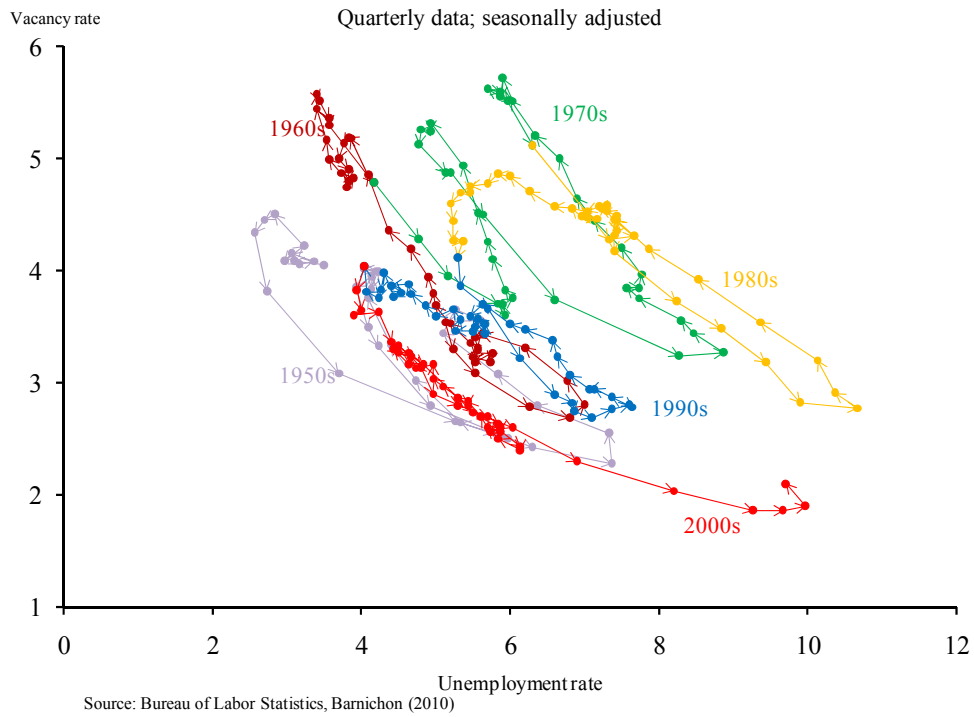


Figure 5. JOLTS-based estimated and shifted Beveridge curves

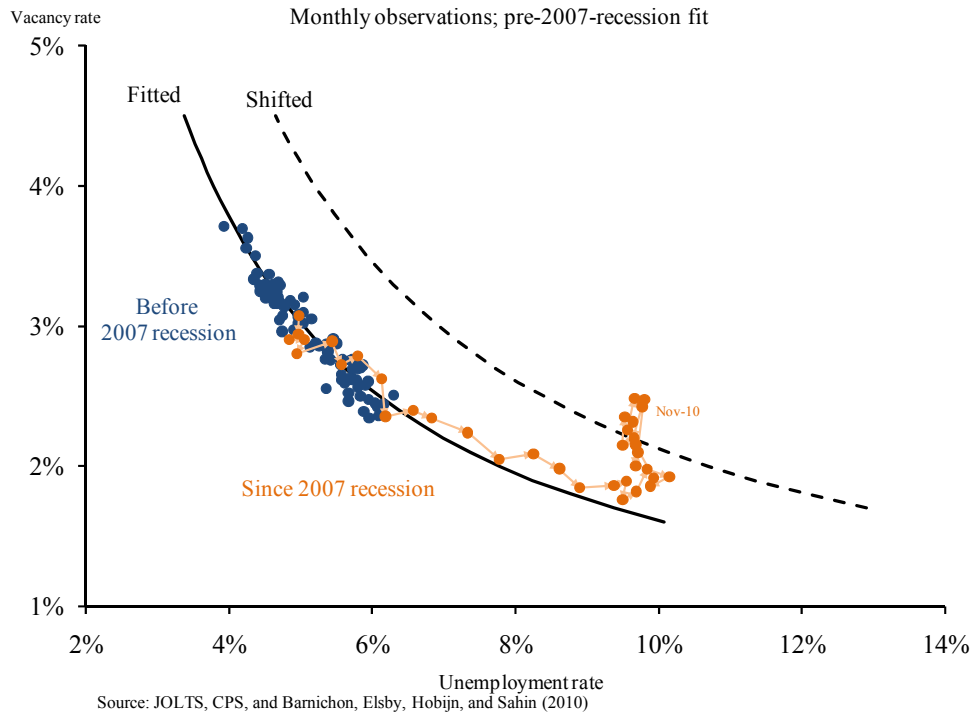


Figure 6. Estimated long-run Job Creation curve

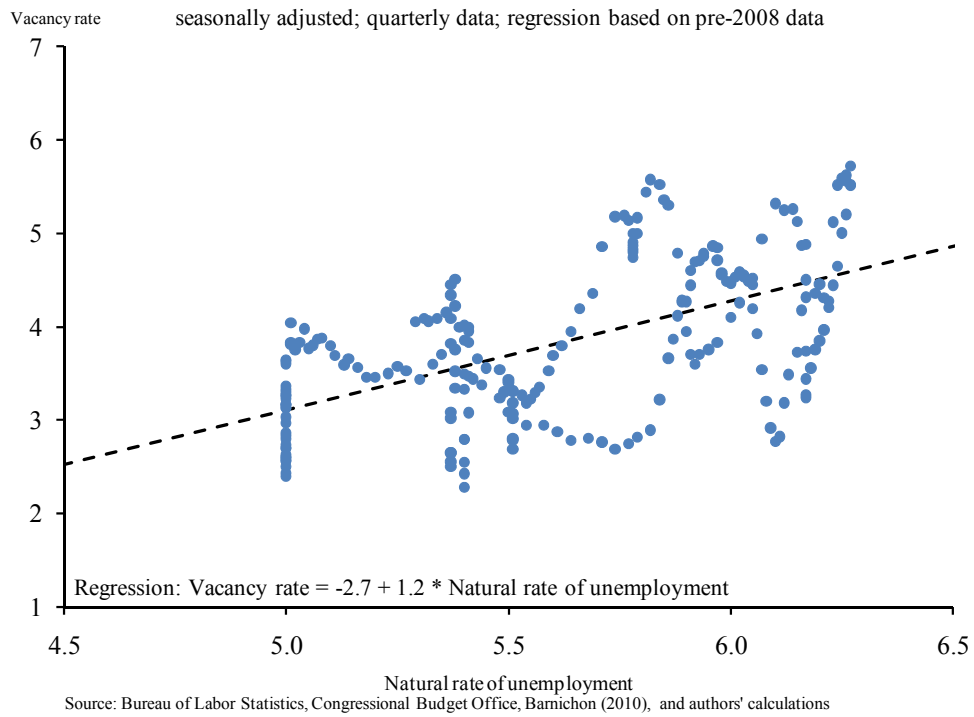


Figure 7. Estimated Job Creation and Beveridge curves

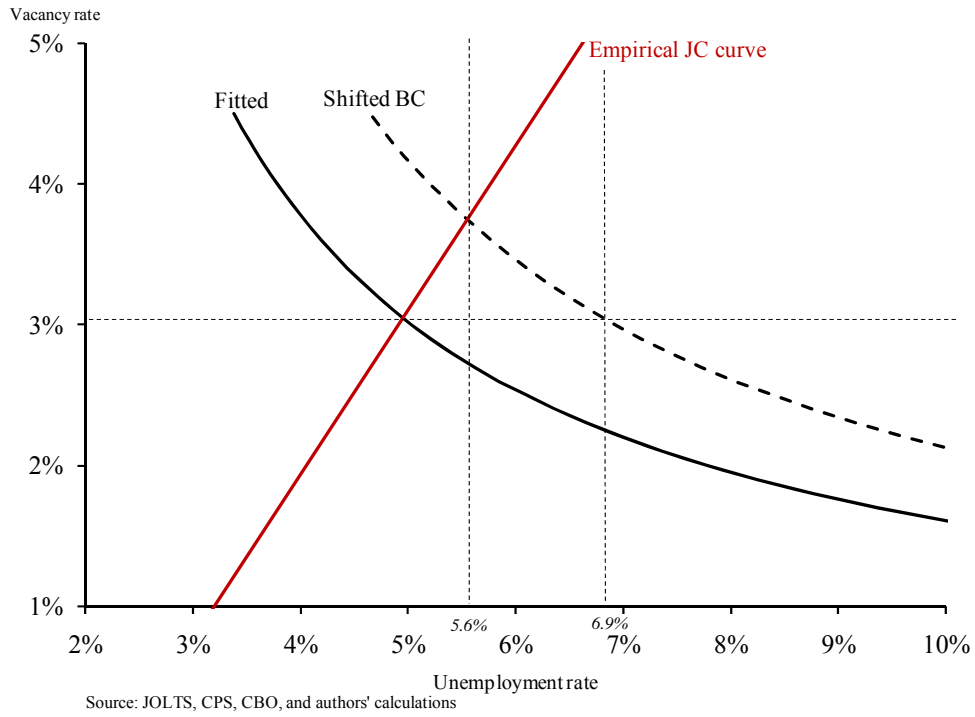


Figure 8. Employment growth and unemployment rate dispersions across industry sectors.

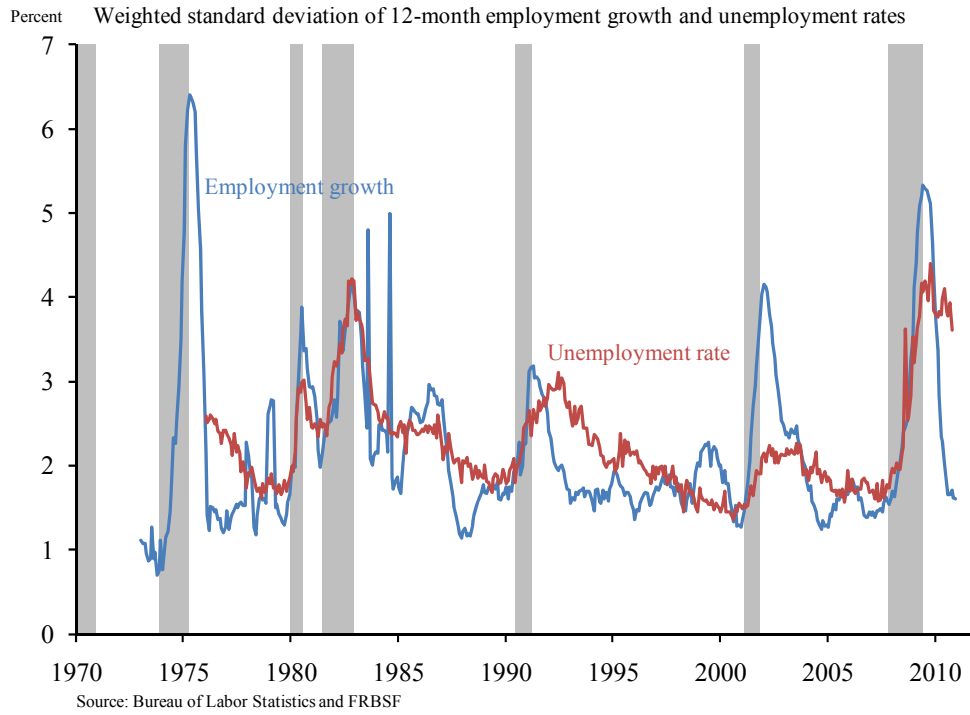


Figure 9. Employment growth and unemployment rate dispersions across states.

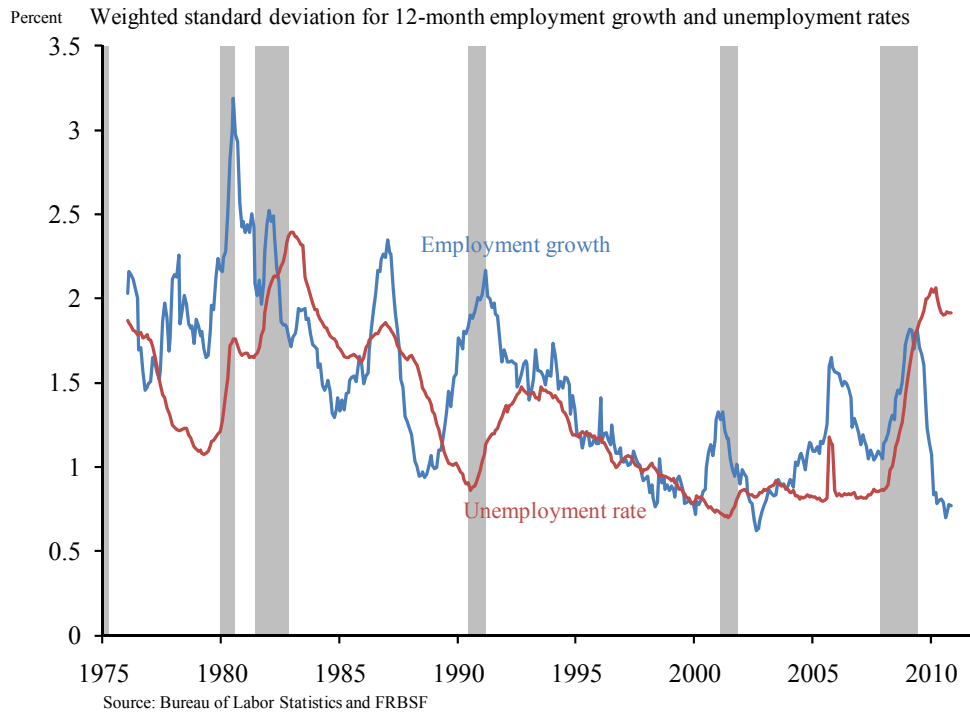


Figure 10. Unemployment rate dispersions across occupations.

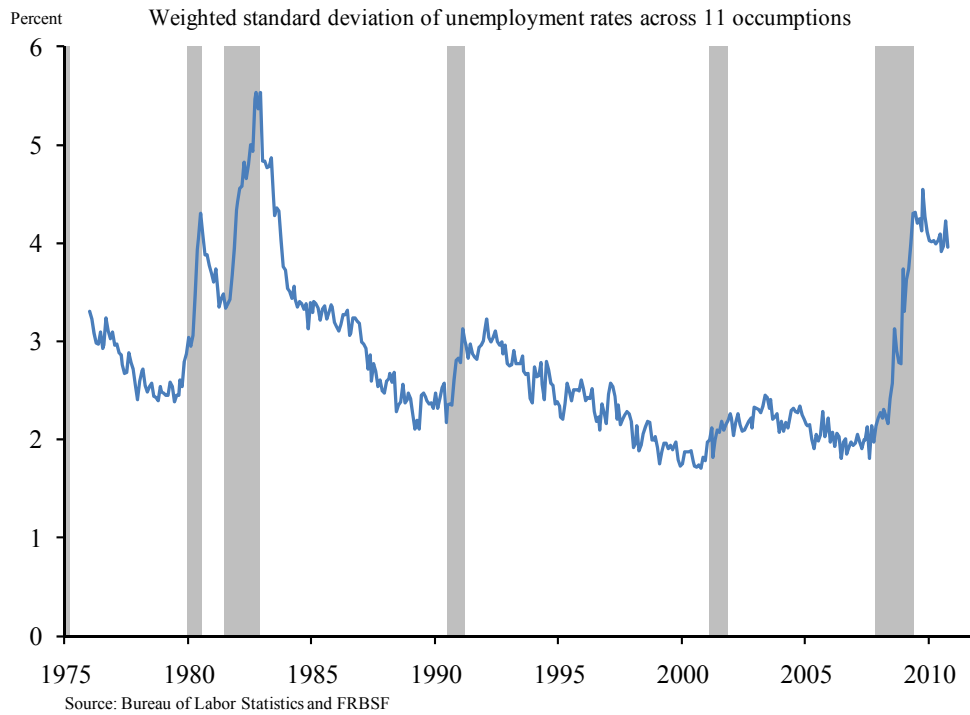


Figure 11. Broad-based shortfall in hiring per vacancy as well as quits rates.

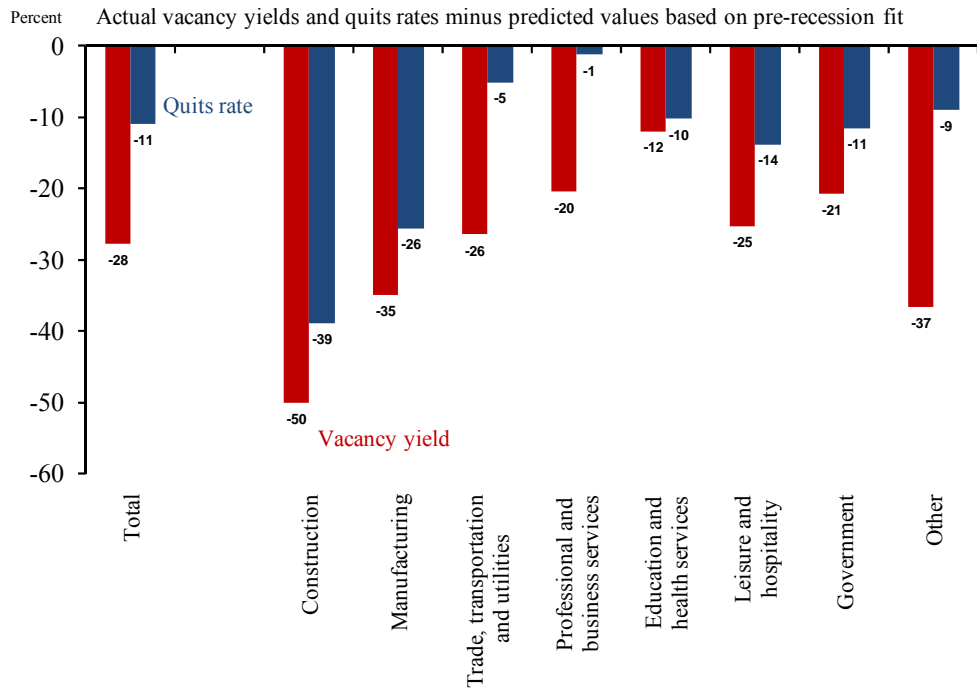


Figure 12. Deceleration in wage growth by industry.

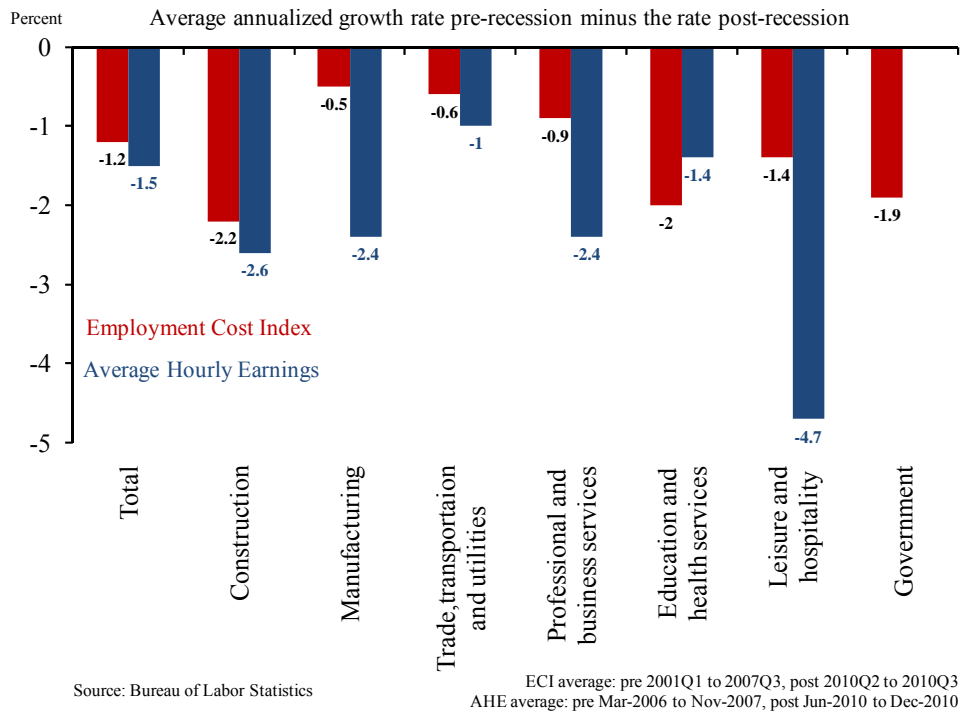


Figure 13. Unemployment duration by reason.

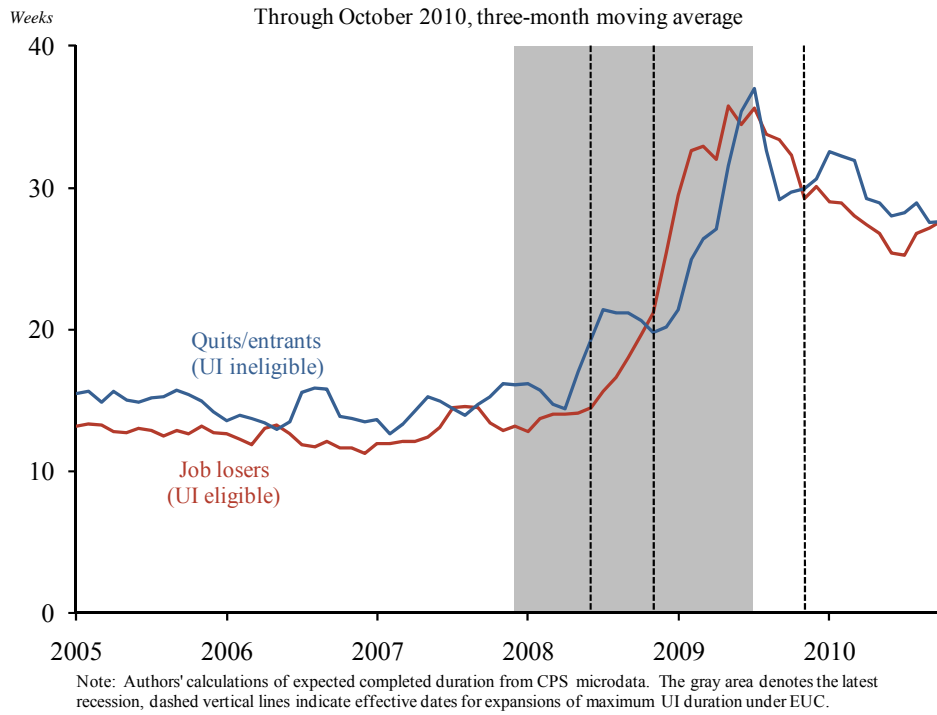


Figure 14. Okun's Law has broken down.

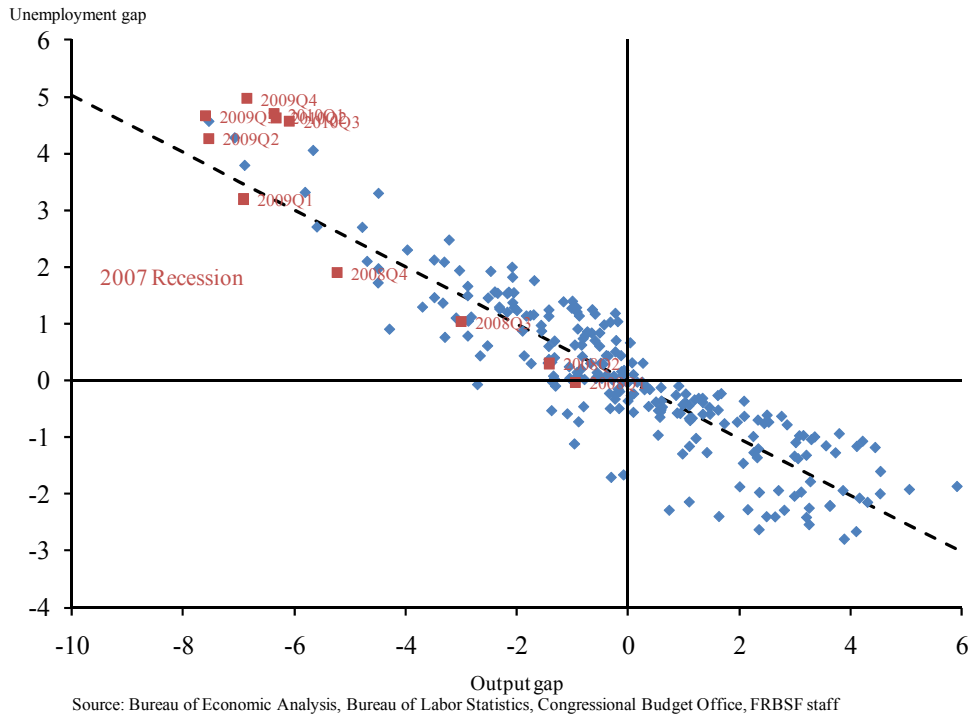


Figure 15. Decomposition of GDP growth since start of 2007 recession.

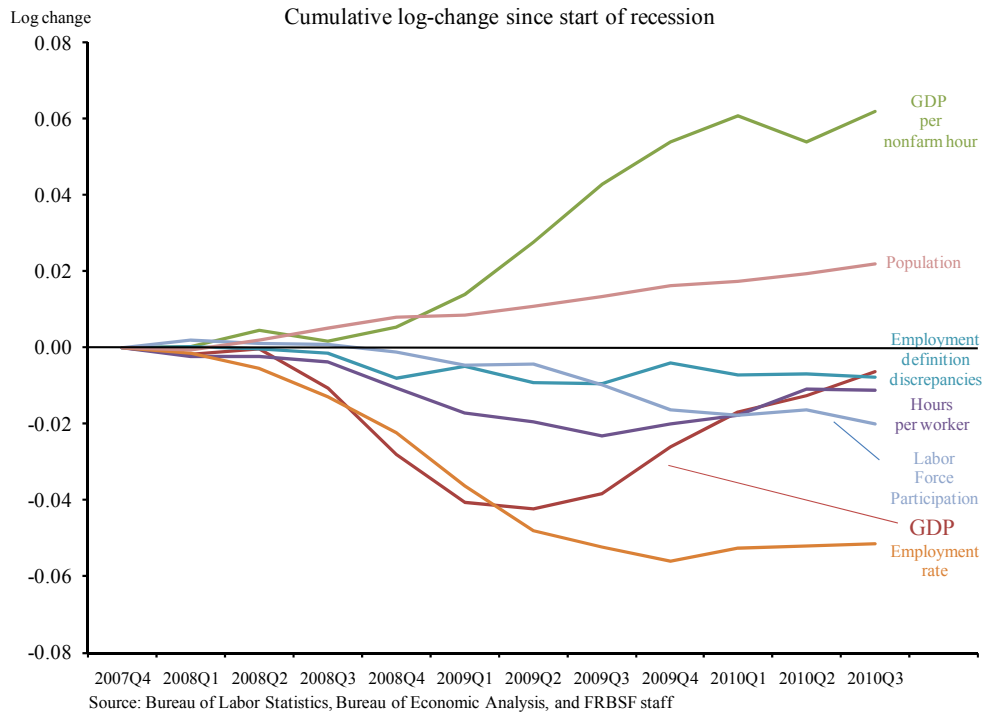


Figure 16. Current situation and factors shifting BC and JC curves.

