FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Which Industries are Shifting the Beveridge Curve?

Regis Barnichon
Federal Reserve Board of Governors

Michael Elsby
University of Michigan and NBER

Bart Hobijn
Federal Reserve Bank of San Francisco

Ayşegül Şahin
Federal Reserve Bank of New York

October 2011

Which Industries are Shifting

the Beveridge Curve?

REGIS BARNICHON, MICHAEL ELSBY, BART HOBIJN, AND AYŞEGÜL ŞAHIN

This version: October 24, 2011.

Abstract

The negative relationship between the unemployment rate and the job openings rate, known as the Beveridge curve, has been relatively stable in the U.S. over the last decade. Since the summer of 2009, in spite of firms reporting more job openings, the U.S. unemployment rate has not declined in line with the Beveridge curve. We decompose the recent deviation from the Beveridge curve into different parts using data from the Job Openings and Labor Turnover Survey (JOLTS). We find that most of the current deviation from the Beveridge curve can be attributed to a shortfall in hires per vacancy. This shortfall is broad-based across all industries and is particularly pronounced in construction, transportation, trade, and utilities, and leisure and hospitality. Construction alone accounts for more than half of the Beveridge curve gap.

Although economic activity in the U.S. economy has grown, be it slowly, since the summer of 2009, the unemployment rate has remained stubbornly high. This continued high level of unemployment is especially puzzling in light of the fact that, during the same period, U.S. employers have started to post substantially more vacancies.[1]

Historically, there has been a tight negative relationship between the unemployment rate and the job openings rate. This relationship is known as the Beveridge curve. However, since the summer of 2009, this relationship seems to have broken down.[2] In August 2011 the unemployment rate was 2.3 percentage points above its level implied by the Beveridge curve.

In this study we decompose the gap between the actual unemployment rate and that implied by the Beveridge curve into different parts using data from the Job Openings and Labor Turnover Survey (JOLTS).

The Beveridge curve can be interpreted as the job openings rate at which the current unemployment rate would be in its flow steady state. This steady state of the unemployment rate is that for which, at the current rates at which workers move between employment, unemployment, and non-participation, the unemployment rate would not change. In order to implement our decomposition, we construct the Beveridge curve by solving a fitted flow-steady-state equation using data on job openings, hires, layoffs and quits from JOLTS as well as data on entry and exit from the labor force from the CPS. The Beveridge curve that we construct in this way fits the pre-2007-recession data very well.

We then use the estimated flow-steady-state equation to derive an approximate additive decomposition of deviations of the unemployment rate from the Beveridge curve into parts due to hires per vacancy, layoffs, quits, as well as labor force entry and exit. We find that the current Beveridge curve gap is almost fully attributable to an unexplained shortfall in the vacancy yield, i.e. the number of hires per vacancy, while a lower than expected quits rate reduces the gap.

We further decompose the Beveridge curve gap to consider which industries account for the unexplained decline in the vacancy yield, as well as for the behavior of the quit and layoff rates. The result of this industry decomposition is that the shortfall in the vacancy yield is widespread across all industries. The vacancy yield deficit is particularly pronounced in construction, manufacturing, trade and transportation, leisure and hospitality, as well as in the industries not classified in JOLTS. From June 2011 through August 2011, the difference between the observed

2

and predicted hires per vacancy in construction alone accounted for more than 0.8 percentage point of the 2.6 percent by which the actual unemployment rate exceeded that implied by the Beveridge curve.

Of course, our decomposition is merely an accounting exercise and does not directly provide any explanations for the deviations of the flow rates from their predicted levels. We discuss some potential explanations as well as how the shift in the Beveridge curve may translate into a higher natural rate of unemployment in the final part of this article.

JOLTS-based Beveridge curve

Due to the high level of worker and job flows, the U.S. labor market has such fast dynamics that it very quickly tends towards its flow steady state.[3] Given this observation about U.S. labor markets, the Beveridge curve is often interpreted as the vacancy rate at which, for a given unemployment rate, the unemployment rate is in its steady state. We also use a similar interpretation in this paper. However, contrary to most studies of the Beveridge curve, which focus on the flow rates derived from labor market status flows, we use the JOLTS hiring, layoffs, and quits rates for defining steady-state unemployment.[4]

The unemployment rate, u_t, is in steady state whenever the growth rate of the labor force, which we denote by $g_t^{(lf)}$, equals the growth rate of employment, denoted by $g_t^{(e)}$.[5] To derive a Beveridge curve from this steady-state condition, we have to relate these growth rates to the vacancy and unemployment rates. We do so by considering the gross flows that underlie these growth rates.

First, the growth rate of the labor force $g_t^{(lf)}$ is given by n_t, the number of people who enter the labor force in a month as a fraction of the number of people in the labor force at the beginning of the month, minus x_t, the number of people who exit the labor force in a month divided by the number of people in the labor force at the beginning of the month. We measure both n_t and x_t from the CPS labor market status flows.

Secondly, the growth rate of employment, $g_t^{(e)}$, can be measured using JOLTS.[6] Employment growth equals hires as a fraction of employment at the beginning of the month, minus quits and layoffs as a fraction of that same employment level.[7] This insight allows us to write the growth rate

of employment in terms of the hires, quits, and layoffs rates reported in JOLTS. We denote the latter two by q_t and l_t.

It is possible to rewrite the hires rate in terms of the job openings rate, v_t, and the number of hires per vacancy, h_t. We do so to be consistent with the prevailing methodology of estimating an empirical matching function, which focuses on the vacancy yield as a function of the ratio of the job openings and unemployment rates, for the construction of the empirical Beveridge curve.[8]

Now that we have measures of the gross flows that drive both $g_t^{(lf)}$ and $g_t^{(e)}$, we introduce the job openings rate and unemployment rate into the steady-state condition for unemployment by estimating how the five flows we measure, i.e. n_t, x_t, h_t, l_t, and q_t, depend on the ratio of the job openings rate to the unemployment rate, the (v/u)-ratio. In a strong labor market there are relatively few unemployed and many vacancies and the (v/u)-ratio is high; the reverse is the case in a weak labor market. Therefore, we use the (v/u)-ratio as a cyclical indicator of labor market tightness in order to capture the "normal" cyclical behavior of the five flows.[9] Specifically, we regress the log-level of each of the flows on the log of the (v/u)-ratio. For these regressions we use monthly seasonally adjusted data that cover the pre-recession sample starting from December 2000, the beginning of JOLTS, until the beginning of the recession. The results of the regressions are reported in Table 1.

The table reveals the following facts. First of all, the vacancy yield moves most closely together with the (v/u)-ratio. Fluctuations in labor market tightness, i.e. the (v/u)-ratio, explain about 88% percent of the fluctuations in the vacancy yield. The estimated elasticity of -0.41 is in line with commonly used models of search frictions in the labor market that assume that the probability of filling a vacancy is decreasing in the (v/u)-ratio. Quits depend negatively on labor market tightness. When there are many job openings, workers are more likely to make job-to-job transitions.[10] This is reflected by the 0.26 elasticity of the quit rate with respect to labor market tightness. The variation in the latter explains two-thirds of the variation in the quits rate. Layoffs as a fraction of employment tend to decrease in tight labor markets. However, the link between labor market tightness and layoffs is less than for the other two employment flows. In fact, layoffs tend to lead movements in the (v/u)-ratio (panel (b), Figure 1).

The three panels in Figure 1 plot the actual and fitted vacancy yields, layoffs rates, and quits rates.[11] The vertical line represents the end of the regression sample. Hence the fitted values to the

right of this line are actual forecasts of these rates based on the pre-recession relationship between these rates and the (v/u)-ratio. Besides the wave of layoffs at the end of 2008 and the beginning of 2009, Figure 1 shows two more large deviations of the actual rates from the fit implied by their historical patterns. The first is that the August 2011 quits rate is about 15% below that predicted based on the level of labor market tightness. Hence, workers hang on to their jobs even more than one would expect based on the current weakness in the labor market. The second, and most profound, deviation is that hires per vacancy are about 30% less than predicted at the current (v/u)-ratio. Hence the JOLTS-based employment growth flows are deviating substantially from their predicted values based on the current degree of labor market tightness.

Such large deviations are not observed for the flows that underlie labor force growth, n_t and x_t. This is largely because, as can be seen from the last two columns of Table 1, fluctuations in labor market tightness do not have any significant explanatory power for these flows.

The estimated flow rate functions now allow us to define the JOLTS-based Beveridge curve as follows. For a given level of job openings rate, v, the Beveridge curve is given by the level of the unemployment rate for which our estimated flow rate equations imply that unemployment is in steady state--i.e. such that the fitted labor force growth rate equals the fitted employment growth rate.

Figure 2 plots the actual and fitted Beveridge curves. We split the observed Beveridge curve into two parts. The blue points are the pre-recession observations based on which the flow rates that underlie the curve are fitted. The orange points are the observations from December 2007 onwards. As can be seen from the figure, the estimated Beveridge curve does not only provide a good fit for the pre-recession observations but also for some of the observations during the recession.

Note that our methodology does not rely on any regression of the job openings rate on the unemployment rate, but instead infers the Beveridge curve using a flow-steady-state unemployment relationship from the job openings rate based on separately fitted flow rates reported in JOLTS.

The real anomaly here is the deviation from the Beveridge curve that occurred during the recovery. Since July 2009 the job openings rate has risen from 1.7% to 2.3%. However, during that same period the unemployment rate has not fallen as much as the Beveridge curve would imply. It actually initially increased from 9.4% to 10.1% and has since come down to 9.1% in August 2011. The result is that, at the August 2011 job openings rate, the actual unemployment rate was 2.3

percentage points higher than the one implied by the Beveridge curve. We refer to this as the Beveridge curve gap. In the rest of this study we decompose this gap into different parts that contribute to it.

Decomposing the Beveridge curve gap

For a given job openings rate, we define the Beveridge curve as the unemployment rate for which our fitted flow rates imply that the unemployment rate is in steady state. However, as we show in Figure 1, we are seeing large deviations of the flow rates from their fitted levels. This is especially true for the vacancy yield and quits rate.

In this section we analyze to what extent these deviations from the fitted flow rates contribute to the Beveridge curve gap. In order to do so, we use an approximate additive decomposition.[12] We use the approximation technique to find the answer to the following type of question.

The August 2011 level of the vacancy yield, h_t, is 27% percent lower than that implied by the estimates reported in Table 1. At the current job openings rate, to what extent is the rise in the steady-state unemployment rate relative to the fitted Beveridge curve explained by the 27% shortfall in hiring?

The answer to this question can be interpreted as the part of the Beveridge curve gap attributable to the current shortfall in the vacancy yield relative to its historical pattern. We can answer this type of question not only for the vacancy yield, but also for the layoffs rate, quits rate, and for the labor force entry and exit rates.

The result is an approximate additive decomposition of the Beveridge curve gap. It decomposes the gap into parts due to deviations of the five actual flow rates from their fitted values as well as a residual part. The residual reflects two main sources of approximation error. The first is that, in order for our decomposition to be additive, we are using a linear approximation. The second is that the actual unemployment rate might not be in steady state.[13]

To smooth out some of the month-to-month fluctuations we report the results of our decomposition in terms of three-month moving averages. The decomposition of the average Beveridge curve gap in the latest three months in our sample is reported in Table 2. On average

from June 2011 through August 2011 the unemployment rate was 2.6 percentage points higher than the level implied by the Beveridge curve.

On the labor force growth side, the contributions of the gross flows offset each other. The Beveridge curve gap is due to the flows that drive net employment growth.

Because both layoffs and quits are below their fitted values we find that they are less than expected and actually contribute negatively to the Beveridge curve gap, by 1.4 and 0.9 percentage points respectively. That is, if layoffs and quits were at their expected levels based on the estimates in Table 1 then more persons would flow into unemployment and this would raise the unemployment rate. The reduction in these flows suppresses the unemployment rate and thus reduces the Beveridge curve gap.

The negative contributions of the layoffs and quits rates are more than offset by the contribution of the shortfall in the vacancy yield. When fewer persons are hired out of unemployment for a given job openings rate more persons remain unemployed and this raises the unemployment rate. This is why the 27% shortfall in the vacancy yield, depicted in panel (a) of Figure 1, translates into a 4.0 percentage point positive contribution to the Beveridge curve gap. Thus, the 2.6 percentage point gap is more than fully accounted for by the unprecedented shortfall in hires per job opening.

Finally, the residual part is 0.9 percentage points. This is partly due to the linear approximation method used and partly due to the unemployment rate being above its flow steady state.

The two panels of Figure 3 show the decomposition of the Beveridge curve gap over time, from the beginning of JOLTS to the latest observation in our sample. Panel (a) of the figure shows the Beveridge curve gap and its employment growth flow determinants. The shortfall in the vacancy yield started before any significant increase in the Beveridge curve gap and even before the spike in the layoff rate signaled the beginning of the major downturn in the labor market. Initially, the reduction in hiring per job opening was mostly offset by a decline in quits. However, during the second half of 2008, at the height of the financial crisis, the shortfall in hires per vacancy increased so rapidly that it was not offset by quits but instead caused a very negative residual.

This can be seen from panel (b) of Figure 3, which plots the contributions of the labor force growth flows and of the residual. This suggests that, during the second half of 2008, the unemployment rate was substantially below its steady-state value. That is, labor market fundamentals were deteriorating so quickly that, in spite of the rapid dynamics of the U.S. labor

market, the observed unemployment rate took about half a year to catch up with the new steady state.

To summarize, we have established that, at the aggregate level, deviations from employment growth flows measured in JOLTS, account for the bulk of the Beveridge curve gap and that the gap is mostly fueled by a shortfall in hires per vacancy. The next step is to consider which industries contribute most to the deviations of the observed vacancy yield, quits rate, and layoffs rate from their fitted values.

Decomposing deviations from fitted flow rates by industry

In order to decompose the deviations from fitted flow rates by industry, we first construct predicted industry level vacancy yields, layoffs rates, and quits rates. Just like for the aggregate flow rates, we estimate the industry-specific flow rates as a function of the ratio of the vacancy rate and the unemployment rate. To control for specific effects in the labor markets within which the industries operate we also include the ratio of job openings in the industry, $v_{i,t}$ to the number of unemployed persons who where last employed in the industry, $u_{i,t}$. Here i indexes the industries.[14]

In particular, we use data for the seven main industries for which JOLTS reports seasonally adjusted job openings, layoffs, and quits. These industries are: (*i*) construction, (*ii*) manufacturing, (*iii*) trade, transportation and utilities, (*iv*) professional and business services, (*v*) education and health services, (*vi*) leisure and hospitality, and (*vii*) government. We construct data for the residual industry, "other", by subtracting the data for the 7 industries from those reported for the total economy.

Mirroring our aggregate analysis, we use the pre-recession sample to fit the flow rates. Table 3 reports the estimates of the parameters for the fitted flow rates by industry. Not surprisingly, since the aggregate flow rates are share-weighted averages of the industry flow rates, the main picture from Table 1 also applies industry-by-industry.

The vacancy yield is most responsive to the degree of labor market tightness, except for the government sector. For each of the private sectors, labor market tightness explains more than two-thirds of the variance of their vacancy yields. Quits also tend to respond quite elastically to the strength of the labor market, except in construction and manufacturing. In all other sectors workers

are more likely to quit during a strong labor market. The strength of the labor market explains between approximately one quarter and one half of the variance of industry quits rates. This fraction is lowest for layoffs. A notable exception is manufacturing.

The results of our decomposition by industry of the deviations of the aggregate vacancy yield, layoffs rate, and quits rate from their fitted values are presented in Table 4. This table is split up in two parts.

Part A of the table shows how much the actual flow rates deviate from their fitted values, both for the total economy as well as for each of the industries. Reported are the average deviations for the last three months in our sample. From column II it can be seen that in all sectors fewer workers quit than would be expected by the current strength of the job market. This shows that the aggregate 13% shortfall in the quits rate is broad-based. Quits are especially low in construction and manufacturing. Column III reveals that in all sectors hires per vacancy are lower than implied by the regression results reported in Table 3. Hires per vacancy are especially low in construction, manufacturing, and the unclassified sectors. The latter include finance and real estate. Finally, the picture for layoffs rates is mixed. The most notable feature of current layoffs rates is the high level of layoffs in manufacturing and construction.

The bottom line of Table 4, and of this paper, is presented in part B. For each sector, it shows how much the deviations of its layoffs rate, quits rate, and vacancy yield from their predicted values contribute to the aggregate Beveridge curve gap. Column VII of the table adds the contributions of these three rates up for each of the industries.[15]

From this part it can be seen that the biggest contributors to the Beveridge curve gap are the vacancy yield deficits in construction, trade and transportation, leisure and hospitality, and "other" sectors. This last category contains finance, insurance, and real estate. Education and health services as well as professional and business services do not seem to contribute much to the Beveridge curve gap. In fact, the latter industry actually reduces the gap. The industry that contributes most to the gap is construction, which is shifting the Beveridge curve right by more than a percentage point.

Interpretation of results

Thus far, we have shown that a broad-based shortfall in the vacancy yield, in particular in construction, is the main culprit behind the current Beveridge curve gap. Here we focus on the potential causes of the low number of hires per vacancy.

The first potential cause is a mismatch between job openings and the unemployed. If the pool of unemployed persons has very different qualifications than those required for the job openings posted then it would be harder to fill these openings relative to other times when this degree of mismatch is less severe. The problem is that currently the only property of a job opening reported in JOLTS is which industry is posting it. Measures of mismatch based on JOLTS data show that industry mismatch initially increased at the onset of the recession but then rapidly reverted to levels only slightly higher than before the recession.[16]

A second possible reason for the shortfall in hires per vacancy is that firms' recruiting intensity declined after 2007, i.e., that firms made less effort (including advertising, screening and wage offers) to fill open vacancies.[17] A substantial number of posted job openings are for replacement hires and given the current level of weak demand and economic uncertainty firms might simply put less effort into recruiting people for such open positions.

Moreover, because many workers are hired without the formal posting of a vacancy, the vacancy yield could decline with a change in the composition of hires. If there is a particular decline in labor demand for jobs with informal hiring, the number of hires per vacancy will decline. This is especially true for jobs in construction, where informal hiring is particularly prevalent. For example, if contractors do not post vacancies to hire craftsmen to work on construction sites but post vacancies to hire bookkeepers, then if there is a lull in building activity and few craftsmen are hired, hires per vacancy will decline because of a change in the composition of hires.

Another possible reason for the shortfall in hires per vacancy might be the search intensity of the unemployed. For example, if extensions of unemployment benefits (UI) reduce the time an average unemployed person spends looking for a job, or if they make workers pickier about which jobs to accept, the extended unemployment coverage program could have lowered hires per vacancy in the 2008-2009 recession. This explanation relies on the effect of (UI) on the incentives for the unemployed to search for and accept job offers.[18]

The potential increase in mismatch is of most concern because it suggests that the shift in the Beveridge curve might be very persistent, potentially leading to an increase in the natural rate of unemployment.[19] For this discussion about the natural rate, several points are important to realize.

First, permanent changes in the constant terms in the regressions reported in Table 1 do not imply a uniform rightward shift of the Beveridge curve. To illustrate this, we have constructed a new Beveridge curve assuming that the average deviations of these intercepts during the last three months in our sample from their estimated values are permanent.

This hypothetical "new" Beveridge curve is plotted in Figure 4. At the 2.3% job openings rate that prevailed in August 2011, the outward shift of the "new" Beveridge curve relative to the fitted historical one is 2.1 percentage points. At a 2.9% percent job openings rate, the average over the pre-2007-recession period, the shift is 1.9 percentage points. At a 4.0% job openings rate the shift is 1.3 percentage points. Also note that the 9.1% unemployment rate in August 2011 was 0.3 percentage points higher than implied by this "new" Beveridge curve. This suggests that the unemployment rate in August 2011 might have been a bit above its flow-steady-state value.

Moreover, one cannot solely use the shift in the Beveridge curve to infer the size of the increase in the natural rate of unemployment. What matters besides the shift in the Beveridge curve is the change in the natural vacancy rate. For the pre-2007-recession period estimates of the natural rate of unemployment are generally around 5%.[20] On the fitted Beveridge curve this coincides with a 3% natural rate of job openings. On the "new" Beveridge curve, this 3% job openings rate coincides with a 6.6% percent unemployment rate. However, most equilibrium models of frictional unemployment[21] suggest that the natural vacancy rate increases when the Beveridge curve shifts out. Suppose that the new natural job openings rate were 3.5% instead of 3%. On the "new" Beveridge curve this would imply a 5.7% natural rate of unemployment rather than the 6.6% associated with the lower vacancy rate.

Hence, one has to be very careful when trying to translate the measured shift in the Beveridge curve in terms of a shift in the natural rate of unemployment. First of all, the current shift in the Beveridge curve is in all likelihood largely temporary. Secondly, even if it was persistent, the shift at the current level of the job openings rate is higher than at any plausible new natural vacancy rate. Finally, there are no good estimates of how the natural job openings rate changes when the

Beveridge curve shifts outward. Such estimates, of course, will become available when we have several more decades of JOLTS data to analyze.

Conclusion

In August 2011, the unemployment rate exceeded the level implied by its historical relationship with the job openings rate by 2.3 percentage points. This is what we call the Beveridge curve gap.

In this paper we constructed a Beveridge curve based on estimated relationships between flow rates reported in JOLTS and the job openings to unemployment ratio, which we used as a measure of labor market tightness. This Beveridge curve fits the pre-2007-recession data remarkably well. Moreover, the estimated flow rates allow us to decompose deviations from the Beveridge curve into parts due to deviations of the job-flow rates from their predicted levels.

Our decomposition reveals that most of the current deviation from the Beveridge curve can be attributed to a shortfall in the vacancy yield, which measures hires per vacancy. This shortfall is broad-based across all industries and is particularly pronounced in construction, transportation, trade and utilities, and leisure and hospitality, and industries not explicitly classified in JOLTS.

Whether this shortfall is due to (*i*) mismatch between job openings and unemployed workers, (*ii*) reduced recruitment effort by employers, (*iii*) a change in the composition of vacancies and hires, or (*iv*) reduced search intensity of unemployed persons, is hard to parse out from the data currently available in JOLTS.

More information about the regional and occupational composition of job openings, hires, and quits would go a long way in helping us distinguish between these issues. Moreover, additional data on search intensity of employers, like for example the number of job offers made, would help us better understand the effort with which they pursue filling the job openings they have.

---

[1] A more detailed discussion of this development is in Michael Elsby, Bart Hobijn and Ayşegül Şahin. "The Labor Market in the Great Recession," *Brookings Papers on Economic Activity*, Spring 2010, 1-48.

[2] See Mark deWolf and Katherine Klemmer. "Job Openings, Hires, and Separations Fall During the Recession," *Monthly Labor Review*, 36-44, May, 2010 for a discussion of JOLTS evidence during the 2007 recession.

[3] See Robert Shimer, "Reassessing the Ins and Outs of Unemployment," *NBER Working Paper No. 13421*, 2007.

[4] Studies that use the labor market status flows to define the Beveridge Curve are, among others, Olivier Blanchard and Peter Diamond, "The Beveridge curve," *Brookings Paper on Economic Activity*, 1:1-60, 1989, and Regis Barnichon and Andrew Figura. "What Drives Movements in the Unemployment Rate? A Decomposition of the Beveridge curve," *Finance and Economics Discussion Series 2010-48*. Washington: Board of Governors of the Federal Reserve System, 2010.

[5] The technical derivations behind this result as well as other results in this article are available in an Appendix in the working paper version: Regis Barnichon, Michael Elsby, Bart Hobijn, and Ayşegül Şahin, "Which Industries are Shifting the Beveridge Curve?," *FRBSF Working Papers 2010-32*, 2010.

[6] Throughout, we do not distinguish between the employment concepts in the Household Survey and the Establishment Survey. JOLTS is based on the latter. We discuss in detail how we deal with the definitional differences between these two concepts in the Appendix.

[7] Beside quits and layoffs, the JOLTS data contain a third class of separations, which includes retirements, deaths, emigration, and other types of separations that are neither considered a layoff or a quit. Because layoffs are not reported by detailed industry in JOLTS, we have added these other separations to layoffs. That is, our measure of layoffs is total separations minus quits.

[8] See Barbara Petrongolo and Christopher A. Pissarides. "Looking into the Black Box: A Survey of the Matching Function." *Journal of Economic Literature*, Vol. 39, No. 2, pp. 390-431, 2001, for a survey of matching function estimates.

[9] JOLTS only started in December 2000. So, we can only use one business cycle, 2001-2007, to infer the "normal" cyclical behavior of the flows.

[10] This is, for example, also documented by Bruce Fallick and Charles A. Fleischman. "Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows," *Finance and Economics Discussion Series 2004-34,* Washington: Board of Governors of the Federal Reserve System, 2004.

[11] For the transformation from the logarithm to the level plotted in the figure we take the exponent of the fitted log flow rate and then correct for Jensen's inequality by assuming the residual is normally distributed.

[12] This decomposition is based on a log-linear approximation of the fitted Beveridge curve.

[13] Other sources of approximation error are the definitional differences between the Household and Establishment Survey employment concepts and the assumption that the parameters in the estimated flow-rate regressions are constant.

[14] For example, how easy it is to hire a nurse does not only depend on how tight the overall labor market is but, more importantly, on how many unemployed nurses there are compared to job openings in health services.

[15] The "Aggregation" line in Table 4 reflects that the weighted sum of the fitted flow rates does not have to equal the aggregate flow rates. The "Composition" line measures how much the aggregate flow rates would have changed due to the change in the industry composition of employment and job openings even if each of the industry flow rates would be equal to its fitted values.

[16] For a detailed analysis of mismatch in the U.S. see Ayşegül Şahin, Joseph Song, Giorgio Topa, and Gianluca Violante. "Measuring Mismatch in the U.S. Labor Market," mimeo, Federal Reserve Bank of New York, 2011.

[17] This explanation is put forward by Steven Davis, Jason Faberman, and John Haltiwanger. "The Establishment-Level Behavior of Vacancies and Hiring," *NBER Working Paper No. 16265*, 2010.

[18] Rob Valletta and Katherine Kuang. "Extended Unemployment and UI Benefits," *FRBSS Economic Letter 2010-12*, 2010, find that the effect of UI on the duration of unemployment is relatively small, however.

[19] An increase in the natural rate due to mismatch has been discussed by, among others, Narayana Kocherlakota, "Inside the FOMC," Speech at Marquette, Michigan, August 17, 2010.

---

[20] For estimates of the natural rate of unemployment see Congressional Budget Office, *The Budget and Economic Outlook*, February 2011.

[21] Like the textbook model in Christopher A. Pissarides, *Equilibrium Unemployment Theory*, Cambridge, MA: MIT Press, 2000.

Appendix: Mathematical and data details

Construction of n_t and x_t from CPS labor market status flows

We construct n_t as the sum of all worker flows in a month from non-participation and from not being part of the civilian working-age population to employment or unemployment. We divide these flows by the size of the labor force within that month. Similarly, we construct x_t as the sum of all worker flows from employment and unemployment to non-participation and to not being part of the civilian working-age population.

Derivation of steady-state condition for the unemployment rate:

Because the labor force in month t, denoted by LF_t, equals the sum of the number of employed, E_t, and the number of unemployed, U_t, the change in the number of unemployed persons can be written as the change in the labor force minus the change in the number of employed persons. That is

$$U_t - U_{t-1} = \Delta U_t = \Delta LF_t - \Delta E_t. \tag{1}$$

Normalizing both sides of this expression by the labor force and using the fact that the unemployment rate, u_t, is the ratio of the number of unemployed persons and the size of the labor force, we can write

$$\frac{LF_t}{LF_{t-1}} u_t - u_{t-1} = \frac{\Delta LF_t}{LF_{t-1}} - \frac{E_{t-1}}{LF_{t-1}} \frac{\Delta E_t}{E_{t-1}}. \tag{2}$$

Defining the growth rates of the labor force and of employment as

$$g_t^{(lf)} = \Delta LF_t / LF_{t-1} \text{ and } g_t^{(e)} = \Delta E_t / E_{t-1} \tag{3}$$

respectively, we can write (2) as

$$\begin{aligned}
\left(1 + g_t^{(lf)}\right) u_t - u_{t-1} &= g_t^{(lf)} - (1 - u_{t-1}) g_t^{(e)} \\
\left(1 + g_t^{(lf)}\right) u_t - \left(1 + g_t^{(lf)}\right) u_{t-1} &= (1 - u_{t-1}) g_t^{(lf)} - (1 - u_{t-1}) g_t^{(e)} \\
\left(1 + g_t^{(lf)}\right) (u_t - u_{t-1}) &= (1 - u_{t-1}) \left(g_t^{(lf)} - g_t^{(e)}\right).
\end{aligned} \tag{4}$$

This simplifies to

$$u_t - u_{t-1} = \frac{1 - u_{t-1}}{1 + g_t^{(lf)}} \left(g_t^{(lf)} - g_t^{(e)}\right). \tag{5}$$

Hence, for the change in the unemployment rate to be zero, that is for unemployment to be in steady state, it must be the case that $g_t^{(lf)} = g_t^{(e)}$. Thus, the unemployment rate is in steady state whenever the growth rate of the labor force equals the growth rate of employment.

Hires as a fraction of the employment level at the beginning of the month.

Denote the number of hires in a month by H_t and the number of job openings reported by V_t. Note that the job openings rate, v_t, is defined as the number of job openings as a fraction of the sum of employment and job openings. The vacancy yield, h_t, is hires per job opening. Given these definitions we can write

$$\frac{H_t}{E_{t-1}} = \frac{E_t}{E_{t-1}} \frac{E_t+V_t}{E_t} \frac{V_t}{E_t+V_t} \frac{H_t}{V_t} = \left(1 + g_t^{(e)}\right) \frac{v_t}{1-v_t} h_t. \tag{6}$$

This gives us hires as a fraction of the employment level at the beginining of the month in terms of the employment growth rate, the job openings rate, and the vacancy yield.

Growth rate of employment in terms of JOLTS flow rates.

The JOLTS layoffs and quits rates are defined as a fraction of the current month's employment level. That is, if L_t is the number of layoffs and Q_t the number of quits, then the layoff rate can be written as $l_t = L_t/E_t$ and the quits rate as $q_t = Q_t/E_t$. Given these definitions the growth rate of employment equals

$$g_t^{(e)} = \frac{E_t-E_{t-1}}{E_{t-1}} = \frac{H_t-L_t-Q_t}{E_{t-1}} = \frac{H_t}{E_{t-1}} - \frac{E_t}{E_{t-1}} q_t - \frac{E_t}{E_{t-1}} l_t = \left(1 + g_t^{(e)}\right)\left(\frac{v_t}{1-v_t} h_t - q_t - l_t\right). \tag{7}$$

Solving the above expression with respect to the employment growth rate yields

$$g_t^{(e)} = \left\{\frac{v_t}{1-v_t} h_t - l_t - q_t\right\} / \left\{1 - \frac{v_t}{1-v_t} h_t + l_t + q_t\right\}, \tag{8}$$

which expresses the growth rate of employment in terms of the job openings rate, the quits rate, the layoffs rate, and the vacancy yield which can all be calculated using data from JOLTS.

Differences between CPS employment and Payroll employment.

The employment concept used to construct the unemployment rate is based on the Current Population Survey while the employment concept used in JOLTS is based on the Establishment Survey. These concepts differ both conceptually as well as in terms of their sampling error. As a

result, the employment growth implied by the Establishment Survey does not always coincide with that on which the unemployment statistics are based. Here, we briefly describe how we account for these differences in our calculation of the fitted Beveridge curve. Just like above, we denote CPS employment as E_t. We denote payroll employment as E_t^*. Because of the definitional discrepancies, $\Delta E_t \approx H_t - L_t - Q_t$. We take care of this approximation error as follows. We write

$$g_t^{(e)} = \left(\frac{\Delta E_t}{E_{t-1}} - \frac{E_{t-1}^*}{E_{t-1}} \frac{\Delta E_t^*}{E_{t-1}^*} \right) + \left(1 + g_t^{(e)} \right) \frac{E_t^*}{E_t} \left(\frac{v_t}{1-v_t} h_t - q_t - l_t \right). \tag{9}$$

When we define relative size of the employment measures and the adjusted difference in the two employment growth rates measures respectively as

$$z_t = \frac{E_t^*}{E_t} \text{ and } d_t = \left(\frac{\Delta E_t}{E_{t-1}} - z_{t-1} \frac{\Delta E_t^*}{E_{t-1}^*} \right) \tag{10}$$

then we use that

$$g_t^{(e)} = \frac{d_t + z_t \left(\frac{v_t}{1-v_t} h_t - q_t - l_t \right)}{1 - z_t \left(\frac{v_t}{1-v_t} h_t - q_t - l_t \right)}. \tag{11}$$

We then solve for the Beveridge curve for the particular values $z_t = \bar{z}$ and $d_t = \bar{d}$, which are their sample average over the pre-recession period.

Log-linear decomposition of the Beveridge curve gap.

Given the regressions

$$lnf_t = \widehat{m}^{(f)} + \widehat{\alpha}^{(f)} ln \left(\frac{v_t}{u_t} \right) + \varepsilon_t^{(f)}, \text{ where } f = n, x, h, l, \text{ or } q, \tag{12}$$

the results of which are reported in Table 1, the fitted Beveridge curve that we consider defines the unemployment rate u^* as an implicit function of the job openings rate, v, through the fitted steady state condition

$$0 = \widehat{m}^{(n)} \left(\frac{v}{u^*} \right)^{\widehat{\alpha}^{(n)}} - \widehat{m}^{(x)} \left(\frac{v}{u^*} \right)^{\widehat{\alpha}^{(x)}} - \frac{\bar{d} + \bar{z} \left(\frac{v}{1-v} \widehat{m}^{(h)} \left(\frac{v}{u^*} \right)^{\widehat{\alpha}^{(h)}} - \widehat{m}^{(l)} \left(\frac{v}{u^*} \right)^{\widehat{\alpha}^{(l)}} - \widehat{m}^{(q)} \left(\frac{v}{u^*} \right)^{\widehat{\alpha}^{(q)}} \right)}{1 - \bar{z} \left(\frac{v}{1-v} \widehat{m}^{(h)} \left(\frac{v}{u^*} \right)^{\widehat{\alpha}^{(h)}} - \widehat{m}^{(l)} \left(\frac{v}{u^*} \right)^{\widehat{\alpha}^{(l)}} - \widehat{m}^{(q)} \left(\frac{v}{u^*} \right)^{\widehat{\alpha}^{(q)}} \right)}. \tag{13}$$

The aim of our decomposition is to figure out how this implicit function would change when the $\hat{m}$'s change. This can be done through the application of the implicit function theorem. This yields that

$$
\left[\hat{\alpha}^{(n)}\frac{\hat{n}}{u^*} - \hat{\alpha}^{(x)}\frac{\hat{x}}{u^*} - \frac{\bar{d} + \bar{z}\left(1 + \hat{g}_t^{(e)}\right)}{1 - \bar{z}\left(\frac{v}{1-v}\hat{h} - \hat{l} - \hat{q}\right)}\right](u - u^*) \approx
$$
$$
\hat{n}\left(\frac{m^{(n)} - \hat{m}^{(n)}}{\hat{m}^{(n)}}\right) - \hat{n}\left(\frac{m^{(x)} - \hat{m}^{(x)}}{\hat{m}^{(x)}}\right) \qquad (14)
$$
$$
-\frac{\bar{d} + \bar{z}\left(1 + \hat{g}_t^{(e)}\right)}{1 - \bar{z}\left(\frac{v}{1-v}\hat{h} - \hat{l} - \hat{q}\right)}\left[\frac{v}{1-v}\hat{h}\left(\frac{m^{(h)} - \hat{m}^{(h)}}{\hat{m}^{(h)}}\right) - \hat{l}\left(\frac{m^{(l)} - \hat{m}^{(l)}}{\hat{m}^{(l)}}\right) - \hat{q}\left(\frac{m^{(q)} - \hat{l}^{(q)}}{\hat{m}^{(q)}}\right)\right].
$$

All variables denoted with a $\widehat{}$ in this equation refer to their values on the fitted Beveridge curve.

This linear approximation allows us to write the deviation of the unemployment rate from the Beveridge curve, $(u - u^*)$, as a linear function of the percentage deviations of the actual flow rates from their fitted values. These percentage deviations measure $\left(m^{(f)} - \hat{m}^{(f)}\right)/\hat{m}^{(f)}$ for $f = n, x, h, l,$ or q.

Decomposition of deviation from fitted flow rates by industry.

Here we derive the decomposition for the vacancy yield. After the derivation we briefly discuss how it can also be applied to the layoffs and quits rates. The most important thing to realize for this decomposition is that the aggregate vacancy yield is a share-weighted average of the industry-specific vacancy yields, where the shares are the industry's share in total vacancies.

For this derivation we denote the aggregate vacancy yield again by h and its fitted value by $\hat{h}$. We use i as the industry index and h_i for the industry-specific vacancy yield and $\hat{h}_i$ for its fitted value. The share of industry i in total job openings is denoted by s_i and its sample average over the pre-recession period is $\bar{s}_i$. This allows us to write

$$
h = \sum_i s_i h_i, \qquad (15)
$$

such that

$$
\begin{aligned}
h - \widehat{h} &= \sum_i (s_i - \bar{s}_i) h_i + \sum_i \bar{s}_i \left(h_i - \widehat{h}_i \right) + \sum_i \bar{s}_i \left(\widehat{h}_i - \widehat{h} \right) \\
&= \sum_i (s_i - \bar{s}_i) h_i - \sum_i (s_i - \bar{s}_i) \widehat{h}_i + \underbrace{\sum_i \bar{s}_i \left(h_i - \widehat{h}_i \right)} + \underbrace{\left(\sum_i s_i \widehat{h}_i - \widehat{h} \right)} \\
&= \underbrace{\sum_i (s_i - \bar{s}_i)(h_i - h)}_{\substack{\text{Actual composition} \\ \text{effect}}} - \underbrace{\sum_i (s_i - \bar{s}_i) \left(\widehat{h}_i - \widehat{h} \right)}_{\substack{\text{Fitted composition} \\ \text{effect}}} + \underbrace{\sum_i \bar{s}_i \left(h_i - \widehat{h}_i \right)}_{\substack{\text{Industry-specific} \\ \text{effects}}} + \underbrace{\left(\sum_i s_i \widehat{h}_i - \widehat{h} \right)}_{\text{Aggregation effect}}
\end{aligned}
\tag{16}
$$

Hence the deviation of the vacancy yield from the fitted vacancy yield can be decomposed into four parts: The first measures the difference between the current vacancy yield and that which would have been observed if the distribution of vacancies across industries was constant at its pre-recession average. The second term reflects the change in the fitted vacancy yield in case one corrects for the deviation of the current cross-industry distribution of vacancies from its pre-recession average. The third term reflects the contribution of each of the industries to this difference due to their actual vacancy yields deviating from their historical average. The final part is the aggregation error that reflects that the vacancy-share-weighted fitted vacancy yields do not exactly aggregate to the aggregate fitted vacancy yield.

The only thing we used for this decomposition is that the aggregate vacancy yield is a share-weighted average of the industry-specific vacancy yields. This is also true for layoffs and quits rates. The aggregate layoffs and quits rates are employment-share-weighted averages of the industry-specific ones. Hence, we can apply a similar decomposition to those flow rates.

Tables and figures

Table 1. Fitted flow rates as a function of the (v/u)-ratio.

Sample size: n=84 *Dec 2000- Nov 2007*	Employment growth, JOLTS-based			Labor force growth	
	Vacancy yield	Layoffs rate	Quits rate	Entry	Exit
Dependent variable	$\ln(h_t)$	$\ln(l_t)$	$\ln(q_t)$	$\ln(n_t)$	$\ln(x_t)$
Constant	0.03	-4.13	-3.71	-3.25	-3.28
	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)
$ln(v_t/u_t)$	-0.41	-0.13	0.26	-0.03	-0.04
	(0.02)	(0.03)	(0.02)	(0.02)	(0.02)
R^2	0.88	0.18	0.65	0.03	0.06
$\hat{\sigma}$	0.04	0.06	0.04	0.04	0.03

Note: Standard errors in parentheses

Table 2. Aggregate decomposition of Beveridge curve gap

	Average for June 2011 – August 2011	
	Beveridge curve gap	2.6
	Employment growth, JOLTS-based	
(i)	Vacancy yield, h_t	4.0
(ii)	Layoffs, l_t	-1.4
(iii)	Quits, q_t	-0.9
	Labor force growth	
(iv)	Entry, n_t	0.5
(v)	Exit, x_t	-0.6
(vi)	Residual	0.9

Note: Numbers do not add up due to rounding

Table 3. Fitted flow rates by industry.

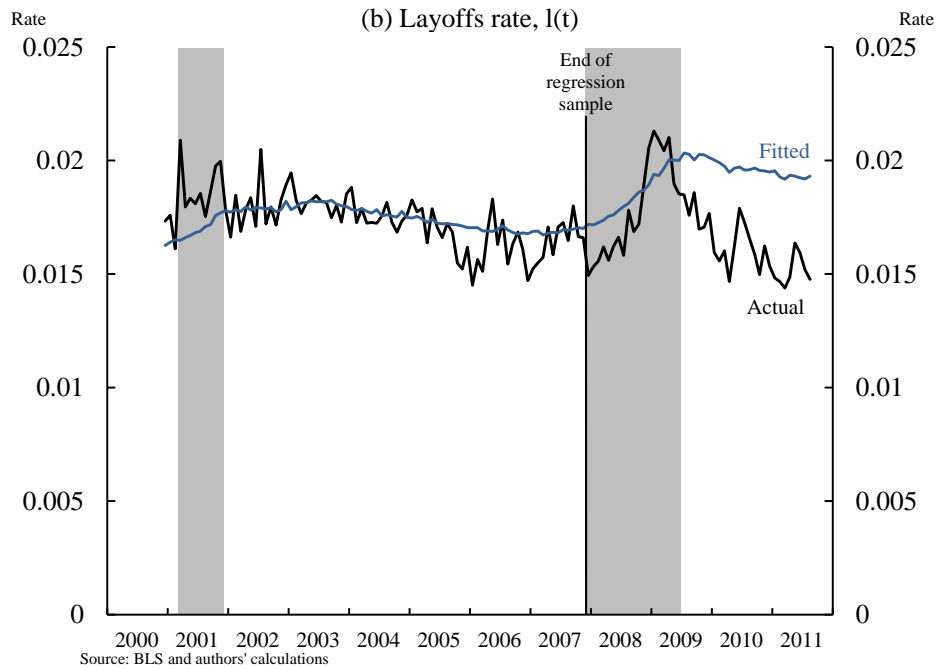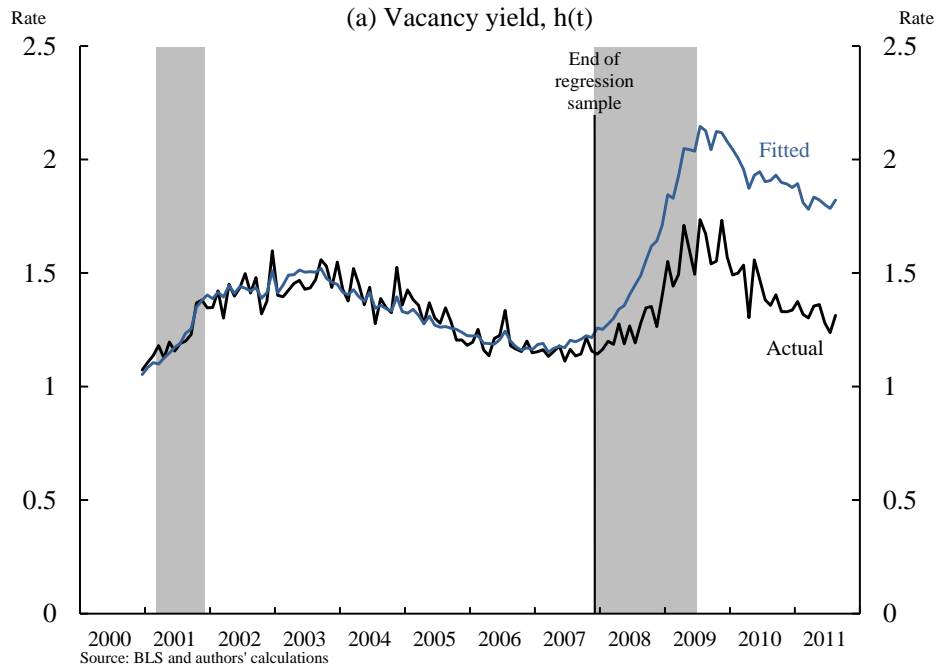|	Construction	Manufacturing	Trade, transportation and utilities	Professional and business services	Education and health services	Leisure and hospitality	Government	Other
	Sample: Dec-2000 through Nov-2007, $n = 84$							
Layoffs rate, $ln(l_t)$								
constant	-3.42	-4.37	-4.15	-3.73	-4.73	-3.86	-4.83	-4.45
	(0.08)	(0.04)	(0.03)	(0.06)	(0.06)	(0.05)	(0.04)	(0.09)
$ln(v/u)$	-0.21	0.96	-0.11	-0.26	-0.18	-0.39	0.07	0.15
	(0.13)	(0.13)	(0.15)	(0.10)	(0.08)	(0.22)	(0.08)	(0.21)
$ln(v_i/u_i)$	0.00	-0.69	-0.04	0.11	0.13	0.19	-0.15	-0.22
	(0.09)	(0.07)	(0.13)	(0.09)	(0.08)	(0.21)	(0.09)	(0.18)
R^2	0.12	0.53	0.13	0.11	0.06	0.10	0.04	0.05
$\hat{\sigma}$	0.14	0.12	0.10	0.12	0.12	0.15	0.11	0.14
Quits rate, $ln(q_t)$								
constant	-3.66	-4.07	-3.52	-3.29	-3.98	-2.91	-4.79	-3.85
	(0.09)	(0.03)	(0.02)	(0.05)	(0.04)	(0.02)	(0.03)	(0.05)
$ln(v/u)$	-0.17	-0.16	0.40	0.55	0.14	0.31	0.14	0.26
	(0.15)	(0.09)	(0.11)	(0.09)	(0.05)	(0.10)	(0.06)	(0.13)
$ln(v_i/u_i)$	0.13	0.30	-0.08	-0.38	0.09	-0.02	0.16	-0.04
	(0.10)	(0.05)	(0.09)	(0.08)	(0.05)	(0.09)	(0.07)	(0.11)
R^2	0.02	0.53	0.55	0.31	0.34	0.51	0.35	0.27
$\hat{\sigma}$	0.16	0.09	0.07	0.11	0.07	0.07	0.09	0.08
Vacancy yield, $ln(h_t)$								
constant	-0.04	-0.19	0.20	0.43	-0.15	0.31	-0.18	-0.70
	(0.08)	(0.02)	(0.02)	(0.04)	(0.03)	(0.02)	(0.03)	(0.05)
$ln(v/u)$	0.06	-0.07	0.31	0.38	0.13	0.14	0.09	0.25
	(0.12)	(0.07)	(0.10)	(0.07)	(0.04)	(0.09)	(0.05)	(0.12)
$ln(v_i/u_i)$	-0.75	-0.37	-0.71	-0.64	-0.44	-0.60	-0.31	-0.75
	(0.08)	(0.04)	(0.08)	(0.06)	(0.04)	(0.08)	(0.05)	(0.10)
R^2	0.80	0.87	0.81	0.67	0.68	0.79	0.35	0.82
$\hat{\sigma}$	0.13	0.07	0.06	0.08	0.06	0.06	0.07	0.08

Note: Standard errors are in parentheses

Table 4. Industry decomposition of deviations of aggregate flow rates from fitted values and Beveridge curve gap.

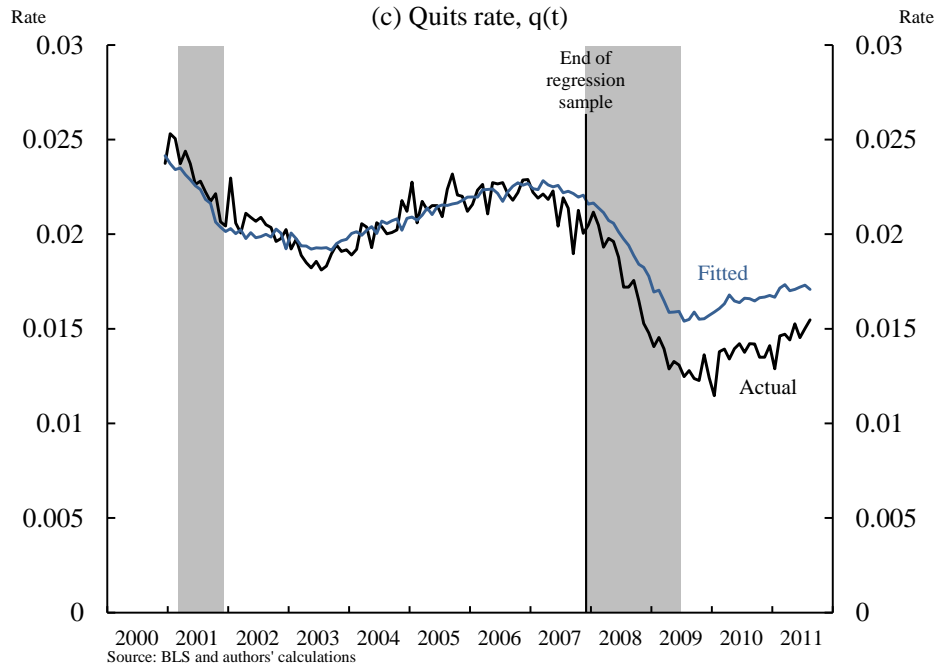|	Average: June 2011 – August 2011						
	Part A *Industry-specific deviation from fit*			Part B *Contribution to Beveridge curve gap*			
	I	II	III	IV	V	VI	VII
	l_t	q_t	h_t	l_t	q_t	h_t	Total
Total	-20	-13	-29	-1.8	-1.1	3.9	1.8
Aggregation	-	-	-	-0.1	0.0	-0.1	-0.2
Composition	-	-	-	-0.1	0.1	-0.4	-0.4
Industry deviations							
Construction	8	-43	-44	0.1	-0.2	0.9	0.8
Manufacturing	15	-31	-40	0.1	-0.2	0.4	0.3
Trade, transportation and utilities	-29	-8	-34	-0.4	-0.1	1.0	0.5
Professional and business services	-24	-5	-12	-0.4	-0.1	0.3	-0.2
Education and health services	-12	-11	-16	-0.1	-0.1	0.2	0.1
Leisure and hospitality	-34	-11	-25	-0.3	-0.1	0.6	0.2
Government	-9	-1	-22	0.0	0.0	0.2	0.1
Other	-20	-21	-46	-0.1	-0.2	0.9	0.6

Note: Numbers don't add up to totals due to rounding. Part A reports the percentage deviation of the individual job flow rates from their fitted values. Part B reports the industry decomposition of the Beveridge curve job-flows part of the Beveridge curve gap in percentage points of the unemployment rate. The composition effect is measured relative to the average employment and vacancy distribution from December 2000 through November 2007

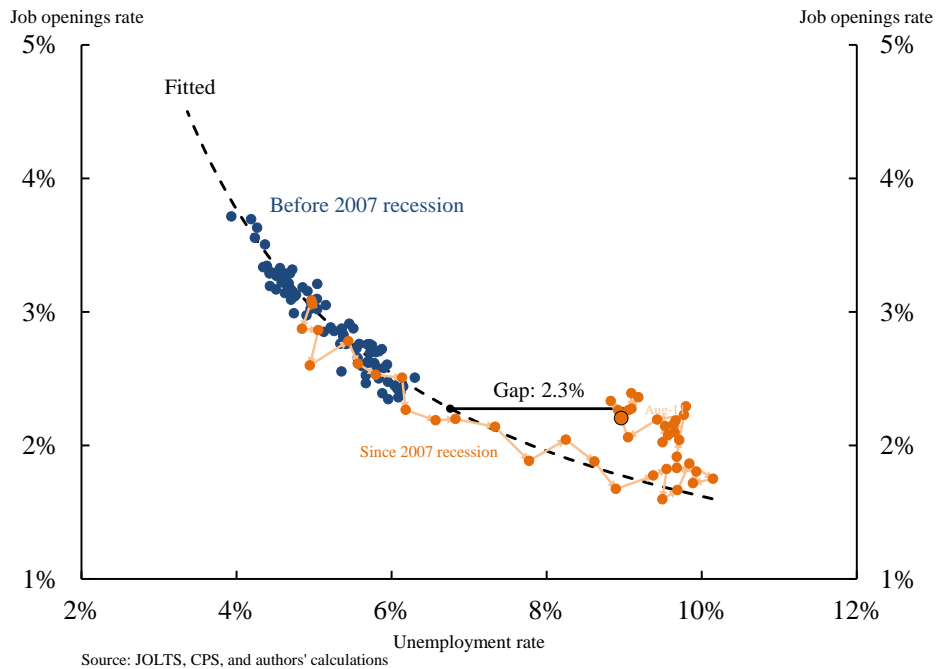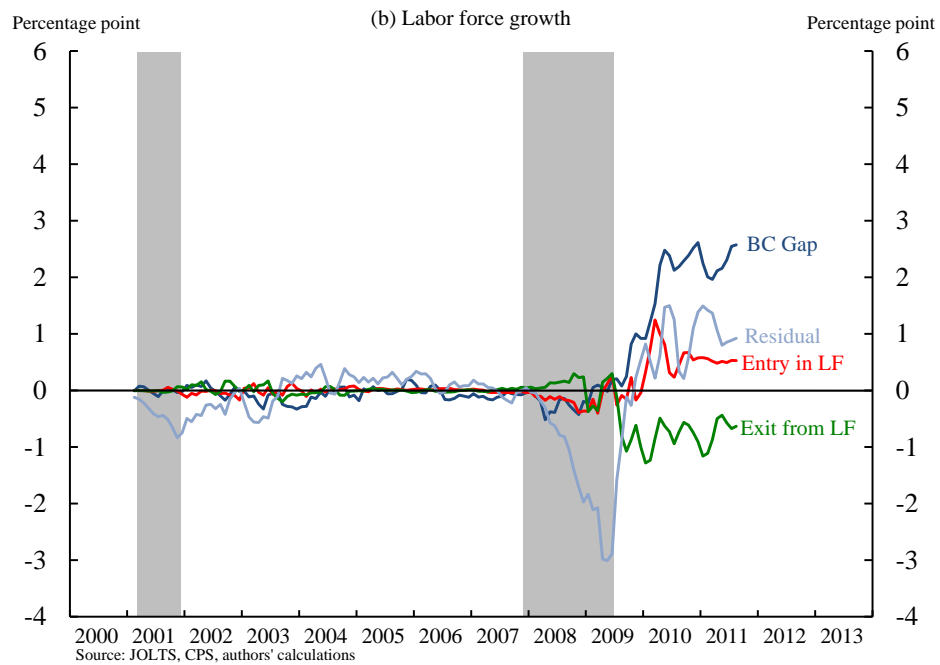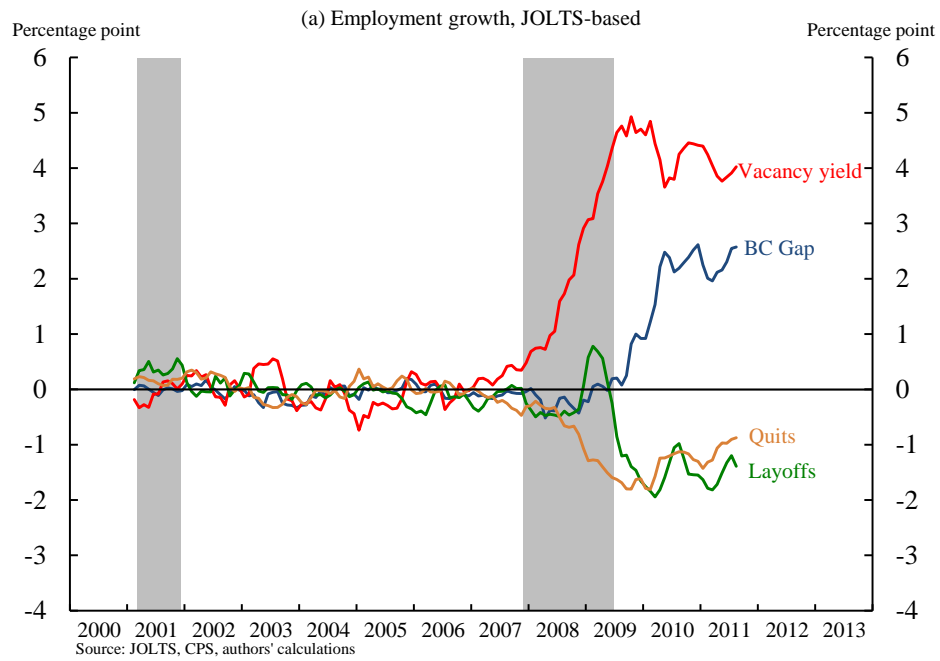**Figure 1.** Fitted and forecasted employment growth flows from JOLTS.

(a) Vacancy yield, h(t)

Source: BLS and authors' calculations

(b) Layoffs rate, l(t)

Source: BLS and authors' calculations

Figure 1 (continued). Fitted and forecasted employment growth flows from JOLTS

(c) Quits rate, q(t)

Source: BLS and authors' calculations

Figure 2. Actual and fitted JOLTS-based Beveridge curve.

Source: JOLTS, CPS, and authors' calculations

24

Figure 3. Beveridge curve gap decomposition over time, February 2001 – April 2011

(a) Employment growth, JOLTS-based

Source: JOLTS, CPS, authors' calculations

(b) Labor force growth

Source: JOLTS, CPS, authors' calculations

Note: all data are plotted in 3-month moving averages.

25

Figure 4. "New" Beveridge curve if current percentage deviations from fitted flow rates are permanent

Source: JOLTS, CPS, and authors' calculations