

Which industries are shifting the Beveridge curve?

According to JOLTS data, the failure of the unemployment rate to improve much despite growth in the economy is attributable to a shortfall in hires per vacancy in all industries, especially construction; this shortfall is what is causing the current labor market's shift in the Beveridge curve, which measures the negative relationship between the unemployment rate and the job openings rate

Regis Barnichon,
Michael Elsby,
Bart Hobijn,
and
Ayşegül Şahin

Although economic activity in the U.S. economy has grown, albeit 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 March 2012 the unemployment rate was 2.8 percentage points above its level implied by the Beveridge curve. The Beveridge curve can be interpreted as the job openings rate at which the current unemployment rate would be in its flow steady state. A flow steady state, so named because the Beveridge curve involves the measurement of flows from one labor force status (employed, unemployed, or not in the labor force) to another, occurs when these flows do not cause a change in the unemployment rate.

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). 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 Current Population Survey (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 attributable to hires per vacancy, layoffs, and 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 in order to consider which industries account for the unexplained decline in the vacancy yield as well as for the behavior of the quits and layoffs 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, and leisure and hospitality, as well as in the industries not classified in JOLTS. From January 2012 through March 2012, the difference between the observed and predicted

Regis Barnichon is a researcher at CREI; Michael Elsby is professor of economics at the University of Edinburgh; Bart Hobijn is senior research advisor at the Federal Reserve Bank of San Francisco, part-time visiting professor of economics at VU University Amsterdam, and a research fellow at the Tinbergen Institute; and Ayşegül Şahin is an assistant vice president at the Federal Reserve Bank of New York. Email: rbarnichon@crei.cat, mike.elsby@ed.ac.uk, bart.hobijn@sf.frb.org, and aysegul.sahin@ny.frb.org. The views expressed in this article solely reflect those of the authors and are not necessarily those of the Federal Reserve Bank of New York, the Federal Reserve Bank of San Francisco, or the Federal Reserve System as a whole.

hires per vacancy in construction alone accounted for 0.5 percentage point of the 2.8 percentage points 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

Because of the high levels 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 use a similar interpretation in this article. 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 hires, layoffs, and quits rates for defining steady-state unemployment.⁴

The unemployment rate, u_t , is in a steady state whenever the growth rate of the labor force, which we denote by $g_t^{(l)}$, 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, $g_t^{(l)}$, is given by n_t , calculated as 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 , calculated as 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 data.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 $g_t^{(l)}$ 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 logarithm of each of the flows on the logarithm of the v/u ratio. For these regressions we use monthly seasonally adjusted data that cover the prerecession 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 closely together with the v/u ratio. Fluctuations in labor market tightness, i.e., variations in the v/u ratio, explain about 89 percent of the fluctuations

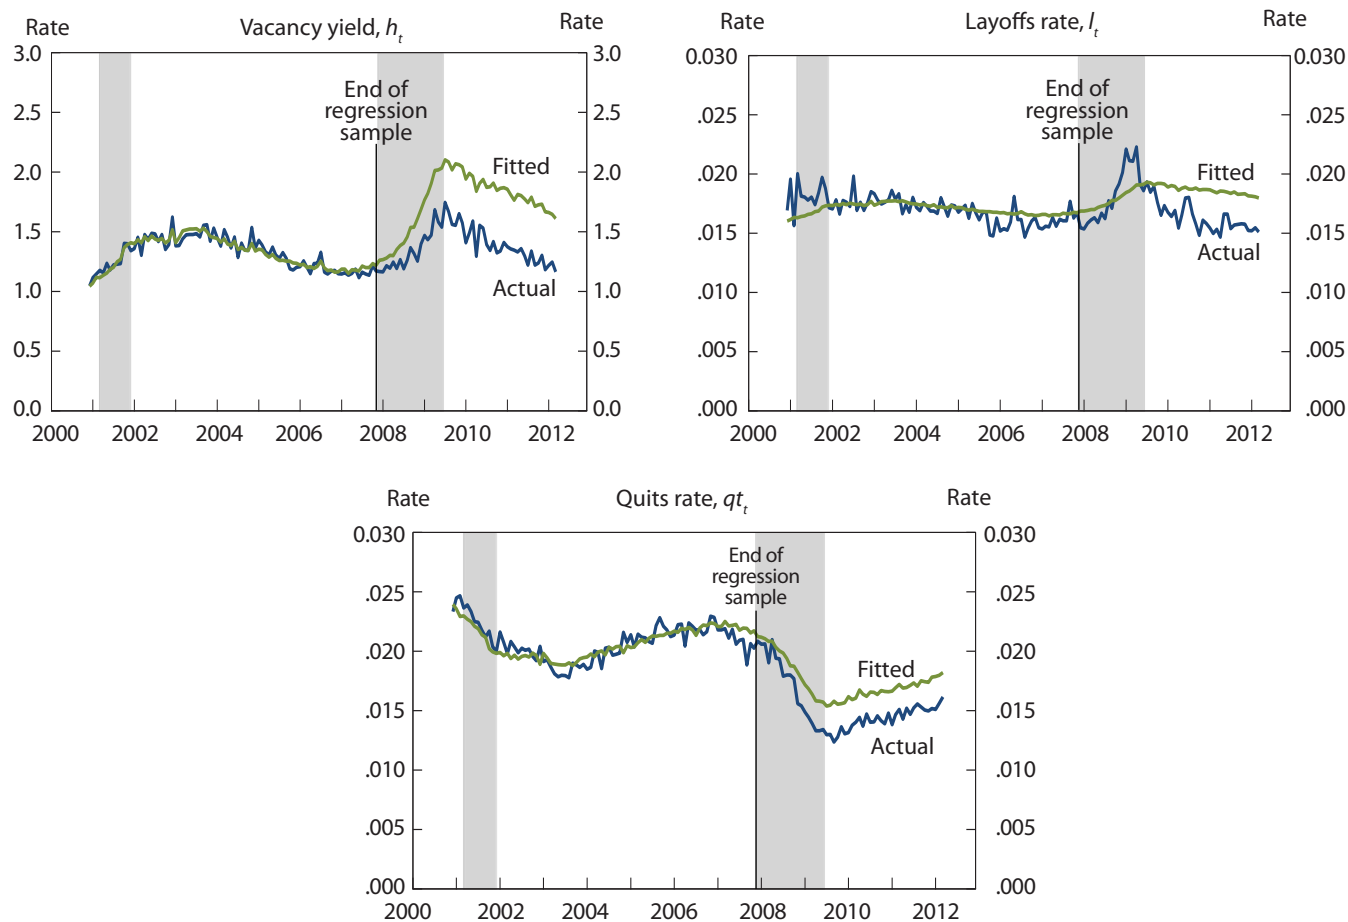
Table 1. Fitted flow rates as a function of the v/u ratio, December 2000–November 2007

Item	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	0.02 (.01)	-4.14 (.02)	-3.72 (.01)	-3.25 (.01)	-3.28 (.01)
$\ln(v_t/u_t)$	-.41 (.02)	-.11 (.03)	.26 (.02)	-.03 (.02)	-.04 (.02)
R^2	.89	.15	.67	.03	.07
$\hat{\sigma}$.04	.06	.04	.04	.03

NOTES: Standard errors are in parentheses. Sample size $n = 84$.

SOURCE: U.S. Bureau of Labor Statistics and authors' calculations.

Chart 1. Actual and fitted and forecasted employment growth flows from JOLTS



NOTE: Shaded areas represent recessions as designated by the National Bureau of Economic Research.

SOURCE: U.S. Bureau of Labor Statistics and authors' calculations.

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 decreases as the v/u ratio rises. 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 quits 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, as shown in the second panel in chart 1, 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.

The three panels in chart 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 prerecession relationship between these rates and the v/u ratio. Besides the wave of layoffs at the end of 2008 and the beginning of 2009, chart 1 shows two more large deviations of the actual rates from the fit implied by their historical patterns. The first is that the March 2012 quits rate is about 13 percent below that predicted on the basis of the level of labor market tightness. Hence, workers hang onto their jobs even more than one would expect on the basis of the current weakness in the labor market. The second, and most profound, deviation is that hires per vacancy are about 38 percent less than predicted at the current v/u ratio. Hence, the JOLTS-based employment growth flows are deviating substantially from their predicted values, which are 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 meaningful 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 the job openings rate, v , the Beveridge curve is determined by the level of the unemployment rate for which our estimated flow rate equations imply that unemployment is in its steady state—i.e., the fitted labor force growth rate equals the fitted employment growth rate.

Chart 2 plots the actual and fitted Beveridge curves. We split the observed Beveridge curve into two parts. The blue points are the prerecession observations on the basis of which the flow rates that underlie the curve are fitted. The orange points are the observations from December 2007 onwards. As can be seen, the estimated Beveridge curve does not only provide a good fit for the prerecession 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 on the basis of separately fitted flow rates reported in JOLTS using a flow-steady-state unemployment relationship from the job openings rate.

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 percent to 2.7 percent. However, during that same period, the unemployment rate has not fallen as much as the Beveridge curve would imply. The jobless rate actually initially increased from 9.5 percent at the National Bureau of Economic Research-determined end of the recession in July 2009 to 10.0 percent in October 2009 and has since come down to 8.2 percent in March 2012. The result is that, at the March 2012 job openings rate, the actual unemployment rate was 2.8 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 the various 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 chart 1, we are seeing large deviations in 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 perform the analysis, we use an approximate additive decomposition.¹² We use the approximation technique to find the answer to the following type of question.

The March 2012 level of the vacancy yield, h_t , is 38 percent lower than that implied by the estimates reported in table 1. At the March 2012 job openings rate, to what extent is the rise in the steady-state unemployment rate relative to the fitted Beveridge curve explained by the 38 percent 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 attributable to deviations of the five actual flow rates from their fitted values and to 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

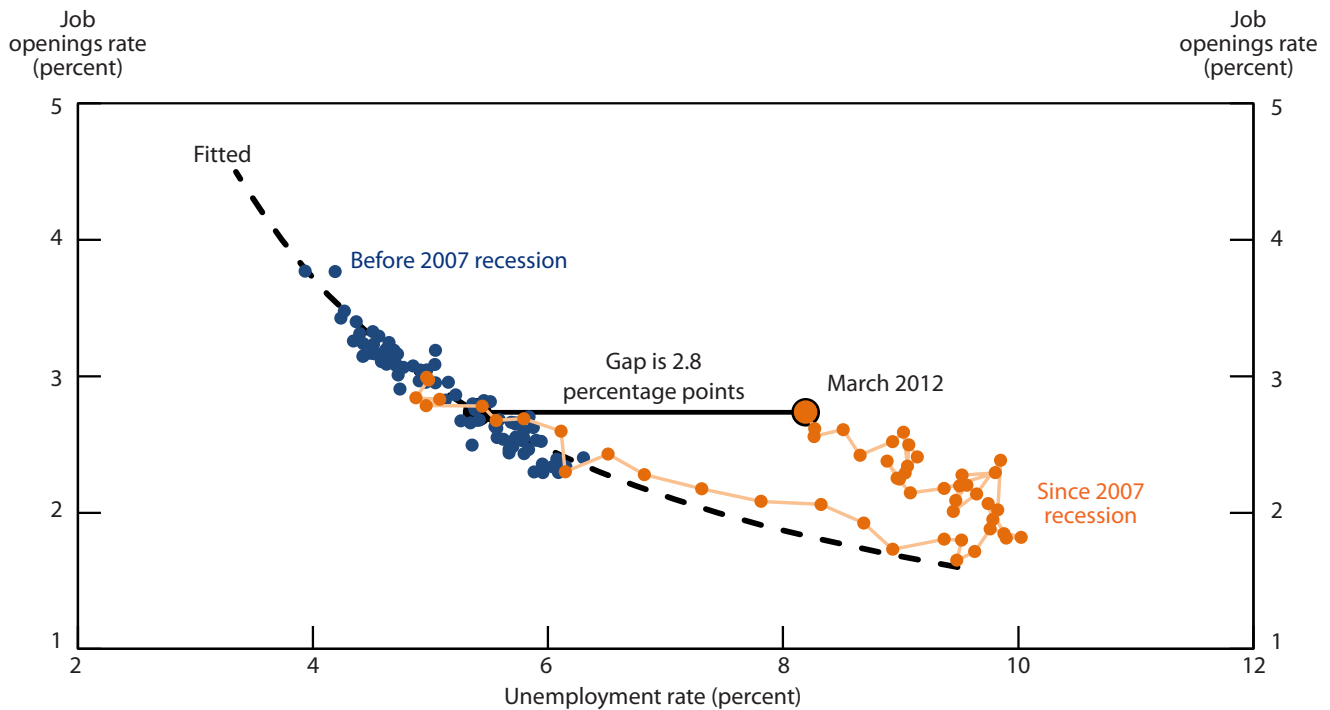
To smooth out some of the month-to-month fluctuations, we report the results of our decomposition in terms of 3-month moving averages. The decomposition of the average Beveridge curve gap in the latest 3 months in our sample is reported in table 2. On average from January 2012 through March 2012, the 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 lessen the Beveridge curve gap, by 0.8 percentage points each. 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 per-

Chart 2. Actual and fitted JOLTS-based Beveridge curve



SOURCES: JOLTS, CPS, and authors' calculations.

sons remain unemployed and this raises the unemployment rate. This is why the shortfall in the vacancy yield, depicted in the first panel of chart 1, translates into a 3.0 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 is 1.1 percentage points. This residual 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 chart 3 show the decomposition of the Beveridge curve gap over time, from the beginning of JOLTS to the latest observation in our sample. The first panel shows the Beveridge curve gap and its employment-growth flow determinants. The shortfall in the vacancy yield started before the substantial increase in the Beveridge curve gap and even before the spike in the layoffs 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 the second panel of chart 3, which plots the contributions of the labor force growth flows and of the residual. The large negative residual 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

Table 2. Aggregate decomposition of Beveridge curve gap, January–March 2012 averages

Item	Percentage points
Beveridge curve gap	2.6
Employment growth, JOLTS-based	
Vacancy yield, h_t	3.0
Layoffs, l_t	–.8
Quits, q_t	–.8
Labor force growth	
Entry, n_t	.5
Exit, x_t	–.3
Residual	1.1

NOTE: Numbers do not add up to totals because of rounding.
SOURCE: U.S. Bureau of Labor Statistics and authors' calculations.

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 were last employed in the industry, $u_{i,t}$. The i denotes which industry is being indexed.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 (1) construction, (2) manufacturing, (3) trade, transportation and utilities, (4) professional and business services, (5) education and health services, (6) leisure and hospitality, and (7) government. We construct data for the residual industry, “other,” by subtracting the data for the seven industries from those reported for the total economy.

Mirroring our aggregate analysis, we use the prerecession sample to fit the flow rates. Table 3 reports the estimates of the parameters for the fitted flow rates by industry. Not surprisingly, because 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 the measure that is the most responsive to the degree of labor market tightness, except for the government sector. For each of the private industries, 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 industries, 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. The responsiveness to labor market tightness 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 into two parts.

Part A shows how much the actual flow rates deviate from their fitted values, both for the total economy and 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 industries, fewer workers quit than would be expected from the current strength of the job market. This column shows that the aggregate 13-percent shortfall in the quits rate is broad-based. Quits are especially low in construction and manufacturing. Column III reveals that, in all industries, 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 industries. The latter include finance and real estate, two industries that were particularly hard hit by the recession. 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 article, is presented in part B. For each industry, the table 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 up the contributions of these three rates for each of the industries.¹⁵

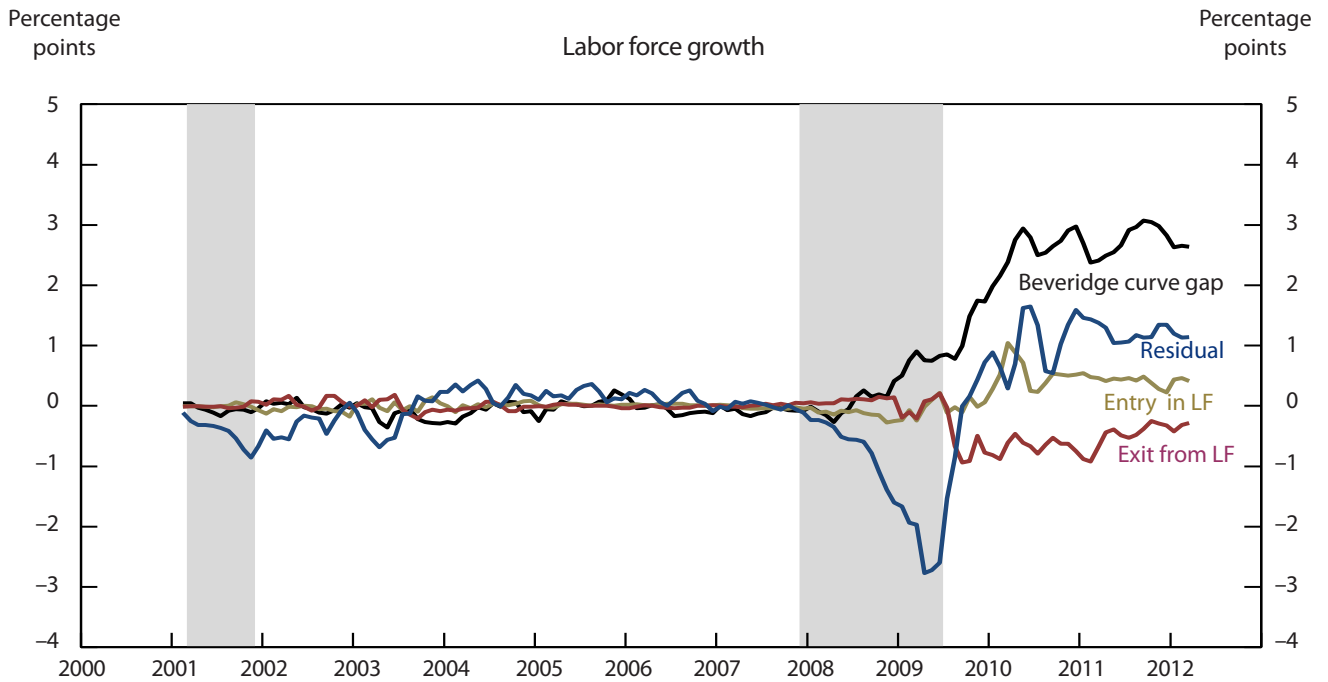
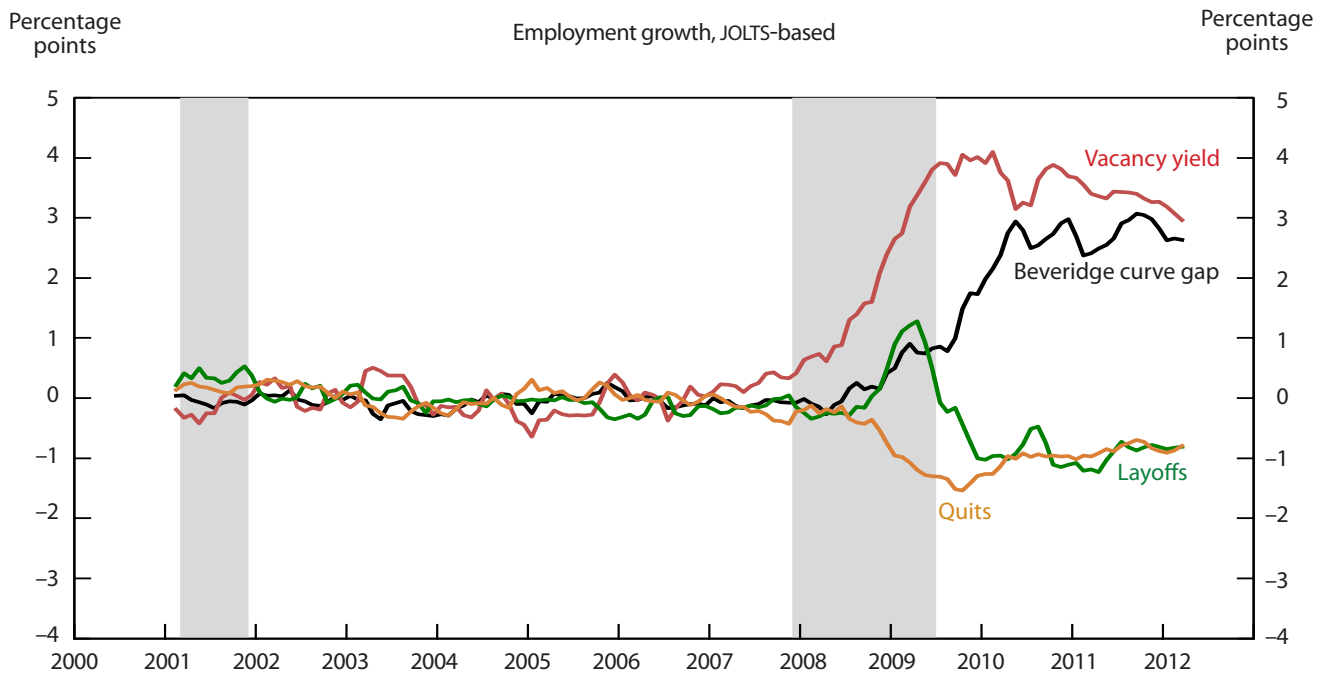
From this part it can be seen that the biggest contributors to the Beveridge curve gap are the vacancy yield deficits in construction, trade and transportation, leisure and hospitality, and “other” industries. This last category contains finance, insurance, and real estate. Education and health services 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 slightly. 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 open-

Chart 3. Beveridge curve gap decomposition over time, February 2000–March 2012



NOTE: Shaded areas represent recessions as designated by the National Bureau of Economic Research.

SOURCE: U.S. Bureau of Labor Statistics and authors' calculations.

Table 3. Fitted flow rates by industry, December 2000–November 2007

Item	Construction	Manufacturing	Trade, transportation, and utilities	Professional and business services	Education and health services	Leisure and hospitality	Government	Other
	Layoffs rate, $\ln(l_i)$							
Constant	-3.48 (.08)	-4.36 (.03)	-4.14 (.03)	-3.70 (.06)	-4.76 (.07)	-3.85 (.05)	-4.88 (.03)	-4.41 (.09)
$\ln(v/u)$	-.23 (.12)	.97 (.11)	.00 (.15)	-.18 (.10)	-.18 (.08)	-.46 (.20)	.12 (.07)	-.01 (.23)
$\ln(v_i/u_i)$	-.02 (.08)	-.68 (.06)	-.10 (.13)	.04 (.09)	.14 (.09)	.31 (.19)	-.19 (.08)	-.06 (.19)
R^2	.18	.58	.09	.07	.05	.09	.07	.02
$\hat{\sigma}$.14	.10	.09	.12	.12	.15	.10	.15
Quits rate, $\ln(q_i)$								
Constant	-3.64 (.08)	-4.07 (.03)	-3.53 (.02)	-3.31 (.05)	-4.01 (.04)	-2.92 (.02)	-4.84 (.03)	-3.89 (.05)
$\ln(v/u)$	-.07 (.13)	-.10 (.09)	.35 (.11)	.51 (.09)	.13 (.05)	.31 (.09)	.14 (.06)	.29 (.14)
$\ln(v_i/u_i)$.11 (.09)	.27 (.05)	-.04 (.09)	-.35 (.08)	.08 (.05)	-.02 (.09)	.16 (.06)	-.08 (.11)
R^2	.04	.56	.56	.29	.34	.53	.41	.25
$\hat{\sigma}$.14	.08	.07	.11	.07	.07	.08	.09
Vacancy yield, $\ln(h_i)$								
Constant	-.02 (.07)	-.20 (.02)	.21 (.02)	.44 (.04)	-.18 (.03)	.31 (.02)	-.21 (.02)	-.70 (.05)
$\ln(v/u)$.01 (.12)	-.04 (.07)	.37 (.10)	.38 (.07)	.10 (.04)	.16 (.08)	.07 (.05)	.20 (.12)
$\ln(v_i/u_i)$	-.71 (.08)	-.39 (.04)	-.75 (.09)	-.65 (.06)	-.41 (.04)	-.61 (.08)	-.31 (.05)	-.71 (.10)
R^2	.81	.87	.80	.66	.68	.79	.40	.83
$\hat{\sigma}$.13	.07	.06	.09	.05	.06	.07	.07

NOTES: Standard errors are in parentheses. Sample size is $n = 84$.

SOURCE: U.S. Bureau of Labor Statistics and authors' calculations.

ings and the unemployed. If the pool of unemployed persons has very different qualifications from those required for the job openings posted, then it would be harder to fill these openings relative to other times when the degree of mismatch is less severe. The problem with using JOLTS data to assess the mismatch 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; that is, 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 these 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 decline in labor demand for jobs that typically are filled through 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 effort an average unemployed person puts into 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 during

and after the 2007–2009 recession. This explanation relies on the effect of UI on the incentives for the unemployed to search for and accept job offers.¹⁸

The possibility that the amount of mismatch has increased 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 Several points are important to realize when discussing the natural rate.

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 from their estimated values during the last 3 months in our sample are permanent.

This hypothetical “new” Beveridge curve is plotted in chart 4. At the 2.7-percent job openings rate that prevailed in August 2011, the outward shift of the “new” Beveridge curve relative to the fitted historical one is 2.4 percentage points. At a 3.5-percent job openings rate, the shift is 1.8 percentage points. At a 4.3-percent job openings rate, the shift is 1.5 percentage points. Also note that the 8.2-percent unemployment rate in March 2012 was 0.3 percentage points higher than implied by this “new” Beveridge curve. This suggests that the unemployment rate in March 2012 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 period before the 2007–2009 recession, estimates of the natural rate of unemployment are generally around 5 percent.²⁰ On the fitted Beveridge curve, this coincides with a 3-percent natural rate of job openings. On the “new” Beveridge curve, this 3-percent job openings rate coincides with a 7.1-percent unemployment rate. However, most equilibrium models of frictional unemployment²¹ suggest that the natural vacancy rate increases when the Beveridge curve shifts out. Suppose that the new natural job openings rate were 3.5 percent instead of 3 percent. On the “new” Beveridge curve, this would imply a 6.1-percent natural rate of unemployment rather than the 7.1 percent 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, because a large part of the mismatch in the labor market seems to be temporary and because UI extensions are set to expire, the current shift in the Beveridge curve is in all likelihood largely temporary. Secondly, even if it were 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.

Table 4. Industry decomposition of deviations of aggregate flow rates from fitted values and Beveridge curve gap, January–March 2012

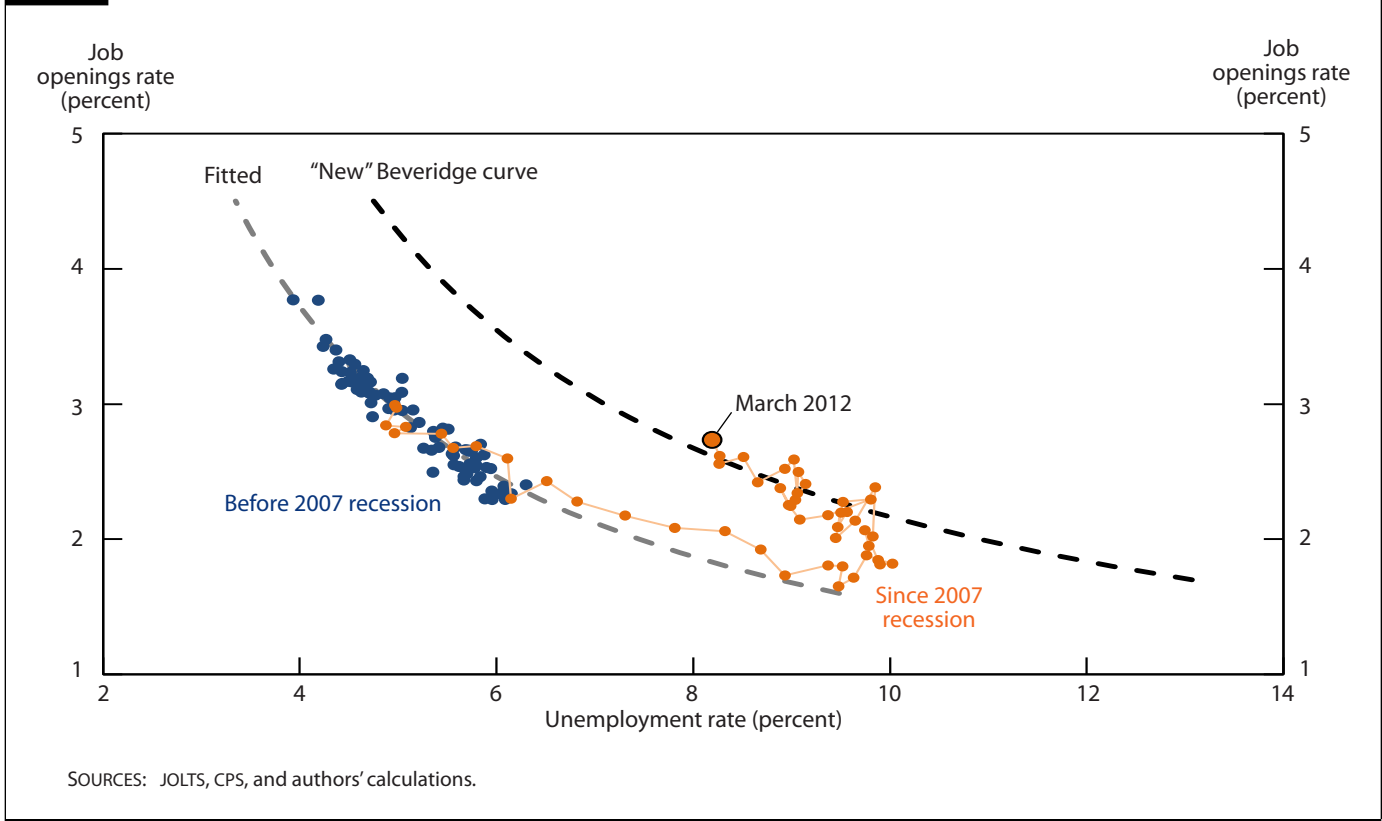
Industry	Part A			Part B			
	Industry-specific deviation from fit			Contribution to Beveridge curve gap			
	l_t	q_t	h_t	l_t	q_t	h_t	Total
Total	-16	-13	-26	-0.8	-0.8	3.0	1.4
Aggregation	—	—	—	-2	.1	-.1	-
Composition	—	—	—	.0	.0	-.3	-.2
Industry deviations							
Construction	-7	-37	-42	.0	-.1	.7	.5
Manufacturing	6	-35	-40	.0	-.2	.3	.1
Trade, transportation and utilities	-22	-9	-31	-2	-.1	.7	.4
Professional and business services	-18	-15	-15	-2	-.2	.3	-.1
Education and health services	-3	-4	-17	.0	.0	.2	.2
Leisure and hospitality	-17	-15	-20	-.1	-.2	.4	.1
Government	0	0	-17	.0	.0	.1	.1
Other	-11	-16	-39	-.1	-.1	.6	.4

NOTE: Numbers do not add up to totals because of 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.

SOURCE: U.S. Bureau of Labor Statistics and authors' calculations.

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



IN THIS ARTICLE 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—in March 2012, the unemployment rate exceeded the level implied by its historical relationship with the job openings rate by 2.8 percentage points—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, leisure and hospitality, and industries not explicitly classified in JOLTS.

Whether this shortfall is due to (1) mismatch between job openings and unemployed workers, (2) reduced recruitment effort by employers, (3) a change in the composition of vacancies and hires, or (4) reduced search intensity of unemployed persons is difficult to parse 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 among these issues. Moreover, additional data on search intensity of employers, such as information on the number of job offers made, would help us better understand the effort with which they pursue filling the job openings they have. □

NOTES

ACKNOWLEDGMENTS: We would like to thank Mary Daly, Pieter Gautier, and Rob Valletta as well as the participants of the Bureau of Labor Statistics JOLTS Symposium for their comments and suggestions. We are grateful to Brian Lucking, Joseph Song, and Ted Wiles for their outstanding research assistance. A draft of this article was presented at the BLS JOLTS Symposium on December 10, 2010. This

article uses JOLTS data through March 2012.

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, pp. 1–48.

² See Mark deWolf and Katherine Klemmer, “Job openings, hires, and

separations fall during the recession,” *Monthly Labor Review*, May 2010, pp. 36–44, for a discussion of JOLTS evidence during the 2007 recession.

³ See Robert Shimer, “Reassessing the Ins and Outs of Unemployment,” National Bureau of Economic Research, Working Paper No. 13421, September 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*, No. 1, 1989, pp. 1–60, 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*, Board of Governors of the Federal Reserve System, August 2010.

⁵ The technical derivations behind this result as well as other results in this article are available in an appendix in Regis Barnichon, Michael Elsby, Bart Hobijn, and Ayşegül Şahin, “Which Industries are Shifting the Beveridge Curve?” Federal Reserve Board of San Francisco, Working Paper 2010–32, June 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 this article’s appendix.

⁷ Besides quits and layoffs, the JOLTS data contain a third class of separations, which includes retirements, deaths, emigration, and other types of separations that are considered neither a layoff nor 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*, June 2001, pp. 390–431, for a survey of matching function estimates.

⁹ JOLTS 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 in “Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows,” *Finance and Economics Discussion Series 2004–34*, Board of Governors of the Federal Reserve System, May 2004.

¹¹ For the transformation from the logarithm to the level plotted in the chart, 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 Survey 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 with 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 because of the change in the industry composition of employment and job openings even if each of the industry flow rates were equal to its fitted values.

¹⁶ For a detailed analysis of mismatch in the United States, see Ayşegül Şahin, Joseph Song, Giorgio Topa, and Gianluca Violante, “Measuring Mismatch in the U.S. Labor Market,” manuscript, Federal Reserve Bank of New York, 2011.

¹⁷ This explanation is put forward by Steven Davis, Jason Faberman, and John Haltiwanger in “The Establishment-Level Behavior of Vacancies and Hiring,” National Bureau of Economic Research, Working Paper No. 16265, 2010.

¹⁸ However, Rob Valletta and Katherine Kuang, in “Extended Unemployment and UI Benefits,” *FRBSF Economic Letter 2010–12*, April 19, 2010, find that the effect of UI on the duration of unemployment is relatively small.

¹⁹ An increase in the natural rate because of mismatch has been discussed by, among others, Narayana Kocherlakota, in “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 and x_t as the sum of all worker flows in a month from nonparticipation 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 during that month. Similarly, we construct x_t as the sum of all worker flows from employment and unemployment to nonparticipation 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 persons, E_t , and the number of unemployed persons, U_t , the change in the number of unemployed can be written as the change in the labor force minus the change in the number of employed persons. That is,

$$U_t - U_{t-1} = \Delta U_t = \Delta LF_t - \Delta E_t \quad (1)$$

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

$$\frac{LF_t}{LF_{t-1}} u_t - u_{t-1} = \frac{\Delta LF_t}{LF_{t-1}} - \frac{E_{t-1}}{LF_{t-1}} \frac{\Delta E_t}{E_{t-1}} \quad (2)$$

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

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

respectively, we can write (2) as

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

This simplifies to

$$u_t - u_{t-1} = \frac{1 - u_{t-1}}{1 + g_t^{(lf)}} (g_t^{(lf)} - g_t^{(e)}) \quad (5)$$

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

The number of hires in a month is denoted by H_t and the number of job openings reported is denoted 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

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

This gives us hires as a fraction of the employment level at the beginning 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 layoffs 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

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

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

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

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 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:

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

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

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

we can write the growth rate of the number of employed workers as

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

We then solve for the Beveridge curve for the particular values $\bar{z}_t = z$ and $\bar{d}_t = d$, which are their sample average over the prerecession period.

Log-linear decomposition of the Beveridge curve gap

Given the regressions

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

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

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

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

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

All variables denoted with a $\hat{}$ 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 $(m^{(f)} - \hat{m}^{(f)}) / \hat{m}^{(f)}$ for $f = n, x, h, l, \text{ or } q$.

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, \hat{h}_i for the industry-specific vacancy yield, and \hat{h}_i for the fitted value of the industry-specific vacancy yield. The share of industry i in total job openings is denoted by s_i , and its sample average over the prerecession period is \bar{s}_i . This allows us to write

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

such that

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

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 were constant at its prerecession 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 prerecession average. The third term reflects the contribution of each of the industries to this difference because of 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.

This decomposition is derived using the fact 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 in the sense that 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.